Book of Abstracts
International Conference on

Industrial Digital Twins for Smart Manufacturing

(ICIDTSM 2026)

20-22 July 2026 | IC&SR, IIT Madras

EDITORS

N. Arunachalam | G. L. Samuel | Sivasrinivasu Devadula

Organized by

Manufacturing Engineering Section
Department of Mechanical Engineering
Indian Institute of Technology Madras
Chennai, Tamil Nadu – 600036

Editors

N. Arunachalam
G. L. Samuel
Sivasrinivasu Devadula

Book of Abstracts

International Conference on

Industrial Digital Twins for Smart Manufacturing

ICIDTSM 2026 (20-22 July, 2026)

Organized & published by

Manufacturing Engineering Section
Department of Mechanical Engineering
Indian Institute of Technology Madras
Chennai, Tamil Nadu – 600036

ICIDTSM

2026

Table of Contents

Indian Institute of Technology Madras

Chennai, India · Established 1959 · Institute of Eminence

Founded in 1959 with technical assistance from West Germany, the Indian Institute of Technology Madras is India's highest-ranked engineering institution. Its 620-acre green campus, carved from Guindy National Park, is home to about 550 faculty and 8,000 students, and to a rare continuity between academia and industry: fundamental research, corporate R&D and deep-tech entrepreneurship proceed within walking distance of one another, in the midst of a protected forest.

NIRF Rank 1	431	567	620 acres
Engineering — 10 consecutive years	patents filed, FY 2025–26	startups incubated	campus

RESEARCH & ACADEMIA

Ranked No. 1 in Engineering in India's NIRF rankings for ten consecutive years, No. 1 Overall for seven, and No. 2 among research institutions, IIT Madras pioneered NPTEL (2003), which opened IIT classrooms to the nation. Recent additions include the School of Sustainability (2023) and the School of Innovation and Entrepreneurship (2025).

INDUSTRY COLLABORATION

The IIT Madras Research Park, India's first university-based research park, hosts more than 250 R&D companies and laboratories and 15 Centres of Excellence, with partners ranging from DRDO and BHEL to Saint-Gobain. Its success led to 50 further research parks being planned under the Startup India initiative.

INNOVATION & ENTREPRENEURSHIP

The IITM Incubation Cell, India's leading deep-tech incubator, has nurtured 567 startups collectively valued at over ₹74,100 crore, incubating more than 100 new ventures in a single year. Its portfolio includes Ather Energy (electric vehicles; the cell's first IPO), Agnikul Cosmos (3D-printed rocket engines), The ePlane Company (electric air taxis), Mindgrove Silicon (India's first commercial SoC), HyperVerge (AI-based identity verification) and Detect Technologies (industrial AI and IoT for asset integrity and safety).

THE CAMPUS & ITS BIODIVERSITY

The campus preserves one of India's rarest habitats — tropical dry evergreen forest. Blackbuck, spotted deer, paradise flycatchers and some 100 bird species live among the lecture halls, protected by the student-run Prakriti wildlife club: a daily reminder that technology here has grown with nature, not over it.

Conference Information

Name of Conference	International Conference on Industrial Digital Twins for Smart Manufacturing (ICIDTSM 2026)
Date	20 - 22 July, 2026
Organizer	Indian Institute of Technology Madras, Chennai - India
Venue	IC&SR Hall
Website	https://ge.iitm.ac.in/icidtsm-2026
Email	icidtsm2026@gmail.com
Official Language	English

Venue Layout: IC&SR Hall

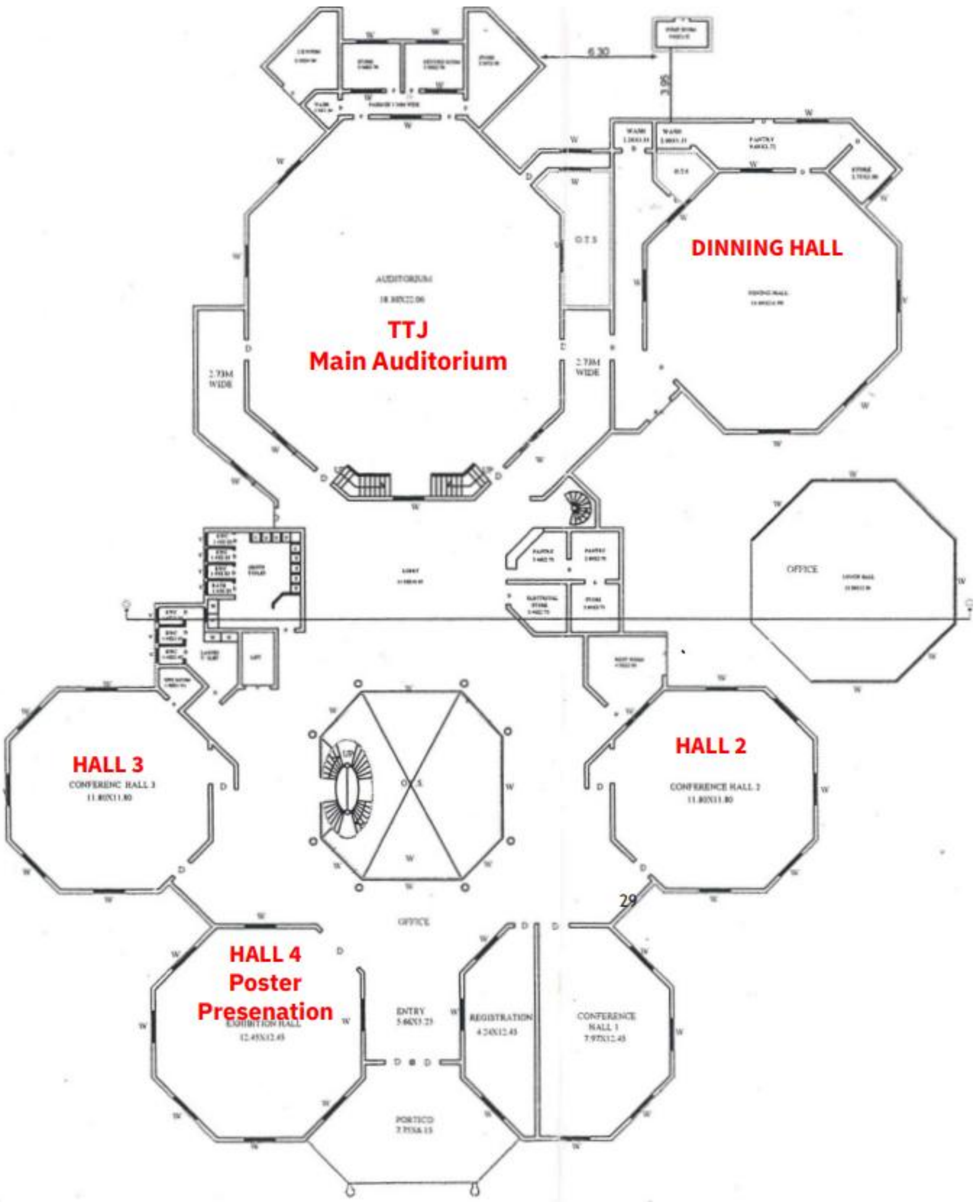

Disclaimer

About the Conference

The International Conference on Industrial Digital Twins for Smart Manufacturing (ICIDTSM 2026) will be held from 20–22 July 2026 at the Indian Institute of Technology Madras, bringing together global leaders from academia, industry, research laboratories, and technology enterprises to explore the next evolution of intelligent manufacturing systems.

ICIDTSM 2026 is designed as a high-impact international platform to examine how emerging technologies are reshaping industrial ecosystems — from connected machines and autonomous robotics to AI-driven decision systems and fully integrated Digital Twin environments. The conference will focus on bridging fundamental research, applied engineering, and scalable industrial deployment, enabling meaningful dialogue between scientists, system architects, policymakers, and manufacturing practitioners.

As global industries transition from Industry 4.0 toward more resilient, human-centric Industry 5.0 paradigms, the need for real-time intelligence, interoperability, sustainability, and autonomous adaptability has never been greater. This conference aims to address these challenges by fostering interdisciplinary collaboration and encouraging innovative solutions that integrate sensing, computation, control, intelligence, and cyber-physical synchronization into cohesive smart manufacturing ecosystems.

Participants can expect keynote lectures by distinguished experts, technical paper presentations, industry panels, technology demonstrations, and opportunities for strategic research and industrial partnerships.

About the Editors

Dr. N. Arunachalam received his B.E. in Mechanical Engineering from the College of Engineering, Guindy (1994-98), Anna University, and his M.E. (Manufacturing) from MIT, Anna University (1999-2001). His Ph.D. is from the Indian Institute of Technology Madras, Chennai, in 2010. He worked as a Research Engineer at General Electric, GRC, Bangalore, India, for three years from 2011 to 2013 before joining IIT Madras. His research interests center around High Performance Manufacturing, which includes Industry 4.0 Implementation and Practices Sensors, Signal processing (Feature extraction & Selection), Model building using machine learning techniques, Decision making models, Development of Physics/Data driven, Predictive models to assess the process induced damage in grinding, milling, EDM and 3-D printing processes, Prognostics and health management of machine tools (PHMMT), Advanced micro and nano level polishing process development, and Metrology in Production Lines.

Dr. G.L. Samuel is a distinguished Professor in the Department of Mechanical Engineering at the Indian Institute of Technology Madras, Chennai, India, where he has been on the faculty since 2005. He received his PhD from IIT Madras in 2001 and subsequently completed a Post-Doctoral Fellowship at Kyungpook National University, South Korea (2003–2005). He leads the Centre of Excellence for Advanced Laser Material Processing and Surface Engineering (CALM-IITM). His research spans Micro Manufacturing, Advanced Laser Material Processing (including femtosecond and ultra-short pulse lasers), Metrology and Computer-Aided Inspection, Biomedical Devices (microfluidics and lab-on-chip), and Artificial Intelligence in Manufacturing, including Digital Twins and Industry 4.0. He holds multiple patents, has published over 81 international journal papers, and has received prestigious recognitions including the Fulbright-Nehru Fellowship at Texas A&M University (2025) and the Erasmus Mundus European HERITAGE Scholarship (2013).

Dr. Sivasrinivasu Devadula received his Ph.D. from the Indian Institute of Technology Madras, Chennai, India, in 2008. He worked as a Postdoctoral Research Fellow at Nottingham University, UK, for 3½ years (2008-2011). Following it, he worked as a Senior Research Fellow at KTH-Royal Institute of Technology, Sweden, for 3½ years (2011-2015) before joining IIT Madras. His research interests include, in particular, machining of 3D surfaces on advanced engineering materials (superalloys, shape memory alloys, high entropy alloys, diamond, engineering ceramics, engineering composites), abrasive waterjet machining, novel tool path strategies, optimization, process modelling and simulation, machine tools, and in general – conventional and unconventional machining processes, machine tool metrology, micro machining.

Organizing Committee

1. **Organizing Committee:**
   a. Patron: Prof. V. Kamakoti, Director, IIT Madras
   b. Convener: Prof. N. Arunachalam
   c. Co-convener: Dr. Sivasrinivasu Devadula
   d. Members:
      1. Prof. Amitava Ghosh, ME, IITM
      2. Dr. Anil Meena, ME, IITM
      3. Dr. Anuj Kumar Tiwari, ME, IITM
      4. Prof. C. Balaji, ME, IITM
      5. Prof. P. Chandramouli, ME, IITM
      6. Dr. Charchit Kumar, ME, IITM
      7. Prof. K. Hariharan, ME, IITM
      8. Prof. Krishnan Balasubramanian, ME, IITM
      9. Prof. P.V. Manivannan, ME, IITM
      10. Prof. Murugaiyan Amirthalingam, MME, IITM
      11. Prof. Palaniappan Ramu, ED, IITM
      12. Dr. Piyush Shakya, ME, IITM
      13. Prof. Ramesh Babu N (Retd.), ME, IITM
      14. Prof. G.L. Samuel, ME, IITM
      15. Prof. Sathyan S, ME, IITM
      16. Prof. Somashekhar S. Hiremath, ME, IITM
      17. Prof. Sundararajan Natarajan, ME, IITM
      18. Prof. Sushanta Kumar Panigrahi, ME, IITM

2. **International Advisory Committee**
   1. Members:
      1. Prof. Dr.-Ing. habil. Barbara S. Linke, Department of Mechanical and Aerospace Engineering, University of California, Davis, California, USA
      2. Prof. Peter Gorm Larsen, Department of Electrical and Computer Engineering - Software Engineering & Computing Systems, Aarhus University, Denmark
      3. Prof. Satish Bukkapatnam, Department Head, Industrial & Systems Engineering, Texas A&M University, Texas, USA
      4. Prof. Xun W Xu, Department of Mechanical and Mechatronics Engineering, The University of Auckland, New Zealand

3. **National Advisory Committee:**
   a. Members:
      1. Prof. Amit Rai Dixit, IIT Dhanbad
      2. Prof. Apurbba Kumar Sharma, IIT Roorkee
      3. Prof. Asim Tewari, IIT Bombay
      4. Prof. Bhupesh Kumar Lad, IIT Indore
      5. Prof. J Ramkumar, IIT Kanpur

6. Prof. Kesavan D, IIT Palakkad
7. Prof. M Ravisankar, IIT Tirupati
8. Prof. M. Ramji, IIT Hyderabad
9. Prof. N Venkaiah, IIT Tirupati
10. Prof. Palani I A, IIT Indore
11. Prof. Ramesh Singh, IIT Bombay
12. Prof. S Aravindan, IIT Delhi
13. Prof. Santhakumar Mohan, IIT Palakkad
14. Prof. Sudip Misra, IIT Kharagpur
1. Prof. Surjya Kanta Pal, IIT Kharagpur

Sponsors

Office of Global Engagement, IIT Madras

Dassault Systemes

Meera Laser Solutions

Anusandhan National Research Foundation

Aktis Engineering Soltuions

National Instruments

VVDN Technologies

Barani Hydraulics Pvt. Ltd.

Keynote Speakers

Mr. Aritra Kumar Ghosh, VP, Mercedes-Benz, India

Mr. Aritra Ghosh is the Vice President of Manufacturing Engineering and IT Hub for Operations at Mercedes-Benz Research and Development India, based in Bengaluru. In this role, he leads organisational transformation and digital strategy at scale across manufacturing engineering and IT operations. He has extensive experience in automotive manufacturing engineering, having previously served as Vice President of Manufacturing Engineering at Mercedes-Benz R&D India since April 2023 before taking on the expanded IT operations mandate in May 2025. A thought leader in smart manufacturing and Industry 4.0, he has been a speaker at leading industry forums including the AUTO SCM Summit 2024, where he represented Mercedes-Benz on topics of supply chain innovation and sustainable manufacturing. He brings deep expertise in manufacturing digitalization, operational transformation, and the application of advanced technologies in large-scale automotive production environments.

Prof. Asim Tewari, IIT Bombay

Prof. Asim Tewari is a Professor in the Department of Mechanical Engineering and a faculty member of the Centre for Machine Intelligence and Data Science (C-MInDS) at IIT Bombay. He holds a B.Tech. from IIT Kanpur and M.S. and Ph.D. degrees from the Georgia Institute of Technology, Atlanta. With over two decades of experience in corporate R&D and national research laboratories, his research focuses on mathematical models for microstructural mechanics and applied machine learning. He has over 100 international journal and conference publications and fourteen patents (ten international, four Indian), has graduated thirteen Ph.D. students, and has secured over 15 million USD in research funding. At IIT Bombay, he leads a Cyber-Physical Systems and Data Analytics group working on smart manufacturing, IoT, and digital factory solutions, and has established state-of-the-art facilities including an Advanced Machining Excellence Cell and a 4D X-ray microscopy laboratory.

Prof. Mohit Law, IIT Kanpur

Prof. Mohit Law is a Professor of Mechanical Engineering at the Indian Institute of Technology Kanpur. Before joining academia, he worked as a machine tool designer in industry before pursuing graduate studies in the United States and Canada, followed by postdoctoral research in Germany. His use-inspired research spans intelligent manufacturing, machining dynamics, digital twins, machine tool metrology, vision-based measurement, and vibration control. He has authored nearly 100 publications, is an Associate Member of CIRP, and serves as an Associate Editor of the CIRP Journal of Manufacturing Science and Technology. A recipient of two distinguished teaching awards and several best paper awards, he actively collaborates with industry through research, consultancy, and technology transfer, with a strong focus on educating engineers for next-generation manufacturing.

Mr. Navin Manaswi, Department of Telecommunications (DoT), Government of India.

Mr. Navin Manaswi is an IIT Kanpur alumnus with over 12 years of experience in Artificial Intelligence, Generative AI, and emerging technologies. He is a recognized Google Developers Expert in AI and has represented India at the UN ITU-T in Geneva on global AI, Digital Twins, and Cybersecurity standards. He is a published author of three books — on Deep Learning (Springer US), GANs, and Generative AI and has authored over 70 articles on AI, ML, and Deep Learning. He has led AI initiatives for Google, Digital India, and Startup India, and has delivered end-to-end AI solutions across telecom, insurance, smart cities, and Industrial IoT. A sought-after keynote speaker and corporate trainer in Generative AI, he is currently engaged in building global standards for AI and Digital Twins.

Prof. Peter Gorm Larsen, Aarhus University, Denmark

Peter Gorm Larsen is a professor at the Department of Electrical and Computer Engineering, in the section for Software Engineering & Computer Systems, at Aarhus University in Denmark. He currently leads the AU Centre for Digitalisation, Big Data, and Data Analytics (DIGIT) and the AU Centre for Digital Twins. After receiving his M.Sc. in Electronic Engineering at the Technical University of Denmark in 1988, he worked in industry before returning to complete an industrial Ph.D. in 1995. In his industry career, as a development engineer and manager, he gave industrial courses all over the world on formal methods and saw the technology he developed applied in areas as diverse as secure message processing for options trading and near-field communication firmware. He returned to academia in 2005. His prime research goal is to improve the development of complex mission-critical applications by implementing and applying well-founded technologies, in particular the design of robust tools that help engineers leverage models in early design stages to reduce overall product development risk, especially in the areas of Cyber-Physical Systems and Digital Twins. He currently coordinates the Horizon Europe project RoboSAPIENS.

Prof. N. Ramesh Babu, IIT Madras

Prof. N. Ramesh Babu is the V Balaraman Institute Chair Professor in the Department of Mechanical Engineering at IIT Madras and Secretary of the Advanced Manufacturing Technology Development Centre (AMTDC) at IIT Madras Research Park. He received his Ph.D. from IIT Madras in 1990 for his pioneering work on laser dressing of grinding wheels. His research interests include manufacturing technology, grinding, abrasive waterjet machining, sheet metal fabrication, laser beam machining, CNC, PLCs, robotics, and process modeling and simulation of manufacturing systems. He has over 61 international journal papers and 104 national/international conference publications. He served as Head of the Department of Mechanical Engineering, IIT Madras (2017–2020) and as a Member of the IIT Madras Board of Governors. He is also a Member of the Academic Council of Anna University and the Development Council for Machine Tools, Ministry of Heavy Industries, Government of India.

Prof. Ramesh Singh, IIT Bombay

Ramesh Singh is G. K. Devarajulu Chair Professor in the Department of Mechanical Engineering at IIT Bombay. His research interests are laser additive manufacturing, ultra-high-speed micromachining, and finite element simulation of manufacturing processes. He is actively involved in the design and development of special-purpose machines for the Indian industry and strategic sector. He received his PhD from the Georgia Institute of Technology in 2007. He has published more than 220 international journal articles and peer-reviewed conference publications. He has been awarded 4 Indian and 2 US patents. He has guided 19 PhDs and over 80 Master's students. He received the prestigious Swarnajayanti Fellowship in Engineering Sciences in 2015 and the SME NAMRI Outstanding Paper Award at Virginia Tech in 2016. He is the Associate Editor of the ASME Journal of Micro-Nano Science and Engineering and the Transactions of the IISE. He serves on the Editorial Board of Nature Scientific Reports, JMST Advances, and the International Journal of Precision Technology. He is a fellow of ASME and serves as the President of the International Institute of Micro-Manufacturing for 2025-26.

Prof. Sam Anand, University of Cincinnati

Sam Anand is a Professor of Mechanical Engineering and Director of the Center for Global Design and Manufacturing at the University of Cincinnati (UC). He also serves as the Founding Director of Siemens Simulation Technology Center, Co-Director of the Industry 4.0/5.0 Institute, and Lab Director of the Digital Futures Smart Manufacturing Lab at UC. Prof. Anand received an M. Eng (Distinction) degree in Mechanical Engineering from the Indian Institute of Science in 1984. Subsequently, he received an MS and Ph.D. in Industrial Engineering from Pennsylvania State University in 1987 and 1990, respectively, and has worked in all areas of Advanced Manufacturing and Digital Manufacturing for the past 35 years. His research expertise includes Digital Design and Manufacturing, Smart Manufacturing, Modeling and Optimization of Additive Manufacturing, Industrial Internet of Things (IIoT), and Augmented/Mixed Reality applications in Manufacturing and Healthcare. He has graduated 80 MS and 15 Ph.D. students, published over 150 technical papers, and received approximately $25M in research grants and contracts from industry, state, and federal agencies, including the US National Science Foundation, the US Department of Labor, and the US Department of Defense. His recent grants include OH BWC-sponsored contracts on Generative AI-based Human Digital modeling for injury mitigation, US Army-sponsored DMDII and MxD contracts focusing on IIoT, Digital Manufacturing, and Additive Manufacturing, where he has partnered with Siemens, Boeing, GE Global Research, Raytheon, and Faurecia. Industry sponsors of his research include Raytheon Technologies, Honeywell, P&G, GE Aviation, GE Power & Water, Ford Motor Company, Eaton, Stanley Tools, and Lexmark. He has previously served as the Associate Editor of the ASME Journal of Manufacturing Science and Engineering and as the Associate Editor of the SME Journal of Manufacturing Systems.

Prof. Santhakumar Mohan, IIT Palakkad

Santhakumar Mohan is a Professor in the Department of Mechanical Engineering and Head of the Centre for Robotics and Control at the Indian Institute of Technology (IIT) Palakkad. He received his Ph.D. in Robotics and Control from IIT Madras in 2010. After serving as an Assistant Professor at NIT Calicut and a Postdoctoral Fellow at the Korean Advanced Institute of Science and Technology (KAIST), Republic of Korea, he joined IIT Indore in 2012 before moving to IIT Palakkad. He has also held visiting faculty and research positions at IISc Bangalore, KAIST, RWTH Aachen (Germany), ECN France, and BSTU Russia. His research focuses on service and field robotics, particularly underwater and wheeled mobile robots, parallel robotic platforms, and assistive and rehabilitation robotics. He has received several prestigious international fellowships and awards, including the Outstanding Young Scientist Award from the Korea Robotics Society (2014), the Alexander von Humboldt Fellowship, the EMARO+ Fellowship, the SAVI Fellowship, the World Class University Fellowship, and the Brain Korea 21 Fellowship.

Prof. Mohan has published more than 150 research papers, authored one book, edited four books, completed 18 externally funded research projects, and is currently leading 3 sponsored projects. He holds 12 granted Indian patents and 5 published Indian patent applications. He has served as Technical Editor of IEEE/ASME Transactions on Mechatronics and as Associate Editor of IEEE Robotics and Automation Letters, and currently serves on the editorial boards of several leading robotics and control journals, including Robotica and Intelligent Service Robotics. He is an Adjunct Professor at IIT Madras, a Faculty Fellow of the IIT Palakkad Technology IHub Foundation (IPTIF), and the founding Director of IPTIF, the IIT Palakkad Technology Research and Innovation Park (TRIP), and the Technology Innovation Foundation of IIT Palakkad (TECHIN).

Prof. Satish Bukkapatnam, Texas A&M University

Prof. Satish T. S. Bukkapatnam is the Rockwell International Professor and Head of the Department of Industrial & Systems Engineering at Texas A&M University, and Director of the TEES Institute for Manufacturing Systems. He earned his Ph.D. in Industrial & Manufacturing Engineering from Pennsylvania State University in 1997. His research focuses on smart manufacturing, harnessing high-resolution nonlinear dynamic information from wireless MEMS sensors, and improving monitoring and prognostics of ultraprecision and nanomanufacturing processes. He is a Fellow of SME (2018) and IISE (2016), and a recipient of the Fulbright-Tocqueville Distinguished Chair Award (2019–20), the IISE Technical Innovation Award (2014), and the IISE Hamed K. Eldin Outstanding Young IE Educator Award (2012). His research has yielded five pending patents in smart manufacturing and process monitoring.

Dr. Suchetana Chakraborty, IIT Jodhpur

Dr. Suchetana Chakraborty is an Associate Professor in the Department of Computer Science and Engineering at the Indian Institute of Technology Jodhpur, Rajasthan, India. She received her M.Tech and Ph.D. degrees in Computer Science and Engineering from IIT Guwahati, India, in 2011 and 2014, respectively, and has previously held faculty positions at IIIT Guwahati and BITS Pilani, Hyderabad Campus. Her research broadly spans Ubiquitous Computing, IoT, Edge Computing, Wireless Sensor Networks, Wearables, and Ambient Sensing for smart living applications. She was a Visiting Fellow at the School of

Computer Science and Engineering, UNSW Sydney, Australia, from June to September 2024, under the International Research Mobility Grant. At IIT Jodhpur, she has established the UbiSys Research Lab, which currently hosts around 20 students and project staff, including postdoctoral researchers and senior research associates. Over a 15-year career, Dr. Chakraborty has published more than 80 research papers in leading international journals and conferences and has led or co-led multiple high-value research projects funded by ANRF, DST, MeitY, DRDO, TIH, and industry partners. Her contributions have been recognized with several prestigious honors, including the Google Academic Research Award (2024), the Ad Hoc Networks Best Paper Award (2024), and the Honorable Mention Award at CHI 2026. Earlier accolades include the NIXI Fellowship (2014), the TCS Research India Fellowship (2011), the TUM Postdoc Mobility Fellowship (2014), and the Moonshot Award (2021).

She served as Associate Dean of Research and Development at IIT Jodhpur from 2023 to 2025 and has been an Area Editor of Ad Hoc Networks (Elsevier) for the past five years. She is actively engaged with the broader research community, serving on the Technical Program or Organizing Committees of leading conferences such as ICDCN, IEEE PerCom, ICBC, MobileHCI, CHI, LCN, COMSNETS, IWQoS, and WWW, and as a reviewer for reputed journals including IEEE TNSM, TMC, TCC, and TOSN, Sensors, Elsevier JPDC, JSA, Super Computing, PMC, FGCS, and Blockchain Research & Applications, among others.

Dr. Vignesh Vishnudas Shanbhag, NORCE, Norway

Dr. Vignesh V. Shanbhag received his M.Tech. in Manufacturing Engineering from the Vellore Institute of Technology (VIT), India, in 2015, and a joint Ph.D. in Manufacturing Engineering from Deakin University, Australia, and the Indian Institute of Technology Madras, India, in 2019. He has over 10 years of research experience in condition monitoring, signal processing, and data-driven diagnostics. He has received two Innovation Awards for his work on developing novel non-destructive inspection techniques for hydraulic cylinders and lithium-ion batteries. Currently, his research focuses on condition monitoring, gaze tracking, remanufacturing, and demanufacturing.

Mr. Aritra Ghosh is the Vice President of Manufacturing Engineering and IT Hub for Operations at Mercedes-Benz Research and Development India, based in Bengaluru. In this role, he leads organisational transformation and digital strategy at scale across manufacturing engineering and IT operations. He has extensive experience in automotive manufacturing engineering, having previously served as Vice President of Manufacturing Engineering at Mercedes-Benz R&D India since April 2023 before taking on the expanded IT operations mandate in May 2025. A thought leader in smart manufacturing and Industry 4.0, he has been a speaker at leading industry forums including the AUTO SCM Summit 2024, where he represented Mercedes-Benz on topics of supply chain innovation and sustainable manufacturing. He brings deep expertise in manufacturing digitalization, operational transformation, and the application of advanced technologies in large-scale automotive production environments.

Publication Opportunities

ICIDTSM 2026 is pleased to be partnering with **ASTM International: Smart and Sustainable Manufacturing Systems** to offer authors the opportunity to submit research presented at the conference.

ASTM International: Smart and Sustainable Manufacturing Systems journal fosters transdisciplinary research that crosses the boundaries of information science, systems engineering and engineering design, manufacturing, data science, and product life cycle with a focus on how to make manufacturing systems smarter and sustainable.

All authors with accepted submissions *(either oral or poster)* will be invited to submit their work for consideration of publication in ASTM International: Smart and Sustainable Manufacturing Systems. Be aware that submitted manuscripts will undergo a peer review process before publication.

Message from the Editors

It is our great pleasure to present the Abstract Book of the International Conference on Industrial Digital Twins for Smart Manufacturing (ICIDTSM 2026), organized by the Department of Mechanical Engineering, Indian Institute of Technology Madras, from 20–22 July 2026.

Industrial Digital Twins have emerged as a transformative technology for realizing intelligent, connected, and autonomous manufacturing systems. By integrating advanced sensing, artificial intelligence, robotics, cyber-physical systems, data analytics, simulation, and cloud technologies, digital twins are enabling unprecedented capabilities in real-time monitoring, optimization, predictive maintenance, and autonomous decision-making. ICIDTSM 2026 provides an interdisciplinary platform to bring together researchers, academicians, industry professionals, and students to exchange ideas, present innovative research, and explore future directions in smart manufacturing.

The conference received an enthusiastic response from researchers across academia and industry, with high-quality submissions spanning a broad spectrum of topics, including Industrial Digital Twins, Robotics and Automation, Artificial Intelligence and Machine Learning, Cyber-Physical Systems, Industrial Internet of Things (IIoT), Advanced Manufacturing, Additive Manufacturing, Machine Vision, Smart Metrology, Digital Manufacturing, and Sustainable Manufacturing. Each submitted abstract underwent a rigorous peer-review process conducted by the Technical Program Committee and expert reviewers to ensure the quality and relevance of the work presented.

This Abstract Book represents the collective efforts of researchers contributing to the advancement of manufacturing science and engineering. We hope that the technical presentations, keynote lectures, panel discussions, and interactions during the conference will foster meaningful collaborations, inspire innovative research, and strengthen partnerships between academia, industry, and government organizations.

We sincerely thank all the authors for sharing their valuable research contributions, the reviewers for their dedicated efforts in maintaining the scientific quality of the conference, the session chairs and keynote speakers for their support, and the organizing committee, volunteers, sponsors, and partner organizations for making ICIDTSM 2026 a success.

We wish all participants an engaging, productive, and memorable conference experience at IIT Madras.

Editors

Prof. N. Arunachalam
Prof. G.L. Samuel
Dr. Sivasrinivasu Devadula

Program Plan

Day 1: Monday, 20 July 2026

Time	Session / Event
09:30 AM - 10:15 AM	Inaugural Session
10:15 AM - 11:00 AM	**Inaugural Talk:** Prof. Satish Bukkapatnam (Texas A&M University) Title: Architecture and validation of digital twins for smart manufacturing Chair: Prof. GL Samuel
11:00 AM - 11:30 AM	High Tea Break
11:30 AM - 12:00 PM	Keynote 1: Prof. N. Ramesh Babu Title: Smart Manufacturing - Opportunities and Challenges Chair: Prof. Sivasrinivasu Devadula
12:00 PM - 12:30 PM	Keynote 2: Prof. Mohit Law Title: Everyone wants intelligent manufacturing. Who will teach it? Chair: Prof. Sivasrinivasu Devadula
12:30 PM - 01:30 PM	Oral Presentation: **Parallel Sessions** Oral Session 1 (Sensing & Communication Technologies) at TTJ Auditorium, IC&SR Oral Session 2 (Computation Metrology & Machine Vision) at Hall 3, IC&SR Oral Session 3 (AI for Vision, Safety & Intelligent Monitoring) at Hall 2, IC&SR
01:30 PM - 02:30 PM	Lunch Break
02:30 PM - 03:00 PM	Keynote 3: Prof. Santhakumar Mohan Title: Role of Robots in Industry 4.0/5.0 towards Smart Manufacturing Chair: Dr. Anuj Kumar Tiwari
03:00 PM - 04:30 PM	Oral Session 4 (Advanced Manufacturing & Additive Technologies) at TTJ Auditorium, IC&SR Oral Session 5 (Robotics & Automation Systems) at Hall 3, IC&SR Oral Session 6 (Digital Twins for Robotics & Autonomous Manufacturing) at Hall 2, IC&SR
04:30 PM - 04:45 PM	High Tea Break
04:45 PM - 05:30 PM	Panel Discussion 1: Digital Twin Readiness - Maturity and Industrial Impact. Moderator: Prof. GL Samuel Panel Members: Prof. Satish Bukkapatnam, Prof. N. Ramesh Babu, Prof. Mohit Law, Dr. Vignesh Vishnudas Shanbhag

Day 2: Tuesday, 21 July 2026

Time	Session / Event
09:30 AM - 10:00 AM	Keynote 4: Prof. Sam Anand Title: Digital Twins for Products, Processes, and Humans in Manufacturing Chair: Prof. Sushanta Kumar Panigrahi
10:00 AM - 10:30 AM	Keynote 5: Dr. Suchetna Chakraborty Title: From Sensor to Simulation: Trustworthy Sensing for Reliable Digital Twins in Smart Manufacturing Chair: Dr. Rahul Yadav
10:30 AM - 11:00 AM	High Tea Break
11:00 AM - 11:30 AM	Keynote 6: Prof. Asim Tewari Title: Agentic AI - the Next Frontier of Digital Twins Chair: Dr. Charchit Kumar
11:30 AM - 12:30 PM	Oral Session 7 (AI & ML for Smart Manufacturing) at TTJ Auditorium, IC&SR Oral Session 8 (Physics-based Modeling & Simulation) at Hall 3, IC&SR Oral Session 9 (Condition Monitoring, Prognostics & Predictive Maintenance) at Hall 2, IC&SR
12:30 PM - 01:00 PM	Keynote 7: Dr. Vignesh Vishnudas Shanbhag (NORCE, Norway) Title: Remanufacturing and Demanufacturing Chair: Prof. K Hariharan
01:00 PM - 01:30 PM	Sponsored Company Presentations
01:30 PM - 02:30 PM	Lunch Break
02:30 PM - 03:00 PM	Keynote 8: Prof. Ramesh Singh Title: Development of Laser DED-based Autonomous Robotic Restoration Process Chair: Prof. Amitava Ghosh
03:00 PM - 04:30 PM	Oral Session 10 (Cybersecurity) at TTJ Auditorium, IC&SR Oral Session 11 (Intelligent Manufacturing Systems) at Hall 3, IC&SR Oral Session 12 (Industrial Digital Twins & Autonomous Factories) at Hall 2, IC&SR
04:30 PM - 04:45 PM	High Tea Break

04:45 PM - 05:30 PM	Panel Discussion 2: Closing the Loop - Why Digital Twins Struggle to Enable Real-Time Control and Autonomous Decisions. Moderator: Prof. K. Hariharan Panel: Prof. Sam Anand, Prof. Asim Tewari, Prof. Santhakumar Mohan, Dr. Suchetana Chakraborty
05:30 PM - 06:00 PM	Group Photo
06:00 PM onwards	Cultural Event followed by Gala Dinner

Day 3: Wednesday, 22 July 2026

Time	Session / Event
09:30 AM - 10:00 AM	Keynote 9: Prof. Peter Gorm Larsen Title: Experience with Establishing Digital Twins Related to Manufacturing Chair: Dr. Anil Meena
10:00 AM - 10:30 AM	Keynote 10: Mr. Aritra Kumar Ghosh (Mercedes Benz) Title: Digital Twins in Manufacturing Chair: Prof. Somashekhar S. Hiremath
10:30 AM - 11:00 AM	High Tea Break
11:00 AM - 12:30 PM	Oral Session 13 (Digital Twins for Smart Infrastructure, Energy and Industrial Applications) at TTJ Auditorium
12:30 PM - 01:00 PM	Sponsored Company Presentations
01:00 PM - 02:00 PM	Lunch Break
02:00 PM - 02:30 PM	Keynote 11: Mr. Navin Manaswi (DoT, GoI) Title: Leveraging Energy Digital Twins to Achieve Net-Zero and Reduce Industrial Power Costs in India Chair: Dr. Piyush Shakya
02:30 PM - 03:15 PM	Panel Discussion 3: From Digital Twins to Living Systems - Self-Learning, Self-Healing, and Adaptive Manufacturing. Moderator: Prof. Palaniappan Ramu Panel: Prof. Peter Gorm Larsen, Prof. Ramesh Singh, Mr. Navin Manaswi, Mr. Aritra Kumar Ghosh (TTJ, IC&SR)
03:15 PM - 03:30 PM	High Tea Break
03:30 PM - 04:30 PM	Poster Presentation (Theme: Digital Twin for Smart Manufacturing) at HALL 4, IC&SR
04:30 PM - 05:15 PM	Valedictory Session and Prize Distribution

Presentation Schedule

Oral Presentation Abstract List

Oral Presentation Session 1
Theme: Sensing & Communication Technologies
Date: 20 July 2026 | Time: 12:15 PM - 01:30 PM |Venue: TTJ AUDITORUM, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3590	Vamsi Mallareddy (IIT Madras)	Design and Development of Test Platform for Electro-mechanical Characterisation of Triboelectric Sensors
3177	Shashi Bhushan Gunjan (IIT Madras)	xviiiRTD-Based Embedded Thermal Sensing and Displacement Estimation System for CNC Machine Tool Accuracy Improvement
2621	Prerit Sunil Shah (BITS Pilani Hyderabad Campus)	xviiiA PCB Structural Health Monitoring Digital Twin Framework Integrating MPU-6050 IMU Sensing and Machine Learning for Real-Time Vibration Fatigue Life Prediction
2572	Dhakshinesh Anandh (UPES, Bidholi, Dehradun)	VINS-Fusion-SAM: IMU-Gated Semantic Masking for Robust Visual-Inertial Odometry on Edge AI
2568	Leema Christy Devasagayam (Schneider Electric Ltd)	AI-Driven Software-Defined Architecture for Intelligent Power Quality Monitoring and Control
2898	Amgad Hassan Alamin Hassan (Delhi Technological University)	Safety Constrained Energy-Latency Optimization for LoRaWAN in Remote Petroleum Field Monitoring

Oral Presentation Session 2
Theme: Computation Metrology & Machine Vision
Date: 20 July 2026 | Time: 12:15 PM - 01:30 PM | Venue: Hall 3, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3934	Ramunivagaira Gnanaprakash (IIT Tirupati)	Proximity-Aware Human-Robot Collaboration: A Dynamic Safety Framework for Industrial Environments
3086	Navodita Singh (IISER, Bhopal)	Digital Twin for Efficient Robotic Scanning using RGB-D Sensing
3055	Ayyappan Selvaraj (VIT University Chennai)	Visibility-Aware YOLOv11 for Robust Real-Time Underwater Object Detection in Autonomous Surveillance
3048	Ramkumar K (PSG College of Technology, Coimbatore)	A smart vision system to detect operation's anomalies in assembly lines
3165	Mohamed Farook K (IIT Tirupati)	Real-Time Vision-Based Casting Defect Detection
2703	Prashant Priyadarshi (BIT Sindri, Dhanbad)	A novel framework towards reshaping the CNN topology

Oral Presentation Session 3
Theme: AI for Vision, Safety & Intelligent Monitoring
Date: 20 July 2026 | Time: 12:15 PM - 01:30 PM | Venue: Hall 2, IC&SR

Abstract ID	Presenter (Affiliation)	Title
4812	Vignesh M (NITTTR, Chennai)	Gesture-Based AI Safety Monitoring and Fault Detection for Machine Operators Using Edge AI and Human Pose Estimation
3135	Lavanya A (NIT Puducherry)	Stability-Based Degradation Trend Modeling for RUL Prediction of Broken Rotor Bar Faults

3140	Rajadurai S (Meenakshi Sundararajan Engineering College)	CLIP-Based Zero-Shot Visual Twin for Manufacturing
3072	Suryajith S S (College of Engineering Trivandrum)	Real-Time Distortion Monitoring During Welding of Structures Using Image Processing Techniques
3114	Abhishek Shrivastava (IIT Madras)	Investigation of Failure and Progressive Wear in Carbide Micro-Drills via Acoustic Emission Signal Analysis

Oral Presentation Session 4
Theme: Advanced Manufacturing & Additive Technologies
Date: 20 July 2026 | Time: 03:15 PM - 05:00 PM | Venue: TTJ Auditorium, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3820	Vipesh Singh (IIT Madras)	Acoustic Levitation and Passive Rotation of Planar Objects Using Perforated Vibrating Plates
3755	M Amarnath (IIITDM, Jabalpur)	Enhancement in Mechanical Properties and Microstructural Investigation of Aluminium Alloy 2024 using Microwave Hybrid Processing
3092	Eshan Amalnerkar (ESTECO Software India Pvt. Ltd.)	Towards Automated Structural Performance Optimization: Multi-Objective Design of a BIW Front Subframe with Coupled Shape Morphing and Stress Constraints
2899	D Mallikarjuna Reddy (VIT, Vellore)	Geometric Influences on Compressive Performance of Additively Manufactured Gyroid Cellular Architectures
3168	Aravindh Raj B R (Indian Institute of Information Technology)	Mechanical Performance of PETGCF-Based Sandwich Lattice Metamaterials Fabricated Using Fused Filament Fabrication
3167	Shishir Kasturi (VNRVJIET)	Process Parameter Optimization of PETG-CF 3D Printing Using Machine Learning
2404	K. Prabhakar Sriram (Delhi Technological University (DTU-DCE))	Six Sigma Application to Precision Instrument Manufacturing Facility for Defect Reduction: A Case Study
2647	Rajalakshmi K (Sri Sai Ram Engineering College)	IIoT-Based Smart Steam Management for Deodorization Efficiency in Refineries

Oral Presentation Session 5
Theme: Robotics & Automation Systems
Date: 20 July 2026 | Time: 03:15 PM - 05:00 PM | Venue: Hall 3, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3778	Sakshi Jha (The Neotia University)	Design and Conceptual Development of a Trunk-Mounted End-Effector for Autonomous Rubber Tapping
2655	Dr. J. Ronald Aseer (BIT Puducherry)	Lightweight Robotic Gripper Finger Design Using a Generative Approach
2628	Rituraj Bhattacharjee (MIT, Manipal)	Design and Experimental Evaluation of an Autonomous Water Surface Cleaning Robot
2623	Rajashekhar V S (IIT Kanpur)	Module Fabrication and Mixed Reality-Based Control of a Soft Snake Robot
3014	Anitha N (IIIT, Tiruchirappalli)	Performance Analysis of Discrete-Time Sliding Mode Control with Field-Oriented Current Regulation for PMSM Servo Drive under Variable Load in Robotics Application
3109	Puja Appasaheb Thorat (Sanjivani University, Kopargaon)	Development of an Autonomous Mobile Robot for Warehouse Application: A Review
3154	Ritvik Ravi (IIT Tirupati)	Real-Time Ergonomic Risk-Driven Human-Cobot Task Reallocation via Edge AI: An Industry 5.0 Framework with Economic Impact Quantification

Oral Presentation Session 6
Theme: Digital Twins for Robotics & Autonomous Manufacturing
Date: 20 July 2026 | Time: 03:15-05:00 PM | Venue: Hall 2, IC&SR

Abstract ID	Presenter (Affiliation)	Title
4687	M Selvakumar (IIT Madras)	Fidelity Mesh: Quantifying the Effect of URDF Geometric Fidelity on Autonomous Mobile Robot Digital Twin
4468	Monica Srinivas (VIT)	A Lightweight Web-Based Dashboard for Real-Time Monitoring and Control of Autonomous Mobile Robot Digital Twins
3117	Bhuvanesh P S (Madras Institute of Technology)	Design and Development of a Digital Twin-Based Autonomous Mobile Robot with Real-Time Path Adaptation
2625	Uttam Agarwal (BITS Pilani Hyderabad Campus)	Development of a Real-Time UAV Digital Twin for Cyber-Physical Monitoring and Visualization
3041	Dr. Velmurugan C (IIIT, Tiruchirappalli)	Development of a Real-Time Digital Twin Framework for Robotic Material Handling Systems in an Industrial Production Line
3090	Sreejith S (Cambridge Institute of Technology, Bengaluru)	A Predictive Digital Twin for GNSS-Denied UAV Navigation

Oral Presentation Session 7
Theme: AI & ML for Smart Manufacturing
Date: 21 July 2026 | Time: 12:15 PM - 01:30 PM | Venue: TTJ Auditorium, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3207	Rakesh Verma (IIT Madras)	Real-Time FPGA Edge Integration for High-Fidelity Chatter Control
3200	Harshvardhan Kharat (IIT Madras)	Adaptive Maintenance of Wind Turbine Gearboxes Using Digital Twin-Inspired Simulation
3170	Abhishek Punia (IIT Guwahati)	Physics-Guided Neural Network for Cutting Force Prediction in Turning
3150	Rakesh Verma (IIT Madras)	A Three-Layer Cyber-Physical System for a Boring Tool to Monitor and Predict Chatter During Machining
3775	Karthik M (Vidyavardhaka College of Engineering)	A Hybrid Contextual Reasoning Architecture for Fully Offline Conversational AI on Resource-Constrained Edge Devices
2664	Dr. Rahul Kumar Ray (FLAME University)	SMRITI: A Smart Manufacturing Reasoning Interface for Twin Integration using LLM-Enhanced Explainable AI

Oral Presentation Session 8
Theme: Physics-based Modeling & Simulation
Date: 21 July 2026 | Time: 12:15 PM - 01:30 PM | Venue: Hall 3, IC&SR

Abstract ID	Presenter (Affiliation)	Title
4338	Akash K (IIT Madras)	Development of a Physics-Based Digital Twin Framework for the CNC End Milling Process
3178	Patel Kunj Ketankumar (IIT Madras)	Physics-Based Digital Twin Development for Polymer Fused Filament Fabrication 3D Printing
3196	Sneha Anand (IIT Madras)	Adaptive Robotic GMAW Welding Simulation with DDPG Reinforcement Learning, Physics-Based Bead Geometry, and Gap-Adaptive Control for Marine Hull Fabrication
3120	Mourya Konathala (IIT Madras)	Firmware-Native Industrial Digital Twin for Thermo-Kinetic Geometric Drift Estimation and Compensation in Smart Fused Deposition Modeling
2990	Mohammed Abdul Razzaq (VIT)	A Comprehensive Study on Streamlining Backward Facing Step Flow Simulation Using the PyAnsys Framework
2890	Kiran Wakchaure (Sanjivani University)	Development of Digital Twin Framework for Compound Pendulum Experimentation empowered with IoT and Augmented Reality

Oral Presentation Session 9
Theme: Condition Monitoring, Prognostics & Predictive Maintenance
Date: 21 July 2026 | Time: 12:15 PM - 01:30 PM | Venue: Hall 2, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3604	Mahesh Babu N (IIT Madras)	Adaptive Optimization of Surface Roughness in CNC Milling using Experimental Machining Data
4458	Syed Ismail Sultan (IIT Kanpur)	Development of a Real-Time Digital Twin Framework for Autonomous Monitoring and Control of a Motorized Spindle
3121	R. Dhanalakshmi (MIT, Anna University, Chennai)	Smart Manufacturing Approach For Rotating Shaft Fault Detection Using Digital Twin And Vibration Analysis
3156	Jaikrishna R (Meenakshi Sundararajan Engineering College)	Smart Factory Predictive Maintenance Hub: An Edge-to-Cloud Architecture for RUL Prediction and Fault Classification in Rotating Industrial Machinery
2560	Abdelhalim Mohamed Hassan Ali (Delhi Technological University)	Physics-Informed Digital Twin for Bearing RUL Prediction with Bi-LSTM
2698	Dominic Sanju James R (BITS, Hyderabad)	Dual-Fault Diagnosis in Single-Stage Spur Gear Systems

Oral Presentation Session 10
Theme: Cybersecurity
Date: 21 July 2026 | Time: 03:15 PM - 05:00 PM | Venue: TTJ Auditorium, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3164	Palak Aggarwal (Doon University)	Robustness Analysis of Federated Learning Defenses Against Local Model Poisoning Attacks
3003	Harish Mrithyunjayan S (Bharathiar University)	A Digital Twin-Based Cybersecurity Framework for G-code Anomaly Detection in Additive Manufacturing Systems
3015	Siddesh R (Bharathiar University, Coimbatore)	AI-Based Cybersecurity Framework for Anomaly Classification in an Autonomous Factory Using a Confidence Weighted Voting Ensemble of SALSTM, TCN, and DBN
2893	Deepa Jose (KCG College of Technology)	Multi-Sensor Orchestrated Cyber Physical System Interactive Framework using Digital Twin for Secure Vehicular Networks
3146	Thulasi M Santhi (NIT Tiruchirappalli)	Digital Twin-Integrated Attack Detection Framework for Adaptive Cruise Control in Cyber-Physical Vehicle Networks
3163	Deepti Joshi (Doon University)	A stable and reliable defense mechanism for detecting label-flipping attacks in federated learning
3162	Ayush Maindola (Doon University)	A Secure and Robust Federated Learning Framework for Detecting and Mitigating Anonymous Free-Rider Attacks
2518	Nagharajhan K P K (Independent Researcher)	A DPDP-Compliant Edge AI Human Digital Twin Using Physics-Driven ODE and Adaptive Reinforcement Learning

Oral Presentation Session 11
Theme: Intelligent Manufacturing Systems
Date: 21 July 2026 | Time: 03:15 PM - 05:00 PM | Venue: Hall 3, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3472	Swathy S (Cochin University of Science and Technology)	Dynamic Beam Shaping using Coherent Beam Combining
3246	Ankush Goyal (Bennett University)	Capsule-Enhanced Generative Adversarial Networks for Financial Time Series: A Comparative Study of RNN Architectures and Mamba
3119	G. Devasena (IIIT, Tiruchirappalli)	Impact of Multimodal LLMs and Agentic AI in Smart Manufacturing - a systematic review
3123	Brijith Jacob (Cochin University of Science and Technology)	Developing a Generative DTDL Framework for Smart Manufacturing

3149	Avantika V (Meenakshi Sundararajan Engineering College)	Tiny ML-Compatible Lightweight Physics-Informed Digital Twin for Edge Systems: A Simulation Study with Exploratory Federated Unlearning
3136	Anurag Nandkumar Bansode (IIIT, Tiruchirappalli)	Designing a Distributed Ledger-Enabled Trust-Aware Federated Learning Framework for Auditable Backdoor Detection in Collaborative Intrusion Detection Systems
3132	Abirami J (Meenakshi Sundararajan Engineering College)	Autoencoder-Driven Prediction-Error-Triggered Digital Twin Synchronization on TinyML Edge Nodes for Bandwidth-Efficient IIoT Predictive Maintenance
3066	Ali Mohamed Ali Elsharef (Delhi Technological University (DTU))	Low-Power AI Inference Accelerator: MAC Unit Optimization via Operand Isolation and Switching Activity Reduction

Oral Presentation Session 12
Theme: Industrial Digital Twins & Autonomous Factories
Date: 21 July 2026 | Time: 03:15 PM - 05:00 PM | Venue: Hall 2, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3180	Anirban Tudu (IIT Madras)	A Real-Time Virtual Reality-Based Digital Twin Framework for Thermal Error reduction of Machine Tool Spindles
3118	Senthil Sivakumar M (IIIT Tiruchirappalli)	Integration of 3D Printing, AI, and IoT for Smart Manufacturing towards Automation for Industry 4.0: A Review
2701	Vijayakumar (SRM Institute of Science and Technology)	Orchestrating Industrial Complexity: A Proposed Human-Swarm Interaction Framework for Agentic Business Process Digital Twins
2979	Sharon Blessy (IIIT Tiruchirappalli)	Phased Integration of 5G Technology in a Digital Twin-Based Autonomous Factory
3166	Suraj Kumar (IIT Bombay)	Development of a Real-Time Digital Twin Framework for Robotic Material Handling Systems in an Industrial Production Line
3115	Ashish Mehta (IIT Madras)	From P&IDs to Knowledge Graphs: Enabling Process Intelligence through Digitization
3158	N V Jatin Pochiraju (SRM Institute of Science and Technology)	Metaverse-Based Digital Twin of a Smart e-Yantra Laboratory for Industrial IoT and Smart Manufacturing Applications

Oral Presentation Session 13
Theme: Digital Twins for Smart Infrastructure, Energy and Industrial Applications
Date: 22 July 2026 | Time: 11:30 AM - 12:30 PM | Venue: TTJ Auditorium, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3100	Dhivya K (VIT)	ABSA-GraphRec: Perception-Aware Constraint Learning via LLM-Enhanced Aspect Sentiment for Graph-Based Tourism Recommendation
2709	Samuel Mathias Menezes (Padre Conceicao College of Engineering)	Digital Twin of Mechanical Machines using Unity
2663	Annie Alex (VIT, Chennai)	Synergizing Multi-Source Energy Management: An Intelligent, Unified Architecture for Next-Generation Electric Vehicles
2659	Richie Garg (Schneider Electric)	AI-Driven Testpoint Prediction and Digital Twin Integration for Calibration Process Optimization in Smart Metering Devices
2656	Harish Muthu S (BITS, Hyderabad)	A Digital Twin Framework for Vibration-Based Condition Monitoring of a Single-Stage Spur Gearbox
2606	Shriiraam Satish (Bharathidasan Institute of Management)	A Distributed Sensor Network-Based Digital Twin for Real-Time Environmental Monitoring and Intelligent Decision Support in Indoor Workspaces
2605	Chakhyusa Saisrita Mittra (VIT, Chennai)	Digital Twin-Enabled Smart Grid for Hybrid Renewable Energy Integration with Self-Powered Sensing and Advanced Storage

Poster Presentation Abstract List

Theme: Digital Twin for Smart Manufacturing

Date: 22 July 2026 | Time: 03:30 PM - 04:30 PM | Venue: Hall 4, IC&SR

Abstract ID	Presenter (Affiliation)	Title
3733	Hamsashree S (NIT Puducherry)	Deep Learning Enabled Microstructure-Property Correlation in A356 Aluminium Alloy: Focus on Silicon Refinement and Grain Globularity by GIPUC Method
3458	Clement Rufus Raj A (Central University of Tamil Nadu)	Leveraging Aluminum 6061 for a Sustainable circular economy via powder Metallurgy by Integrating Machine Learning
3229	Shiwam Shorya Sharma (IITM Pravartak Technologies Foundation)	Cloud-Enabled Digital Twin Framework of an Industrial Robot for Seamless Integration and Task Automation in Autonomous Factories
3228	Manikala Prem Sainadh (NIT Puducherry)	Machine Learning-based Fault Diagnostics for Industrial Robots using Accelerometer Signals
3204	V Aruna (IIT Madras)	Automated Painting of Aerospace Components: An end-to-end pipeline from STL geometry to coordinated robot execution
3198	Arshiya Singh (VIT, Chennai)	Bubble Detection and Parameter Optimisation for Sterile Vial Filling
3159	Kalaiyarasan R (NIT Tiruchirappalli)	Tribological Behavior of Wire Arc Additive Manufactured Super Duplex Stainless Steel using Laser Shock Peening
3151	Srihari Prasad (IIT Madras)	AI-Based Predictive Maintenance Framework for Wind Turbine Using Real-Time SCADA Data
3107	Harish Muthu S (BITS Pilani, Hyderabad)	Development and Implementation of ROS 2 Architecture for Differential Drive Indoor AMRs
2674	Porrkodi M (Central University of Tamil Nadu)	Additive Manufacturing of Casting Tooling
3079	Jeevika Kannan (Meenakshi Sundararajan Engineering College)	Multimodal Tutor Kit for Blind Children
2557	G. Anitha (R.M.D Engineering College)	Smart Prescription-Based Automated Medicine Dispenser
3101	Loganathan, Velayutham (TTANI.AI)	AI Research - Autonomous Cement Plant Operations
2498	Ajithkumar A (Shinkai-labs Pvt. Ltd.)	Multi-Domain Digital Twin Framework for Real-Time Monitoring and Object Identification Across Air, Sea, and Ground Domains

List of Abstracts

Abstract ID: 3590

Design and Development of Test Platform for Electro-mechanical Characterisation of Triboelectric Sensors

Vamsi Mallareddy*, Tadesse Wasie, Nayana K V, Nithish S, Charchit Kumar

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India
**me24d401@smail.iitm.ac.in*

Abstract

The global demand for digital twins in smart manufacturing has escalated the need for a new type of autonomous, self-powered sensors. Traditional sensors, which bridge the gap between physical assets and their digital counterparts, require a continuous power supply. They are often powered by batteries that require regular charging and replacement after use. Along with various environmental concerns, batteries are unreliable for remote sensing and continuous monitoring. As smart manufacturing shifts toward a more sustainable approach, the integration of energy-harvesting technologies with sensors, such as triboelectric-based sensors, has emerged as a promising solution.

In this regard, a new technology, Triboelectric Nanogenerators (TENGs), has gained significant attention from researchers. TENG can utilise ambient mechanical energy from industrial environments and convert it to usable electrical signals. TENGs mainly operate on the combined principles of contact electrification and electrostatic induction during the repeated interaction between tribolayers. When two different triboelectric materials come into contact, charge transfer occurs mainly due to differences in their work functions. The transferred charges induce a potential that drives electron flow through the external circuit, producing the highly sensitive signals required for sensing applications. TENG devices typically operate in two fundamental modes: vertical contact-separation mode and lateral-sliding mode.

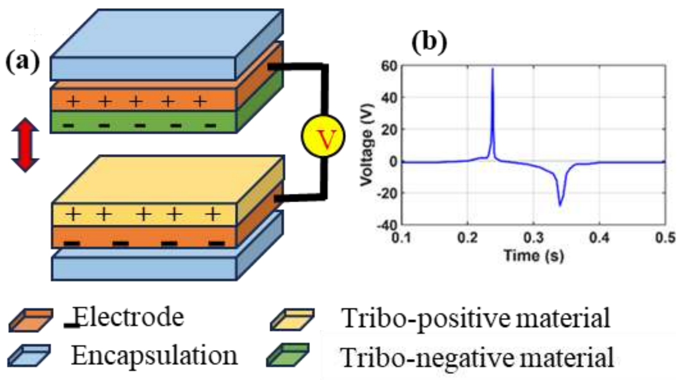

Figure 1. (a) Simplified schematic showing a TENG device operation under contact-separation mode. and (b) a representative plot of output voltage characteristics from a TENG device.

To rigorously evaluate triboelectric sensor functionality and performance, systematic testing under controlled conditions is essential. However, no standardised testing protocols currently exist for self-powered sensors. This research addresses the gap by developing an indigenous electromechanical test platform that precisely controls critical parameters, including applied load, displacement, and frequency. The rig is also equipped with a universal contact alignment mechanism to ensure flat-flat contact between the tribolayers. Multiple folding mechanisms are designed and developed using additive manufacturing. Various polymeric tribolayers are fabricated using a replica moulding technique and assembled within the folding mechanisms. Each device is further analysed to replicate realistic operating conditions while maintaining effective contact area and proper structural interface alignment between the tribolayers.

The fabricated sensing devices tested on this platform demonstrated a significant increase in voltage output with increasing applied pressure. The observed high sensitivity and efficiency indicate that triboelectric sensors have strong potential as suitable replacements for conventional sensor technologies in industrial monitoring applications, including pressure, strain, and low-frequency vibration measurements.

Abstract ID: 3177

RTD-Based Embedded Thermal Sensing and Displacement Estimation System for CNC Machine Tool Accuracy Improvement

Shashi Bhushan Gunjan[1]*, Ridha Sham[2], D.S. Srinivasu[1]

[1]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India*
[2]*Department of Electrical Engineering, National Institute of Technology Puducherry, Karaikal, India*
**gunjanshashi@gmail.com*

Abstract

Thermal deformation is a dominant source of dimensional error in precision CNC machine tools, because heat generated by spindle bearings, servo drives, ball screws, guideways, cutting action, coolant, and ambient variations creates non-uniform temperature fields within the machine structure.- During prolonged operation, the controller continues to execute the tool path with respect to nominal calibration conditions, whereas the physical tool center point (TCP) gradually shifts due to structural expansion and bending. This work addresses the sensing-control gap by developing an embedded RTD-based thermal sensing and displacement-estimation architecture for real-time CNC thermal-error compensation.

The proposed system follows an indirect compensation strategy in which distributed temperature measurements are converted into TCP displacement estimates. The sensing layer uses Class A Pt100 resistance temperature detectors located at thermally sensitive regions of the CNC structure. The RTDs are interfaced with LTC2497 delta-sigma ADCs via a ratiometric excitation circuit, so that reference-voltage drift is rejected by measuring the resistance ratio rather than the absolute voltage. The embedded processor is an ARM Cortex-M7 microcontroller that acquires ADC data via I2C, computes RTD resistance, applies the Callendar-Van Dusen relation for temperature linearization, and executes locally regression-based displacement inference (Figure 1). The architecture also separates the analog and digital domains and incorporates galvanic isolation to ensure safer connections to industrial controllers.

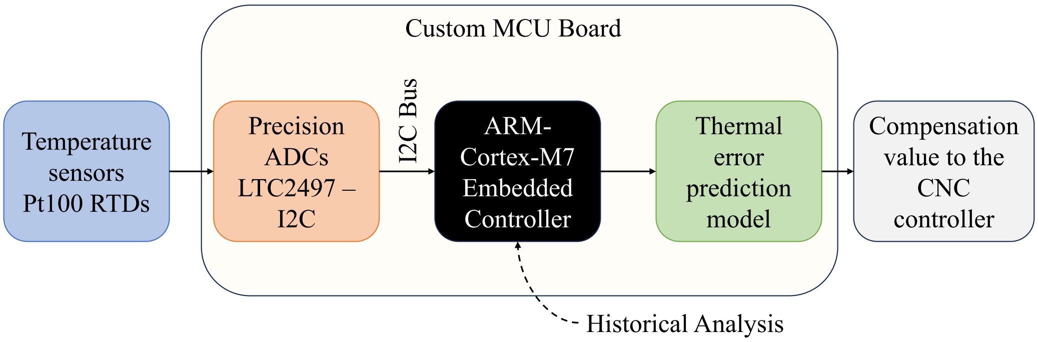

Figure 1. RTD-based thermal sensing and embedded displacement estimation architecture for CNC thermal compensation.

For model development, six temperature features were considered from representative structural locations, including spindle housing, column top, column base, X-axis guideway, Y-axis guideway, and lead-screw region. Calibration data were obtained by recording the temperature vector and corresponding TCP displacement during machine warm-up and cool-down transients. Two regression models were evaluated: multiple linear regression (MLR), which assumes additive thermal influence, and second-order polynomial regression, which augments the input space with quadratic and interaction terms. The polynomial formulation is suitable for machine tools because the thermal field is spatially coupled and exhibits nonlinear transient behavior; at the same time, it remains computationally lightweight enough for embedded execution on the Cortex-M7 processor's floating-point unit.

The assembled prototype was validated at the hardware, sensing, and model levels (Figure 2). Power-rail verification showed stable 5 V, 3.3 V, and 3.8 V outputs from a 12 V industrial input. I2C communication with the ADCs was verified at 100 kHz with correct acknowledgment and data framing. Temperature measurement was compared with a calibrated Fluke 1524 reference thermometer across 25-60 °C. The maximum observed temperature error was ±0.11 °C, and the static noise floor was approximately ±0.07 °C. These values are adequate for machine-tool thermal compensation, where sub-degree thermal changes can produce micron-level TCP drift. In model evaluation, the MLR model achieved an R2 score of 0.87, a mean absolute error of 4.2 µm, and a maximum error of 12.1 µm. The polynomial regression model improved performance to an R2 score of 0.94, a mean absolute error of 2.1 µm, and a maximum error of 6.8 µm, confirming the benefit of including nonlinear and interaction effects in thermal displacement prediction.

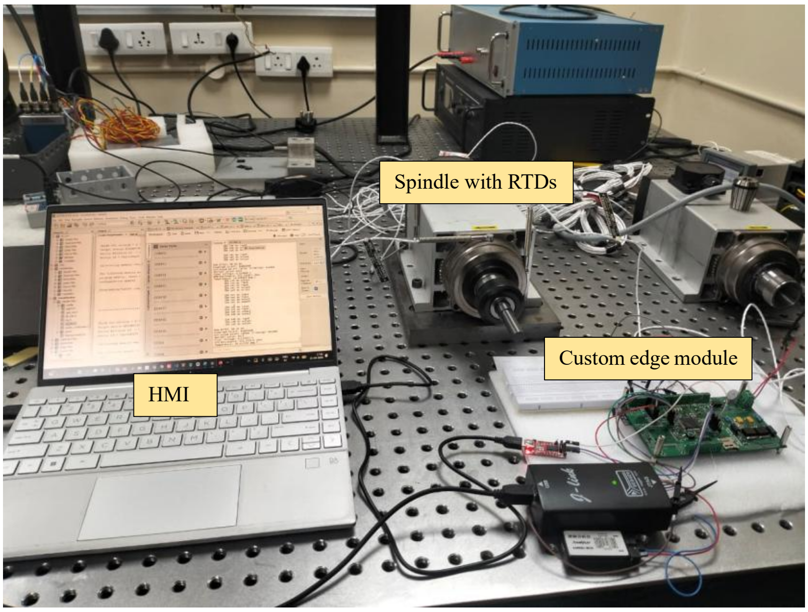

Figure 2. Experimental setup for temperature measurement using RTD and the proposed edge module.

The results demonstrate that a low-latency embedded thermal sensing module can bridge raw thermal measurements and actionable compensation information for CNC machines. The work contributes to industrial digital twins and smart manufacturing by providing a scalable cyber-physical layer that streams temperature states, estimates thermal drift, and supports future closed-loop compensation or digital-twin-based model updating. Future work will focus on integration with commercial CNC/PLC interfaces, sensor-placement optimization, compensation validation during cutting trials, and adaptive model updating under varying production conditions.

Keywords: CNC machine tools; thermal error compensation; Pt100 RTD; embedded sensing; edge computing; smart manufacturing.

Abstract ID: 2621

A PCB Structural Health Monitoring Digital Twin Framework Integrating MPU-6050 IMU Sensing and Machine Learning for Real-Time Vibration Fatigue Life Prediction

Prerit Sunil Shah[1]*, Arpit Kapoor[1], Aryan Karas[1], Ayan Das Sarkar[2], Sarakapalli Sai Krishna Nikhitha[2], Anbumani P[2], G.R. Sabareesh[1], Sudha Radhika[2]

[1]*Department of Mechanical Engineering, BITS-Pilani Hyderabad Campus, India*
[2]*Department of Electrical Engineering, BITS-Pilani Hyderabad Campus, India*
**f20220996@hyderabad.bits-pilani.ac.in*

Abstract

Real-time structural health monitoring of PCBs under sustained vibration loading is critical for aerospace, automotive, and defence applications. The conventional approaches (finite element analysis (FEA) or destructive testing) provide no continuous health monitoring information. This work attempts to address that gap with a hardware-in-the-loop digital twin framework based on an MPU-6050 MEMS inertial measurement unit bonded at the centre of mass of the base of a power supply PCB mounted on an electrodynamic shaker. Six DOF inertial data (tri-axial acceleration and gyroscope) is collected at 50Hz via an Arduino Uno to a host PC, where a complementary filter reconstructs spatial orientation and a six-element feature vector is passed to a dual-stage machine learning pipeline. A Random Forest regressor forecasts next-step vertical acceleration while the Isolation Forest flags anomalous vibration states; the absolute prediction error supports a divergence-based health index that decreases monotonically towards failure without relying on FEA-generated training data. A PyQt5 interface provides live orientation plots, a 3D CAD/STL viewer, and continuous session logging. Controlled excitation tests at 1G, 3G, and 5G were done to benchmark the framework against established FEA-derived baselines.

Keywords: Digital Twin, PCB, MPU-6050, IMU, I2C, Vibration Fatigue, Machine Learning, Structural Health Monitoring, Random Forest, Isolation Forest.

Introduction

PCB assemblies in aerospace, defence, and automotive systems endure sustained broadband vibration that accumulates fatigue damage in solder joints over time. Conventional reliability assessment relies on destructive testing or computationally intensive Finite Element Analysis (FEA), neither of which provides real-time health monitoring capability. The digital twin framework integrating physics models, live sensor data, and ML surrogates can offer an alternative for continuous structural monitoring.

The prior work of the present group established an adaptive CatBoost–MOA-Hu-Washizu model achieving R2 = 0.9982 for the prediction of fatigue life of single-component PCBAs under vibration, and a Hybrid Computational Intelligence Framework (HCIF) using Random Forest that achieved R2 = 0.89 and MAE = 5.1 on multicomponent PCBAs. Both approaches relied on FEA-generated datasets. The present work addresses this gap by replacing simulated signals with real-time MEMS IMU measurements from a power supply PCB on an electrodynamic shaker for hardware-in-the-loop digital twin operation.

Methodology

Experimental Setup

A commercially available power supply PCB carrying a transformer, capacitors, inductor, and cooling fan is mounted to an aluminium fixture clamped to an electrodynamic shaker (Fig. 2). Excitation levels of 1G, 3G, and 5G are induced (20–2000Hz).

Sensing and Data Acquisition

The MPU-6050 (±16g accelerometer; ±2000◦/s gyroscope) is bonded by an adhesive at the centre of mass of the power supply board, and the IMU communicates via I2C at 400kHz. An Arduino Uno reads all six DOF registers at 50Hz and forwards data at 115200baud to a host PC, where pyserial logs time-stamped frames to CSV/XLSX files.

Signal Processing and Feature Engineering

Raw accelerometer counts are converted to physical units by the full-scale sensitivity factor. Roll (φ) and pitch (θ) are recovered with a complementary filter (α = 0.98):

$$\phi_k = \alpha\left(\phi_{k-1} + \dot{\phi}, \Delta t\right) + (1-\alpha)\phi_{accel} \tag{1}$$

A six-element feature vector is constructed at each sample: **x**

$$\mathbf{x} = \left[a_x, a_y, a_z, \|\boldsymbol{a}\|, |a_z|, \phi\right]^{\mathsf{T}}, \|\boldsymbol{a}\| = \sqrt{a_x^2 + a_y^2 \ + a_z^2} \tag{2}$$

All features are standardised with the Standard Scaler before being input to the model.

ML Pipeline and Health Assessment

A Random Forest regressor (N_{trees} = 100) predicts $a^{z(t+1);}$ an Isolation Forest (ρ = 0.10) simultaneously flags anomalous vibration states in the feature space. The divergence metric

$$\varepsilon = \left|\, a_z^{real} - \hat{a}_z \right| \tag{3}$$

is monitored continuously. The composite health index

$$\mathrm{H} \ = 100\left(1 - \frac{\varepsilon}{\varepsilon_{max}}\right) \tag{4}$$

provides a monotonically decreasing fatigue state indicator similar to the D = 1 Miner–Palmgren damage criterion. After ≥ 50 samples are collected, both models auto-retrain and are serialised via joblib into a per-PCB model library for cross-session deployment.

Framework Overview

The four-layer architecture (Physical Setup & Sensing → Data Acquisition & Transmission → Data Processing → Digital Twin Application) is illustrated in Figure 1. The PyQt5 GUI is shown in Figure 3.

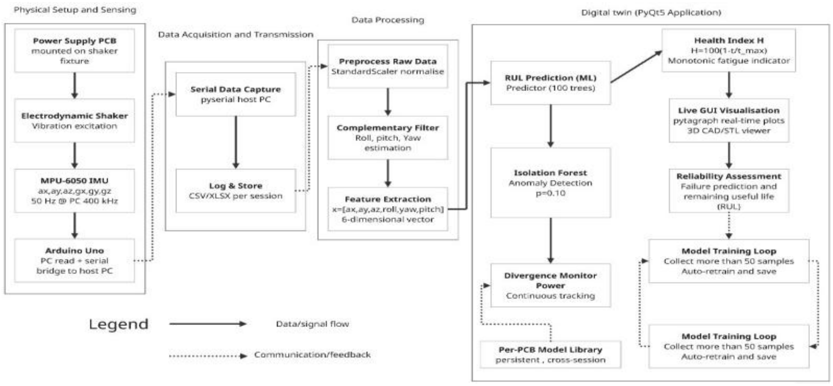

Figure 1. Proposed Digital Twin framework showing four layers: Physical Setup & Sensing, Data Acquisition & Transmission, Data Processing, and Digital Twin Application (PyQt5). Solid arrows: data/signal flow; dashed arrows: communication/feedback.

Experimental setup and GUI

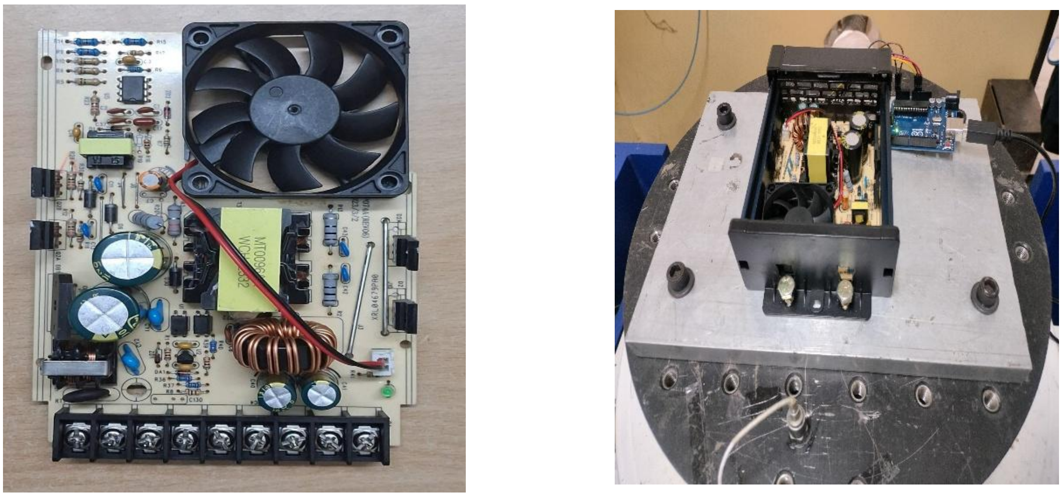

(a) (b)

Figure 2. Experimental test setup. (a) Power supply PCB, (b) PCB on electrodynamic shaker.

Results and discussion

The electrodynamic shaker subjects the PCB to 1G, 3G, and 5G random vibration (20–2000Hz), logging IMU data at each level until visible solder joint failure or a predefined limit. Health-index trajectories, anomaly flags, and divergence spikes will be correlated and compared against Steinberg three-band damage fractions (68.3%, 27.1%, 4.3% at 1σ–3σ) used in previous work. Framework accuracy will be compared with FEA-derived baselines of R2 = 0.9982 (CatBoost–MOA–HuWashizu) and R2 = 0.89 (HCIF–Random Forest).

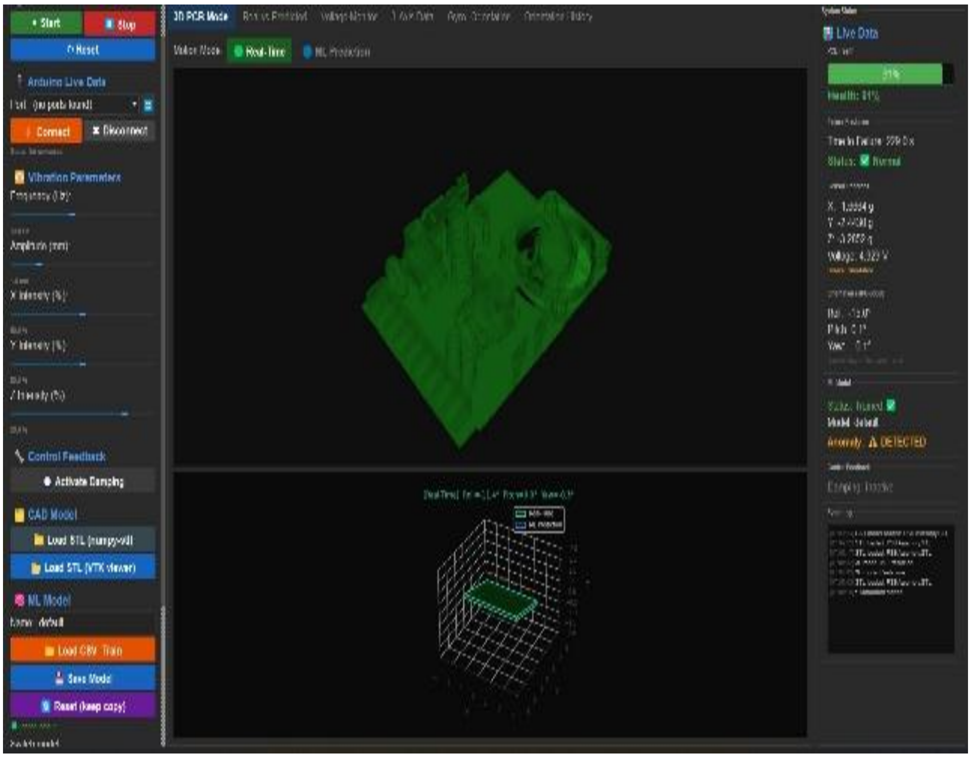

Figure 3. PyQt5 Digital Twin GUI: 3D CAD/STL viewer, real-time orientation plot, and System Status panel with live PCB health index, sensor readings, and ML model status.

Conclusions

A cost-effective hardware-in-the-loop digital twin framework for PCBA structural health monitoring has been developed and is ready for experimental validation. Key contributions will be: (1) a six feature IMU derived inertial vector replacing FEA-computed inputs; (2) a dual-stage ML pipeline combining Random Forest regression and Isolation Forest anomaly detection; (3) a persistent per-PCB model library for cross-session deployment; (4) a normalised composite health index providing a continuous fatigue state indicator. These advances will make the framework ready to deploy on physical assemblies without finite element preprocessing.

Abstract ID: 2572

VINS-Fusion-SAM: IMU-Gated Semantic Masking for Robust Visual-Inertial Odometry on Edge AI

Agrima Jain[1], Dhakshinesh Anandh[1], J Dhiviya Rose[1]*, S Christalin Nelson[1], Muthukumar K A[2]

[1]*School of Computer Science, UPES, Dehradun, India*

[2]*Department of Computer Science and Engineering, SSN College of Engineering, Chennai, Tamil Nadu, India*

**dhiviyarj@ddn.upes.ac.in*

Abstract

Visual-Inertial Odometry (VIO) systems based on nonlinear sliding-window optimization achieve high pose estimation accuracy in static environments, but their performance degrades severely in the presence of dynamic objects, texture-less regions, and extreme lighting. These conditions generate geometrically inconsistent feature tracks that heavily corrupt bundle adjustment. To address this, we present VINS-Fusion-SAM, an extension of the VINS-Fusion framework that integrates the Semantic Segment Anything Model (SAM) directly into the feature tracking front-end. By generating dense segmentation masks, our system constrains optical-flow tracking and Harris corner candidate proposals to strictly static, geometrically reliable regions, successfully suppressing outlier feature tracks before they enter the estimator.

A critical challenge for SLAM-SAM integration is the severe high GPU inference cost associated with running foundation models at camera frame rates on resource-constrained Edge AI devices. We overcome these limitations by deploying MobileSAM over an asynchronous, decoupled node architecture in ROS 2. To strictly maintain real-time throughput on a shared-memory edge platform (Nvidia Jetson Orin Nano), we introduce two key strategies: aggressive prompt-grid decimation to reduce GPU memory footprint and an IMU-Gated Adaptive Segmentation mechanism. By continuously monitoring the trace of the IMU pre-integration rotation covariance, the system accurately predicts high-uncertainty motion conditions. When rapid rotational motion occurs, costly GPU inference is pre-emptively skipped, and the previous semantic mask is warped forward using an affine approximation of the IMU relative rotation. We evaluate the proposed system utilizing the EuRoC MAV dataset inside a ROS 2 Humble environment. Through quantitative evaluation employing Absolute Trajectory Error (ATE) metrics derived via the evo toolkit, our system demonstrates significant improvements in trajectory robustness while preserving the computational bandwidth necessary for real-time edge robotics.

Keywords: Visual-Inertial Odometry; Segment Anything Model; SLAM; Edge AI; ROS 2; IMU-Gated Segmentation; MobileSAM; Real-Time Robotics.

Introduction

State-of-the-art VIO systems such as VINS-Fusion[1] achieve competitive trajectory accuracy under controlled conditions by tightly coupling visual feature tracks with IMU pre-integration in a sliding-window bundle adjustment backend. However, the core assumption of scene stationarity is routinely violated in real-world deployments: dynamic foreground objects corrupt the reprojection error landscape, large texture less surfaces yield sparse and drift-prone feature populations, and abrupt illumination changes degrade optical-flow reliability. These failure modes are well-documented in the SLAM literature[2] but remain inadequately addressed by geometry-only outlier rejection methods such as RANSAC, which cannot distinguish semantically dynamic features from genuine static-scene outliers. Foundation models for visual segmentation, particularly the Segment Anything Model (SAM)[3] and its lightweight derivative MobileSAM[4], offer a principled mechanism for semantic scene decomposition. However, deploying such models within a real-time VIO pipeline introduces non-trivial computational constraints that have not been systematically addressed for edge AI hardware.

Methodology

Semantic Front-End Integration. VINS-Fusion-SAM inserts a segmentation mask generation stage between raw image input and the Harris corner detector. MobileSAM is invoked on each incoming frame to produce a dense binary foreground mask that labels pixels belonging to semantically dynamic object classes. Both optical-flow tracking of existing landmarks and new Harris corner candidate proposals are gated by this mask, restricting feature populations exclusively to static, geometrically stable regions. This semantic filtering eliminates the root cause of estimator contamination rather than reactively rejecting corrupted tracks post-hoc. IMU-Gated Adaptive Segmentation (IGAS). The principal computational bottleneck is sustained GPU inference at camera frame rates. IGAS addresses this by monitoring the trace of the IMU pre-integration rotation covariance matrix as a lightweight proxy for platform motion intensity. When the covariance trace remains below a calibrated threshold — indicating smooth, predictable motion — MobileSAM inference proceeds normally. When rapid rotational motion is detected, inference is pre-emptively skipped, and the previous frame's semantic mask is propagated forward via an affine warp derived from the IMU relative rotation estimate. This ensures the segmentation mask remains geometrically consistent with the new frame without incurring GPU inference latency, preserving real-time throughput on the Nvidia Jetson Orin Nano.

System Architecture. The full pipeline is implemented as an asynchronous, decoupled node architecture in ROS 2 Humble, enabling the segmentation node and the VIO estimator to operate on independent execution threads without mutual blocking. Prompt-grid decimation reduces the MobileSAM input resolution aggressively to minimize GPU memory allocation while retaining sufficient spatial resolution for reliable dynamic-object boundary delineation. The combined system runs end-to-end on the shared-memory architecture of the Jetson Orin Nano without requiring inter-process data copies between the segmentation and tracking modules.

Results and Discussion

The proposed system is evaluated on the EuRoC MAV benchmark dataset[5] within a ROS 2 Humble simulation environment. Trajectory accuracy is assessed using Absolute Trajectory Error (ATE) metrics computed via the evo evaluation toolkit, enabling direct comparison with the unmodified VINS-Fusion baseline across all benchmark sequences. VINS-Fusion-SAM demonstrates consistent and measurable improvements in trajectory robustness across sequences characterized by dynamic foreground objects, textureless regions, and aggressive lighting variation, precisely the conditions under which the baseline estimator suffers the most severe feature track corruption. Tracking failures observed in baseline runs are suppressed by the semantic front-end gate, yielding cleaner feature populations that reduce drift accumulation over extended trajectories. The IGAS mechanism successfully prevents localization breakdown during high-rotation segments: by warping the prior mask forward rather than triggering GPU inference mid-maneuver, the system avoids the latency spikes that would otherwise stall the feature tracker and destabilize the estimator state. Real-time throughput is maintained throughout all evaluation sequences on the Jetson Orin Nano, confirming that the computational overhead introduced by semantic masking is fully absorbed by the asynchronous, decoupled architecture and prompt-grid decimation strategy. These results demonstrate that foundation model-class semantic intelligence can be made compatible with the strict latency and memory budgets of embedded edge robotics without compromising odometry accuracy.

Conclusion

VINS-Fusion-SAM presents a practical and deployable approach to semantically robust visual-inertial odometry on resource-constrained edge platforms. By integrating MobileSAM into the VIO feature tracking front-end and introducing the IMU-Gated Adaptive Segmentation mechanism, the system delivers improved trajectory robustness in challenging dynamic environments while preserving the real-time operational constraints required for embedded robotics. The asynchronous ROS 2 architecture and prompt-grid decimation together ensure that the computational demands of foundation model inference do not compromise estimator throughput on the Jetson Orin Nano. The proposed framework applies to a wide range of engineering systems, including UAVs, autonomous mobile robots, and vision-guided robotic manipulators, where robust visual-inertial localization and efficient use of onboard computational resources are crucial for reliable real-time operation. Future work will focus on extending evaluation to custom indoor dynamic datasets, incorporating instance-level semantic labeling for richer scene understanding, and exploring adaptive prompt-grid strategies that dynamically balance segmentation resolution against available GPU headroom.

Abstract ID: 2568

AI-Driven Software-Defined Architecture for Intelligent Power Quality Monitoring and Control

Leema Christy Devasagayam[1, 2]*, Dr Ajitha Shenoy K B[1], Prateek Jain[2]

[1]*Manipal Academy of Higher Education, Manipal, India*
[2]*Schneider Electric India Pvt. Ltd., Bangalore, India*
**leema.v2004@gmail.com*

Abstract

The global integration of renewable energy sources, distributed generation, and loads that demand high power densities is also causing major transformation in modern electrical power systems. The developments have turned the need to monitor power quality (PQ) into an engineering challenge and not an operational requirement. The traditional hardware-based monitoring systems were designed with the power grid in mind, which is centralized and predictable, but not true in today's dynamic and heterogeneous power world. The paper proposes the software-defined architecture of the proposed system that brings together the concepts of artificial intelligence (AI), machine learning (ML), and edge computing to guarantee power quality monitoring and management in modern power systems with intelligence and in real-time.

Methodology

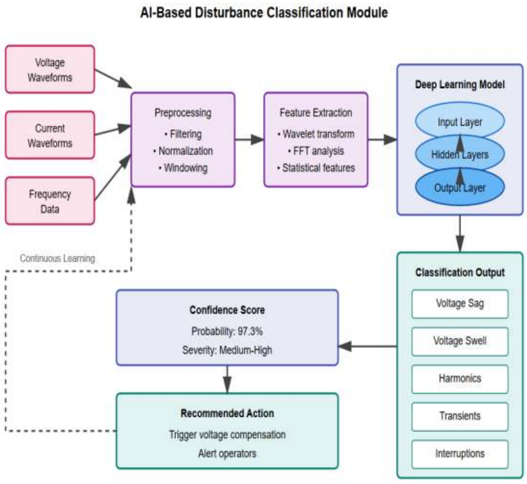

Figure 1. AI-Based Disturbance Classification Module illustrating the overall pipeline of signal input to deep learning classification to suggested control actions.

The suggested architecture embraces a three-layer software-defined architecture. The monitoring points are equipped with smart sensors, reconfigurable power electronic converters, and edge computing devices known as the Hardware Layer. It includes the Software Layer with powerful Deep Learning algorithms such as Convolutional Neural Networks (CNNs) to classify a disturbance in real time, predictive control models, and cybersecurity modules. The Management Layer enables system-wide optimization, resource allocation, policy enforcement, and operator interfaces. Edge intelligence takes the raw voltage, current, and frequency signals to the edge of the system to allow a swift response in under a millisecond without centralized cloud support. Features in power waveforms can be automatically detected by the CNN and grouped into a number of different types of disturbances, such as voltage sags, voltage swells, harmonics, and transients, without requiring any engineering features to be programmed into the CNN. Predictive ML algorithms can be used to predict disturbances minutes to hours before they occur to prevent them, based on past behaviour. The legacy and modern infrastructure are interoperable thanks to communication standards and protocols like IEC 61850, MQTT, OPC UA, and DNP3. A multi-layer cybersecurity system with network segmentation and continuous threat monitoring is constructed to protect critical grid resources through the incorporation of an ML-based anomaly detection.

Results and Discussion

The system has been proven to be effective in real-world implementation in over 50 deployments, including isolated communities with a 2.3 million square kilometre area and high-density urban microgrids. The CNN-based disturbance classifier claimed an accuracy of more than 97 percent and was also faster and more multi-classable than other signal-processing techniques, such as wavelet transforms and probabilistic neural networks. The edge processing meant there were no sub-cycle delays, while there were no bottlenecks in the cloud latency due to centralized architectures. A 5th of the cost of the hardware-based solutions employed at the baseline was saved on the cost of power quality management, and the system availability was 99.9%. The predictive control algorithms predicted voltage sags and voltage harmonics and pre-emptively made parameter changes that reduced the stress on the equipment and process and prevented any interruptions. The architecture exhibited smooth scalability of building-sized microgrids up to utility-scale systems without any hardware changes. To ensure the system was integrated into existing SCADA, protection relay, and energy meter systems, a protocol gateway middleware was used, enabling a phased approach with no need for complete replacement of the existing system. Cybersecurity stress tests ensured resilience to data injection attacks and denial-of-service situations, where anomaly detection, using ML, detects threats in real time. Based on the above findings, it can be concluded that the software-defined architectures based on artificial intelligence are a revolutionary change in the direction of power quality management. The proposed system leverages deep learning, edge intelligence, and software-defined control, which eliminate the limitations of reactive, hardware-heavy monitoring and deliver the adaptive, proactive, and interoperable power grid capabilities needed. Future directions will be federated learning over distributed installations to jointly improve models, integrating digital twins into predictive maintenance, and investigating ways to use quantum computing to optimize power flows.

On the manufacturing shop floor, this real-time PQ monitoring and control capability directly translates to operational resiliency. Early detection of sags, harmonics, and transients helps prevent nuisance trips and premature wear to sensitive automation equipment such as variable-frequency drives, PLCs, robotic controllers, and CNC machines, ensuring continuous production and tighter process tolerances to avoid unplanned downtime and mid-process voltage events. Recent examples of using a cyber-physical digital-twin for real-time fault protection and intelligent energy monitoring using smart meters, both presented in, illustrate how such an edge-synchronized digital replica can be integrated directly into a plant's electrical infrastructure; meanwhile, generative-AI-based approaches to predictive fault management highlight the potential for disturbances to be predicted and mitigated before they become disruptive to production. The software-defined, edge-based architecture allows for incremental deployment with existing SCADA systems and energy-management solutions, seamless integration with digital twin and predictive-maintenance initiatives, and robust scaling capabilities from single production lines to plant-wide and multi-site deployments with minimal hardware upfront costs, making it ideal for Industry 4.0 and smart manufacturing applications.

Keywords: Artificial Intelligence, Machine Learning, Power Quality Monitoring, Software Defined Systems, Edge Computing, Deep Learning, Intelligent Control.

Abstract ID: 2898

Safety Constrained Energy–Latency Optimization for LoRaWAN in Remote Petroleum Field Monitoring

Amgad Hassan*, Kaustubh Singh

Electronics & Communication, Delhi Technological University, Bawana RD, New Delhi, India
**amgad_hassa_ec22b1_36@dtu.ac.in*

Abstract

Safety, environmental protection and equipment reliability are important aspects of continuous monitoring in remote petroleum extraction facilities. Such installations are usually in remote areas where wired communication equipment cannot be used or is too costly. LoRaWAN technology, with its capabilities of long-range communication, reduced power consumption, and low deployment cost, has proved to be a viable solution for these Industrial Internet of Things (IIoT) applications. Under the ETSI 1% duty-cycle regulation, though, there is a fundamental trade-off between communication energy usage and emergency alert latency. Lower transmissions have advantages in battery life, but can also result in a delay in reporting critical information, like pressure surges, gas leaks, or equipment problems. The currently available LoRaWAN parameter selection strategies can mostly be found to optimize either energy or communication performance and do not expressly account for the high latency requirements of safety-critical petroleum monitoring applications.

The monitoring problem in a petroleum field is formulated as a cross-layer optimization problem under practical communication constraints with the goal of minimizing the overall energy consumption of the network while adhering to a maximum allowed emergency reporting delay that is crucial for the safety of the petroleum field. The proposed framework takes into account the propagation properties, the medium access behavior, the duty cycle restrictions imposed by regulations, and the reliability of communication and combines them into an overall optimization model. Industrial wireless propagation is modeled as the log-normal shadowing model to account for the metallic infrastructure and large-scale fading often seen in petroleum facilities, and channel access is modeled as the unslotted Pure ALOHA protocol, which is used to estimate collision probability and the retransmission behavior. Unlike adaptive heuristic or learning-based optimization methods, a deterministic constrained search is used to find the combination of the most optimal spreading factor and transmission interval that meets the various constraints of the application, while minimizing energy consumption, including latency, coverage, and reliability requirements, as well as regulatory constraints.

A custom MATLAB simulation framework based on analytical models for the LoRaWAN physical layer and MAC layer was used for performance evaluation. In real industrial operating conditions, the spreading factor, the distance of communication, the time interval between each communication, and the number of nodes were varied systematically to generate the spreading data. Analytically computed packet airtime, communication latency, duty-cycle utilization, collision probability, and transmission energy for each feasible operating configuration. The generated dataset was then evaluated to find a deployment point that was feasible and to determine the operating point that meets the specified safety requirements. The results show that the proposed framework sets the deployment limits for the safety-critical LoRaWAN monitoring in a remote petroleum environment practically. The feasibility analysis shows that under the ETSI 1% duty-cycle regulation, the emergency alert latency of less than 180 s, at a communication distance of 1 km, cannot be realized without increasing the density of the gateways. The optimized configuration reduces the expected latency for an alert to be issued in an emergency by ~41 s over a conventional fixed SF12 deployment, and lowers the energy consumed per alert by over

80% at high node densities due to lower airtime and fewer retransmissions. It is also the energy-latency trade-off that shows that optimizing the operating point always occurs on the boundary of the latency constraint, which means that the optimal energy consumption is achieved by operating at the minimum required transmission rate to meet the required safety limit.

It also reveals that the optimized parameters in the cross-layer approach can consistently outperform the conventional static LoRaWAN parameters in terms of communication efficiency, emergency responsiveness, regulatory compliance, and communication reliability. The proposed framework offers a quantitative guideline for the deployment of the large-scale LoRaWAN-based petroleum monitoring system under strict safety and regulatory requirements and a practical parameter selection strategy. The results offer a practical basis for the implementation of energy-efficient industrial monitoring networks and a basis for future applications of the Industrial Digital Twin for remote petroleum infrastructure.

Abstract ID: 3934

Proximity-Aware Human-Robot Collaboration: A Dynamic Safety Framework for Industrial Environments

R. Gnana Prakash*, Govind. N. Sahu

Smart Manufacturing and Automation Laboratory, Department of Mechanical Engineering, Indian Institute of Technology Tirupathi, Andhra Pradesh – 517619

gnanaprakash@iittp.ac.in

Background and motivation

The widespread adoption of collaborative robots under the Industry 5.0 paradigm demands dynamic, real-time responses to human presence without reliance on physical safety fencing. ISO/TS 15066 codifies Speed and Separation Monitoring (SSM), always mandating a protective separation distance between every moving robot part and any detected human. Existing approaches represent two inadequate extremes: complete robot shutdown upon any workspace intrusion, which preserves safety at the cost of throughput, or expensive safety-rated sensor arrays that remain economically impractical for small and medium-sized enterprises (SMEs). Recent RGB-D and neural-network-based approaches bridge this gap; notable works have demonstrated an 11% task-efficiency gain over blanket-stop baselines using combined SSM and Power and Force Limiting (PFL) on a KUKA LBR, and achieved an end-to-end response latency of 0.806 s, within the ISO/TS 15066 temporal budget. This work presents a fully implemented and experimentally validated ROS 2 framework deployed on a real industrial cobot, in which the robot transitions to a paused state upon human entry into the defined safety envelope and resumes trajectory execution once the workspace is cleared, thereby addressing the reproducibility and embedded-throughput gaps in prior art by offloading inference to a host GPU to eliminate it as a latency bottleneck.

Experimental setup and methodology

The experimental platform for Proximity-Aware Human-Robot Collaboration integrates five key components, as illustrated in Figure 1. The JAKA ZU5 cobot exposes a ROS 2 Control-compatible hardware interface, continuously streaming joint encoder positions as joint state feedback. Depth perception is provided by the ZED 2i stereo camera, whose passive stereo design inherently avoids infrared interference with surrounding industrial sensors. Human detection is performed by a YOLOv8n model running on an NVIDIA RTX 5050 GPU leveraging Tensor Cores, achieving inference throughput that matches the camera frame rate. Consequently, the dominant latency contribution shifts to the ROS 2 communication layer, which remains below 100 ms throughout operation.

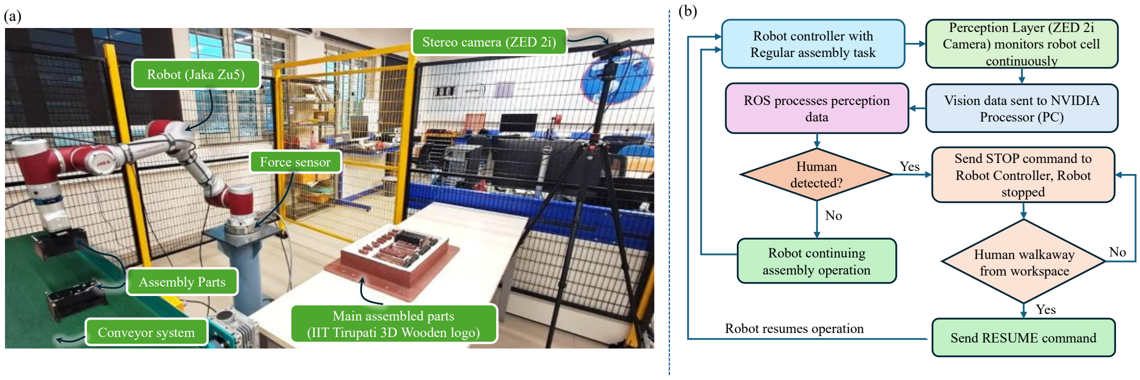

Figure 1. (a) Experimental setup, (b) flowchart showing proposed methodology for Proximity-Aware HRC.

The proposed methodology for Proximity-Aware HRC is structured into four sequential layers: the Perception Layer, Detection Pipeline Layer, Safety Node Layer, and Motion Execution Layer, each of which is described briefly below.

(i) Perception layer: The ZED 2i passive stereo RGB-D camera provides a 110°×70° field of view, a depth range of 0.2–20 m, and an integrated IMU for vibration compensation. Its passive stereo triangulation principle eliminates infrared projection interference with co-located industrial sensors. The ZED SDK publishes per-frame-aligned RGB and depth images directly into the ROS 2 ecosystem. **(ii) Detection pipeline layer:** YOLOv8n inference is executed on the host PC via CUDA on an NVIDIA RTX 5050 GPU (Ada Lovelace architecture), leveraging dedicated Tensor Cores for INT8/FP16 acceleration. Offloading inference to the host GPU eliminates the throughput bottleneck characteristic of embedded solutions, ensuring that detection latency does not dominate the overall system response time. A Boolean person-detected signal is published to the detection topic at the camera frame rate**. (iii) Safety node layer:** The safety node subscribes to the detection topic and transitions the system state accordingly. Upon receiving a True signal, it publishes a pause command to the /robot_pause topic and logs a diagnostic warning. Upon workspace clearance, a resume command is issued. The default system state at startup is RUNNING, with no pause condition imposed**. (iv) Motion execution layer:** A threaded ROS 2 executor processes /robot_pause callbacks in a dedicated background thread while the JAKA Zu5 trajectory executes in the main thread, preventing deadlocks and nested spin race conditions. A joint-state readiness guard ensures valid encoder data is available prior to the first motion goal, eliminating violent initialisation transients. A non-blocking sleep mechanism maintains callback responsiveness during inter-move delays. Upon a pause event, the robot holds its current joint configuration precisely; upon resumption, execution continues from the held state without re-homing.

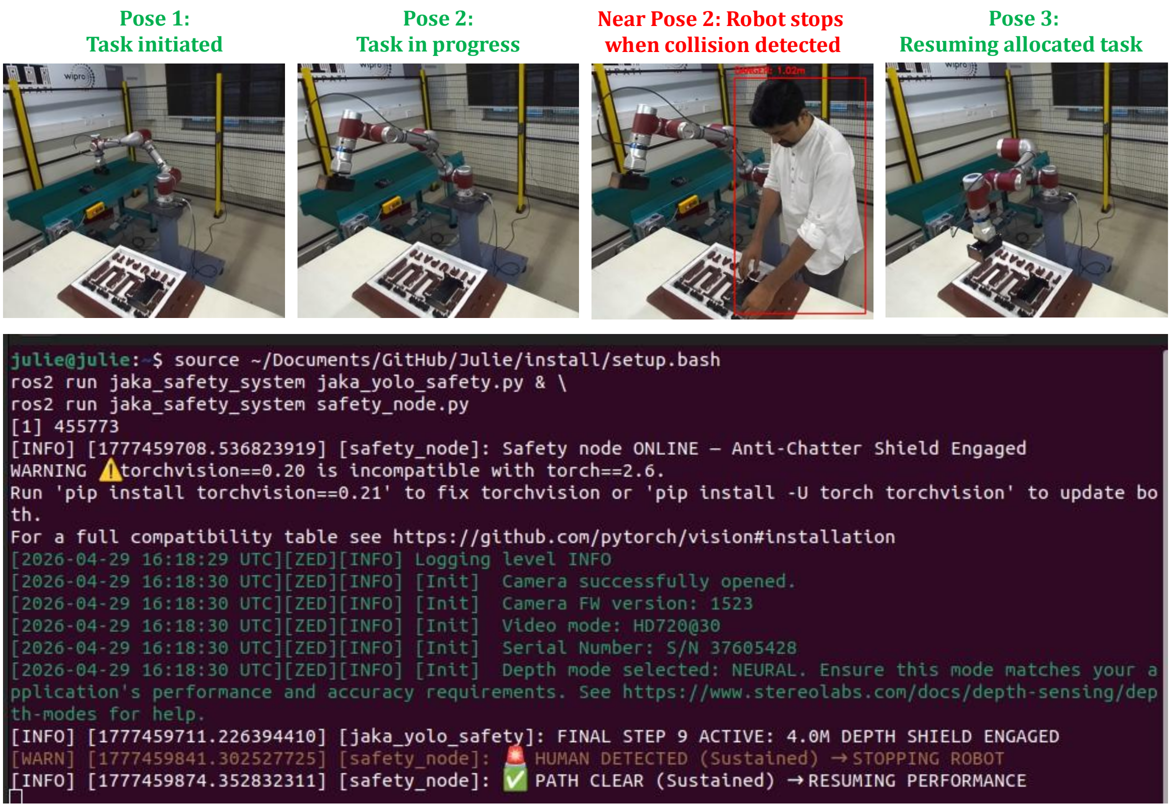

Figure 2. Results showing assembly progress and stoppage of the robot while human intervention.

4. Results & Discussion

The proposed proximity-aware stop-and-resume framework enables safe human-robot collaboration on the JAKA Zu5 manipulator through four sequential operational states. Pose 1 (Task Initiated) demonstrates the robot entering the pick-and-place trajectory with the safety node's RUNNING-default policy ensuring zero startup penalty. Pose 2 (Task in Progress) shows continuous trajectory tracking across all six joints, with the GPU-accelerated YOLOv8n perception pipeline maintaining real-time throughput without frame loss. Near Pose 2 (Robot Stops When Collision Detected) illustrates the critical safety mechanism: upon human intrusion, the perception pipeline immediately triggers a pause command, halting the robot deterministically, with all 40 stop events executed correctly with zero deadlocks or race conditions. Pose 3 (Resuming Allocated Task) demonstrates seamless task resumption without re-homing once the workspace is cleared, with all 30 resume transitions executing flawlessly. Three critical architectural mechanisms enabled this performance: MultiThreadedExecutor-based callback handling, wait_for_joint_state guard preventing initialization jerks, and non-blocking ros_sleep() preserving responsiveness. Terminal logs confirm consistent "HUMAN DETECTED → STOPPING ROBOT" and "PATH CLEAR → RESUMING PERFORMANCE" events, demonstrating deterministic, low-latency behavior that balances safety with operational efficiency; future work will extend toward ISO/TS 15066-compliant Sensitive Speed Monitoring for adaptive speed reduction based on depth-measured separation distance.

Abstract ID: 3086

Digital Twin for Efficient Robotic Scanning using RGB-D Sensing

Navodita Singh*, Antriksh Gangwar, Hiba K. T, Rushikesh Bhosale, Dr. Sundaram Muthu

Data Science and Engineering, Indian Institute of Science Education and Research, Bhopal, India
**navodita25@iiserb.ac.in*

Abstract

Robotics is widely developed and extensively used across industrial applications, enabling automation in manufacturing and inspection processes. Robotic arm-based scanning systems are commonly employed for tasks such as surface inspection and 3D reconstruction. These systems typically operate using predefined, structured scanning trajectories. Conventional robotic scanning systems often rely on predefined scanning trajectories with limited simulation-driven validation and optimization, making it difficult to balance surface coverage, reconstruction quality, and scanning efficiency. To overcome these limitations, a digital twin framework is developed to enable virtual generation, evaluation, and iterative optimization of candidate scanning trajectories prior to physical deployment. Performance metrics obtained from simulation are used as feedback to refine trajectory parameters and guide trajectory selection before execution.

In this work, we focus on a digital twin framework of a UR7e robotic work cell integrated with an Intel RealSense D435i sensor. The main objective is to develop a simulation-feedback-driven scanning framework that iteratively optimizes trajectory parameters to maximize surface coverage while reducing operational time before physical deployment. The proposed digital twin models the complete work cell, including surrounding walls, the mounting table, a 3D-printed camera mount, the RGB-D sensor, and the target object to be scanned, as shown in Figure 1(a), which illustrates the corresponding digital twin environment, while Figure 1(b) shows the real robotic setup. The UR7e robot provides a reach of 850 mm with ±0.03 mm repeatability, while the D435i offers RGB-D sensing with up to 1280×720 depth resolution and a range of approximately 0.2–10 m, enabling accurate and dense 3D data acquisition. This synchronized digital twin enables real-time visualization, simulation-driven validation, and performance analysis of scanning operations prior to execution.

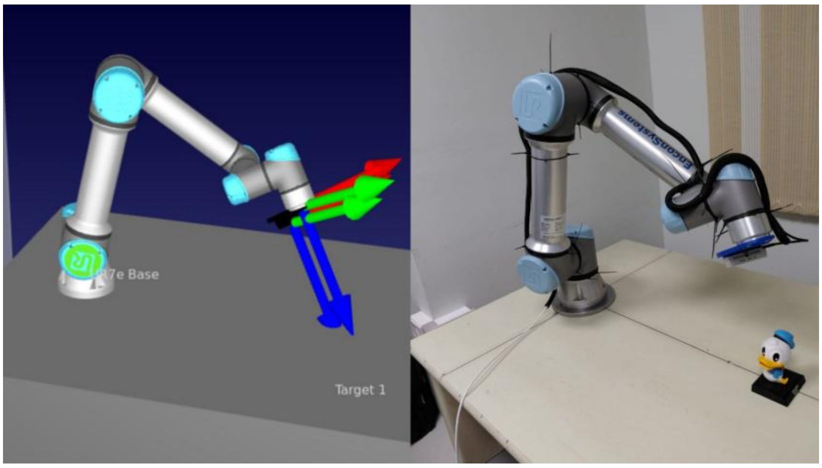

(a) (b)

Figure 1: (a) Digital twin environment, (b) corresponding real robotic setup

Candidate scanning trajectories $\{\tau_i\}$ (zigzag, spiral, circular, polar) are parameterized in Cartesian space and mapped to joint space via inverse kinematics: $q(t) = f^{-1}(x(t))$, subject to kinematic and workspace constraints. Feasibility is enforced through collision checking and reachability within the digital twin. Each trajectory is discretized and evaluated in simulation using metrics including surface coverage, reconstruction error, execution time, and trajectory length. These simulation results provide iterative feedback for refining trajectory parameters and automatically selecting the trajectory that offers the best trade-off between coverage, reconstruction quality, and execution efficiency prior to deployment.

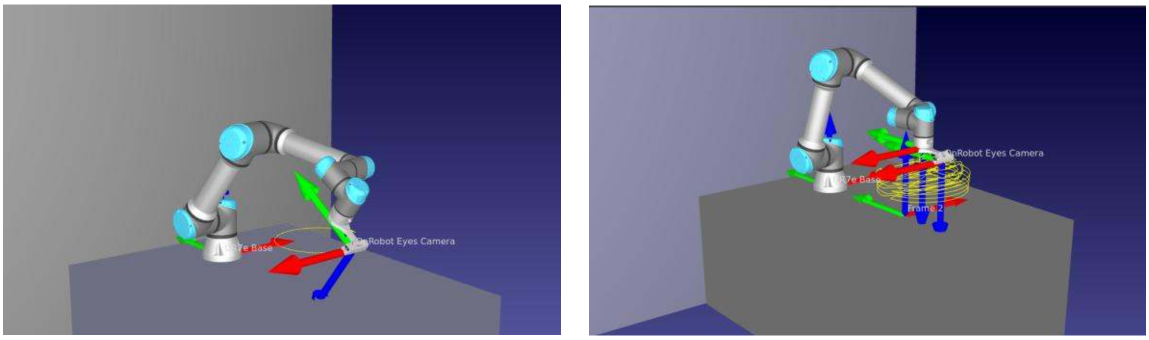

(a) (b)

Figure 2. (a). Circular scanning trajectory generated for the UR7e robotic arm, (b) Spiral scanning trajectory generated for the UR7e arm.

Circular scanning, illustrated in Figure 2(a), is described by the parametric path:

$$x(t) = xc + r\,cos(2\pi t), y(t) = yc + r\,sin(2\pi t), z(t) = z0, t \in [0,1] \tag{1}$$

where (x_c,y_c) is the object center, r is the scanning radius, and z_0 is the constant scanning height. Spiral scanning, illustrated in Figure 2(b), is defined as:

$$x(t) = xtarget + R\,cos(2\pi Nt), y(t) = ytarget + R\,sin(2\pi Nt), z(t) = zstart + Ht, t \in [0,1] \tag{2}$$

where (x_{target},y_{target}) represents the object center, R is the spiral radius, N is the number of revolutions, H is the total vertical scanning height, and z_{start} is the initial scanning height.

Trajectory selection is formulated as a constrained multi-objective optimization problem:

$$\tau^* = argmax\_\{\tau \in T\}\,[\alpha C(\tau) - \beta T(\tau) - \gamma L(\tau) - \delta \Phi(\tau)] \tag{3}$$

where C(τ) denotes normalized surface coverage, T(τ) execution time, L(τ) trajectory length, and Φ(τ) penalties for collision or kinematic infeasibility. The weights α, β, γ, and δ define the trade-off between coverage maximization and execution efficiency. Performance is evaluated using reconstruction error with respect to the reference CAD model, relative rotation error (RRE), and relative translation error (RTE).

Results and Discussion

A comparative assessment of zigzag, circular, spiral, and polar trajectories shows clear trade-offs between coverage and efficiency. The zigzag trajectory provides high coverage because of its systematic back-and-forth pattern, but it requires longer execution due to frequent directional changes. The circular trajectory offers smoother and faster motion, although its predominantly planar nature limits coverage. The polar and spiral trajectories provide a better balance by enabling gradual and continuous spatial exploration. Based on the iterative simulation-feedback process, the spiral trajectory is refined by optimizing its geometric parameters and identified as the preferred scanning strategy, providing an improved balance between surface coverage, motion smoothness, reconstruction quality, and computational efficiency.

Application

CAD-based Inspection: The reconstructed point cloud obtained from the optimized scanning trajectory is aligned with the reference CAD model to enable geometric inspection. Deviations between the scanned data and the CAD model are analyzed for defect detection and dimensional verification, supporting automated quality assessment in manufacturing workflows.

The 3D reconstruction is generated from RGB-D images using the Open3D pipeline, which integrates depth maps into a volumetric representation (e.g., TSDF – Truncated Signed Distance Function) and extracts a dense point cloud or mesh. Specifically, synchronized RGB and depth frames are first converted into RGB-D images, followed by camera pose estimation and frame integration into a global model using volumetric fusion or odometry-based registration. This process enables accurate and scalable 3D scene reconstruction suitable for inspection tasks.

The proposed framework demonstrates how simulation feedback within a digital twin can be used to iteratively optimize candidate scanning trajectories before execution, improving scanning performance while reducing reliance on manual trajectory selection. The approach enhances scanning performance through pre-deployment validation and offers a scalable solution for intelligent inspection in smart manufacturing environments.

Abstract ID: 3055

Visibility-Aware YOLOv11 for Robust Real-Time Underwater Object Detection in Autonomous Surveillance

Ayyappan S*, Ilavendhan A

Vellore Institute of Technology, School of Computer Science and Engineering, Chennai, India
**ayyappan.2024@vitstudent.ac.in*

Abstract

This study presents a YOLOv11-based framework augmented with a Dynamic Convolutional Attention Block (DCAB) for real-time underwater object detection in Autonomous Underwater Vehicles (AUVs) surveillance. The framework is evaluated using Precision, Recall, mAP@0.5 and mAP@0.5–0.95 across four typical underwater imaging conditions: clear, turbid, darkened and low-contrast. Trained on the UPRC2020 dataset (17,345 images, four marine target categories), the proposed detector attains 84.8% mAP@0.5 overall and retains 75.8 – 83.4% mAP@0.5 across all four degradation regimes, with a worst-case relative drop of only 9.1%. On a single NVIDIA RTX 3050 GPU, the model achieves 98.3 FPS (10.18 ms per frame), approximately 3.3 × the conventional 30 FPS real-time threshold required for live AUV deployment. Precision is essentially condition-invariant (80.2–81.6%), indicating that the DCAB block preserves discriminative features against scattering and illumination loss while suppressing false positives in degraded imagery. The proposed detector is therefore appropriate for real-time marine surveillance on resource-constrained AUV platforms.

Keywords: Underwater Object Detection, YOLOv11, Dynamic Convolutional Attention, AUVs, Real-Time Vision, Visibility-Aware Detection, Marine Surveillance.

Introduction

Marine robots, environmental monitoring, and offshore inspection systems all rely on robust underwater object detection. Detection accuracy is sharply reduced in underwater imagery due to wavelength-dependent attenuation, scattering, restricted illumination, and dynamic-range compression. Although YOLO-family architectures have substantially improved detection speed and robustness on standard benchmarks, there is still a lack of comprehensive evaluation across varied underwater degradation conditions. To address this gap, this work introduces a Dynamic Convolutional Attention Block (DCAB) inserted at the detection feature scales of YOLOv11 and rigorously evaluates the resulting detector across four physically motivated visibility regimes, clear, turbid, darkened, and low-contrast, for the deployment of AUVs in real-time.

Methodology

The proposed workflow comprises image pre-processing, multi-scale feature extraction, DCAB-based refinement, and visibility-aware evaluation. The dataset is UPRC2020 (17,345 ROV-captured images, four marine targets: echinus, holothurian, scallop, starfish), resized to 640 × 640 and split 85% / 6% / 10% into 14,706 training, 980 validation, and 1,659 test images. YOLOv11 provides multi-scale features through Cross Stage Partial (CSP) connections; a Dynamic Convolutional Attention Block (DCAB) is inserted at the P3, P4, and P5 detection scales for spatial and channel-wise refinement. DCAB combines an input-conditioned dynamic 3 × 3 convolution that blends K = 4 candidate kernels via SoftMax attention, CBAM-style channel attention, and 7 × 7 spatial-attention maps; a residual link keeps it channel-preserving and adds 3.13 M parameters (5.72 M total).

$$\mathrm{F} = \mathrm{CSP}(\mathrm{I}) \tag{1}$$

$$\mathrm{F}' = \mathrm{DCAB}(\mathrm{F}) \quad (2)$$

$$\mathcal{L} = \mathcal{L}_{\mathrm{cls}} + \mathcal{L}_{\mathrm{obj}} + \mathcal{L}_{\mathrm{bbox}} \quad (3)$$

Training uses COCO-pretrained YOLOv11n weights, AdamW (lr 1×10^{-3}), batch 8, and AMP with a 150-epoch budget and patience-30 early stopping on validation mAP@0.5. Augmentation comprises Mosaic, flip, HSV jitter, and underwater noise injection ($p = 0.3$), Jaffe-style attenuation, multiplicative darkening, and dynamic-range compression. All training and benchmarking use a single NVIDIA RTX 3050 (8 GB) GPU.

Result and Discussion:

Table 1 reports detection performance across the four visibility regimes. The detector attains 83.4% mAP@0.5 on clear imagery and retains 75.8 – 79.5% mAP@0.5 across the three degraded regimes, with a worst-case relative drop of only 9.1%. Precision is essentially condition-invariant (80.2 – 81.6%), with most degradation absorbed by Recall; the model rarely hallucinates targets even when visibility is severely reduced. On the held-out validation split, the model attains 84.8% mAP@0.5 and 49.7% mAP@0.5–0.95, with per-class mAP@0.5 ranging from 92.2% (echinus) to 73.5% (holothurian). The detector runs at 98.3 FPS (10.18 ms/frame) at 640 × 640 on an RTX 3050, approximately 3.3 × the conventional 30 FPS real-time threshold, confirming suitability for live AUV deployment.

Table 1. Per-condition detection performance of Visibility-Aware YOLOv11 + DCAB on the UPRC2020 test split.

Condition	Prec. (%)	Rec. (%)	mAP@0.5 (%)	mAP@0.5–0.95 (%)
Clear	81.6	77.1	83.4	48.8
Turbid	80.4	71.9	79.5	45.6
Darkened	80.2	67.5	75.8	42.6
Low-contrast	80.2	69.1	77.3	44.2
Overall (val)	81.6	79.3	84.8	49.7

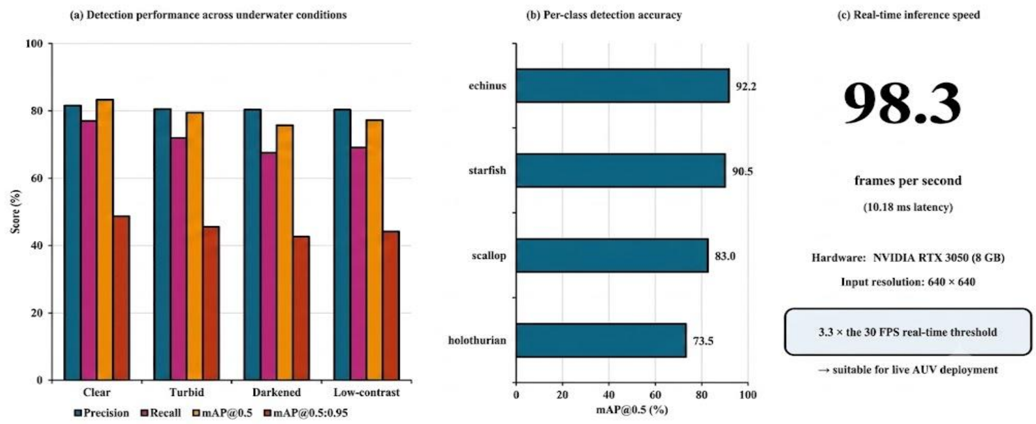

Figure 1. Visibility-Aware YOLOv11 + DCAB — performance summary. (a) Per-condition Precision / Recall / mAP. (b) Per-class mAP@0.5. (c) Real-time inference speed.

Conclusion:

Visibility-Aware YOLOv11 + DCAB is a capable real-time underwater object detector, achieving 84.8% mAP@0.5 overall and 98.3 FPS on an NVIDIA RTX 3050, with 75.8 – 83.4% mAP@0.5 retained across four physically motivated visibility regimes. The condition-invariant Precision (80.2 – 81.6%) confirms that DCAB-based attention preserves discriminative features under visibility degradation. Future work targets baseline ablation against YOLOv5/v8/vanilla-YOLOv11.

Abstract ID: 3048

AssemblyGuard: A Smart Edge AI Vision Prototype for Single-Class Bolt Count Verification in Manual Assembly Lines

Ramkumar Kanthasamy[1]*, Sujeet Subramani[1], Devanandh Yuvaraj[1], AkashRai Selvaraj[1], Ganesh Kumar Nithyanandam[1], Suriya Sundaramoorthy[2]

[1]*Department of Mechanical Engg., PSG College of Technology, Coimbatore, India*

[2]*Department of Computer Science & Engg., PSG College of Technology, Coimbatore, India*

*ram24kj@gmail.com

Abstract

This paper presents AssemblyGuard, a student-team-developed edge AI vision prototype for assembly-state verification in manual assembly stations where full robotic automation is difficult or uneconomical, such as large or low-volume assemblies in wind-turbine gearboxes, excavators and heavy equipment. The current work is deliberately restricted to a single-product, single-class bolt-count verification problem using a rectangular fixture with six predefined M12 bolt positions. A YOLOv8n detector was trained on 370 annotated bolt-assembly images and exported to ONNX for deployment on an NVIDIA Jetson Nano using ONNX Runtime CPU inference. The pipeline combines OpenCV image acquisition, 640 x 640 letterbox preprocessing, ONNX inference, Non-Maximum Suppression, bolt counting, 15-frame temporal smoothing, and persistence-based anomaly confirmation. Cloud services are not used for AI inference; Telegram Bot API and an AWS EC2 Node.js dashboard are used only for alerting and monitoring. The detector achieved 98.52% overall accuracy, 99.03% precision, 99.48% recall, and 99.30% mAP@0.5 after 150 epochs. Operational validation across 567 controlled trials produced 521 correct predictions, corresponding to 91.88% operational accuracy. All observed errors were single-bolt undercounts caused mainly by reflection or partial occlusion. A 5-hour stability test verified sustained headless operation, Telegram command handling, anomaly alerts and dashboard communication. AssemblyGuard is therefore positioned as a low-cost laboratory prototype for monitorable edge-based bolt-count verification in manual assembly; larger datasets, latency measurement, GPU/TensorRT acceleration, and multi-class/parameterized configurations remain future work.

Keywords: Edge AI, YOLOv8, ONNX Runtime, Jetson Nano, Bolt Count Verification, Manual Assembly Line Monitoring, Computer Vision, Temporal Smoothing, Assembly Anomaly Detection.

Introduction

Manual assembly remains necessary in large, customized, or low-volume manufacturing because component size, product variation, and fixturing complexity often make full robotic automation impractical. Examples include wind-turbine gearbox assembly, excavator assembly, and heavy-equipment stations where operators must follow fastening sequences while supervisors require reliable assembly-state visibility. Fastener omissions can affect safety and reliability. This study addresses a constrained verification task: checking whether a six-bolt fixture reaches the expected bolt count during manual assembly. The aim is not to claim a universal inspection system, but to demonstrate a low-cost edge-AI monitoring pipeline with local inference, temporal decision stability and remote alerting.

Literature Survey

Deep-learning-based vision inspection has been applied to assembly defects. Lin et al. studied ratchet-wrench assembly inspection covering missing, misplaced, foreign and extra parts. Shen et al. proposed a YOLOv5-based PCBA defect detector and dataset for assembled electronic components. Assembly-line anomalies have also been studied using historical tabular inspection data and root-cause analysis. Edge computing is relevant because it reduces dependence on remote cloud processing and can support lower-latency industrial decision-making. Knaak et al. showed that temporal information and low-power embedded implementation can improve defect monitoring. Compared with these works, AssemblyGuard has a narrower scope but focuses on a Jetson-based, camera-driven pipeline for single-class bolt-count verification with temporal smoothing, persistent anomaly confirmation, and remote monitoring.

Methodology

The proposed pipeline contains nine stages. Frames are acquired from a USB webcam using OpenCV and resized to 640 x 640 pixels with letterbox padding. Pixel values are normalized and rearranged into CHW format for ONNX Runtime. The YOLOv8n ONNX model is executed with the CPU Execution Provider at a confidence threshold of 0.35. Outputs are refined using Non-Maximum Suppression with IoU 0.40, and surviving detections are counted as bolt instances. A 15-frame sliding window smooths the count using statistical mode, with mean-based fallback when required. The smoothed count is compared with the reference value of six: count = 6 is OK, count < 6 is MISSING, and count > 6 is EXTRA. Alerts are generated only when the anomaly persists across the confirmation window and are then sent through Telegram and the AWS dashboard.

Experimental Setup

The hardware setup consists of an NVIDIA Jetson Nano, a USB webcam at 1280 x 720 pixels, a power supply, and a rectangular metal fixture with six M12 bolt positions. The approximate deployment cost is INR 17,300. The software stack includes Python, OpenCV, YOLOv8n exported to ONNX, ONNX Runtime, Telegram Bot API, and a Node.js Express dashboard on AWS EC2. The model was trained for 150 epochs on 370 annotated images collected under varied bolt configurations, lighting, and viewing orientations. Operational validation consisted of 567 controlled trials across bolt counts from 0 to 6, including fully assembled, partially assembled, and over-assembled states. The fixture is a laboratory representation and is not yet a wind-turbine gearbox or excavator validation.

Results and Discussion

The detector showed stable convergence over 150 epochs. Box loss decreased from 1.55 to 1.10, classification loss from 0.85 to 0.42, and distribution focal loss from 1.11 to 0.92. Precision and recall stabilised near 0.990 by epoch 50. The final model achieved the metrics shown in Table 1.

Table 1. Model and operational validation summary.

Metric	Value
Overall Accuracy	98.52%
Precision	99.03%
Recall	99.48%
mAP@0.5	99.30%
F1-score	0.992
Operational Trials	567
Correct Predictions	521
Operational Accuracy	91.88%
Continuous Stability Test	5 hours
Maximum Error	1 bolt

During 567 controlled trials, the system correctly predicted the bolt count in 521 cases and produced 46 incorrect predictions. The errors were mainly due to metallic reflections or partial occlusion near fixture corners. No large-count error or unsafe false OK condition was observed in the recorded validation set. Telegram logs confirmed online status reporting, cloud IP update, missing-bolt alert delivery, /status, /logs, and clean shutdown. In the 5-hour headless stability test, the system continued monitoring, generated repeated anomaly alerts, handled commands, and communicated with the AWS dashboard/heartbeat interface. This is reported as an endurance/stability test, not as a substitute for statistically balanced 1000-cycle validation.

Conclusion

AssemblyGuard demonstrates the feasibility of using YOLOv8-based object detection, ONNX Runtime edge inference, temporal smoothing, and Telegram/AWS monitoring for single-class bolt-count verification on a six-bolt manual assembly fixture. The system achieved high model-level detection metrics and 91.88% operational accuracy across 567 controlled trials. The present scope is intentionally limited to one product, one component class, and one fixture geometry. It should be read as laboratory evidence for a monitoring pipeline before translation to large manual assembly stations such as wind-turbine gearboxes, excavators, and other heavy equipment. Future work will expand class/state balance, add synthetic/augmented data, measure inference time and end-to-end latency under load, explore GPU/TensorRT acceleration and extend the method to multi-class, positional and parameterized assemblies.

Abstract ID: 3165

Real-Time Vision-Based Casting Defect Detection

Mohamed Farook K*, R· Gnana Prakash, Dr Govind N· Sahu

Smart Manufacturing and Automation Laboratory, Department of Mechanical Engineering, Indian Institute of Technology Tirupati, Andhra Pradesh – 517619

**ermdfarook@gmail.com*

Background and motivation

In the foundry industry, quality inspection of a cast part plays a critical role in ensuring product reliability, process stability, and customer satisfaction. Traditionally, visual inspection has been performed by human operators due to its flexibility and ability to interpret complex surface variations. However, this approach is inherently subjective and susceptible to variability arising from factors such as operator skill level, fatigue, and environmental conditions. Visual inspections carried out during night shifts often exhibit reduced accuracy due to decreased alertness and inconsistent judgment, leading to higher defect escape rates and increased customer rejections. Moreover, the lack of standardized acceptance criteria further amplifies inconsistency in decision-making across different inspectors. These challenges not only affect product quality but also contribute to rework, scrap generation, and increased operational costs. Therefore, there is a growing need for reliable, consistent, and scalable inspection systems that can complement human expertise while minimizing variability in industrial quality control processes. While addressing this gap, we propose real-time vision-based casting defect detection.

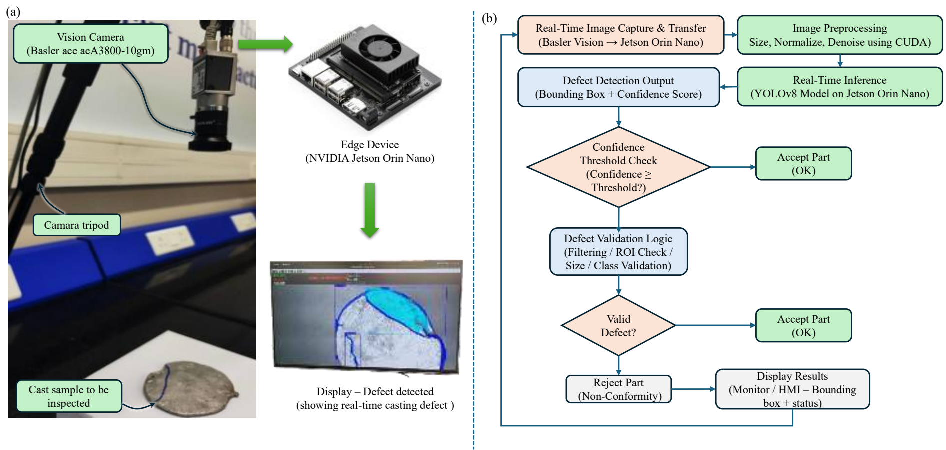

Figure 1. (a) Experimental setup, (b) Proposed methodology for real-time vision-based casting defect detection system.

Experimental setup

The hardware system was assembled sequentially as shown in Figure 1(a). First, the Jetson Orin Nano Developer Kit was initialised by connecting it to a monitor and flashing JetPack 6.x to enable CUDA and AI support. Next, the Basler acA3800-10gm camera was interfaced with the Jetson via a GigE Ethernet cable using PoE for both power and data transmission. The required drivers, including the Basler pylon SDK and pylon library, were then installed, and camera parameters such as exposure, gain, and resolution were configured using the pylon Viewer. Finally, the network interface was optimised by assigning a static IP address and enabling Jumbo Frames to ensure efficient high-resolution image transfer at the desired frame rate.

System Architecture and Implementation

The casting defect detection system is organised in a modular and structured manner to ensure scalability, maintainability, and ease of industrial deployment, as illustrated in the workflow diagram presented in Figure 1(b). The architecture is decomposed into functionally distinct components, each addressing a specific phase of the development pipeline.

A centralized configuration module governs all key system parameters, including dataset paths, model settings, and training hyperparameters, ensuring consistency across all stages of the pipeline. Dataset acquisition and initial exploration are managed through a dedicated data collection component, which facilitates efficient retrieval and preliminary analysis of defect class distributions. Preprocessing and augmentation constitute a critical stage in which raw images undergo cleaning, annotation, and geometric and photometric transformations to enhance model generalization and robustness against real-world variation. The model training component leverages the YOLOv8 architecture — a state-of-the-art single-stage object detector — applying transfer learning and fine-tuning strategies to achieve optimal defect localization performance.

Quantitative evaluation is performed by a dedicated assessment module that computes standard detection metrics, including precision, recall, F1-score, and mean Average Precision (mAP), enabling rigorous benchmarking against baseline methods. For real-time industrial deployment, a camera-integrated inference module interfaces the trained model with the Basler industrial camera via the Pylon SDK, enabling live defect detection directly on the production line with minimal latency.

Reproducibility is ensured through explicit dependency management, while comprehensive documentation — covering environment setup, configuration guidelines, and execution procedures — supports adoption and extension by future practitioners. This structured, modular design facilitates efficient development, systematic testing, and seamless deployment of the automated casting defect detection system within industrial manufacturing environments.

Results and discussion

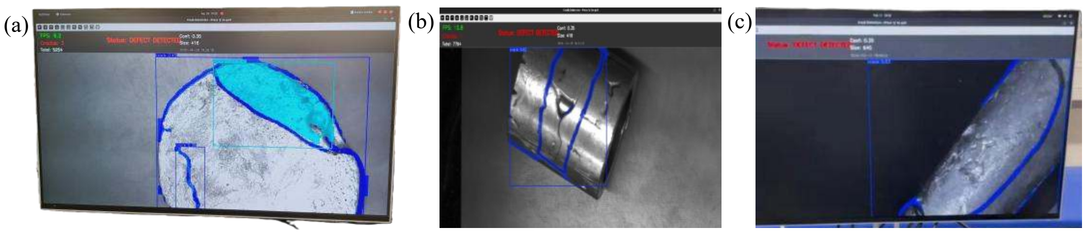

Figure 2. Detected casting defects (a) crack, (b) cold shut, and (c) shrinkage.

The proposed vision-based defect detection system demonstrates strong performance under both test and real-time industrial conditions, achieving a precision of 0.94, recall of 0.91, mAP@0.5 of 0.93, and F1-score of 0.92. It operates in real time on the Jetson Orin Nano Super at 30–60 fps with ~35 ms latency, meeting the requirements of a 10 fps pipeline, while CUDA acceleration ensures efficient edge processing. The system accurately detects dimensional defects with consistency and integrates human-informed acceptance criteria to align with practical industrial standards, reducing false rejections and improving adaptability. Compared to manual inspection, it minimizes human error, reduces inspection time, and retains domain expertise, while remaining scalable through YOLOv8 and pylon SDK integration. However, performance is limited by the 10 fps Basler camera and depends on dataset quality, indicating scope for higher-speed imaging and dataset expansion. Overall, the approach provides a cost-effective, scalable, and industry-ready solution for automated defect detection. Moreover, some tested output images are shown in Figure 2.

Abstract ID: 2703

A novel framework towards reshaping the CNN topology

Prashant Priyadarshi*, Dr. Pooja Sharma, Avinash Kumar,

Department of CSE & IT, BIT Sindri, Dhanbad, Jharkhand, India
**prashant6850@gmail.com*

Abstract

Radiomics has been established as a technique that can facilitate the extraction of quantitative information from imaging studies, in order to describe objectively the phenotypic characteristics of tumors that cannot be easily assessed using conventional methods of analysis. On the other hand, Convolutional Neural Network (CNN) models have attained excellent results in medical imaging studies, but interpretation of the features learned by such networks usually remains challenging. The establishment of the quantitative connection between deep learning and radiomic features thus becomes crucial in order to increase the reliability of AI-driven clinical decision support systems. In this study, a model based on radiomics is proposed to study the connection between deep learning and geometric radiomic features of computed tomography (CT) images of patients with Non-Small Cell Lung Cancer (NSCLC).

For experimental assessment, the NSCLC Radiomics dataset was utilized, which is publicly accessible and consists of CT images along with respective tumor segmentation masks. Preprocessing of the CT images involved selection of five representative slices per patient image set and normalization of their intensities and resizing them to 224 × 224 pixels. The descriptive statistics of radiomic features revealed great variation in the morphology of tumors analyzed in the patients (see Table 1). Five ImageNet pre-trained CNN architectures, namely ResNet18, VGG16, DenseNet121, EfficientNet-B0, and MobileNetV2, were used as feature extractors. Four different architectural modifications, Fully Connected (FC) Projection, Multi-Layer Feature Extraction, Spatial Pyramid Pooling (SPP), and Channel Attention, are used for improving the interpretability of features. Three activation statistics (mean, standard deviation (SD), and maximum (Max)) were calculated for CNN feature maps, and try to find the relation with radiomic descriptors using Spearman correlation and a regression equation.

As can be observed from the experimental results in Tables 2(a) and 2(b), the baseline CNNs' regression skills are enhanced with the proposed architecture changes. The DenseNet121 and FC Projection performed the best overall model, $R^2 = 0.273$, while SPP had a positive effect on ResNet18, EfficientNet-B0, and MobileNetV2. This indicates that different CNN architectures must be correspondingly processed to improve the features. Furthermore, correlation analysis between the CNN activation statistics and the volume of the bounding box also revealed that the volume is highly correlated, and far more than any other geometric parameters.

Table 1. Dataset statistics.

Feature	Mean	SD	Max
Mask Volume (Voxels)	1,593,867	518,395	2,780,483
Tumor Volume Ratio	0.01034	0.00535	0.05402
Bounding Box Volume (Voxels)	37,661,667	13,536,048	89,351,370

Regression Model

CNN activation statistics were modeled using this linear regression.

$$A_{pred} = \beta_0 + \beta_1 V_m + \beta_2 F_r + \beta_3 V_b \quad (1)$$

where A_{pred} refers to the predicted CNN activation. V_m, F_r, and V_b refers to mask volume, tumor volume ratio, and bounding box volume, respectively, and β_0, β_1, β_2, and β_3 are the regression coefficients estimated using linear regression

The R^2 values one can achieve in the present study are moderate, but shown to be consistent with the aim of the study, as only three geometric radiomic descriptors were utilized for regression, whereas CNN activations further encode texture, intensity, and high-level semantic information. By analyzing the proposed architectural changes, we showed that the conservation of both spatial context and learning compact feature representations brings substantial improvements in interpretability for CNN architectures without altering the backbone CNNs.

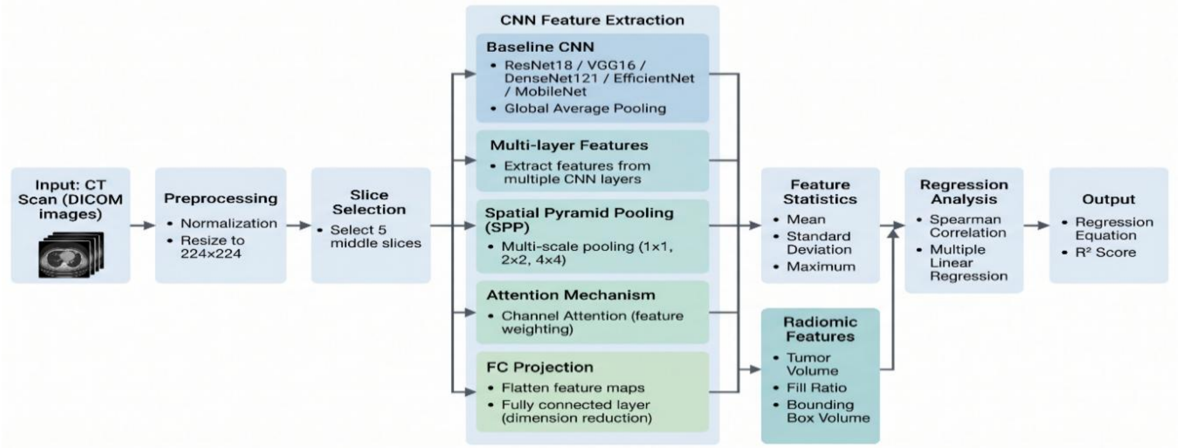

Figure 1. The proposed framework that integrates baseline CNN models with architectural changes such as FC projection, multi-layer extraction, SPP, and attention.

In other words, the overall framework provides quantitative associations between the radiomic descriptors and CNN feature representations, which helps to explain deep learning models used for medical image analysis. They show the power of SPP to improve spatial representation at multi-scale and the effectiveness of FC Projection in DenseNet121. The proposed methodology could be applied in the context of explainable computer-aided diagnosis, radiomics-guided clinical decision support, prediction of treatment response, and future imaging-based Digital Twin systems for personalized healthcare.

Table 2. (a) CNN Architecture Summary

Model	**Architecture**	**Best**	**R^2**
ResNet18	Residual Learning	SPP	0.157
VGG16	Sequential Convolution	ML	0.057
DenseNet121	Dense Connectivity	FC	0.273
EfficientNet-B0	Compound Scaling	SPP	0.057
MobileNetV2	Inverted Residual	SPP	0.089

Table 2. (b) Comparison of R^2 values across baseline and modified CNN architectures

Model	Baseline	Modified Architectures			
		FC	Multilayer	SPP	Attention
ResNet18	0.045	0.045	0.137	0.157	0.144
VGG16	0.018	0.026	0.057	0.01	0.018
DenseNet121	0.057	0.273	0.058	0.021	0.038
EfficientNet-B0	0.041	0.021	0.019	0.057	0.042
MobileNetV2	0.066	0.056	0.042	0.089	0.065

Dataset Note: The NSCLC Radiomics dataset currently available is public and consists of CT images with masks of the tumor segmentations. Thus, the study in this paper concentrates on spatial radiomic features. We are trying to get temporal imaging data to extend this framework.

Abstract ID: 4812

Gesture-Based AI Safety Monitoring and Fault Detection for Machine Operators Using Edge AI and Human Pose Estimation

Vignesh M[1], Saravana Perumaal S[2]*

[1]*PG Scholar, Dept. of Mechanical Engineering, NITTTR, Chennai*
[2]*Associate Professor, Dept. of Mechanical Engineering, NITTR, Chennai*
**vigneshmanoharan32@gmail.com*

Abstract

Industrial safety remains a critical concern in manufacturing environments, particularly in machining operations where unsafe operator behaviour can result in accidents, equipment damage, and production downtime. Conventional safety systems primarily focus on machine condition monitoring and physical safeguarding mechanisms, with limited emphasis on real-time assessment of operator actions. This work presents a Gesture-Based AI Safety Monitoring and Fault Detection System for Machine operators using the Edge AI platform. The proposed framework employs real-time human pose estimation to monitor operator posture and gestures and evaluates them against a predefined Allowed Motion Envelope (AME). Safety violations such as hand intrusion into restricted zones, prolonged head-down posture, excessive body leaning, and body bending are automatically detected. The system computes Standard Operating Procedure (SOP) compliance scores, ergonomics scores, and overall operator safety scores while generating automated safety reports. Experimental validation conducted in a laboratory environment demonstrated successful real-time monitoring with a safety performance of 81.05%. The proposed solution offers a low-cost, scalable, and privacy-preserving approach for enhancing workplace safety through edge AI-enabled operator monitoring.

Introduction

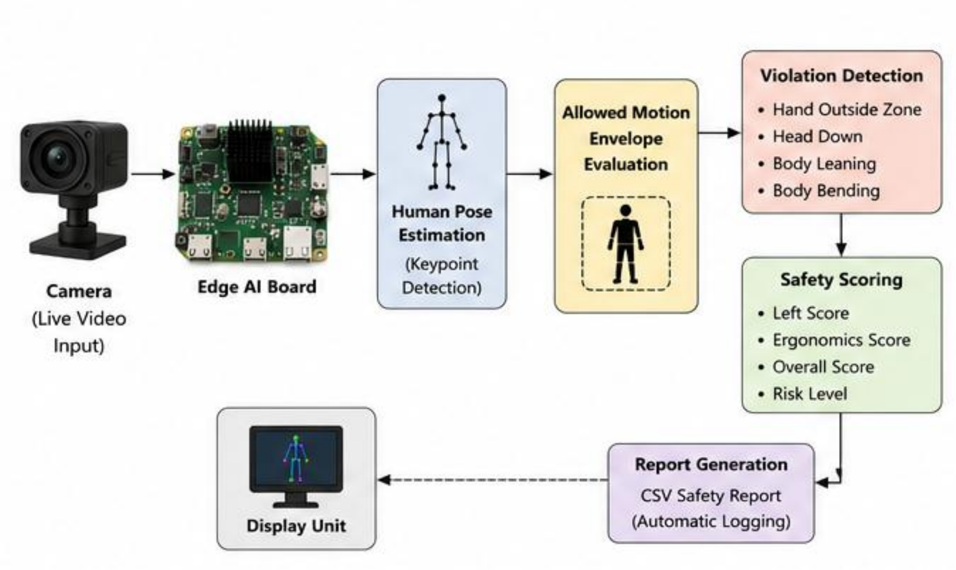

Figure 1. System architecture of the proposed safety monitoring framework.

The increasing adoption of Industry 4.0 technologies has significantly enhanced manufacturing productivity and automation. However, operator-related accidents continue to remain a major challenge in machining environments. Unsafe practices such as entering hazardous machine zones, improper posture, distraction, and non-compliance with safety procedures contribute to workplace injuries and operational risks.

Figure 2. Human pose estimation and key point detection.

Recent advancements in Artificial Intelligence (AI), Computer Vision, and Edge Computing have enabled real-time human activity monitoring without requiring wearable devices. Human pose estimation techniques can identify and track body keypoints, allowing continuous evaluation of operator movements and posture. Despite these developments, most existing systems focus on activity recognition or surveillance applications and do not provide integrated safety assessment tailored for industrial environments.

This study proposes an edge AI-based framework for real-time operator safety monitoring using the edge AI platform. The system evaluates operator behaviour against predefined safety rules and automatically generates quantitative safety reports for risk assessment and decision-making.

Methodology

The proposed system consists of six primary modules:

1. Video Acquisition Module
2. Human Pose Estimation Module
3. Allowed Motion Envelope Evaluation Module
4. SOP Violation Detection Module
5. Ergonomics Assessment Module
6. Safety Report Generation Module

Figure 3. System architecture of the proposed safety monitoring framework.

A camera continuously captures operator movements near the machine. The edge AI processor performs real-time human pose estimation to extract skeletal keypoints including head, shoulders, elbows, wrists, hips, knees, and ankles. The extracted keypoints are evaluated against an **Allowed Motion Envelope**

(AME) representing safe operating boundaries around the machine. Deviations from these boundaries are classified as safety violations.

SOP Violations

- Hand Outside Safe Zone
- Head Down Violation

Ergonomics Violations

- Excessive Body Leaning
- Excessive Body Bending

Safety scores are computed using the frequency and severity of detected violations.

$$\text{Overall Safety Score} = \frac{\text{SOP Score} + \text{Ergonomics Score}}{2} \quad (1)$$

Based on the calculated score, operators are categorized as:

All observations are automatically stored in CSV-based safety reports for future analysis.

Results and Discussion

Experimental validation was conducted in a Machine laboratory environment using the edge AI platform. Multiple operator activities were performed to simulate both safe and unsafe working conditions. The experimental results demonstrate that the proposed framework can effectively identify unsafe operator behaviours in real time. The system successfully detected violations related to SOP compliance and ergonomics while operating entirely on an edge AI device without cloud dependency. The low computational requirements and real-time performance make the framework suitable for deployment in small- and medium-scale manufacturing industries.

Conclusion

This paper presented a Gesture-Based AI Safety Monitoring and Fault Detection System for Machine operators using the Edge AI platform. By combining human pose estimation, safety-rule evaluation, and automated reporting, the proposed framework enables real-time monitoring of operator behaviour and workplace safety. Experimental results validated the effectiveness of the system in detecting unsafe actions and generating quantitative safety assessments.

Abstract ID: 3135

Stability-Based Degradation Trend Modeling for RUL Prediction of Broken Rotor Bar Faults

Lavanya Aruldass[1]*, Suvendu Mohanty[2], Revathi S[1], Piyush Shakya[2]

[1]*NIT Puducherry, Karaikal, 609609*

[2]*IIT Madras, Chennai, 600036*

**lavanya4241@gmail.com*

Abstract

Induction Motor (IM) plays a vital role, accounting for 70-80% of total industrial electrical energy consumption and serving as a prime mover for critical machinery. These motors are extensively employed in industrial applications due to their advantages, including cost-effectiveness, rugged construction, self-starting capability, and low maintenance requirements. Despite these advantages, these motors are susceptible to various types of failure due to thermal stress, mechanical cyclic fatigue, electromagnetic loading, and environmental contamination; each involves a distinct physical mechanism requiring dedicated treatment. This work restricts its scope to the Broken Rotor Bar (BRB) fault, one of the most insidious failure modes, accounting for 5-10% of reported industrial motor failures; the other modes above are noted only for motivation and are left for dedicated future work. The presence of BRB disrupts the symmetrical nature of the rotor circuit, resulting in asymmetrical magnetic flux, speed oscillations, torque pulsation, and localized thermal hotspots in adjacent rotor bars. Neglecting these faults can lead to cascading failures and render the motor inoperable. Due to the stochastic nature of BRB initiation and gradual progression, there exists a need for condition monitoring and data-driven prognostic approaches.

Recent research focuses on fault detection via signal processing. Motor Current Signature Analysis (MCSA) analyzes sideband frequency components in the current spectrum; vibration, acoustic emission, and stray flux signals are similarly used to identify fault signatures. These methods detect and assess fault severity but lack prognostic capability to quantify the degradation trajectory and predict Remaining Useful Life (RUL). An Industry 4.0 environment needs a framework for transitioning from corrective to predictive maintenance.

This work uses Prognostics and Health Management (PHM) via pole migration of the system. The pole shift quantifies degradation, a larger shift indicating greater rotor bar deterioration. This health indicator (HI) models degradation trends, providing reliable RUL prediction without large-scale dataset dependencies.

The transfer function of the induction motor is given by:

$$G(s) = \frac{k}{LJS^2 + (RJ + BL)S + (RB + k^2)}$$

Here K is the motor constant, L and R the electrical inductance and resistance, J and B, the mechanical inertia and viscous damping coefficient, and s the Laplace operator. The system poles are:

$$s_{12} = \frac{-a \pm \sqrt{a_1^2 + 4a_2a_0}}{2a_2}$$

Rotor resistance R is a degradation-sensitive parameter: BRB progressively increases R, directly affecting a_1 and a_0, while a_2 remains constant.

The framework assumes: (i) a linear, time-invariant, lumped-parameter equivalent circuit linearized about a nominal operating point — valid for small perturbations, not large-signal transients such as direct-on-line starting ; (ii) BRB severity is represented as an equivalent rise in R, the parameter most sensitive to rotor-circuit asymmetry ; (iii) a second-order (two-pole) reduction that neglects faster electrical transients, justified when the electrical time constant is much smaller than the mechanical one, as for comparable reduced-order motor models ; and (iv) L, J, B, K remain constant while only R varies, decoupling electrical from mechanical degradation. These hold near rated speed/load; validity may degrade under saturation, variable-speed, or heavy-load conditions.

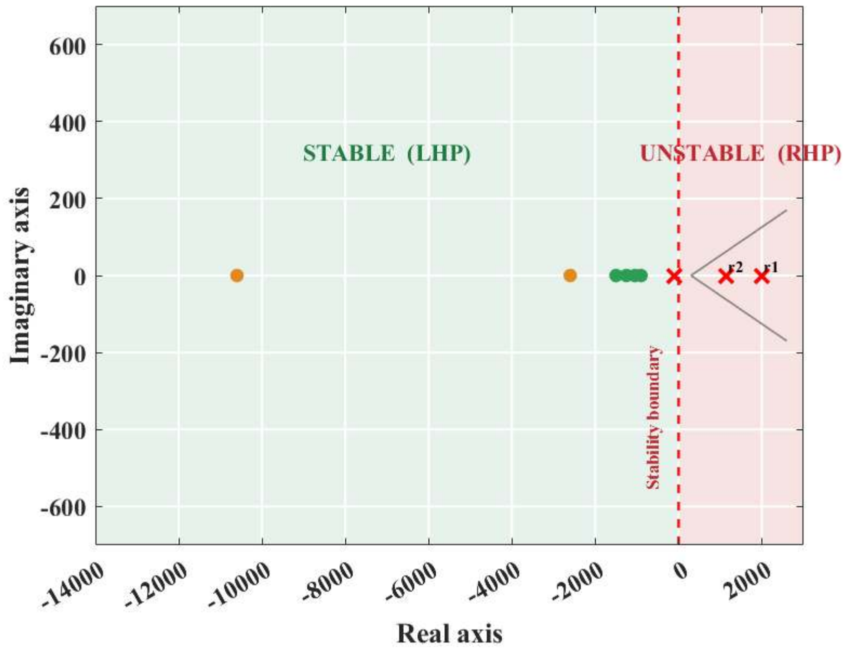

Figure 1. S-plane representation of pole migration under increasing rotor resistance.

Figure 1 shows the dominant pole in the complex plane as the rotor resistance increases with the number of broken bars. At low severity, poles sit in the left half-plane (stable, high damping). At medium severity, they migrate rightward but remain stable. At the third level, the dominant pole crosses the stability boundary ($\sigma = 0$), marking the onset of instability and unsafe operation. This gradual RHP shift ties rotor degradation to system dynamics, motivating the HI. As noted above, these three levels are discrete checkpoints on the underlying continuous R(t) trajectory used for RUL extrapolation.

This study is analytical/simulation-based: pole locations in Figure 1 come from the closed-form transfer function under parametrically varied R (a physics-based surrogate for BRB severity), not from measured current signals, and no experimental rig or motor nameplate data is used at this stage. Future work will validate the framework on an instrumented motor with artificially seeded BRB faults, obtaining K, L, R, J, B via standard tests (no-load, locked-rotor, DC resistance; IEEE Std 112) and estimating $R(t)$ online via recursive least squares or Kalman filtering from stator voltage/current to validate simulated pole/HI trends against real fault progression.

Abstract ID: 3140

CLIP-Based Zero-Shot Visual Twin for Manufacturing

Rajadurai S*, Kanishk H, Sandhya C, Damodharan K, Karthick S

Meenakshi Sundararajan Engineering College, Anna University, Kodambakkam, Chennai-24
**rajadurais825@gmail.com*

Abstract

Visual quality inspection in modern smart manufacturing systems faces severe challenges due to scarce labelled defect data, high annotation costs, and the inability of conventional deep learning models to generalise to unseen defect categories. Current DT frameworks for visual inspection are based on supervised models that require thousands of labelled images per defect class and require complete retraining when new product variants or defect types appear. These constraints lead to lengthy deployment cycles, expensive operational costs, and brittle inspection pipelines that fail under dynamic manufacturing conditions. A critical gap that challenges the promise of autonomous and adaptive smart manufacturing systems is the lack of ability to detect novel defects in real time.

Figure 1 shows a simple working flow of a CLIP-based defect detection system with a digital twin. Firstly, in the sensing layer, images are captured by the industrial camera and the images are resized and normalized in the preprocessing step. Then, in the inference layer, the CLIP model transforms both images and text descriptions into embeddings and compares them to identify defects. The detected results are then fed into the digital twin module where the system is updated in real time, including the classification of the defect and the generation of heatmaps. On the other hand, the knowledge layer stores the defect information by embeddings and keeps it updated. Finally, the results are displayed through visual outputs, alert and dashboard in the output layer for easy monitoring.

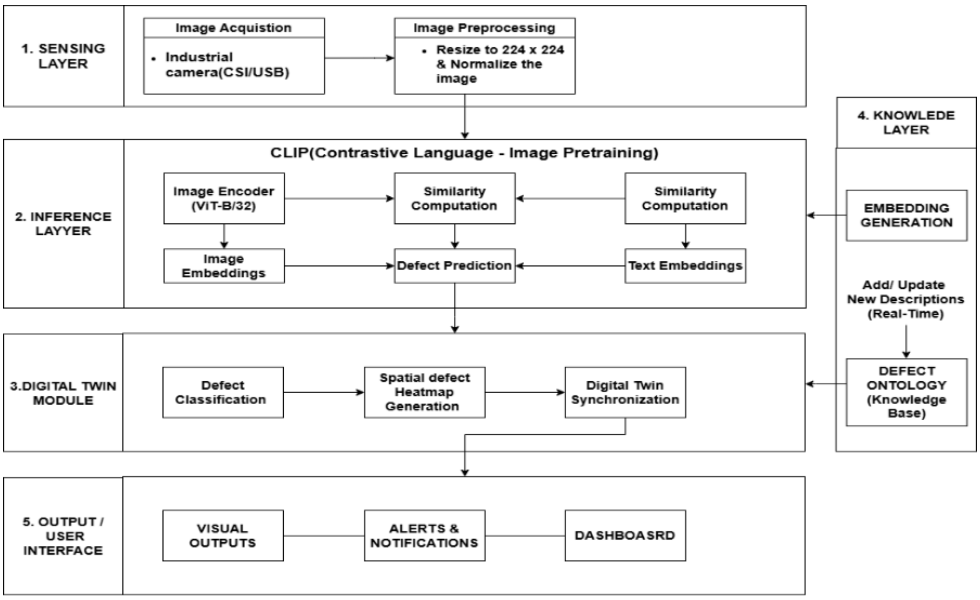

Figure 1. Architecture diagram of a clip-based digital twin system for automated defect detection.

In this paper, we propose a novel VisualTwin framework, a CLIP-based zero-shot visual digital twin that offers open-vocabulary defect intelligence for smart manufacturing systems without the need of labelled training data. The VisualTwin model describes the manufacturing inspection system with a three-layer architecture, which comprises of a Sensing Layer, an Inference Layer and a Knowledge Layer. The Sensing Layer acquires real-time product surface images through industrial cameras. The Inference Layer is based on a Contrastive Language-Image Pretraining (CLIP) vision-language model to calculate the cosine similarity between visual embeddings and textual defect descriptors. The Knowledge Layer is comprised of a living natural language defect ontology that operators can update in real time by simply adding new text descriptions such as “circular crack on weld seam” or “surface delamination near edge” without any model retraining.

To address the issue of false-positive classification in continuous ontology updates, the presented framework incorporates cosine-similarity confidence thresholding, ontology consistency validation, and Grad-CAM based spatial heatmap verification. Predictions with confidence below a pre-defined threshold are flagged for operator verification before they are added to the knowledge base of the Digital Twin, thus ensuring that the defect classification is reliable and consistent.

Together, these three components form a novel zero-shot visual digital twin system that operates in real-time and adapts continuously to evolving manufacturing requirements. The main contribution of this work is the introduction of a visual twin intelligence module, VisualTwin, which exploits the shared vision-language embedding space of CLIP to perform defect classification with zero labelled examples. The measurement method evolves from static supervised inspection to dynamic open-vocabulary detection with real-time DT synchronisation. The system relies on a living knowledge base that updates defect ontologies through natural language that facilitates early novel defect detection and prevents unobserved anomalies from escalating to system-wide quality failures. The proposed system comprises a self-updating DT layer that produces spatial defect heatmaps and temporal anomaly logs for predictive quality control at the process level.

The MVTec AD industrial anomaly detection benchmark and cross-domain generalisation tests on metal, wood, leather and ceramic substrates are used to evaluate the system performance, to comprehensively assess the defect detection capability and zero-shot transferability.

The VisualTwin framework represents significant advances towards visual digital twin technology that enables manufacturing inspection systems to operate through zero-shot and open-vocabulary intelligence approaches, and introduces novel contributions through five research areas: (i) CLIP-based vision-language embedding for zero-shot defect classification, (ii) incorporation of a living natural language defect ontology that enables real-time knowledge updates, (iii) spatial defect heatmap generation and temporal anomaly logging within the digital twin, (iv) cross-domain generalisation across heterogeneous product substrates without fine-tuning, and (v) edge-deployable inference pipeline using lightweight ViT-B/32 variants for real-time manufacturing deployment. These contributions together form a new intelligent visual digital twin system that operates in real time.

Our proposed manufacturing inspection framework demonstrates effective zero-shot defect detection and strong cross-domain generalization for industrial surface defect images from MVTec AD dataset. By leveraging the pre-trained CLIP vision-language model, the framework eliminates the need for large labelled datasets and repeated model retraining, making it suitable for detecting previously unseen defect categories. We introduce VisualTwin, a framework for intelligent visual inspection that combines zero-shot learning, open-vocabulary defect recognition, and Digital Twin integration. This allows a cost-effective, scalable and practical solution to the problem of labelled-data scarcity in an industrial quality control for the next-generation smart manufacturing systems.

Compared with the state-of-the-art supervised anomaly detection methods, the proposed VisualTwin framework can achieve effective zero-shot defect detection without the requirement of large labelled datasets and repeated retraining of the model. The framework uses a pre-trained CLIP model, which reduces training time and computational cost, and allows deployment on standard computing systems. This makes the proposed approach a cost-effective, scalable and practical solution for real-time smart manufacturing applications.

Abstract ID: 3072

Real-Time Distortion Monitoring During Welding of Structures Using Image Processing Techniques

Suryajith S S, Lijo Johnson, Joy Varghese V M

Department of Mechanical Engineering, College of Engineering Trivandrum, Thiruvananthapuram, Kerala, India

*suryajithsssathy@gmail.com

Abstract

Welding-induced residual stresses and distortions pose significant challenges to structural integrity in manufacturing industries, particularly in high-precision applications such as aerospace and automotive sectors. Traditional assessment methods like X-ray diffraction and ultrasonic testing suffer from limitations including high costs, safety concerns, limited accessibility, and computational inefficiency. This study presents an innovative, non-destructive image processing technique for quantifying welding deformations and residual stresses in Gas Tungsten Arc Welding (GTAW) components, offering a cost-effective alternative to conventional approaches.

The methodology involves laser-engraving a grid pattern (comprising 360 cells with 5×5 mm and 10×10 mm sections) on low-carbon steel plates (120×120×3 mm). Pre-weld and post-weld high-resolution images (4000×4000 pixels) are captured using a Canon 1500D camera. A custom MATLAB application with an intuitive graphical user interface (GUI) was developed to automate the complete analysis workflow. The algorithm implements bottom-hat morphological filtering with a 10-pixel disc-shaped structuring element for optimal grid extraction, followed by adaptive thresholding and connected component analysis for segmentation. Grid cells are identified via centroid detection and bounding box extraction, with one-to-one correspondence established between pre-weld and post-weld states using Euclidean distance minimization.

Deformation quantification is achieved by comparing cell dimensions before and after welding, calculating longitudinal and transverse strains. Residual stresses are subsequently estimated using 2D linear-elastic stress-strain relationships incorporating Young's modulus (210 GPa) and Poisson's ratio (0.29) for low-carbon steel. The application successfully identified all 360 grid cells in pre-weld images and matched 286 cells post-weld, with unmatched cells attributed to severe distortion or obscuration near the weld zone.

For validation, a three-dimensional thermo-mechanical finite element model was developed using ANSYS APDL. The simulation employed a modified Goldak double-ellipsoidal moving heat source (80 A current, 16 V voltage, 5 mm/s velocity, 85% efficiency) with temperature-dependent material properties sourced from JMatPro. Beer's empirical heat transfer coefficient was applied for accurate cooling simulation. The sequential uncoupled approach involved transient thermal analysis followed by static structural analysis using SOLID90 and SOLID186 elements with bilinear isotropic hardening.

Results demonstrate strong qualitative agreement between image-based measurements and FEA predictions. Both methods consistently identified tensile residual stresses (peak ~300-450 MPa) concentrated along the weld centreline, transitioning to compressive stresses (~-200 to -300 MPa) in the heat-affected zone and base metal regions. Transverse stress components exhibited similar trends with lower magnitudes (~100-200 MPa tensile). The image processing algorithm demonstrated significant improvement over prior grid detection approaches, substantially reducing spurious detections and noise artifacts. Temperature evolution analysis confirmed peak temperatures of

approximately 1500°C at the weld centre, with rapid heating and gradual cooling characteristic of TIG welding processes.

The developed methodology offers particular promise for industrial applications where traditional techniques prove impractical. The non-contact, image-based approach enables rapid in-situ assessment without specialized equipment or radiation safety concerns. The GUI implementation significantly lowers adoption barriers, enabling quality control personnel to perform comprehensive deformation analysis without advanced computational expertise. While some quantitative variations between image-based and FEA results were observed (attributable to 2D simplification, optical artifacts, and linear-elastic assumptions), the strong spatial correlation validates the technique's reliability.

Key limitations include the 2D nature of analysis (neglecting out-of-plane displacements), sensitivity to optical artifacts (glare, uneven illumination), and linear-elastic assumptions that may not capture plastic deformation or phase transformations. Future work will focus on extending to 3D deformation tracking, integrating machine learning for real-time monitoring, and validating across diverse welding configurations and materials. This research advances optical measurement techniques for welding analysis, offering a scalable, cost-effective solution for residual stress evaluation in critical manufacturing sectors.

Abstract ID: 3114

Investigation of Failure and Progressive Wear in Carbide Micro-Drills via Acoustic Emission Signal Analysis

Abhishek Shrivastava*, Amitava Ghosh

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India
**me22d007@iitm.ac.in*

Abstract

Monitoring micro tool performance during operation is critical in precision machining, particularly when working with challenging materials such as austenitic stainless steel 316 (SS316). This study investigates the use of acoustic emission (AE) sensors for real-time, non-intrusive monitoring of carbide micro drills during the micro-machining of SS316. AE sensors can detect high-frequency stress waves generated by tool wear, edge-chipping, and chip formation, providing early indicators of tool degradation and performance changes. The experiments were conducted using uncoated carbide micro drills under controlled drilling conditions. The results reveal that AE signal characteristics evolve systematically with progressive tool wear. At early stages, continuous and low-amplitude AE signals corresponded to stable cutting, while at advanced stages, the emergence of high-frequency bursts (400-500 kHz) and transient spikes indicated chipping and tool breakage. Statistical parameters such as kurtosis and crest-factor increased sharply during failure, confirming their sensitivity as wear indicators. Scanning electron microscopy (SEM) of the drill edges corroborated the AE findings, revealing built-up edge (BUE) formation, flank wear, and edge chipping at corresponding AE transition points. The study demonstrates that AE-based monitoring can serve as a reliable, non-intrusive diagnostic tool for early detection of tool degradation in micro-drilling, paving the way for intelligent process monitoring and predictive maintenance in precision.

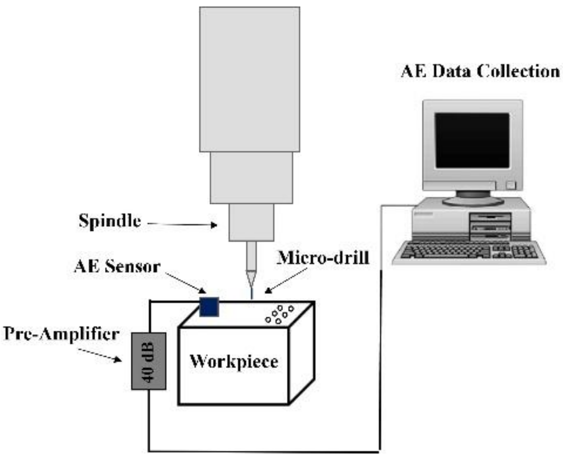

Figure 1. Schematic representation of the experimental setup for AE-based monitoring during micro-drilling.

Methodology

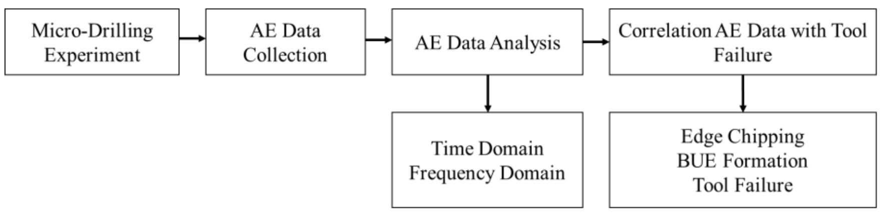

Figure 2. Methodology.

Results and Discussion

This research established a systematic relationship between AE signal characteristics and the progressive wear and failure of uncoated carbide micro-drills during drilling of SS316. The micro-drill failure is predominantly associated with higher frequency ranges (400–500 kHz), which coincided with tool breakage at the 14th hole. Time-domain analysis indicated that AE signals from a fresh tool were stable and continuous, while signals captured during tool breakage became highly unstable and discontinuous. Statistical parameters such as kurtosis and crest-factor showed a sharp increase during tool failure, confirming their reliability as sensitive indicators for micro-drill breakage detection.

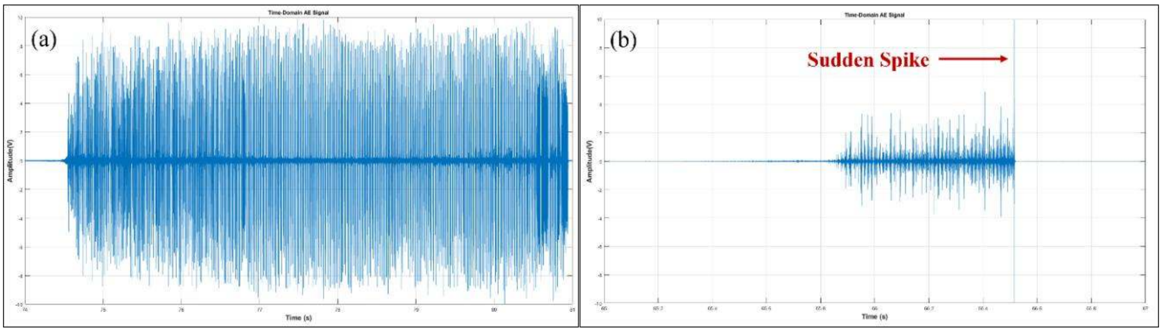

Figure 3. Representative time-domain Acoustic Emission (AE) waveforms during micro-drilling (a) Fresh Tool (b) Broken Tool.

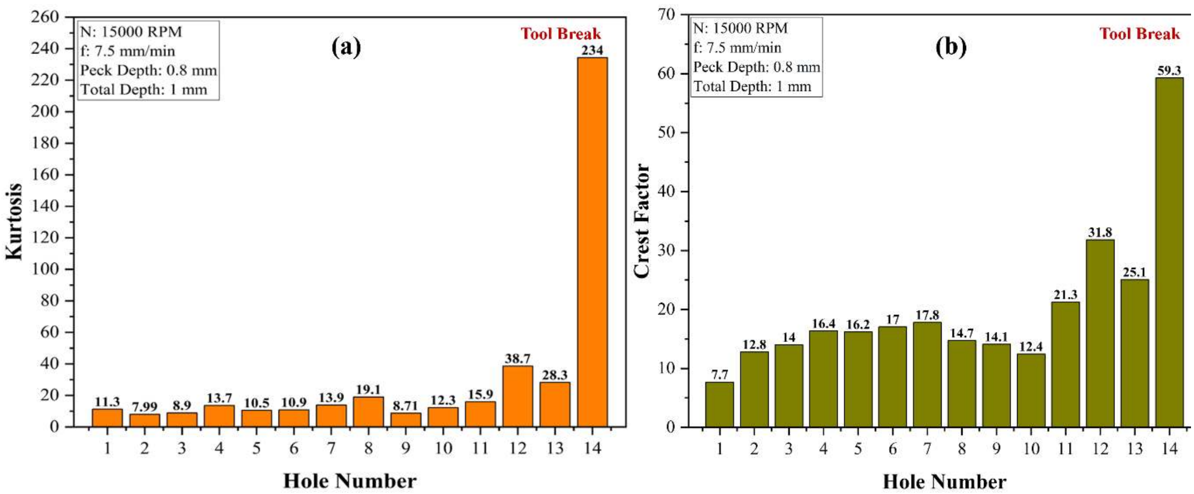

Figure 4. Variation of AE signal (a) Kurtosis (b) Crest Factor with the number of holes during micro-drilling.

Abstract ID: 3820

Acoustic Levitation and Passive Rotation of Planar Objects Using Perforated Vibrating Plates

Vipesh Singh*, Chaitanya Prasad Tudu, Manoj Pandey

Indian Institute of Technology Madras, India
**vipeshsingh95@gmail.com*

Abstract

Near-Field Acoustic Levitation (NFAL) has emerged as a promising technique for the contactless handling of delicate planar objects, offering a route to eliminate mechanical contact in applications such as semiconductor wafer transport, precision manufacturing, and contamination-free material manipulation. Conventional handling methods—conveyors, chucks, and robotic arms—generate particulates through physical contact, which can significantly degrade yield in cleanroom environments; NFAL addresses this by using a high-pressure squeeze film of air to support the object without touching it. This work designs, models, and experimentally investigates an ultrasonic levitation system driven by a 28 kHz Langevin-type piezoelectric transducer and a mechanical horn coupled to a circular radiating plate, with the aim of characterizing both levitation and rotational behaviour within a single, compact hardware platform.

Finite Element Modelling (FEM) of the coupled piezoelectric–structural system was performed to analyse mode shapes and vibration behaviour across the operating frequency range. Eigenfrequency and parametric studies evaluated the effects of horn and plate geometry on displacement amplification and resonance stability, with simulated nodal patterns validated experimentally via Chladni pattern visualization to confirm agreement between the numerical model and physical prototype.

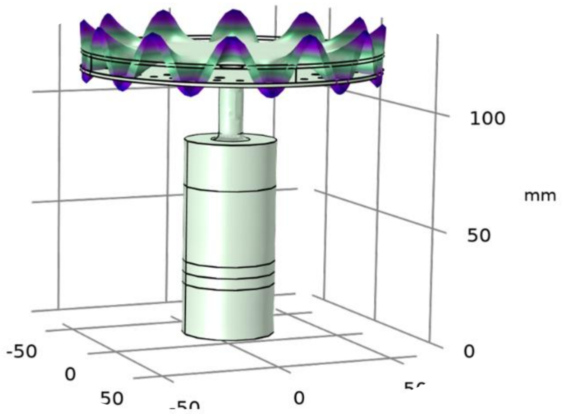

Figure 1. Circumferential waves generated by perforations.

Building on this validated platform, the study further explores perforations on the radiating plate as a means of inducing passive rotation of the levitated object, without requiring auxiliary actuators, motors, or multi-transducer setups. Preliminary experimental results confirm that perforation-based surface design can reliably drive rotational motion using only a single ultrasonic source, suggesting a compact, cost-effective route toward integrated levitation-and-rotation platforms for contactless manipulation. This approach has potential relevance to wafer alignment, inspection, and metrology stations in smart manufacturing systems, where combined positioning and rotation are typically achieved with more complex actuation schemes.

Abstract ID: 3755

Enhancement in Mechanical Properties and Microstructural Investigation of Aluminum Alloy 2024 using Microwave Hybrid Processing

Amrutha P K[1], Mukesh Kumar Roy[1], M Amarnath[2]*, Dinesh Rangappa[3], S Devaraj[4], Pankaj Bhagat[5]

[1]*Department of Physics, Indian Institute of Information Technology Design and Manufacturing, Jabalpur, Madhya Pradesh 482005, India*

[2]*Tribology and Machine Dynamic Laboratory, Department of Mechanical Engineering, Indian Institute of Information Technology Design and Manufacturing, Jabalpur, Madhya Pradesh 482005, India*

[3]*Department of Applied Sciences (Nanotechnology), Visvesvaraya Technological University, Muddenahalli, Chikkaballapura, Karnataka 562101, India*

[4]*Department of Mechanical Engineering, Reva University, Bengaluru, Karnataka 560064, India*

[5]*Department of Mechanical Engineering, Indian Institute of Information Technology Design and Manufacturing, Jabalpur, Madhya Pradesh 482005, India*

**amarnath@iiitdmj.ac.in*

Abstract

To meet the challenging industrial requirements, the mechanical properties of the materials need to be enhanced. This can be achieved by altering the surface properties of materials by using heat treatment methods. In comparison to the conventional heat treatment methods, microwave heat treatment offers specific advantages such as rapid processing time, energy efficiency, and lower environmental hazard. In microwave heat treatment, heat energy is directly transferred to the material due to the interaction of electromagnetic waves with molecules of the test specimen, which leads to volumetric heating. However, for the heat treatment of microwave reflecting materials, viz. metal alloys, a technique known as microwave hybrid heating (MHH) has been used extensively. In the present work, an Aluminium Alloy (AA2024) test specimen is exposed to microwave irradiation, a microwave oven (900 W, 2.45 GHz) is used to enhance the mechanical properties of the test specimen under the above recrystallization temperature. Scanning electron and optical microscopy, X-ray diffractometry techniques were used to evaluate microstructure, grain boundary, and phase transformation details, respectively. Further, density and hardness measurements were also carried out to highlight the significance of MHH in enhancing the mechanical properties of the test specimen.

Keywords: Aluminium alloy, heat treatment, grain refinement, hardness

Methodology

Figure1 depicts the schematic diagram of the MHH setup used in the experimental investigations, which consists of a microwave oven operated at 900 W, 2.45 GHz. In this oven, the AA 2024 specimen is placed in an alumina crucible, embedded with charcoal. The charcoal acts as a susceptor which couples with microwave energy and heats the specimen from the outer surface to the core by conduction mode of heat transfer.

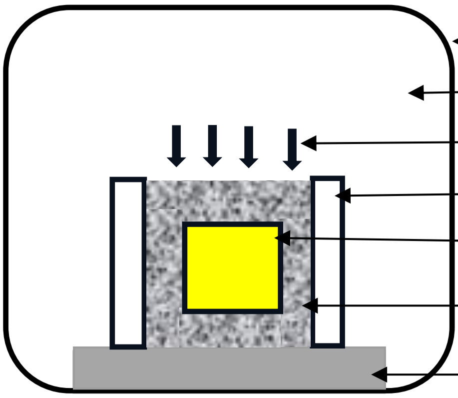

Figure 1. The schematic diagram of a microwave hybrid.

The abovementioned procedure was used by the researchers to enhance the mechanical properties of various metal alloys. The MHH was carried out beyond the recrystallization temperature, from 264°C to 440°C over a period of 600 seconds. Further, the specimen was allowed to cool under the air-cooling method. In order to evaluate the enhancement in surface and mechanical properties, the specimen was subjected to a polishing process by using 400, 600, 800, 1000, and 1200 grit-sized abrasive paper, then alumina powder and 6-micron diamond paste were used for the disc polishing process over a period of 15 minutes, followed by an etching process.

Results and discussion

The effect of MHH on the density of material is measured by using Archimedes' principle. The MHH resulted in hot compaction, which led to better densification. However, the increase in density by 0.3342% was caused by shrinkage and less porosity in the specimen. Further, grain refinement occurred in the specimen due to MHH, resulting in a noticeable increase in the hardness of the material. The hardness measurement of the specimen was performed by using a 1/16-inch diameter ball indenter, Rockwell B scale with a load of 100 kg. Figure2 indicates the hardness and density of the specimen before and after MHH. The hardness values of the specimen subjected to 10 minutes of MHH treatment showed an increase in hardness values of 36 HRB in comparison to the untreated sample, which had a hardness of 35 HRB. The hardness of the sample improved due to an increase in density, grain refinement, and a decrease in porosity after MHH treatment.

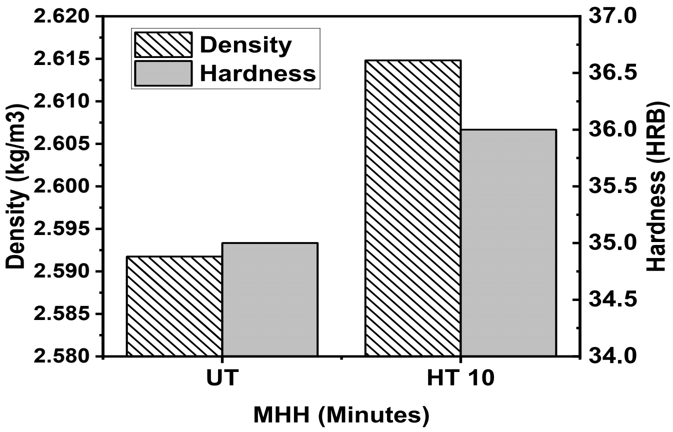

Figure 2. Density and hardness values of AA 2024 material of untreated and MHH.

Abstract ID: 3092

Towards Automated Structural Performance Optimization: Multi-Objective Design of a BIW Front Subframe with Coupled Shape Morphing and Stress Constraints

Eshan Amalnerkar*, Vidit Sharma

ESTECO Software India Pvt. Ltd., Pune, India
**amalnerkar@esteco.com*

Abstract

Automotive Body-in-White (BIW) front subframes must reconcile conflicting targets: low mass for efficiency and electrification, sufficient stiffness and strength to bound dynamic stresses, and high natural frequencies for NVH compliance. Single-objective approaches cannot expose the full Pareto trade-off. This work couples high-fidelity FEA with parametric CAD morphing (Figure 1) in a closed-loop multi-objective evolutionary framework that integrates ANSA (parametric model management and spot-weld realization), MSC Nastran SOL 103 (linear-static and normal-modes analysis), and modeFRONTIER (DOE, surrogate modelling, and Pareto exploration), applied to the subframe (Figure 2).

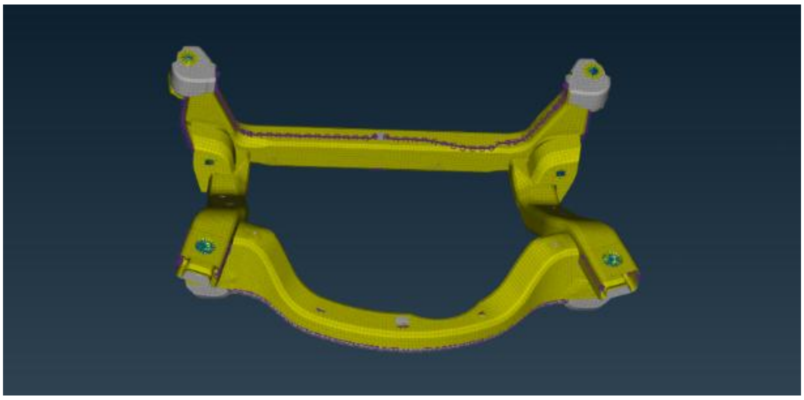

Figure 1. Meshed geometry of the BIW front subframe.

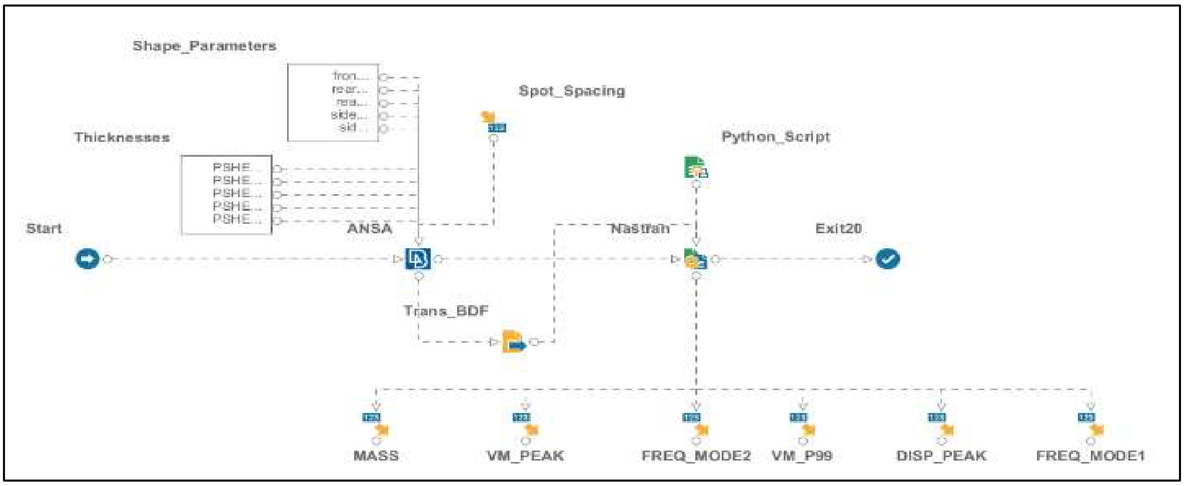

Figure 2. modeFRONTIER closed loop workflow.

Table 1. Optimization formulation and material properties.

Objectives	min MASS; min VM_P99 (99th-pct σvM); max FREQ_MODE1
Constraints	VM_PEAK≤1420 MPa; DISP_PEAK ≤ 26.2 mm
Variables (11)	Thickness×5: 0.5–2.0 mm; Shape morph×5: 0–5.0 mm; Spot-weld pitch: 2–30 mm
Material	Structural steel; E = 210 GPa, ν = 0.3, ρ = 7.85e−9 t/mm3

Approach

Eleven continuous design variables were parameterized in ANSA, comprising five panel thicknesses (PSHELL_700001/2/5/6/7_T, 0.5–2.0 mm), five mesh-morphing parameters (front/rear/side height and rear/side width, 0–5.0 mm) that modify the local section depth and flange geometry without remeshing, and the spot-weld pitch (2–30 mm). All panels were assigned a single linear-elastic steel material (Table 1). Two inequality constraints limited the peak von Mises stress (VM_PEAK) and peak displacement (DISP_PEAK) to their respective baseline values, ensuring that no optimized design compromised structural strength or stiffness. The optimization problem captures the classical lightweight design trade-off between mass, structural integrity, and NVH performance. Mass was minimized to achieve lightweighting, while the 99th-percentile von Mises stress (VM_P99) was minimized instead of the absolute peak stress to reduce the influence of localized mesh-induced stress singularities around spot welds and to provide a smoother response surface for optimization. Simultaneously, the first natural frequency (FREQ_MODE1) was maximized to improve dynamic performance by increasing the separation from typical road and powertrain induced excitation frequencies. Crashworthiness and fatigue performance were beyond the scope of the present linear static and modal analyses.

Optimization and Surrogate Model Validation - Multi-objective optimization was performed using the MOGA-II algorithm with a population of 100 designs evolved over 500 generations. To reduce the computational expense of repeated finite element analyses, response surface metamodels were constructed from a Uniform Latin Hypercube Design of Experiments (DOE). Multiple surrogate techniques, including Adaptive Kriging (AKR), DACE (Kriging), Gaussian Process (GP), Extended Design (ED), Radial Basis Function (RBF), Neural Network (NN), Support Vector Regression (SVR), and polynomial models, were evaluated for each response. The best-performing surrogate was selected based on validation metrics, primarily the coefficient of determination (R^2). The validation results showed that DACE (Kriging) provided the highest predictive accuracy for the displacement (DISP_PEAK) and mass (MASS) responses, Extended Design (ED) was the most accurate surrogate for the frequency responses (FREQ_MODE1 and FREQ_MODE2), and the Neural Network (NN) model achieved near-perfect prediction for the stress responses (VM_P99 and VM_PEAK). All selected surrogate models achieved $R^2 \geq 0.992$, demonstrating excellent agreement with the finite element results and confirming their suitability for the subsequent optimization process. Replacing the MSC Nastran solver with the validated metamodels reduced the evaluation time from approximately 25s per SOL 103 analysis to only a few milliseconds, substantially improving computational efficiency.

To ensure that surrogate approximation errors did not influence the final conclusions, the three highest-ranked Pareto-optimal solutions (Figure 3) were re-evaluated using the original MSC Nastran models. The surrogate predictions agreed closely with the finite element results, with deviations below 2.5% for VM_PEAK and 3.4% for VM_P99, while the largest discrepancies were observed for mass and FREQ_MODE1, reaching approximately 6%. Consequently, all final performance metrics reported in this study are based on the MSC Nastran re-analysis.

Results and Discussion

Three MSC Nastran-validated Pareto-optimal solutions were identified, representing the principal trade-offs between minimum peak stress, maximum first natural frequency, and the best overall compromise (Figure 3). Among these, Design ID 877 was selected as the preferred solution because it achieved the most balanced improvement across all objectives. As summarized in Table 2, the optimized design maintained the structural mass essentially unchanged (+0.05% relative to the baseline) while accepting a modest increase in the 99th-percentile von Mises stress (+2.0%). In return, it achieved a 12.4% reduction in peak von Mises stress and a 3.5% increase in the first natural frequency, while satisfying both displacement and peak-stress constraints. This solution effectively prioritizes the reduction of the critical peak-stress hotspot and enhances the NVH margin, while maintaining acceptable overall structural performance.

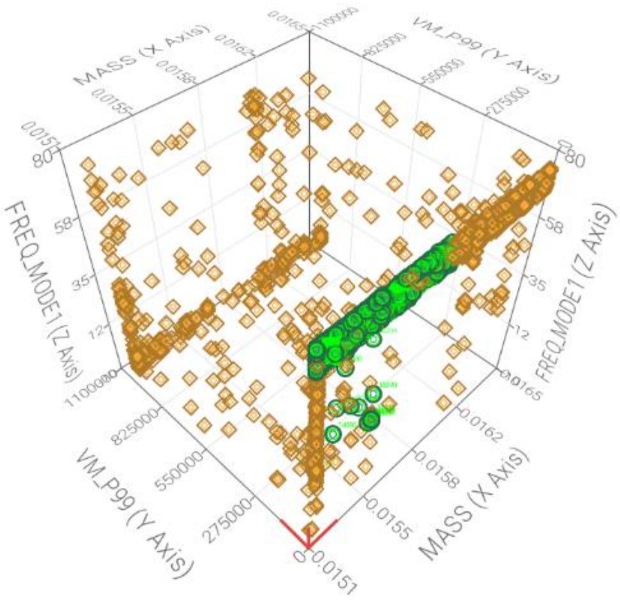

Figure 3. Pareto Frontier for 3 objectives.

Table 2. Baseline vs. optimised performance.

Output	Base	Opt.	% change
MASS (kg)	15.839	15.847	+0.05%
VM_P99 (MPa)	733.3	747.6	+2.0%
VM_PEAK (MPa)	1420	1244	**−12.4%**
FREQ1 (Hz)	68.1	70.6	**+3.5%**
FREQ2 (Hz)	87.8	90.0	+2.5%
DISP (mm)	25.17	25.49	**+1.2%**

Abstract ID: 2899

Geometric Influences on Compressive Performance of Additively Manufactured Gyroid Cellular Architectures

Mallikarjuna Reddy D*, Ayitha Sreenivasulu#, Bandi Isaac, Moturu Uday Kiran, Abhinav V M

School of Mechanical Engineering, Vellore Institute of Technology, Vellore, Tamil Nadu, India
**dmreddy@vit.ac.in, #ayitha.s@vit.ac.in*

Abstract

Gyroid structures are a common type of Triply Periodic Minimal Surface (TPMS) known for their mechanical properties, such as energy absorption and stress distribution. This study investigates the compressive behaviour of gyroid structures produced by Fused Deposition Modelling (FDM). Polymeric gyroid specimens made of PLA material were 3D printed at varying relative densities. Three specimens with different structural configurations were designed using 3D modelling software and fabricated with the QIDI MAX 3D. Compression tests were conducted to assess their mechanical performance under load. Load–extension and stress–strain data were collected to analyse compressive behaviour. Results show that compressive strength and load-bearing capacity increased as the structural configuration changed. The specimens displayed typical cellular material responses, including elastic deformation, plateau collapse, and densification during compression. The sample with the highest geometric parameter achieved the greatest compressive load and strength, indicating enhanced stiffness and stability. These findings underscore the role of geometric design in the mechanical performance of gyroid lattice structures and offer valuable insights for developing lightweight, high-performance components for aerospace, biomedical, automotive, and energy-absorption applications.

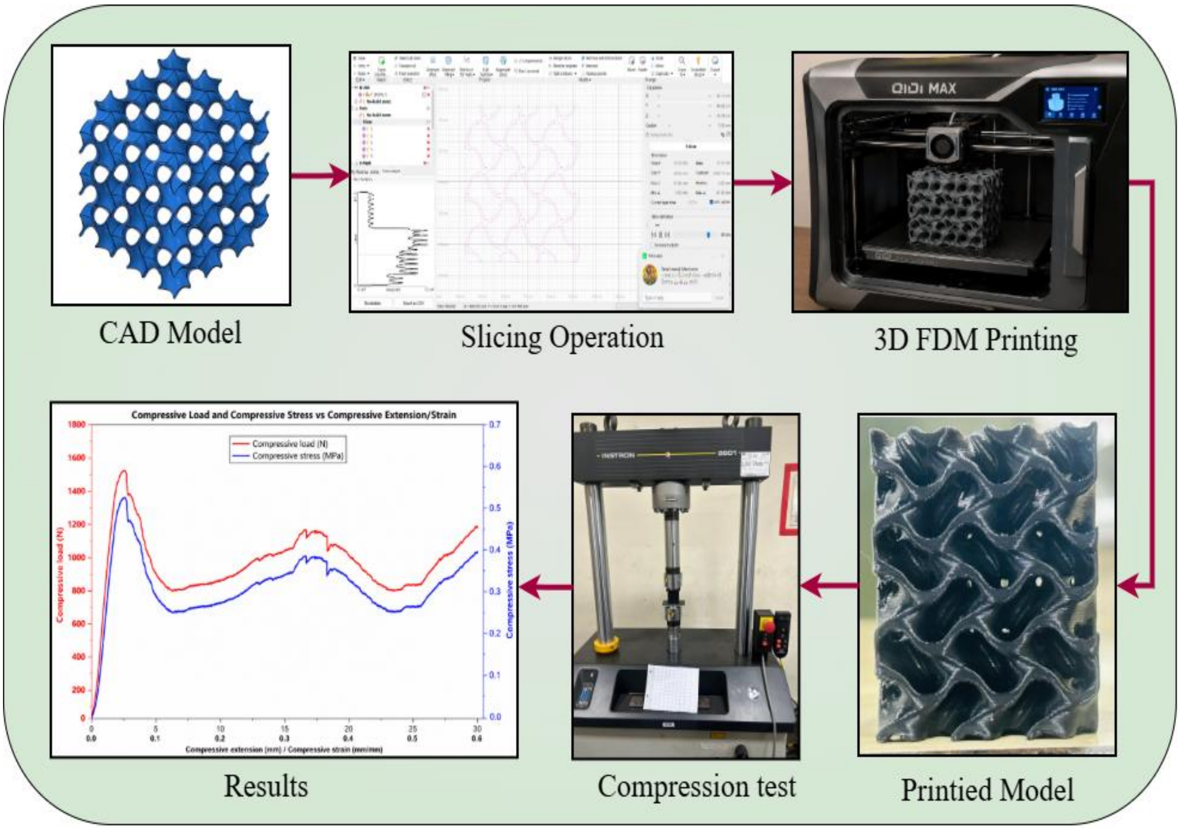

Figure 1. Workflow image of the complete process from CAD design to compression testing and mechanical performance evaluation of FDM-printed gyroid lattice specimens.

As shown in Figure 1, the workflow begins with creating the gyroid lattice CAD model, followed by slicing to generate machine instructions for the FDM printer. The specimen is fabricated using a QIDI MAX 3D printer and subsequently subjected to compression testing using a Universal Testing Machine (UTM). The recorded experimental data are analysed to obtain compressive load–displacement and stress–strain curves, enabling the evaluation of the mechanical performance of the printed gyroid lattice structure.

Printing Parameters:

PLA filament (1.75 mm diameter) was chosen for its printability, dimensional stability, and biodegradability. Gyroid lattice specimens were produced using an FDM 3D printer with a 0.4 mm nozzle. Printing was performed at a nozzle temperature of 240 °C, a build plate temperature of 60 °C, a layer height of 0.2 mm, and a printing speed of 35 mm/s. Specimens were printed vertically, with full cooling applied after the initial layers to ensure surface quality and dimensional accuracy. These parameters were selected to promote consistent layer bonding and reliable mechanical performance during compression testing.

Results and Discussion:

The result graph (as shown in the abstract figure) shows that compressive load (red curve) and compressive stress (blue curve) as functions of compressive extension (mm) and strain (mm/mm) during the specimen's compression test. Under quasi-static compression, the gyroid specimen initially exhibited a linear elastic region, followed by a peak compressive load of approximately 1510 N (0.53 MPa) at a strain of about 0.05 mm/mm. After this peak, the load decreased due to localised buckling and progressive collapse of the lattice cells. A plateau between 0.10 and 0.35 mm/mm strain indicates stable deformation and efficient energy absorption. Fluctuations in load and stress result from the sequential collapse of individual gyroid unit cells. At strains above 0.45 mm/mm, the load increased again as collapsed cell walls made contact, increasing structural stiffness. Overall, the compressive response demonstrates the gyroid lattice structure's superior energy absorption.

Abstract ID: 3168

Mechanical Performance of PETGCF-Based Sandwich Lattice Metamaterials Fabricated Using Fused Filament Fabrication

Aravindh Raj B R*, Velmurugan C

Department of Mechanical Engineering, Indian Institute of Information Technology, Tiruchirappalli, Tamil Nadu, India

**aravibabu8@gmail.com*

Abstract

Introduction

A family of synthetic materials known as metamaterials and lattice structures has mechanical properties that are mostly determined by their geometric arrangement rather than the inherent characteristics of the base material. Energy absorption systems, lightweight applications, and structural elements all make extensive use of these structures. However, when exposed to external loads, traditional lattice constructions frequently experience decreased strength and instability. To overcome these difficulties, sandwich-style constructions with a lattice core and solid outside layers have been created to improve structural integrity while preserving lightweight qualities. The development of additive manufacturing, especially Fused Filament Fabrication (FFF), has made it possible to efficiently produce complicated lattice structures utilising reinforced polymer materials like carbon fiber-reinforced PETG (PETGCF). Three distinct lattice geometries—gyroid structure, octet truss lattice, and hexagonal (honeycomb) lattice—are used in the intermediate region of sandwich lattice metamaterial specimens in this study, which have totally solid top and bottom layers. The mechanical performance of these constructions under tensile, flexural, impact, and hardness testing needs to be experimentally assessed and compared.

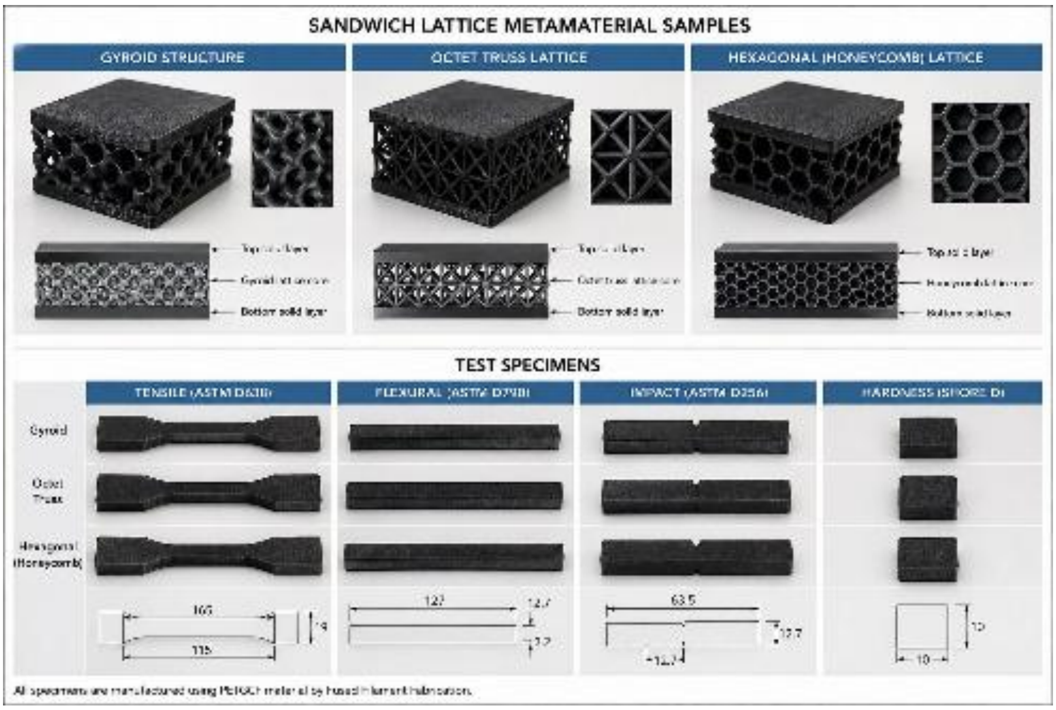

Figure 1. Schematic representation of sandwich lattice metamaterial structure showing solid top and bottom layers with gyroid, octet truss, and hexagonal (honeycomb) lattice cores.

Results and Discussion

Property	Gyroid	Honeycomb	Octet Truss
Tensile Strength (MPa)	44	46	52
Flexural Strength (MPa)	62	58	70
Impact Strength (kJ/m²)	7.5	6.0	4.2
Hardness (Shore D)	80	81	82

In light of its stretch-dominated construction and effective load transfer along strut members, the octet truss lattice demonstrated the highest tensile strength (52 MPa) and flexural strength (70 MPa). Because of this, it can be used in situations where high rigidity and load-bearing capability are necessary. Because of its continuous and smooth geometry, which permits uniform stress distribution and progressive deformation under loading circumstances, the gyroid structure displayed higher impact strength (7.5 kJ/m²) and energy absorption capability. In every test, the hexagonal (honeycomb) structure performed moderately, demonstrating a well-balanced mix of strength and deformation characteristics. Because of this, it can be used in applications where mechanical behaviour must be steady and predictable. Since hardness values are mostly determined by the totally solid exterior layers rather than the internal lattice geometry, they were found to be almost the same for all specimens. By improving load distribution and postponing structural failure, the sandwich arrangement greatly enhanced mechanical performance as compared to traditional lattice constructions. The PETGCF specimens had strong interlayer bonding and homogeneous filament deposition, according to the optical microscopic investigation. There were a few flaws at the interface between the solid layers and the lattice core, suggesting ideal printing circumstances. Because of the intricate geometry of the octet truss construction, small porosity and layer gaps were found. The observed mechanical performance was supported by the honeycomb structure's consistent cell generation and the gyroid structure's more uniform material distribution.

Conclusion

This study shows that the integration of solid outer layers with a lattice core improves the mechanical performance of PETGCF-based sandwich lattice metamaterials made via fused filament fabrication. The findings verify that structural behaviour is significantly influenced by lattice geometry. Tensile and flexural strengths were demonstrated by the octet truss structure, suggesting that it is appropriate for load-bearing applications. On the other hand, because of its continuous geometry and superior energy absorption capacity, the gyroid structure demonstrated superior impact resistance. Honeycomb construction is appropriate for general-purpose applications because it demonstrated balanced mechanical performance in every test. All things considered, these structures exhibit great promise for energy-absorbing and lightweight engineering applications. The optical microscopic analysis showed uniform filament deposition and good interlayer bonding in the PETGCF specimens. Minor porosity was observed in the octet truss structure, while the gyroid and honeycomb structures exhibited more uniform material distribution.

Abstract ID: 3167

Multi-Objective Optimization of PETG-CF 3D Printing Parameters Using Machine Learning

Shishir Kasturi*, Pratheeka Dahagam

Department of Mechanical Engineering, Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering & Technology, Bachupally, Hyderabad, India

**kasturishishir@gmail.com*

Abstract

Fused Deposition Modeling (FDM) of advanced composites like Carbon Fiber-reinforced PETG (PETG-CF) involves complex, non-linear thermodynamic interactions. Traditional statistical optimization techniques struggle to resolve the inherent multi-objective conflict in additive manufacturing: settings that maximize mechanical structural integrity almost inevitably degrade surface finish. This study presents a data-driven Machine Learning framework that transitions from traditional trial-and-error to a mathematically optimized digital twin, integrating Explainable AI and Evolutionary Algorithms.

Methodology

The physical baseline data was generated using a Taguchi L27 Orthogonal Array. Twenty-seven PETG-CF specimens were fabricated across five parameters: Layer Thickness (mm), Infill Density (%), Raster Angle (°), Print Temperature (°C), and Print Speed (mm/s). Each specimen was physically tested for Tensile Strength (MPa) using a Universal Testing Machine and Surface Roughness (Ra, µm) using a digital profilometer. A Random Forest Regressor (ensemble of 100 decision trees) was trained on this empirical dataset to predict the mechanical outcomes of a full-factorial design space comprising 243 virtual combinations. To avoid the "black box" nature of AI, SHAP (SHapley Additive exPlanations) was integrated to dissect the decision trees and verify that the algorithm learned authentic manufacturing physics. Finally, the Non-dominated Sorting Genetic Algorithm II (NSGA-II) was employed over 50 generations to mathematically pinpoint a Pareto Front of optimal compromise settings.

Results and Discussion

The predictive digital twin successfully captured the physical variance of the composite material and the relatively low MAE and RMSE proved the model made consistent, minor errors rather than unsafe predictive outliers. The SHAP analysis visually proved the thermodynamics of the print: Infill Density emerged as the absolute dominant driver of mechanical strength, while Print Temperature and Layer Thickness primarily governed the chaotic flow of carbon fibers, dictating surface finish.

The NSGA-II algorithm mapped a Pareto Front, proving the multi-objective conflict, pushing the printer to its absolute maximum strength, resulting in a 56% degradation in surface quality. To solve this, a Balanced "Golden Ratio" Profile was extracted from the Pareto Front. Physical laboratory validation of this AI-predicted profile yielded an exceptionally durable Tensile Strength and Surface Roughness, perfectly aligning with the digital predictions. Ultimately, this framework provides a scientifically robust, cost-effective methodology for optimizing composite additive manufacturing.

Abstract ID: 2404

Six Sigma Application to Precision Instrument Manufacturing Facility for Defect Reduction: A Case Study

K. Prabhakar Sriram*, Jay Aditya, Harsh Kumar, Prof. Girish Kumar

Department of Mechanical Engineering, Delhi Technological University (Formerly DCE), Delhi, India
**sriramkrovvidi.sk@gmail.com*

Abstract

Manufacturing efficiency and defect reduction remain critical challenges in the small-scale industry sector, where quality failures directly impact product reliability and customer satisfaction. This paper presents a diagnostic case study applying the Six Sigma DMAIC (Define, Measure, Analyse, Improve, Control) methodology and Seven Quality Control (7QC) tools to a small-scale voltmeter manufacturing facility in Delhi, India. The primary objective is to demonstrate how structured quality management tools can systematically identify root causes of defects and generate a targeted improvement action plan applicable to the Indian MSME context.

The DMAIC framework was applied across five structured phases. In the Define phase, a SIPOC (Supplier, Input, Process, Output, Customer) chart was constructed to establish process boundaries and identify critical input-output relationships across the 22-component voltmeter bill of materials. In the Measure phase, defect data were passively collected from existing production quality logs over a 30-day period without intervention, covering a total production volume of 6,340 units across all component types. A total of 869 attribute defects were recorded, yielding an overall defect rate of approximately 13.7%. In the Analysis phase, Pareto Analysis (80/20 rule) was applied to rank defect contributors by impact. In the Improve phase, Ishikawa (fishbone) diagrams were constructed through structured brainstorming sessions with facility management to map root causes across the six M-categories: Man, Machine, Material, Method, Measurement, and Milieu (Environment). The Control phase culminated in the development of a component-specific quality improvement action plan.

Pareto Analysis identified seven Critical-to-Quality (CTQ) components — Glass, Joule, Bobbins, Resistance (15–20 Ω), Coil, Dial, and Zero Adjuster were collectively responsible for approximately 80% of all recorded defects. Ishikawa analysis revealed that defect causation is inherently multi-factorial, spanning all six M-categories, with supplier quality and operator handling practices emerging as dominant causal themes across the CTQ components.

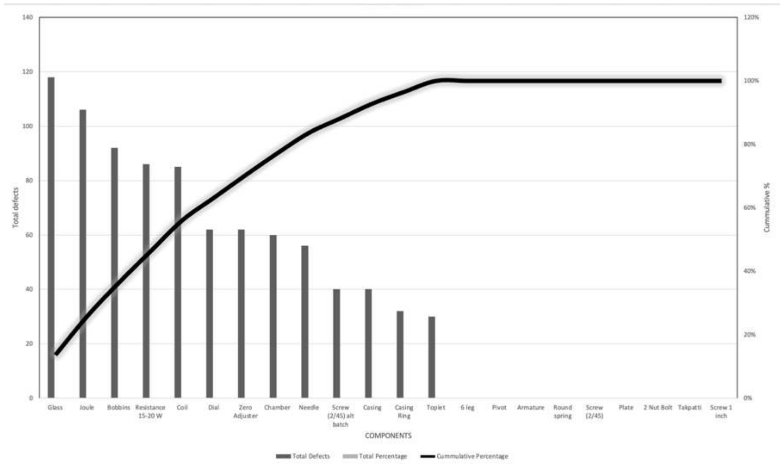

Figure 1. Pareto analysis.

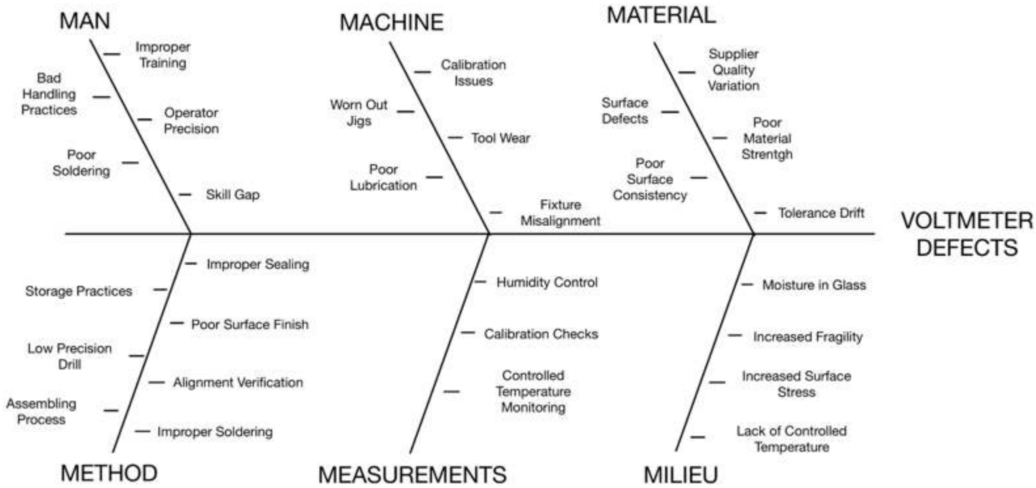

Figure 2. Ishikawa Diagram.

The improvement action plan encompasses component-specific recommendations including: corrosion-resistant spring materials and dimensional AQL controls for the Zero Adjuster; go/no-go gauging and First Pass Yield (FPY) tracking for Bobbins; dielectric strength testing and IPC-J-STD-001 compliance for Coils; ink density and alignment sampling plans for the Dial; humidity monitoring and glass-handling SOPs for Glass; moisture-controlled storage and periodic operator competency assessments for Resistance; and positional tolerance fixturing with CAPA protocols for the Joule. It is acknowledged that the recommendations represent a diagnostic improvement blueprint; implementation and empirical validation of outcomes are identified as directions for future research.

This study reinforces that Six Sigma is a scalable and adaptable methodology whose structured discipline, Pareto-driven CTQ prioritisation, and multi-causal Ishikawa analysis can deliver meaningful quality improvements in precision instrument manufacturing facilities irrespective of organisational scale. The findings demonstrate that rigorous evidence-based quality methodologies are neither exclusive to large enterprises nor dependent on sophisticated infrastructure, and are equally critical and accessible in the MSME manufacturing context.

Keywords: Six Sigma; DMAIC; Defect Reduction; Pareto Analysis; Ishikawa Diagram; 7QC Tools; MSME Manufacturing; Quality Control

Abstract ID: 2647

IIoT-Based Smart Steam Management for Deodorization Efficiency in Refineries

Rajalakshmi K*, Durgadevi S

Sri Sai Ram Engineering College, Affiliated to Anna University, West Tambaram, Chennai, India

*rajieee93@gmail.com

Abstract

The edible oil refining industry is considered to be highly energy intensive, with strict product quality requirements throughout the process. Deodorization is one of the crucial stages in the refining process, where volatile compounds such as free fatty acids, ketones, and pesticides are removed by stripping steam under high temperature and vacuum conditions. Even though steam injection is necessary for the production of neutral, odourless, high-grade edible oil, a limited amount of steam injection is considered a key factor. Too much of steam usage leads to high energy consumption, whereas less injection of steam affects the quality of the product. Thus, it is important to consider the energy loss by continuous monitoring of steam utilization and ensure the product's quality.

Figure 1. Steam Efficiency Monitoring Dashboard.

This paper proposes an IIoT-based steam monitoring system and a Streamlit dashboard for steam consumption in the deodorization process. In the proposed methodology, the historical data of the deodorization unit is received from the PLC and transmitted through the MQTT communication protocol. Further, the data is processed in a Python-based steam efficiency analytics engine. To monitor the steam utilization, the specific steam rate is calculated as the ratio of steam flow to feed flow. By comparing actual steam consumption with target steam consumption, the steam efficiency is calculated. Based on the evaluation, the Streamlit dashboard visualizes the calculated KPI, which helps in remote process monitoring and decision-making.

In contrast to traditional monitoring methods, which rely on manual observation or SCADA systems that display only the raw sensor data, the proposed method gives a continuous KPI evaluation and identifies the steam utilization performance across various operating conditions. Experimental validation was done with 1882 historical data samples of the industrial deodorization process. Figure 1, steam efficiency monitoring dashboard, was successfully developed and also classifies the three different statuses of steam consumption as graphically represented over a time period. This helps in continuous assessment of stripping steam utilization and ensures the connection of the communication protocol. The proposed system provides continuous information on steam consumption and steam efficiency, supporting optimal deodorization operation, helping to maintain product quality, and enabling better decision-making.

Keywords: Deodorization, Stripping steam rate, IIoT, Streamlit dashboard, Industry 4.0, KPI

Abstract ID: 3778

Design and Conceptual Development of a Trunk-Mounted End-Effector for Autonomous Rubber Tapping

Sakshi Jha[1]*, Ajay Krishnan V[2#], Santhakumar Mohan[2]

[1]*The Neotia University, West Bengal*

[2]*Indian Institute of Technology, Palakkad*

**sakshi.jha13024@gmail.com, [#]132314002@smail.iitpkd.ac.in*

Abstract

Rubber tapping is a critical agricultural process involving the controlled removal of bark from Hevea brasiliensis trees to extract natural latex, an essential industrial raw material used in tyre manufacturing, medical products, footwear, adhesives, conveyor systems, and other engineering applications. Despite advances in agricultural mechanization, tapping remains largely dependent on skilled manual labour, requiring a precise cutting angle, controlled bark consumption, and accurate incision depth while avoiding damage to the cambium layer responsible for bark regeneration. Irregular trunk geometry, varying bark thickness, difficult terrain, and the need to operate during early morning hours for maximum latex yield further complicate the process, while repetitive wrist motion exposes tappers to ergonomic strain. A growing shortage of experienced tappers and rising labour costs motivate autonomous tapping solutions capable of improving operational consistency while reducing dependence on manual skill. This work presents the design and simulation of an end effector designed for rubber latex collection. This is an attempt for contribution to autonomy in field of rubber latex collection.

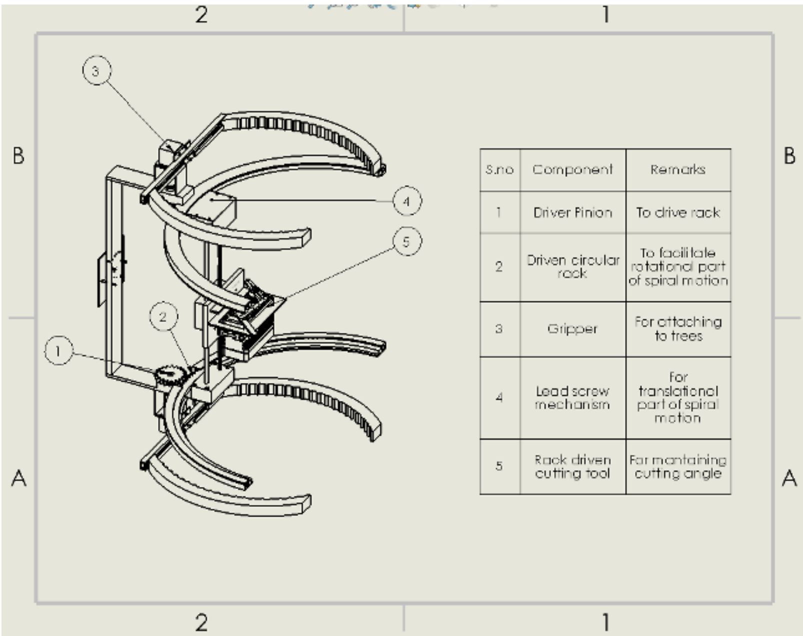

Figure 1. Complete assembly of the autonomous rubber tapping mechanism with labelled subsystems.

Methodology

A rubber-tapping end-effector incorporating a trunk-hugging dual-rack mechanism is proposed to generate a controlled spiral tapping motion around the tree trunk. Rather than relying on a manipulator-only approach that depends heavily on trajectory planning and sensor-based correction, the mechanism is mechanically guided, reducing control complexity while improving repeatability. As shown in Figure

1, the end-effector consists of two semi-circular driven racks (2) positioned around the trunk at different vertical levels, driven by a pinion (1) to generate the rotational component of the spiral motion. A lead-screw mechanism (4) provides the translational component and regulates bark consumption across successive cycles, while guide rods maintain alignment of the moving platform. The cutting tool (5) is mounted on a dedicated carriage and actuated by a servo-driven four-bar linkage, which converts rotational servo motion into guided angular movement of the cutter arm so that the required tapping orientation relative to the trunk surface is maintained throughout traversal, while keeping load off the servo shaft. A gripper (3), actuated by a stepper motor and guided rail system, secures the assembly around the trunk while accommodating moderate variation in trunk diameter, improving stability during cutter traversal.

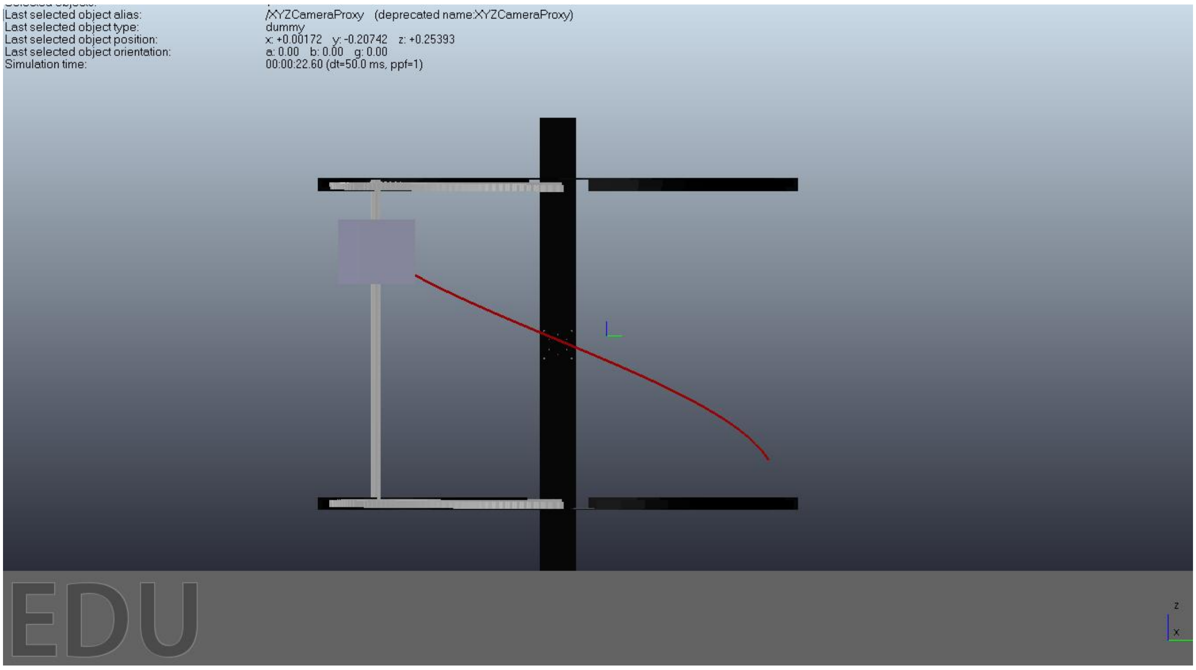

Figure 2. CoppeliaSim kinematic simulation showing the spiral tapping trajectory traced at the cutter endpoint.

The mechanism was modelled in SolidWorks for subsystem integration and motion generation. To verify kinematic feasibility ahead of physical prototyping, two complementary simulation studies were carried out in CoppeliaSim. First, the spiral tapping trajectory generated by coupled rotational and translational actuation was reproduced using a model with a prismatic joint (lead-screw motion) and a revolute joint (rack-and-pinion rotation) driven in coordination, with a tracer dummy attached at the cutter location; the actuation rates were set to match the target 25–30° helix angle, and the resulting path was recorded as a drawn line-strip in the world frame over the simulation run. Second, the servo-driven four-bar cutter-orientation linkage was modelled and actuated within the same Coppelia Sim environment to evaluate link rotations and the angular range of the cutter arm across the actuation cycle.

Results and Discussion

The Coppelia Sim spiral-trajectory simulation produced a smooth, continuous helical trace consistent with the intended 25–30° tapping angle, with no discontinuities over the full simulated traversal, supporting the design assumption that the dual-rack-and-lead-screw arrangement reproduces the manual tapping trajectory without requiring closed-loop correction during a single pass.

The four-bar linkage simulation showed a smooth, bounded angular sweep of the cutter arm appropriate for maintaining cutting orientation as the carriage traverses the curved rack, consistent with its intended role of keeping the load path away from the servo shaft. Together, the two simulations indicate that the rotational (spiral path) and orientation (cutter-angle) requirements of the tapping motion can be decoupled and independently verified, simplifying both mechanical design and future control development.

The CAD design and the two simulation studies demonstrate the feasibility of generating a controlled spiral tapping motion using a mechanically guided approach, with a compact configuration adaptable to varying trunk geometries. The assembly is intended for future integration with a manipulator arm on a mobile platform for autonomous positioning and traversal within a plantation row. Future work will involve fabrication of a scaled or full-size 3D-printed prototype, structural and trajectory-accuracy simulation, integration of depth sensors and machine-vision-based closed-loop control, and field testing to evaluate tapping accuracy, bark consumption, and terrain adaptability.

Abstract ID: 2655

Lightweight Robotic Gripper Finger Design Using Generative Approach

Ayush Kumar, Harsh Kumar Pandey, Satrajeet Kumar, MVA Raju Bahubalendruni, J. Ronald Aseer*

Department of Mechanical Engg, National Institute of Technology Puducherry, Karaikal, India
**ronaldaseer1981@gmail.com*

Abstract

Automation has now become an integral aspect of manufacturing industries to improve efficiency and accuracy in production processes. Robotic grippers form a vital part of automation, providing the precision needed in performing specific tasks. The aim of our research was to develop a design for a four-fingered robotic gripper which would be light and at the same time have the requisite strength.

The first phase involved modeling the robotic gripper design in fusion 360. It allowed us to visualize the different parts forming the robotic gripper and how the entire system works. Once the base design was completed, we used a generative approach to further design the prototype. The inputs used in the generative design included the design space, loads acting on the robot as well as various constraints. From the design results generated, we selected a design that had minimal material yet was stable and robust.

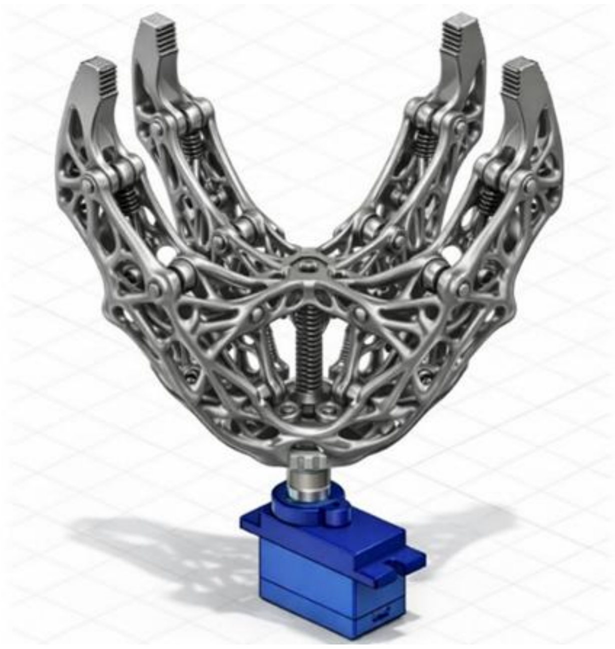

Figure 1. Generated Design.

In order to assess the design further, we conducted finite element analysis through ANSYS workbench. The test helped in analyzing the stress distribution, deformation, and factor of safety in various scenarios. Based on the results obtained from the analysis, it became evident that the optimized design had a better strength-to-weight ratio than the original one.

Lastly, we printed out the models with the help of additive manufacturing technology and performed experiments with different materials including PLA, ABS, and PETG. During the tests, we realized that the PLA material was easily printable while ABS offered increased strength. On the other hand, PETG possessed high flexibility properties. The current study showed the successful design of a lightweight robotic gripper finger utilizing the principles of generative design combined with CAD and simulation

techniques. The main goal of the project was to develop an optimized design of the gripper with minimum use of material while preserving adequate stiffness and structural integrity.

Firstly, the elements of the robotic gripper were designed individually, and their correct fitting, alignment, and functionality were ensured through CAD modeling. Then, the assembly model was developed to confirm its functionality, interaction, and feasibility before proceeding to the optimization stage.In the generative design stage carried out in Autodesk Fusion 360, the design domain, fixed and free areas, loads, and constraints were defined. The goal of optimization was formulated as the minimization of the weight of the design with safety factor of 2.0 and feasible manufacturing constraints. These findings have shown that the application of generative design techniques leads to a significant increase in the stiffness-to-weight ratio of the robotic gripper finger, as it enables removal of extra material from the structure without compromising its strength.

In addition, the comparison of the ABS and Nylon materials in terms of their performance characteristics has shown that the latter material offers increased structural stability, lower displacement, and greater safety margin.

Thus, it has been proved that the combination of generative design and correct material choice is able to dramatically improve the properties and performance of robotic end-effectors. Overall, it can be concluded that the use of generative design together with additive manufacturing methods represents a promising way to produce advanced, high-performance robotic components. In this respect, the method proposed in the paper can be used as a basis for designing future generations of robotic grippers applied to industrial automation, assembly systems, and smart manufacturing facilities.

Abstract ID: 2628

Design and Experimental Evaluation of an Autonomous Water Surface Cleaning Robot

Subhradip Mukherjee[1], Rituraj Bhattacharjee[2]*

[1]*School of Computer Science and Engineering, RV University, 8th Mile, Mysuru Road, Bengaluru, India*
[2]*Department of Instrumentation and Control Engineering, MIT, Udupi - Karkala Road, Manipal, India*
**rituraj.b@manipal.edu*

Abstract

Floating debris accumulation in stagnant water bodies such as ponds and lakes poses serious environmental and public health concerns. Manual cleaning methods are labour-intensive, time-consuming, and unsuitable for continuous operation. This paper presents the design, development, and experimental evaluation of a water surface cleaning bot intended for efficient removal of floating waste in small-scale water bodies. The proposed system employs an Arduino Uno microcontroller, a differential drive propulsion mechanism for surface navigation, and a motor-driven conveyor belt for debris collection. Ultrasonic sensors are integrated for obstacle detection, while a GPS module enables basic location tracking and monitoring. The bot emphasizes low-cost implementation, mechanical simplicity, and real-time responsiveness to ensure ease of deployment and operation. Experimental results demonstrate reliable maneuverability, effective waste collection, and stable floating performance under controlled conditions. The developed prototype offers a practical and sustainable solution for water surface cleaning, with scope for future enhancements toward improved autonomy, scalability, and intelligent navigation. The real-time working of the water bot can be seen in Figure 1.

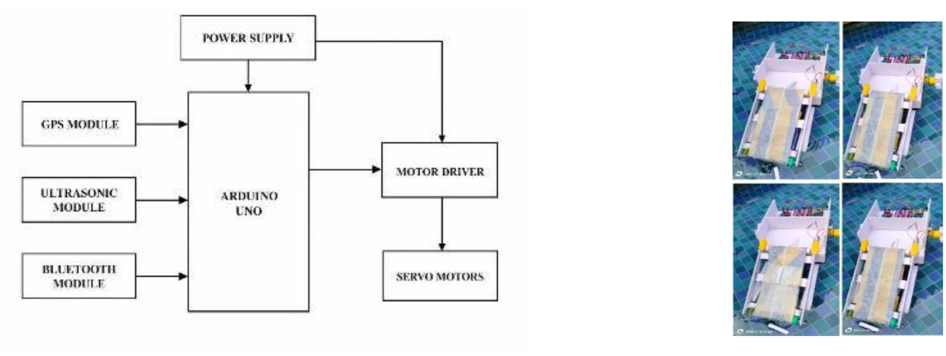

(a) (b)

Figure 1. (a) Hardware architecture of the prototype, (b) Water cleaning in real time.

The methodology is represented as Step 1- Distance measurement and path planning, Step 2- Obstacle detection and avoidance, Step 3- Autonomous conveyor belt garbage collection. The proposed water bot achieved the lowest average cleaning time of 20 minutes, representing a 55.6% reduction compared to manual cleaning (45 minutes) and a 33.3% improvement over the Bluetooth-controlled bot (30 minutes). The approximate implementation cost was also reduced to ₹15,000, compared to ₹20,000 for the Bluetooth-controlled system and ₹25,000 for manual cleaning operations. These results demonstrate that the proposed system provides a more time-efficient and cost-effective solution for autonomous water surface cleaning. Overall, the integration of autonomous navigation and obstacle avoidance significantly enhances operational efficiency while reducing human effort and maintenance costs.

Abstract ID: 2623

Module Fabrication and Mixed Reality-Based Control of a Soft Snake Robot

Rajashekhar V S*, Muhammad Imam Anugrahadi Athaaillah, Aravinth Rajesh, Gowdham Prabhakar

HIVE Lab, Department of Design, Indian Institute of Technology – Kanpur, Uttar Pradesh, India
**raja23@iitk.ac.in*

Abstract

The earthquakes are dangerous and life-threatening for humans, and they occur at unexpected times. The need for robots to be used in earthquake scenarios is being increasingly worked on by researchers around the world. In the user study done by us for the usage of robots for earthquakes, we conclude that robots would play a vital role in these scenarios. Among the robots that can be used for post-earthquake search operations, the soft continuum robots, which can be used as a manipulator and a snake robot, are a good choice. In this work, we present the soft snake robot form of the continuum robot, which can move around in the debris zone. This is shown in Figure 1. The fabrication of the soft modules is shown in Figure 2. The dive system, which is needed to actuate the soft module, was 3D printed, and the process is briefly explained. The initial mixed reality experiments were conducted to control the soft snake robot from the Meta Quest 3 Heads Up display. This would enable teleoperation of the robot in disaster scenarios.

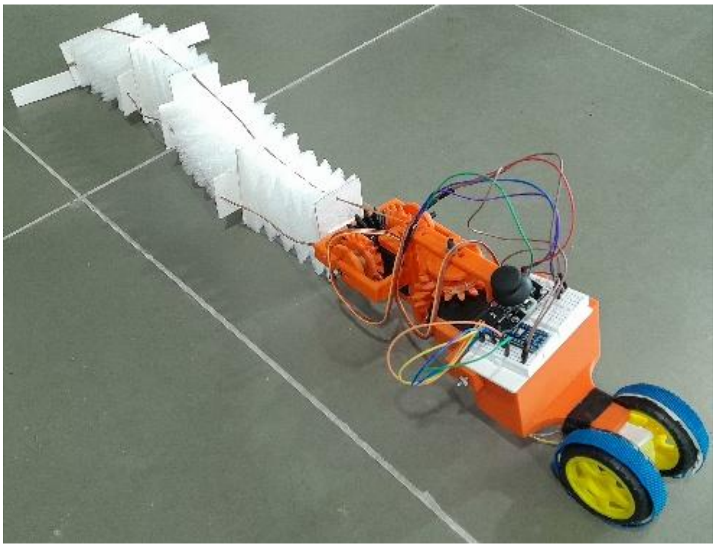

Figure 1. The Soft Snake Robot

Fabrication of Module and Drive System

The modules are made using expanded polyethylene foam (EPE), and the drive system is additively manufactured using polylactic acid (PLA). The reason for using EPE for fabricating the module is due to its lightweight nature, flexible and elastic properties, durability, shock absorption, water resistance, thermal insulating properties, easy to work with and low cost. The polylactic acid (PLA) filament was used to fabricate the drive system parts since it has good tensile strength, Young's modulus, ductility, chemical resistance, and low cost. The fabrication procedure of both the module and the drive system is explained in this section.

Module:

The 2D cross-section of the module is drawn in the FreeCAD^{TM} software in the draft mode. This is shown in Figure 2 (a). It is then fed to the laser cutting machine to cut the template on a 10mm thick acrylic sheet. This can be seen in Figure 2 (b). Then, the template of the module made in acrylic material is removed from the laser cutting machine, as shown in Figure 2 (c). The acrylic template of the module is then placed on the EPE and fixed temporarily, as shown in Figure 2 (d). Then, a paper-cutting knife is used to cut out the module manually in the EPE. This process is shown in Figure 2 (e). Finally, the EPE module is removed, and the acrylic template is removed. The final module is shown in Figure 2 (f). Then the holes for the cables to pass through are made using a paper cutting knife by poking it on the fringes. The entire fabrication process is shown in Figure 2.

Drive System

The gears were modeled using the FreeCAD™ software in the Gear workbench. Then, the pulleys were attached to it. The frames were also designed using the same CAD software. The individual components were exported in the STL (STereoLithography) format. They were imported individually into the Creality slicing software, where the G-Code was generated. The infill density was set at 30%, and the support structure generation was turned on. It was then uploaded to the Ender 5 Plus 3D printer. The material used for printing was PLA filament. It took about 36 hours for all the parts to be 3D printed. The bevel gear with pulley required post-processing since the support material was filled inside the hollow pulley region. The rest of the parts did not need much post-processing. Then, the bevel gear without the pulley was attached to the circular horn of the servo motor. The entire setup was then made, and the cables were connected to the drive system.

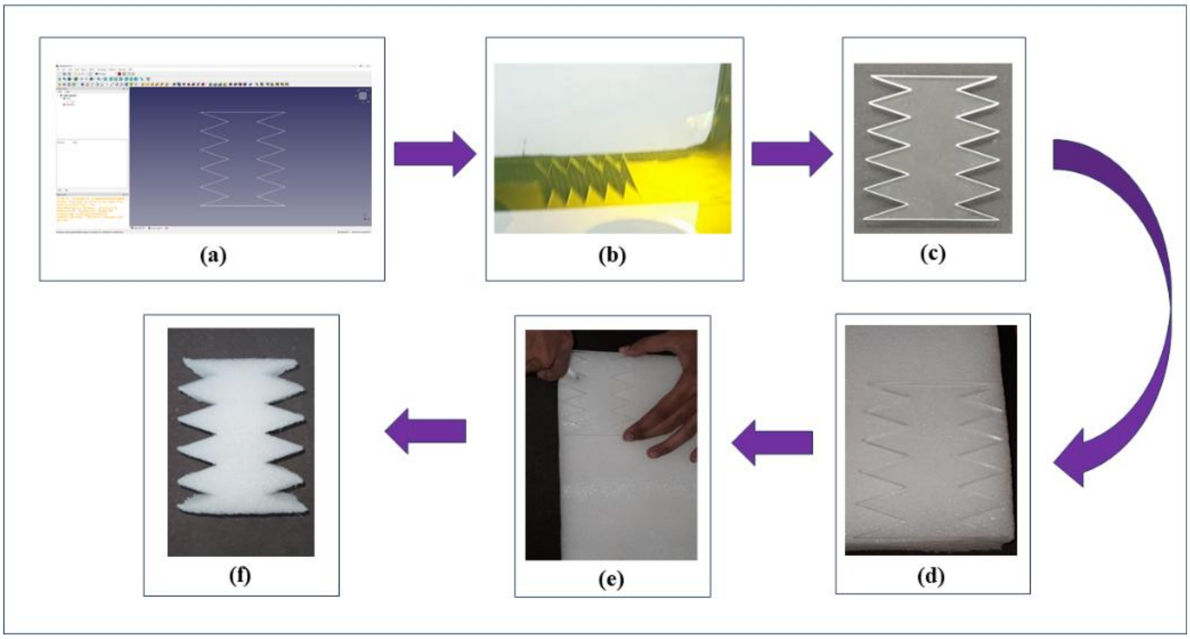

Figure 2: The stages involved in the manufacturing of the module (a) Computer-aided drafting of the module design (b) Laser cutting of the template on the acrylic sheet (c) The fabricated acrylic sheet that is to be used as a template (d) Fixing the template on the EPE temporarily (e) Cutting out the modules using a sharp paper cutting knife (f) The fabricated module.

Mixed Reality-Based Control of the Soft Snake Robot

During the deployment of the soft snake robot in the disaster scenario, teleoperation of the robot would play an important role in controlling the actions of the robot. To enable it, a camera is set up to project the view of the soft snake robot on the computer. An application was created where the two servo motors of the drive system of the soft snake robot were controlled. It was projected on the Meta Quest 3 Heads Up display. Using the hand controllers, the movement of the soft modules was controlled by adjusting the red and green sliders in the application. This is shown in Figure 3. The same setup was tested by connecting to the internet, and the functioning of the robot was tested wirelessly.

Figure 3: The mixed reality experiment in progress. The top left insert shows the point of view of the user through the VR headset. The snake robot that was being controlled is in the bottom left.

Abstract ID: 3014

Performance Analysis of Discrete-Time Sliding Mode Control with Field-Oriented Current Regulation for PMSM Servo Drive under Variable Load in Robotics Application

Anitha N[1]*, Senthil Sivakumar M[1], Velmurugan C[2], Samuel G.L[3]

[1] *Department of Electronics and Communication Engineering, Indian Institute of Information Technology, Tiruchirappalli-620012, Tamil Nadu, India*

[2] *Department of Mechanical Engineering, Indian Institute of Information Technology, Tiruchirappalli-620012, Tamil Nadu, India*

[3] *Department of Mechanical Engineering, Indian Institute of Technology, Madras, Chennai-600036, Tamil Nadu, India*

**anithan@iiitt.ac.in*

Abstract

Servo drive control systems are widely used with Permanent Magnet Synchronous Motors due to their high-performance, high-power density, greater efficiency, accurate torque control, and very high dynamic response characteristics. The servo motor experiences frequent load torque variations due to changes in payload, joint angle transitions, and gravity compensation in robotic applications such as pick-and-place systems, CNC joint drives, and surgical robotic arms. Under such challenging operating conditions, the speed control systems achieve precise speed tracking with minimal overshoot during reference transitions and rapid recovery from load disturbances. The continuous sliding mode control method makes the system follow a predetermined sliding surface, and the closed-loop dynamics become robust against disturbances and parameter uncertainties. The digital implementation of the SMC method, finite sampling rate for microcontroller and DSP-based controllers. The SMC techniques employed in various applications demonstrate slow convergence of the output and are unsuitable for servo motor precision control. A most common issue in the design of a control system is the accurate tracking of periodic reference input and efficient rejection of periodic disturbance using the DT-SMC approach. The vector control FOC technique achieves independent control of the torque and flux components of the servo motor by transforming from the synchronous stationary frame to the synchronous rotating dq reference frame. Vector control provides reliable, predictable speed tracking in steady-state operations, supports the required system dynamics, and offers greater precision with high-performance drives. However, the linear PID is sensitive to disturbances and motor parameter variations, leading to speed overshoots and long settling times during step reference transitions. In a closed loop system, speed control is achieved through accurate tracking of the reference speed with transient responses when the system in its peak and accurate steady state responses.

This paper presents the design, implementation, and performance analysis of a proposed hybrid Discrete Time Sliding Mode Controller (DT-SMC) integrated with Field Oriented Current Regulation (FOCR) for Permanent Magnet Synchronous Motor (PMSM). The Hybrid control approach comprises a discrete sliding surface in the outer speed control loop and in the inner d-axis and q-axis current loops, employs FOCR with a Proportional-Integral-Derivative (PID) controller. The high-bandwidth current loops accurately track the current reference generated from the outer speed controller. In Discrete-Time SMC, the sliding surface technique simultaneously responds to the rate of change and its error, providing nonlinear predictive damping. The mentioned features enable the system to avoid overshooting without requiring anti-wind-up compensators. The control algorithm is developed in discrete time, with finite switching frequencies, and uses a backward difference approximation of the derivative action. The entire PMSM drive system is modeled in MATLAB/Simulink using a three-phase IGBT-based voltage-

source inverter. The behavior of the proposed control scheme is compared with that of conventional Field-Oriented Control (FOC) with a PID speed regulator in the outer control loop. Testing was carried out under four conditions, constant speed and step speed references with no load torque and variable load torque disturbance conditions. The test results demonstrated that the proposed DT-SMC with FOCR achieves higher accuracy, lower overshoot, and faster settling time than traditional FOC with PID control.

The hybrid DT-SMC with FOCR and conventional FOC with PID speed control was developed and simulated using the same PMSM servo drive simulation model created on MATLAB/Simulink. The following are the simulation parameter settings used in the drive system: input voltage is 300V, a three-phase IGBT inverter with Space Vector Pulse Width Modulation (SVPWM), and an operating frequency of 10 kHz. For FOC, inner current loops and outer speed loop controllers are designed using PID controllers. For DT-SMC with FOCR, the inner current loop is controlled using a PID controller, while the outer speed loop control law is designed using a discrete-time sliding surface. Discrete-time equations for the sliding surface $s(k)$ and control law $v(k)$ are $s(k) = cx + \dot{x}$, where x is the error output, and $v(k) = cx + ksgn(s(k))$ are defined as follows: constant parameter values c =200, k =200, relay thresholds ±0.5, and output saturation ±15 A accordingly. The following four test conditions considered for estimation of the drive system's performance are: Constant speed of 1500 rpm with 0 Nm load, Constant speed of 1500 rpm with 2 Nm step at t=0.2s load, Step to 1500 rpm at t=0.2s speed with 2 N·m step at t=0.2s load, Step to 1500 rpm at t=0.2s speed with 0 Nm load.

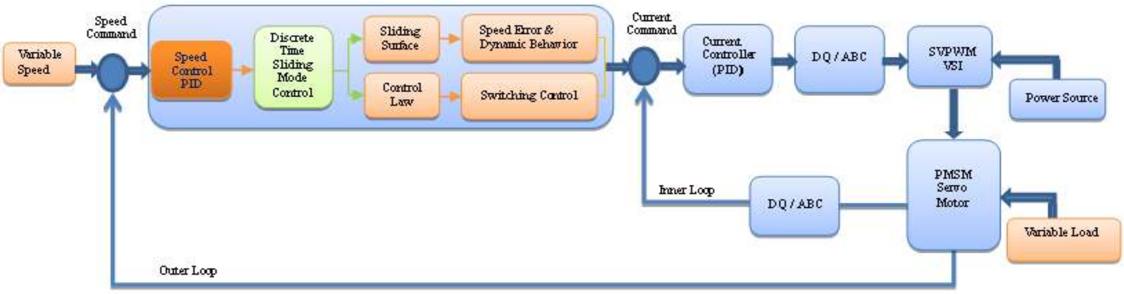

Figure 1. Proposed methodology: DT-SMC with FOCR.

In no-load transient conditions, the DT-SMC controller with no mechanical load showed an overshoot of 0.505% with 3.608 ms and 3.801 ms settling times, respectively, for constant and variable speeds. Speed response increases up to 1500 rpm, without any oscillations and also no deviations from the target value. In contrast, the FOC control displayed an overshoot of 14.368% with 36.695 ms and 36.689 ms settling times, respectively, for constant and variable reference speed. Under the Load disturbance rejection condition, a 2 Nm load is applied with a constant reference speed. Both DT-SMC and FOC give exactly the same performance as in the previous case. For the combined start-up and load condition, a varying speed reference and load are applied. DT-SMC reaches 0.502% overshoot while achieving 10.066 ms settling time within one rising period. FOC gives an overshoot of 6.186%, with a settling time of 202.592 ms. Thus, for all test cases, the proposed hybrid controller DT-SMC-FOCR demonstrates more efficient results compared to the conventional FOC based on three parameters of performance. Firstly, the overshoot, in DT-SMC it remains constant with the value between 0.502% and 0.505%, regardless of operation conditions. On the other hand, the overshoot values obtained from FOC vary greatly between 6.186% for the combined start-up and load conditions and 14.368% for no-load conditions and constant speed references. Secondly, settling time: DT-SMC settles in a very short time ranging between 3.608 ms and 10.066 ms. Whereas, the FOC controller settles between 36.689 ms and 202.592 ms. Thirdly, final speed accuracy. DT-SMC shows excellent results, giving speed in a narrow range of 1492-1501 rpm, which corresponds to 0.5% of the desired speed, even under load. FOC shows the speed range of 1491-1493 rpm, or approximately 0.6% deviation from the desired reference speed. Overall, the DT-SMC technique offers highly enhanced transient performance under all conditions without compromising steady-state accuracy and precise speed tracking control for variable load disturbance on PMSM servo drive.

Abstract ID: 3109

Development of an Autonomous Mobile Robot for Warehouse Application

Puja A. Thorat[1]*, Dr.Ajaykumar Thakur[2], Swaraj Mohapatra[3]

[1]*PhD Scholar, Sanjivani University, Kopargaon, Ahilyanagar, Maharashtra, India*

[2]*Vice Chanceller, Sanjivani University, Kopargaon, Ahilyanagar, Maharashtra, India*

[3]*Director & CEO, VCTR Labs Pvt. Ltd (Vector Robotics), Pune, Maharashtra, India*

* *thoratpujamk@sanjivani.org.in*

Abstract

Warehouse automation is rapidly growing with the adoption of Autonomous Mobile Robots (AMRs), which help in transporting materials efficiently while reducing human effort and operational time. This review paper examines various technologies and methods used in the development of AMRs for warehouse environments. The study mainly focuses on systems developed using ROS and ROS2 frameworks, along with supporting technologies such as Simultaneous Localization and Mapping (SLAM), Navigation Stack (Nav2), LiDAR sensors, computer vision, artificial intelligence, and cloud-based robotics. A review of published work from recent years shows that most AMRs are designed for accurate navigation, obstacle avoidance, and industrial material transport. Although these systems offer good performance, they are often costly, computationally intensive, and mainly suitable for large-scale industries. In addition, many existing robots do not include proper mechanisms for lifting or handling crates, especially when dealing with delicate products such as fruits and vegetables. Some studies are limited to simulation and lack real-world implementation. This paper highlights the need for a low-cost and practical AMR solution for small-scale warehouse applications. Based on the identified research gap, the proposed work aims to develop an affordable ROS2-based autonomous mobile robot equipped with a simple and safe crate-handling mechanism using a back-mounted forklift system. The robot will use economical sensors and efficient navigation techniques to reduce overall cost while maintaining good performance. With an estimated budget of around ₹1.5 lakh, the proposed system can be a useful solution for small-scale any industry like fruit and vegetable warehouses, helping improve logistics, safety, and productivity.

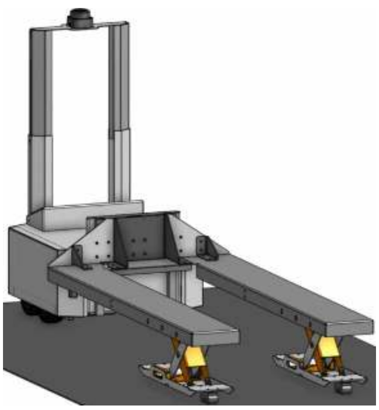

(a)

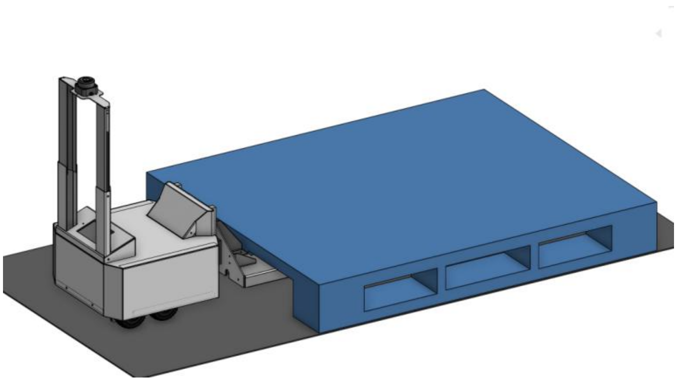

(b)

Figure 1. (a) & (b) Proposed ROS2-based AMR system implementation for small scale ware house.

Autonomous Mobile Robots (AMRs) have become a key technology in modern warehouse automation and Industry 4.0. These robots are capable of navigating dynamically without human intervention using advanced algorithms such as SLAM (Simultaneous Localization and Mapping), path planning, and sensor fusion. The Robot Operating System (ROS) and its advanced version ROS2 provide a flexible and scalable framework for developing AMRs. Over the years (2005–2026), research has evolved from basic ROS navigation systems to highly intelligent ROS2-based systems integrating Artificial Intelligence (AI), cloud computing, and multi-agent coordination. Early systems focused mainly on path planning and mapping, while modern systems emphasize real-time decision-making, robustness, and industrial deployment. Despite advancements, challenges such as high cost, computational complexity, and lack of integrated handling mechanisms still exist. In my AMR system I am going to use Back mounted Forklift mechanism which will carry crates with 15-30kg capacity and Navigation is done by ROS2. In this I have analysed pallet mover mechanism by developing a mechanism analysis app using AI tool and Onshape analysis is helped me in analysing the deformation of the pallet mover with the boundary conditions of actuator speed 7mm/s with capacity of both Linear actuators of 1500N. So that I am able to develop a hardware model of AMR which can be useful for small scale ware house application.

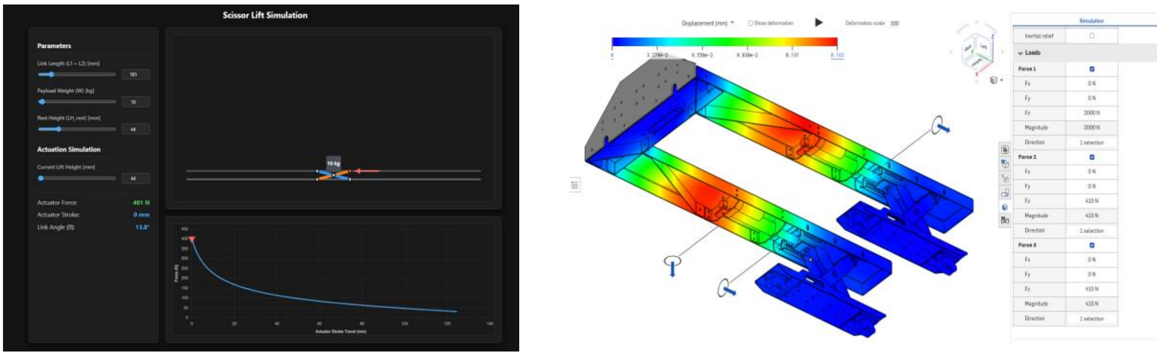

(a) (b)

Figure 2. (a) System mechanism app, and (b) deformation analysis of pallet mover.

Abstract ID: 3154

Real-Time Ergonomic Risk-Driven Human-Cobot Task Reallocation via Edge AI: An Industry 5.0 Framework with Economic Impact Quantification

Ritvik Ravi*, Govind N. Sahu

Smart Manufacturing and Automation Laboratory, Department of Mechanical Engineering, Indian Institute of Technology Tirupathi, Andhra Pradesh – 517619

**me23b044@iittp.ac.in*

Background and Motivation

Work-related musculoskeletal disorders (WMSDs) impose severe health and economic burdens on global manufacturing. Conventional ergonomic risk assessment (ERA) methods, such as RULA and REBA protocols, are reactive and preclude continuous, real-time monitoring in dynamic manufacturing cells. While recent computer vision-based ERA systems show promise, they operate as isolated detection tools: they alert workers but do not close the loop into corrective action. Furthermore, cloud-based ERA inference faces critical bandwidth bottlenecks; a multi-worker cell generates over 400 Mbps, which can exceed the practical uplink capacity of typical industrial floor networks and introduces 20-150ms cloud-round-trip latency that is unsuitable for real-time intervention. This paper proposes an edge-deployed Digital Twin framework that creates a real-time kinematic representation of the worker, integrating 3D human pose estimation, continuous ergonomic risk scoring via a fuzzy-RULA inference engine, and autonomous human-cobot task reallocation to close this loop.

Methodology

The proposed framework comprises four interdependent modules as shown in Figure 1. In this work, a ZED 2i industrial stereo camera and an NVIDIA Jetson Orin Nano running the ZED SDK's HUMAN_BODY_ACCURATE body tracking module have been used. The system is designed to reconstruct each worker's 18-joint skeletal pose at 22-30 fps. Joint angles computed from stereo-derived 3D key points are processed by a Mamdani-type fuzzy inference engine (FIE), resolving the discrete-threshold sensitivity of standard RULA to produce a cumulative Ergonomic Load Index (ELI), representing temporal exposure to ergonomic risk over the work shift. When ELI exceeds a configurable threshold, the system autonomously delegates the highest-ergonomic-demand subtask to an appropriate cobot in the adjacent cell via a ROS2-integrated MQTT scheduler. A Monte Carlo economic model (N = 10,000 iterations) parameterized from published WMSD cost benchmarks and comparable HRC deployment outcomes is used to estimate prospective annual net savings and ROI per automated work cell. By establishing theoretical viability before full-scale deployment, this framework directly advances the Industry 5.0 mandate for human-centric manufacturing, enabling closed-loop ergonomic interventions with quantified economic feasibility.

Each module of the proposed method is described below.

Module 1: Stereo vision and 3D pose reconstruction

Dual industrial stereo cameras acquire synchronized stereo frame pairs. The ZED SDK's HUMAN_BODY_ACCURATE model fuses learned stereo depth estimation with skeletal joint priors to reconstruct 18-joint 3D spatial coordinates in metric space at 22-30 fps, eliminating the need for separate 2D detection and depth-lifting steps required by monocular approaches.

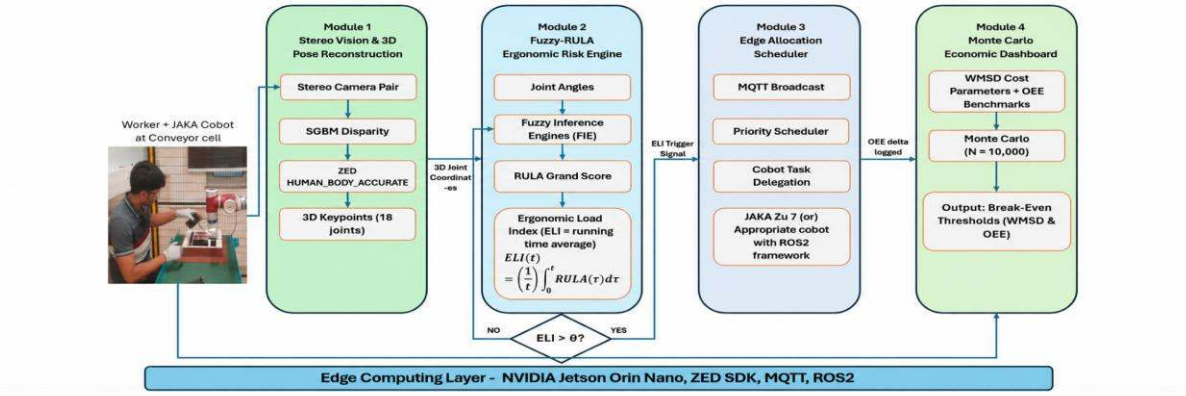

Figure 1. Proposed Industry 5.0 Edge-AI framework architecture.

Module 2: Fuzzy-RULA ergonomic risk engine

Joint angles are mapped to RULA sub-scores via a Mamdani FIE. Continuous membership functions resolve the discrete-threshold sensitivity of rule-based RULA. The per-frame grand score contributes to a running ELI:

$$ELI(t) = \frac{1}{t}\int_0^t RULA\ (\tau)d\tau \tag{1}$$

This integral captures temporal exposure to ergonomic risk, smoothing transient posture fluctuations while preserving cumulative risk trends. When ELI > θ (configurable per policy thresholds), a reallocation trigger is issued.

Module 3: Edge reallocation scheduler

Each workstation edge node broadcasts its ELI via MQTT to a local edge coordinator. The coordinator's rule-based priority scheduler identifies subtasks at the triggered workstation that satisfy kinematic, safety, and task-feasibility constraints and delegates them to the cobot's ROS2 routine. This distributed architecture eliminates cloud dependency, enabling local scaling without overtaxing the factory LAN.

Module 4: Monte Carlo economic dashboard

A parametric economic model is used to assess prospective viability. The WMSD incident rate is modelled as a Beta distribution (α=2, β=5), incident cost as LogNormal (μ = ₹1.8L, σ = ₹0.6L) per literature-informed and illustrative benchmarks, and OEE improvement as Normal (μ = 7%, σ = 2%). Annual net savings (S) per automated work-cell are computed per iteration as:

$$S = Nw * Ir * Rinc * Cwmsd. +, \Delta OEE * Vprod * Hop. - Csys \tag{2}$$

Results and Discussion

Preliminary computational profiling on the JAKA Zu 7 testbed validates the quantitative advantages of the edge-deployment strategy. The TensorRT-optimized pipeline is profiled to support the 17-joint 3D pose reconstruction at 22-30 fps locally on the Jetson Nano. This directly eliminates the 400 Mbps continuous network payload required by cloud-inference architectures and mitigates the 20-150ms round-trip latency, ensuring strictly bounded response times for the fuzzy-RULA engine.

To evaluate deployment feasibility, the parametric Monte Carlo framework (N=10,000) conducted a break-even sensitivity analysis. Assuming a conservative amortized edge-system cost of ₹1.33 lakhs/year (inclusive of dual cameras, an edge GPU, and ROS2 integration), the simulation identifies exceptionally low barriers to capital recovery. The model demonstrates that the system achieves a 1-

year payback period if the autonomous task reallocation prevents a single WMSD incident (LogNormal μ = ₹1.8L) every 16 months, a threshold significantly lower than typical industrial injury frequencies. Alternatively, independent of WMSD cost avoidance, the hardware expenditure is fully justified by an OEE productivity increase of just 0.45% in a standard 4000-hour production cycle. These results mathematically demonstrate that the framework ensures economic resilience even under minimal performance conditions, providing a low-risk pathway for Industry 5.0 adoption.

Abstract ID: 4687

Fidelity Mesh: Quantifying the Effect of URDF Geometric Fidelity on Autonomous Mobile Robot Digital Twin

M Selvakumar*, G L Samuel

Department of Mechanical Engineering, Indian Institute of Technology Madras, India
**me22d053@smail.iitm.ac.in*

Abstract

Digital twin platforms for autonomous mobile robots (AMRs) commonly approximate robot geometry with simplified collision meshes to reduce computational cost, yet the consequences of this simplification for simultaneous localization and mapping (SLAM) accuracy remain uncharacterized in the literature. This paper presents Fidelity Mesh, an experimental framework that quantifies how ten levels of URDF geometric fidelity — from primitive shape approximations (~24 triangles) to Full Fidelity Mesh (~1000,000 triangles) — affect SLAM map quality, navigation safety margins, and computational cost in a ROS Noetic–Unity 2022.3 digital twin of a Doozy AMR 500 differential-drive robot. We hypothesize that geometric under-approximation induces self-occlusion in the simulated LiDAR model, producing ghost returns that corrupt the SLAM occupancy grid specifically in the near-field region surrounding the robot, while far-field map accuracy remains largely mesh-invariant. A full factorial design (10 mesh levels × 3 SLAM algorithms × 10 repetitions, n = 300 runs) combining Unity ray-casting self-occlusion analysis, ground-truth map comparison, and navigation safety trials is proposed to test this hypothesis and identify the minimum mesh fidelity at which additional geometric detail yields diminishing returns.

Introduction

Digital twin architectures that couple a Robot Operating System (ROS) navigation stack with a real-time 3-D engine such as Unity have become a standard tool for developing and validating SLAM and navigation algorithms for autonomous mobile robots (AMRs) before physical deployment. A necessary step in constructing such a twin is supplying a Unified Robot Description Format (URDF) model of the robot, including a collision mesh used by the simulated LiDAR sensor. In practice, the geometric fidelity of this mesh is chosen ad hoc — ranging from coarse bounding-box primitives to full photogrammetric scans — with little guidance on the accuracy this choice costs or buys. This gap matters because the simulated LiDAR ray-casts against the robot's own collision mesh as well as the environment. An under-approximated mesh can create spurious self-intersections ("ghost returns") or, conversely, fail to occlude rays that the physical chassis would actually block. Either error propagates directly into the SLAM occupancy grid in the region immediately surrounding the robot — precisely the region most safety-critical for collision avoidance — while leaving distant map regions, which the robot's own geometry cannot occlude, unaffected. This research makes three contributions: (1) a reproducible ten-level mesh fidelity hierarchy spanning primitive approximation to full fidelity mesh, with a documented construction protocol for each level; (2) a controlled 3×10 factorial experimental design that isolates the effect of mesh fidelity from the choice of SLAM algorithm; (3) a Unity ray-casting self-occlusion methodology that causally links mesh geometry to LiDAR ghost returns, independent of the SLAM pipeline.

Related work

Prior work on AMR digital twins has largely treated mesh geometry as an implementation detail rather than an independent variable, with most ROS–Unity and ROS–Gazebo integrations defaulting to whatever mesh is convenient (typically a coarse collision proxy or a full CAD export) without comparing the two. Separately, the computer-vision and graphics literature on convex decomposition and mesh simplification (e.g., VHACD-style approaches) characterizes geometric fidelity loss in terms of generic shape metrics such as Hausdorff distance, but does not connect that loss to a downstream robotics task such as SLAM accuracy. The self-occlusion question has likewise been studied in the context of physical LiDAR placement and sensor-mount design, but not as a function of the fidelity of the simulated counterpart mesh used in a digital twin. Fidelity Mesh is intended to close this specific gap: a controlled, task-grounded study of mesh fidelity versus SLAM accuracy, safety, and compute cost on a single physical AMR platform.

System architecture and methodology

A. Platform and digital twin pipeline: Experiments use a Doozy AMR 500 differential-drive robot running ROS Noetic, paired with a Unity 2022.3 LTS digital twin connected via the ROS–TCP–Connector bridge. Because untethered navigation over a multi-meter test environment is incompatible with a wired Gigabit Ethernet link, autonomous navigation trials are executed with all sensor topics recorded on board via rosbag; the recording is subsequently replayed into the ROS–Unity pipeline over Gigabit Ethernet for digital-twin synchronization, decoupling navigation-time wireless latency from the synchronization-fidelity measurements reported here. Figure 1 summarizes the resulting pipeline.

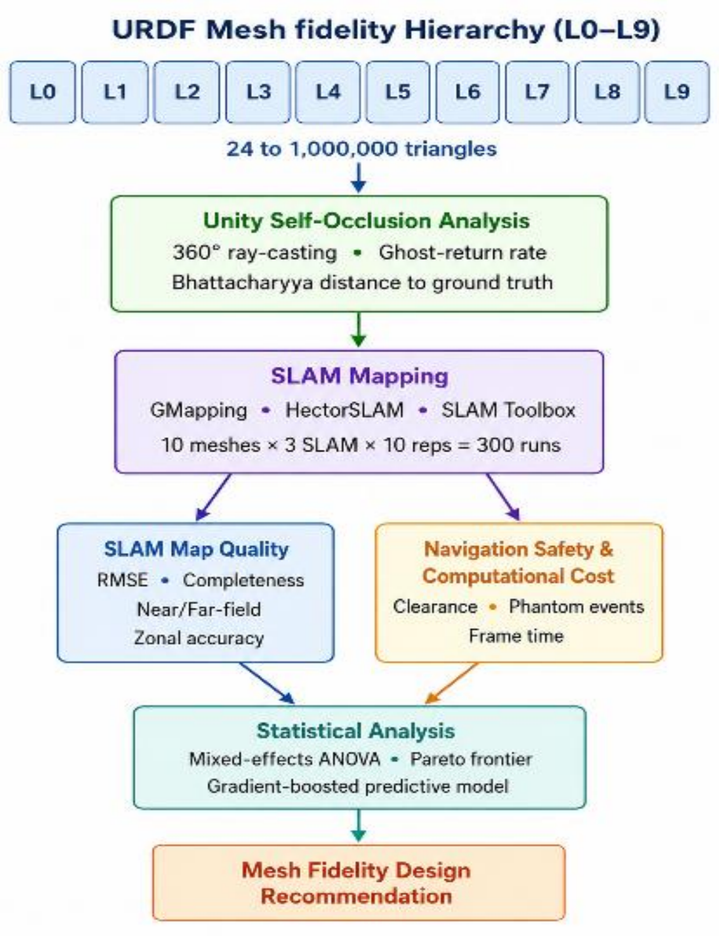

Figure 1. Fidelity Mesh experimental pipeline.

B. Mesh fidelity hierarchy: Ten mesh levels are constructed for the same physical robot, summarized in Table I. Levels L0–L2 are generated programmatically from a measured-dimension YAML specification; L3–L9 are exported from a SolidWorks CAD assembly at two STL deviation tolerances and cleaned in Blender (duplicate-vertex removal, normal correction, optional decimation) and also processed in pymesh library.

Table 1. 10-level URDF mesh fidelity hierarchy used as Factor B of the experimental design.

L0	Box / cylinder primitives	~24
L1	Bounding-box decomposition	~200
L2	VHACD convex decomposition	~2,000
L3	Low-poly CAD export	~15,000
L4	Low-fidelity CAD export	~50000
L5	Low-fidelity CAD export	~100000
L6	Medium-fidelity CAD export	~250000
L7	Medium-fidelity CAD export	~500000
L8	High-fidelity CAD export	~750000
L9	High-fidelity CAD export	~1000,000

Anticipated findings

The design is constructed to test five specific, falsifiable hypotheses, stated here as anticipated findings to be confirmed or refuted once data collection is complete: Near/far-field invariance: far-field RMSE will not differ significantly across mesh levels (ANOVA $p > 0.05$), while near-field RMSE will decrease significantly with increasing fidelity ($p < 0.001$), confirming that the effect is specifically attributable to self-occlusion rather than a general SLAM degradation. Corridor-width safety consequence: L0 will show the largest negative corridor-width error (corridors appearing narrower than reality due to phantom obstacles from oversized primitive geometry), decreasing monotonically with fidelity — connecting geometric approximation directly to an operationally relevant safety margin. Diminishing returns: the marginal near-field RMSE improvement from L2→L3 will be substantially smaller than from L0→L1, with a 90%-of-total-improvement threshold analysis identifying L5 or L6 as the practical minimum recommended fidelity.

Abstract ID: 4468

A Lightweight Web-Based Dashboard for Real Time Monitoring and Control of an Autonomous Mobile Robot Digital Twin

Monica Srinivas[1]*, M Selvakumar[2], Steven John[2], Akash K[2], G L Samuel[2]

[1]*Department of Mechanical Engineering, Vellore Institute of Technology, Vellore, India*
[2]*Department of Mechanical Engineering, Indian Institute of Technology Madras, India*
**monicasrinivas6@gmail.com*

Abstract

Digital Twin technology is a key component in Industry 4.0 and Smart Manufacturing, as it allows the virtual modeling of physical machines for monitoring, visualizing, and analyzing. AMRs are more and more used in manufacturing and logistics settings to streamline material handling and transportation. Digital Twins are a valuable tool for testing and monitoring their navigation, localization and operational capabilities before they are deployed in real-world environments. Most of the current Digital Twin platforms, however, require specific simulation software and powerful computers to operate, making regular monitoring, remote access, and operator interaction difficult.

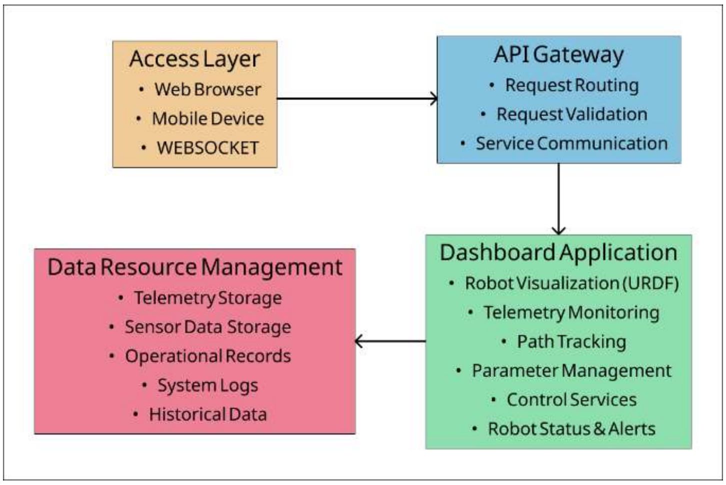

Figure 1. Implemented 4 Layer Architecture of the Web-based AMR Digital Twin dashboard.

To overcome this limitation, this work proposes a lightweight web-based dashboard to monitor and control the AMR digital twin in real-time, using NVIDIA Isaac Sim. The digital twin environment includes the integration with ROS2, realistic LiDAR and camera sensor models, realistic physics simulation with PhysX, and URDF compatibility, all of which accurately represent robot behaviour and have been accelerated by the RTX graphics programming APIs. Because running the simulation environment continuously for monitoring is a high-cost operation, the created dashboard allows users to access the digital twin remotely via a standard web browser without the need for users to directly interact with the simulation software. This way, operators can check the robot from various devices and minimize the hardware needed while increasing the usability of the robot for daily operations.

The dashboard is built with React, TypeScript and Vite. React was chosen for its easily reusable component architecture, Virtual DOM, efficient state management, and support for 3D rendering technologies available in browsers. In addition to WebSocket communication, the dashboard allows for near real-time visualization and only sends robot telemetry, navigation, joint-state and operational status information to the simulated environment, rather than the entire environment, reducing the bandwidth needed. A URDF parser is provided that uses a browser, dynamically creating the robot model that can be used to visualize different configurations of the robot without using pre-rendered graphical assets. This helps to minimize changes to the dashboard and adapt it to various AMR platforms with minimal effort, keeping the same user interface.

The four-layer architecture implemented is shown in Figure 1, which includes the Access Layer, the API Gateway, the Dashboard Application Layer and the Data Resource Management Layer. The Access Layer provides access to the platform independent via the web browser and mobile devices, and API Gateway ensures that the requests are routed and communicated between the different system services. The Dashboard Application Layer offers robot visualization, telemetry monitoring, parameter management, path tracking, robot status monitoring and control functions. The information received from the robots is used to update the URDF model shown in the browser, so that the robot can be dynamically shown in the dashboard. The Data Resource Management Layer will contain telemetry streams, sensor data, operational information, and system logs, which will be used for historical analysis, performance evaluation, and future predictive maintenance applications. The modular approach also enables the integration of further services and monitoring functions with little impact on the current dashboard structure.

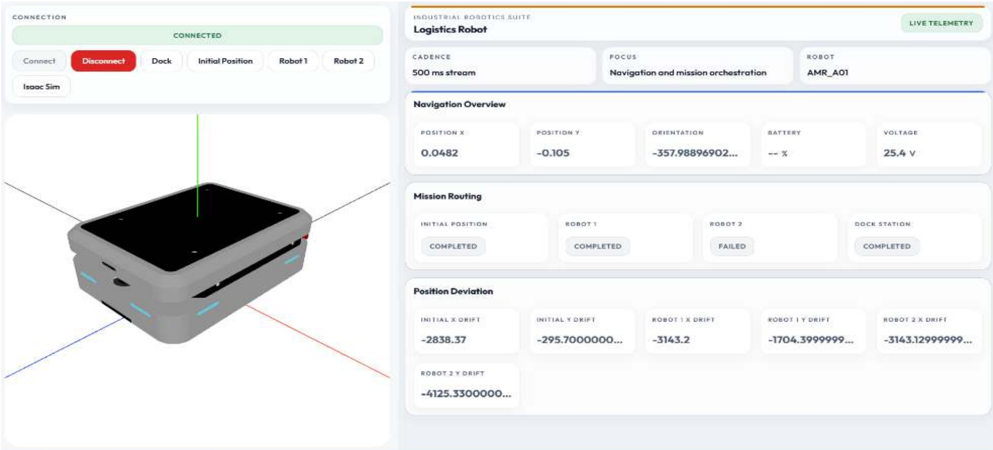

Figure 2. Web Based Dashboard for Real-Time Monitoring and Control of the AMR Digital Twin.

The designed dashboard successfully displays the real-time robot telemetry, navigation information, mission status, battery level, sensor output, and robot movement in a standard web browser, as depicted in Figure 2. Data from LiDAR sensors, maps created using SLAM, odometry data, joint state data and robot status received from the digital twin environment is fed through the communication system and shown in the dashboard in near real-time. The URDF model is implemented in a browser-based application that can reproduce the articulation and movement of a robot and gives an intuitive interface to monitor the robot's operation without the need for continuous execution of the simulation environment. Both for testing and for demonstration, the dashboard enables operators to monitor the robot's behaviour, check the navigation performance and observe the outputs of the robot's sensors from a single interface.

The developed dashboard provides a lightweight, accessible approach to monitoring and controlling AMR digital twins, decreasing reliance on time-consuming simulation software. Future work will involve adding the Authentication, Authorization, and Session Management layer to increase the security of the system, provide multi-user access, and deploy the system in the cloud securely with the help of two technologies: JSON Web Tokens (JWT) and Role Based Access Control (RBAC). The proposed framework also paves the way for the future integration of other smart manufacturing systems like CNC machines and Additive Manufacturing platforms into a single Digital Twin environment, which Multi-AMR fleet management will also be integrated to.

Abstract ID: 3117

Design and Development of a Digital Twin–Based Autonomous Mobile Robot with Real-Time Path Adaptation

Bhuvanesh P S*, Kanisha P, Vaishnavi J, Vainesh V, Arun Prakash C, Pavithra E

Department of Production Technology, Madras Institute of Technology, Anna University, Chromepet, Tamil Nadu 600044, India

*psbhuvanesh2005@gmail.com

Abstract

In today's modern world, Digital Twin technology has emerged as a key technology in contemporary industries and factories by enabling real-time synchronization between physical systems and their virtual counterparts. This study proposes the development of a low-cost, real-time digital twin system for an autonomous mobile robot (AMR) that combines virtual simulation, physical execution, and adaptive navigation. In this system, the Unity engine is used to create a virtual environment, where the paths are generated using NavMesh algorithm and transmitted to an ESP32-based robot execution.

A bidirectional data exchange is enabled by Wi-Fi-based TCP/IP communication system. Unity sends waypoint commands and the robot provides real-time position feedback using encoder-based odometry. This ensures continuous synchronization. An ultrasonic sensor is employed for obstacle detection. This will help to handle dynamic environments. Upon obstacle detection, the robot updates Unity and also triggers dynamic path replanning and execution of the revised path.

Experimental results prove path tracking, real-time synchronization, and improved adaptability to environmental changes. All these results are reliable and will be helpful for the future enhancement. The proposed system highlights the feasibility of cost-effective digital twin implementations for applications in smart manufacturing, warehouse automation, dark factories and intelligent robotic systems.

Introduction

Industry 4.0 has accelerated the adoption of Digital Twin technology in smart manufacturing environments. A digital twin is a virtual representation of a physical system. It continuously updates based on real-time data which it receives. It enables monitoring, simulation, and optimization. In industrial settings, it plays a crucial role in improving operational efficiency, reducing downtime, and allowing predictive maintenance.

Autonomous Mobile Robots (AMRs) are used in applications such as material handling, warehouse automation, and flexible manufacturing systems. However, there remains a challenge in achieving seamless coordination between physical robots and their virtual counterparts. This challenge is very high particularly in low-cost implementations.

This project addresses the above said challenge by developing a real-time digital twin system for an AMR using Unity and an embedded controller. It also focuses on achieving bidirectional synchronization between the virtual and physical environments. Thus, it enables both control and monitoring functionalities.

System Architecture and Methodology

This system consists of three main components: a virtual environment built in Unity, a physical autonomous mobile robot (AMR) controlled by an ESP-32, and a Wi-Fi-based communication system using TCP/IP.

Here is a detailed description of the system:

In Unity, the workspace and robot are modelled. The path planning is carried out using the NavMesh system. The user can select a destination and an optimal collision-free path is generated. This generated path is then sent to the ESP32, which controls the robot's movement by converting the received coordinates into motor commands for differential drive motion. The robot is equipped with motor drivers which executes the movement and wheel encoders to track its movement.

A bidirectional communication setup is used to maintain synchronization between the virtual and physical systems. In one mode, Unity sends the planned path to the robot for execution. In the other mode, the robot is controlled externally and it sends encoder-based position data back to Unity, mirroring its movements. This functionality allows the virtual model to update in real time. The continuous exchange of data creates a closed-loop system which ensures that the digital twin remains aligned with the actual robot during the entire process.

Implementation and Results

This proposed system was tested in a controlled environment. Unity successfully generated navigation paths using NavMesh. These paths were transmitted to the ESP32 for execution through Wi-Fi. The robot proved the ability to follow the paths with reasonable accuracy.

Real-time feedback was received from wheel encoders. This enabled continuous updating of the robot's position in the virtual environment, achieving dynamic synchronization. Minor discrepancies were observed due to factors such as wheel slippage and motor inconsistencies. However, the system maintained overall stability between the digital and physical models.

The results validate the feasibility of developing a low-cost digital twin system capable of real-time interaction and control.

Conclusion

This work presents the design and implementation of a real-time digital twin system for an autonomous mobile robot using Unity and ESP32. It successfully demonstrates bidirectional synchronization, enabling both control and monitoring functionalities. The integration of virtual path planning and real-time feedback establishes a foundation for intelligent and adaptive robotic systems in smart manufacturing environments.

Future Works

- Develop the robot to navigate unknown dynamic environments using SLAM and advanced sensors like LiDAR or depth cameras.
- Integrate IMU and advanced sensors to for better localization and accuracy and smoother movement.
- Implement AI-based path planning using reinforcement learning and dynamic algorithms.
- Extend system for multi-robot (swarm) coordination.

Abstract ID: 2625

Development of a Real-Time UAV Digital Twin for Cyber-Physical Monitoring and Visualization

Uttam Agarwal*, Harshit Singh, Mukunth Maithreyan K, Santosh Alluri, P Anbumani, GR Sabareesh, Sudha Radhika

BITS Pilani Hyderabad Campus, Shameerpet, Hyderabad, India

**f20230699@hyderabad.bits-pilani.ac.in*

Abstract

Industry 4.0 has accelerated the adoption of Digital Twin (DT) technologies for real-time monitoring and optimization of autonomous platforms like unmanned aerial vehicles (UAVs). Digital Twins create synchronized virtual replicas that enhance situational awareness and human-machine collaboration. This extended abstract details a real-time DT framework for a quadrotor UAV, integrating live telemetry, low-latency communication, and 3D visualization into a cohesive cyber-physical system.

Keywords: Digital Twin, UAV, MAVLink, Cyber-Physical Systems, Real-Time Visualization, Robotics, Unity,

Introduction and Motivation

Digital Twins bridge the gap between physical assets and their virtual counterparts. Monitoring UAVs with strictly real-time data is essential to ensure operational safety, avoid collisions in dynamic environments, and enable immediate corrective actions. Without real-time telemetry, the virtual model desynchronizes from reality, making it impossible to detect battery drops or spatial anomalies in time. This makes real-time synchronization a baseline requirement for safe deployment in industrial inspection and logistics tasks.

This work focuses on the implementation of a Digital Twin for a quadrotor UAV, providing real-time visualization of its state and motion. The system establishes a continuous data exchange between the physical UAV and a virtual environment, enabling synchronized representation and analysis.

System Architecture

The proposed system is structured as a three-layer cyber-physical architecture:

- **Physical UAV Layer:** Comprises a quadrotor UAV built on a Pixhawk flight controller, responsible for sensing and generating telemetry data such as position, orientation, velocity, and battery status.
- **Middleware Processing Layer:** Handles telemetry acquisition, parsing, coordinate transformation, and data conditioning using a Python-based interface.
- **Digital Twin Visualization Layer:** Implements the virtual UAV representation in Unity, enabling real-time visualization and interaction based on incoming telemetry data.

The physical system comprises a quadrotor UAV built on a Pixhawk flight controller, which provides telemetry data such as position, orientation, velocity, and battery status through onboard sensing and estimation.

A Python-based middleware, implemented using the pymavlink library, is used to interface with the flight controller and asynchronously parse MAVLink telemetry streams. To ensure smooth and stable data flow, telemetry is rate-controlled in the range of 20–50 Hz. Buffering mechanisms are incorporated to handle intermittent packet loss and maintain continuity in the data stream.

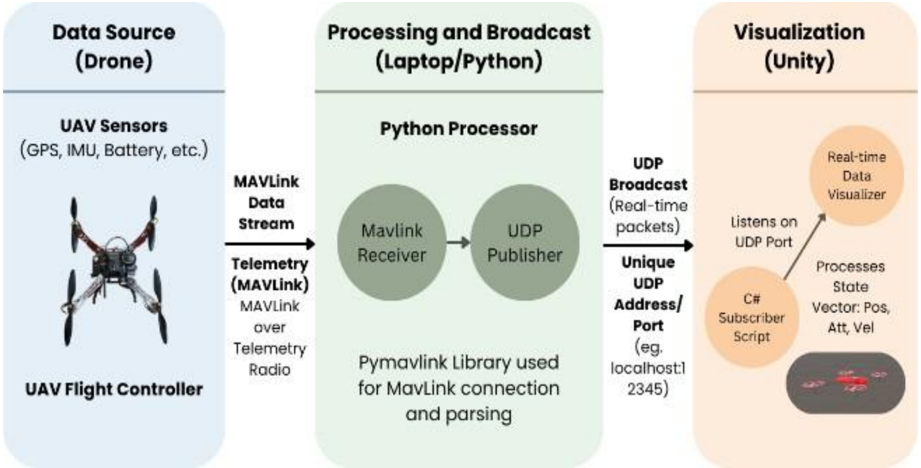

Figure 1. System Architecture.

Processed telemetry data is transmitted using a lightweight UDP communication framework. For standard indoor warehouse and shop floor environments relying on local Wi-Fi, the expected operational distance between the UAV and the remote-control station is typically 100 to 150 meters (Line of Sight). This localized communication layer ensures efficient synchronization between the physical and virtual systems with minimal computational overhead.

Digital Twin Implementation

The Digital Twin is implemented in a Unity-based 3D environment, where a virtual UAV model replicates the behavior of the physical system in real time. A UDP listener within Unity receives the incoming telemetry data, decodes it, and updates the UAV GameObject. The UAV's spatial state is represented using its Transform component, including position and orientation through quaternion-based rotations, ensuring accurate alignment with the physical system.

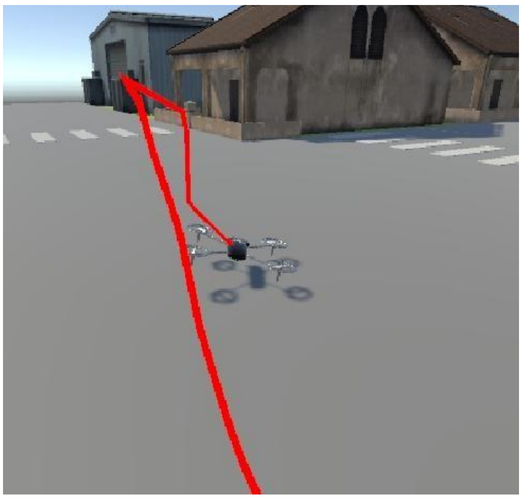

Figure 1. Digital Twin Visualization.

The visualization framework includes trajectory rendering using Unity's LineRenderer component, enabling continuous tracking of the UAV's path during operation. Additional user interface elements provide real-time monitoring of system parameters such as battery status and operational state. The system also supports offline replay using logged CSV telemetry data, allowing post-flight analysis and debugging.

System Performance and Observations

Experimental validation shows the system achieves real-time synchronization with update rates of 10–20 Hz and a baseline latency of 30–50 ms. However, latency is directly impacted by the UAV's distance from the control station. As distance increases or obstacles intervene, signal attenuation causes higher packet loss and communication delays. This increased latency can desynchronize the physical UAV from its DT, leading to inaccurate virtual representations and delayed operator responses. The implemented architecture utilizes buffering to effectively mitigate these distance-based delays and data jitter within the expected operational range, ensuring smooth visualization.

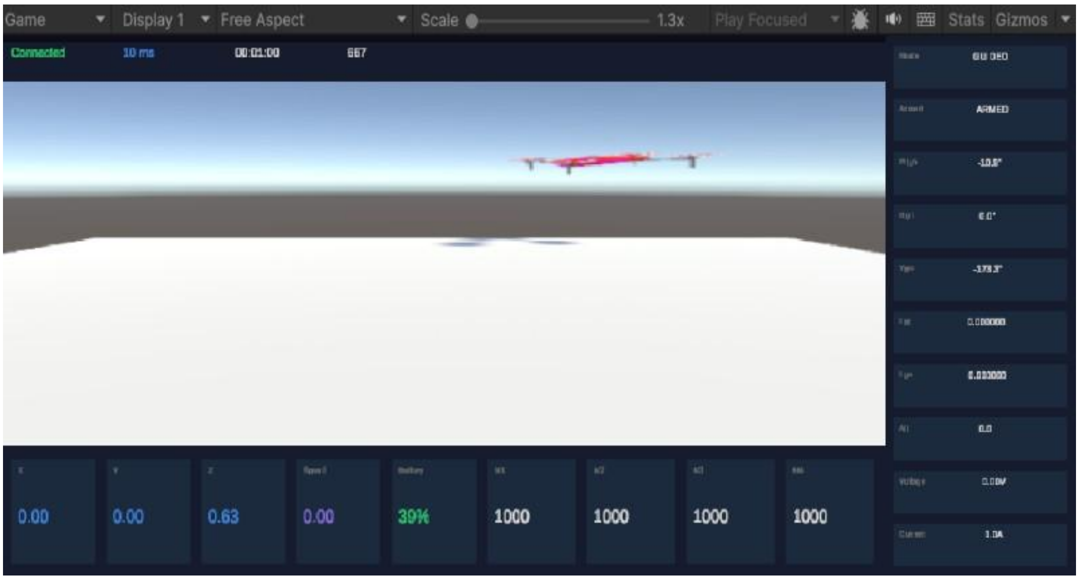

Figure 3. Performance metric.

Applications in Smart Manufacturing

This framework can be effectively implemented in large manufacturing shop floors and warehouses by programming UAVs to follow automated waypoints for high-bay rack inspection and RFID inventory scanning. In these spaces, the digital twin is highly beneficial as it provides a centralized, remote visualization of the entire UAV fleet, allowing operators to monitor flights safely away from hazardous zones. Furthermore, the twin allows operators to optimize flight paths around moving equipment like forklifts and simulate missions before physical execution, thereby guaranteeing safety and enhancing human–machine collaboration.

Conclusion

This work presents a real-time Digital Twin system for UAVs, integrating telemetry acquisition, middleware processing, and 3D visualization within a unified architecture. The system demonstrates the feasibility of creating synchronized cyber-physical representations of robotic platforms using lightweight communication and processing frameworks. The results highlight the potential of Digital Twin technologies in enhancing monitoring, analysis, and future autonomous capabilities in industrial and robotic systems.

Abstract ID: 3041

Development of a Real-Time Digital Twin Framework for Robotic Material Handling Systems in an Industrial Production Line

Velmurugan C[1*], Senthil Sivakumar M[2], Arunachalam N[3], Devasena G[4]

[1]*Department of Mechanical Engineering, Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli-620012, Tamil Nadu, India*

[2]*Department of Electronics and Communication Engineering, Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli-620012, Tamil Nadu, India*

[3]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai – 600036, Tamil Nadu, India*

[4]*Department of Computer Science and Engineering, Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli-620012, Tamil Nadu, India*

[*]*velmuruganc@iiitt.ac.in*

Abstract

A Digital Twin (DT) framework for robotic material handling establishes a real-time virtual representation of physical robotic systems, supporting industrial operations through capabilities such as virtual commissioning, predictive maintenance, and enhanced motion optimization. The transition toward intelligent manufacturing environments necessitates adaptive and data-driven material handling solutions. The rapid evolution of smart manufacturing has significantly increased the demand for intelligent, flexible, and adaptive material handling systems. Robotic technologies such as Automated Guided Vehicles (AGVs) and robotic manipulators have become integral components of modern production environments, enabling efficient intra-logistics and automation. However, conventional control strategies are often static and lack the capability to respond dynamically to real-time changes, leading to inefficiencies such as congestion, idle time, and suboptimal resource utilization. The emergence of the Digital Twin paradigm provides a promising solution to these challenges. By creating a real-time virtual replica of physical systems, digital twins enable continuous monitoring, predictive analysis, and adaptive decision-making. The combination with advanced tool of FlexSim, DT can suggestively improve system visibility and operational performance.

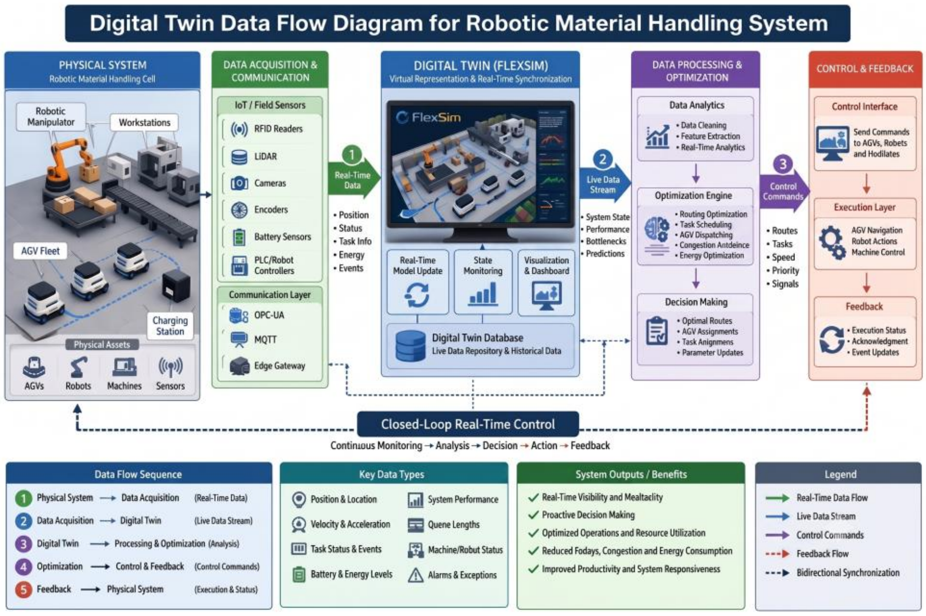

Figure 1. Real time Digital Twin data flow diagram for robotic material handling system.

This study introduces a real-time DT framework for robotic material handling systems developed using FlexSim. The proposed approach establishes a bidirectional connection between the physical system—comprising AGVs and robotic manipulators—and a virtual simulation model as shown in Figure 1. This study offers a real-time DT framework for robotic material handling systems that integrates simulation, real-time data acquisition, and optimization techniques to achieve improved system efficiency and responsiveness. Real-time operational data continuously updates the digital twin, enabling dynamic analysis and optimization of system performance. A closed-loop control strategy is implemented in this work to optimize routing, scheduling, and resource allocation in response to changing system conditions.

The experimental study was conducted to assess the performance of a real-time DT framework for robotic material handling systems. A smart manufacturing cell was modelled using FlexSim, integrating AGVs, robotic manipulators, workstations, and buffer storage units. The system consists of three processing stations—machining, assembly, and inspection—connected through AGV-based material transport. Two robotic manipulators are employed for loading and unloading operations. The layout is designed to replicate a realistic shop-floor environment with multiple routing paths and potential congestion zones. The simulation runs were conducted with the following conditions: Simulation duration- 8h; Warm-up period- 30 min; AGV speed-1.2 m/s; loading/unloading time-20-30 seconds; buffer capacity-10 min; number of AGVs-2 to 6.

A real-time Digital Twin was implemented by linking the physical system with a virtual model developed in FlexSim. System states such as AGV positions, queue lengths, and task statuses were updated at 1-second intervals through API-based communication. A closed-loop control mechanism was adopted, where real-time data continuously updated the simulation model, and optimized decisions—such as AGV dispatching and routing—were fed back to the system.

The experiments were analysed under three scenarios such as (a) Conventional System: Static routing with First-Come-First-Serve dispatching, (b) Simulation-Based Optimization: Offline optimization without real-time updates and (c) Proposed Digital Twin Framework: Real-time adaptive routing and scheduling. The performance of DT framework was evaluated with the key indicators such as: (i) Throughput (units/hour), (ii) Average AGV travel time (minutes), (iii) AGV utilization (%), (iv) Robot idle time (%) and (v) Waiting time (minutes). The results demonstrate a clear performance improvement with the proposed digital twin framework. The results demonstrated the effectiveness of the proposed method compared with the conventional and offline simulation ones. Throughput increased by about 60% (from 105 to 168 units/h), while AGVs' utilization rate improved significantly from 68% to 91%. Moreover, the average travel time decreased notably from 6.5 min to 4.1 min, and waiting time reduced from 5.8 min to 2.7 min. These enhancements were attributed to the benefits brought by the implementation of real-time synchronization, adaptive routing, and dynamic task allocation offered by the integration of a Digital Twin into the system. The experimental results confirms that integration of FlexSim with DT pointedly augments the efficiency of robotic material handling systems. The proposed framework enables real-time adaptability, improved resource utilization, and reduced system delays.

Abstract ID: 3090

Closing the Loop: A Predictive Digital Twin for GNSS-Denied UAV Navigation

Sreejith S[1]*, Tejaswini S A[1], Arpitha Gowda N[2], Gnana Vasavi G[2], Dr. Priyanka Desai[1]

[1]*Department of Information Science and Engineering, Cambridge Institute of Technology, K R Puram, Bengaluru, Karnataka, India*

[2]*Department of Electrical and Electronics Engineering, Cambridge Institute of Technology, K R Puram, Bengaluru, Karnataka, India*

**worksreejith005@gmail.com*

Abstract

Autonomous Unmanned Aerial Vehicles (UAVs) operating in GNSS-denied environments face significant challenges in reliable localization, navigation stability, and mission assurance. Conventional approaches based on Simultaneous Localization and Mapping (SLAM) and Visual-Inertial Odometry (VIO) provide real-time state estimation but remain inherently reactive and lack predictive validation mechanisms. This work proposes a closed-loop digital twin–driven framework that integrates real-time SLAM-based localization with a dynamically synchronized virtual replica to enable predictive simulation and mission validation.

The proposed system is implemented using a ROS2-based modular architecture integrated with PX4 offboard control. The navigation pipeline incorporates 3D SLAM and Extended Kalman Filter (EKF)-based sensor fusion combining vision, inertial, and optical flow inputs to achieve robust state estimation. A deterministic finite state machine governs mission-level autonomy, enabling structured navigation, target detection, and system monitoring.

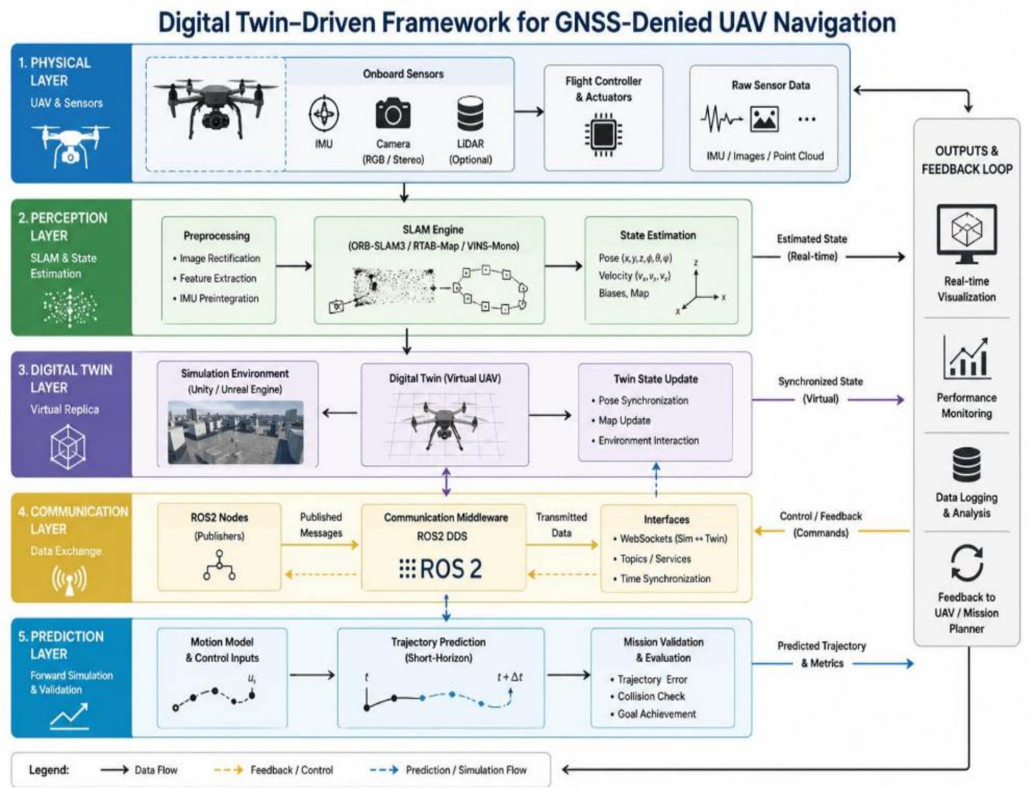

Figure 1. Architecture of the proposed closed-loop digital twin framework integrating SLAM, EKF-based state estimation, ROS2 communication, predictive trajectory forecasting, and mission validation for autonomous UAV navigation in GNSS-denied environments.

The digital twin is realized as a high-fidelity virtual environment that continuously mirrors the UAV state through low-latency communication. A bidirectional synchronization mechanism ensures consistent state alignment between physical and virtual systems. Unlike conventional SLAM-based navigation and passive digital twin frameworks, the proposed system enables predictive interaction, where the digital twin performs short-horizon trajectory forecasting using a nonlinear UAV kinematic state-transition model integrated with PX4 control inputs.

The predictive module estimates future UAV states over a defined time horizon, enabling proactive validation of navigation decisions and early detection of potential trajectory deviations. This capability enhances mission reliability and allows adaptive response in dynamic environments.

Experimental evaluation is conducted in a GNSS-denied simulation environment using AirSim integrated with ROS2, where the proposed framework demonstrates improved localization consistency, reduced trajectory deviation, and enhanced mission validation compared with conventional reactive SLAM-based navigation and PX4 offboard control without predictive digital twin feedback. Quantitative analysis using Absolute Trajectory Error (ATE) and prediction accuracy metrics confirms the effectiveness of the proposed framework. The system enables accurate short-horizon trajectory prediction while maintaining stable real-time performance.

The proposed framework improves localization robustness, predictive mission validation, and operational reliability compared with conventional reactive UAV navigation, demonstrating its potential for safety-critical applications including disaster response, infrastructure inspection, and autonomous exploration. Future work will focus on real-world deployment, multi-UAV coordination, and learning-based trajectory prediction to further enhance system performance.

Abstract ID: 3207

Real-Time FPGA Edge Integration for High-Fidelity Chatter Control

Rakesh Verma[1], Sanjayan N [2], Sivasrinivasu Devadula[1]*

[1]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India*
[2]*Department of Computer Science and Engineering, Indian Institute of Information Technology, Tiruchirappalli, India*
**devadula@iitm.ac.in*

Abstract

Machine tool productivity, particularly in operations involving flexible structures like boring bars, is fundamentally limited by regenerative chatter. While Hardware-in-the-Loop (HiL) simulators provide a safe, non-destructive platform to emulate cutting forces and test active damping strategies, they rely on instantaneous and accurate displacement data. Obtaining this displacement data by doubly integrating acceleration signals in the real-time (RT) software layer is computationally taxing and highly sensitive to external noise and DC offsets. Correcting these offsets in software requires multiple high-pass and low-pass digital filters, which severely increase the phase delay and distort the control loop. Conversely, employing physical analog hardware integrators to bypass the software burden introduces massive inherent phase lags that act as negative damping and artificially lower the stability limits of the simulated machine tool.

The experimental setup (Figure 1) consists of an electrodynamic shaker used to introduce controlled disturbances into the system, with the workpiece securely held in a clamping vice. The resulting structural response is measured and fed through a signal conditioning stage, including an amplifier, before being processed by a PC for real-time control and analysis. This configuration enables precise emulation of dynamic cutting conditions and facilitates validation of the proposed control architecture.

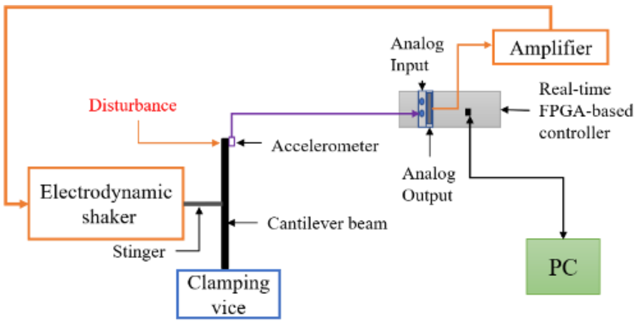

Figure 1. Schematic of the experimental setup for real-time FPGA-based control

Effective edge control is critical in regenerative chatter suppression, as even minimal latency or phase distortion in the feedback loop can destabilize the system and significantly reduce achievable stability limits. To address this challenge, this paper proposes a novel architecture that implements the double integration algorithm directly within the Field Programmable Gate Array (FPGA) fabric. By moving this critical computation to the edge, the proposed system entirely bypasses the latency and filter-heavy processing of the RT software layer while simultaneously eliminating the severe phase distortions native to analog integrator circuits. This FPGA-based edge integration allows for highly deterministic, low-latency digital signal processing. The resulting high-fidelity displacement data prevents the artificial distortion of the chatter boundaries and provides a highly robust, ultra-fast feedback loop. Consequently, this architecture significantly enhances the real-time edge control of regenerative chatter, allowing active vibration suppression strategies such as Direct Velocity Feedback (DVF) to operate at much higher sampling frequencies and broader bandwidths without degradation.

Abstract ID: 3200

Adaptive Maintenance of Wind Turbine Gearboxes Using Digital Twin-Inspired Simulation

Harshvardhan Kharat*, Pranav M N, Amindra Khilla, Piyush Shakya

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai 600036, India
**me23b174@smail.iitm.ac.in*

Abstract

Wind turbine drivetrain reliability remains a major operational challenge, even though wind power is one of the fastest-growing renewable energy sources worldwide. Since the gearbox is a demanding and failure-prone component, it accounts for a large share of unscheduled downtime and maintenance costs. A single malfunction can cost over $100,000 in repairs and cause weeks of lost energy production. Conventional maintenance techniques, such as reactive (fix-on-failure) or calendar-based replacement, don't consider the machine's actual condition, often leading to unnecessary costs or avoidable failures. Although vibration-based condition monitoring has progressed, most research relies on physics-only, non-data-driven models or single-axis signal analysis. At present, no comprehensive framework exists for wind turbine gearboxes that integrates physics-informed structural simulation, AI/ML-based prognosis, and multi-domain signal processing within a digital twin paradigm.

The proposed three-tier digital twin-inspired approach is used for monitoring the adaptive gearbox health. $[a_x(t), a_y(t)] = x(t)T$ is the biaxial vibration acceleration vector of a three-stage planetary gearbox with a gear ratio of 80:1 at sampling rate f_s. The formula is $f_{GMF} = N_t \cdot f_{shaft}$, where N_t is the number of gear teeth, and f_{shaft} is the shaft rotation frequency. Remaining Useful Life (RUL) ranges from 0 to 1, with RUL = 0 indicating failure and RUL = 1 indicating health. The process has three steps: (1) Multi-domain signal analysis: To identify fault indicators, the tangential (a_x) and radial (a_y) channels are analyzed using FFT, PSD, STFT, and CWT. (2) Data-driven prognosis: A two-layer LSTM network (128 and 64 units, dropout, Adam optimizer) uses ten statistical features (RMS, standard deviation, kurtosis, skewness, peak amplitude, among others) extracted over 2048-timestep sliding windows for continuous RUL prediction. A composite Health Index (HI) is also developed for threshold-based maintenance. (3) Physics-informed FEA: Reaction forces and torsional compliance under applied moments are obtained using an ANSYS finite element gear mesh simulation as a structural validation layer. The main innovation is combining physics-informed simulation with LSTM-based prognosis—a digital twin-inspired method not yet demonstrated for wind turbine gearboxes. To improve prediction accuracy and generalization, future work will evaluate various AI/ML models. The reported RMSE and MAE are influenced by the small dataset used, and future work will use a larger run-to-failure dataset for better validation and more reliable results.

Key results from signal analysis, LSTM prognosis, and FEA structural simulation are summarised in Tables 1 and 2 below.

Table 1. Signal Analysis & Fault Identification Findings.

Method	Key Finding
FFT	By FFT analysis from 1100 RPM gear speed, Gear Mesh Frequency (GMF) is 306.8 Hz, which is the first significant peak for a 17-tooth gear. Structural resonance harmonics are observed at 4000, 8000, 16000 and 24000 Hz, with 16000 Hz harmonic amplitude being 4 times greater than the GMF amplitude, which indicates significant gear degradation through increased mesh impact forces.
PSD	Dominant low-frequency energy (0–1000 Hz) in a_x → shaft/bearing rotation. Sideband clusters around GMF in ay (radial channel) indicate amplitude modulation → gear eccentricity or distributed wear. The Gear eccentricity causes sideband clusters around the GMF in the radial channel (a_y) to exhibit distributed wear or amplitude modulation.
STFT	Horizontal energy bands at ≈ 4000 Hz and ≈ 8000 Hz are time-stationary over the entire observation window. The absence of vertical streak impulses rules out localised defects such as pitting and spalling. The distributed fault signature is compatible with wear or misalignment.
CWT	The shaft-rate amplitude fluctuates because of periodic energy bursts in the a_x direction at around 500 Hz. Subtle secondary impulsive excitations are produced by randomly distributed bursts in the 3–5 kHz range. A pattern that is consistent with either gear misalignment or eccentricity (periodic load change).

Table 2. LSTM Prognosis & FEA Simulation Performance Metrics.

LSTM-Based RUL Prognosis	
Mean Absolute Error (MAE)	0.0229
Mean Squared Error (MSE)	0.0009
Health Index — Warning threshold	HI < 0.3 (triggered at ~window 310)
Health Index — Critical threshold	HI < 0.1 (triggered at ~window 430)
RUL tracking fidelity	Full degradation curve (0.93 → 0) tracked; spike at index ~1400 recovered
Physics-Informed FEA Simulation (ANSYS)	
Load–deformation behaviour	Linear elastic (Hookean); no plastic deformation observed
Tangential force F_y range	25 N → 249 N (primary load transfer; dominant force component)
Torsional compliance (slope)	Z-rotation 0.01° → 0.099° over 100–1000 N·mm; confirms elastic stiffness

Abstract ID: 3170

Physics-Guided Neural Network for Cutting Force Prediction and Chatter Identification in Turning

Abhishek Punia*, Rinku Kumar Mittal

Department of Mechanical Engineering, Indian Institute of Technology Guwahati, Guwahati, Assam, India
[*]a.punia@gmail.com

Abstract

Accurate prediction of cutting forces and chatter identification is necessary for process optimization, improved productivity, and the development of intelligent machining systems. Traditionally, mechanistic time-domain models have been widely used to accurately predict cutting forces and facilitate chatter identification. These models accurately capture the underlying nonlinear dynamic physical mechanisms involved in cutting processes, such as the effect of regeneration on chip thickness and structural dynamics. Hence, these models generalize very well across different cutting conditions. However, small discrepancies occur between the simulation predictions and the experimentally measured responses. This difference mainly arises from simplifications and assumptions, as well as from neglected factors such as tool wear, changes in structural dynamics during machining, thermal effects, material variability, and tool nose radius, in these mechanistic models. The challenge of integrating all practical effects to achieve more precise responses can be addressed through the utilization of data-driven models based on machine learning. However, purely data-driven machine learning models require large amounts of data, lack interpretability, and generalize poorly to unseen cutting conditions and setups. This work proposes a physics-guided neural network model for the accurate and reliable prediction of cutting forces and chatter, leveraging the advantages and addressing the limitations of both mechanistic time-domain simulation models and purely data-driven models.

The proposed model learns the discrepancy between the mechanistic time-domain simulated cutting force and the experimentally obtained cutting force by learning the residual between the two forces. However, the major challenge in learning directly from force signal is that simulated and experimental signals are generally not synchronized due to differences in spindle phase, initial conditions, and transient behavior. Instead of learning directly from raw force signals, the proposed framework first converts the simulated cutting-force response into a compact physics state representation that describes the static and dynamic behavior of the machining process. Since these descriptors summarize the overall characteristics of the force signal rather than individual time samples, they naturally eliminate the synchronization requirement while preserving the dominant machining physics.

The proposed framework consists of four stages. A mechanistic time-domain simulation is first performed to generate the dynamic cutting force response. At each simulation step, the instantaneous uncut chip thickness is calculated using the regenerative surface profile from previous spindle revolutions. The cutting force components are then computed using nonlinear cutting force coefficients, while the dynamic response of the machine-tool system is represented by the second-order equation of motion,

$$\ddot{x} + 2\zeta\omega_n\dot{x} + \omega_n^2 x = \frac{\omega_n^2}{k}F(t)$$

solved by using a fourth-order Runge–Kutta (RK4) integration method. The updated tool displacement is subsequently used to modify the chip thickness in the next time step, thereby capturing the regenerative effect responsible for chatter. This iterative procedure produces the simulated cutting force signals and corresponding vibration response for the specified machining conditions.

In the second stage, the simulated force signal is transformed into a compact physics state representation using a set of physically meaningful descriptors that characterize both the static cutting load and the dynamic response of the machining process. This low-dimensional representation preserves the essential information required for force prediction and stability analysis while significantly reducing the sensitivity to temporal misalignment between simulation and experiment. In the third stage, the physics state representation, together with the machining parameters, is provided as input to a physics-guided neural network. Since the mechanistic simulator already captures the dominant cutting physics, the network learns only the descriptor discrepancy between simulation and experiment.

Finally, the corrected physics state is used for two downstream tasks. First, the corrected descriptors are employed to estimate the experimental cutting-force characteristics. Second, the corrected dynamic descriptors provide physically meaningful stability indicators that can be supplied to conventional machine-learning classifiers, such as Support Vector Machines (SVMs), Random Forests, or similar algorithms, to distinguish stable, transitional, and chatter machining conditions.

The proposed framework combines the interpretability of mechanistic modeling with the adaptability of machine learning while requiring only limited experimental data. By learning only the small discrepancy between simulation and experiment, the framework improves cutting-force prediction, provides reliable stability indicators, and eliminates the need for direct synchronization of simulated and experimental force signals.

Abstract ID: 3150

A Three-Layer Cyber-Physical System for a Boring Tool to Monitor and Predict Chatter During Machining

Ananthakrishna Ayankalath Thekkepat [1], Rakesh Verma [2], Sivasrinivasu Devadula[2]*

[1]*IITM Pravartak Technologies Foundation, IITM Research Park, Chennai, India*

[2]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India*

**devadula@iitm.ac.in*

Abstract

The boring operations are inherently susceptible to self-excited vibrations, commonly known as chatter, which can degrade surface finish, accelerate tool wear, and potentially damage machined components. The real-time insight into the dynamic behaviour of such tools remains limited. On the other hand, to support the present digital transformation aimed at realising Factories of the Future, it is necessary to adopt Smart manufacturing concepts on the shopfloor.

This paper presents a Cyber-Physical System (CPS) framework for a boring tool, with the objective of real-time monitoring and prediction of machining vibrations at the Tool Centre Point (TCP). The proposed CPS architecture addresses the gaps highlighted by enabling continuous tracking of vibration signatures and early detection of chatter phenomena. The developed system is structured within the CPS architecture (Figure 1), comprising three primary layers: physical, edge, and cloud layers. The physical layer consists of the damped boring bar integrated with appropriate sensing elements, such as accelerometers or vibration sensors, along with a data acquisition system. This layer captures high-frequency vibration data directly from the machining environment under varying cutting conditions. The acquired sensor signals are transmitted to the edge layer, where initial data processing and analysis are performed. In this work, the edge layer is implemented on a local computing system using Simulink, enabling real-time signal conditioning, filtering, and feature extraction. Advanced signal processing techniques are employed to derive meaningful indicators of tool vibration and stability, such as amplitude, frequency content, and stability metrics. By processing data at the edge, latency is minimized, allowing faster detection of undesirable dynamic behaviour during machining. The processed information is subsequently communicated to the cloud layer, which serves as the central hub for data visualization, storage, and higher-level analytics. The cloud platform provides an intuitive interface for monitoring the vibration characteristics of the boring bar in real time. It also incorporates predictive capabilities for chatter detection based on predefined thresholds or learned patterns. In addition, the system generates alerts to notify operators of imminent instability and offers recommendations for parameter tuning, such as adjustments in spindle speed or feed rate, to maintain stable cutting conditions. The digital twin (DT) in the cyber layer of the proposed CPS architecture not only mirrors the physical state of the damped boring bar but also enhances decision-making through predictive insights. The DTs have emerged as a cornerstone of modern smart manufacturing systems, enabling seamless integration between physical assets and their virtual counterparts. By providing a dynamic, real-time digital representation of physical systems, digital twins facilitate continuous monitoring, predictive analysis, and remote control of manufacturing operations. This capability is particularly valuable in high-precision machining processes, where tool performance and stability directly influence product quality, productivity, and operational efficiency.

The proposed framework demonstrates the potential to improve machining performance, reduce downtime, and extend tool life by enabling proactive control strategies. Furthermore, the modular nature of the system allows for scalability and adaptation to other machining processes and tool configurations.

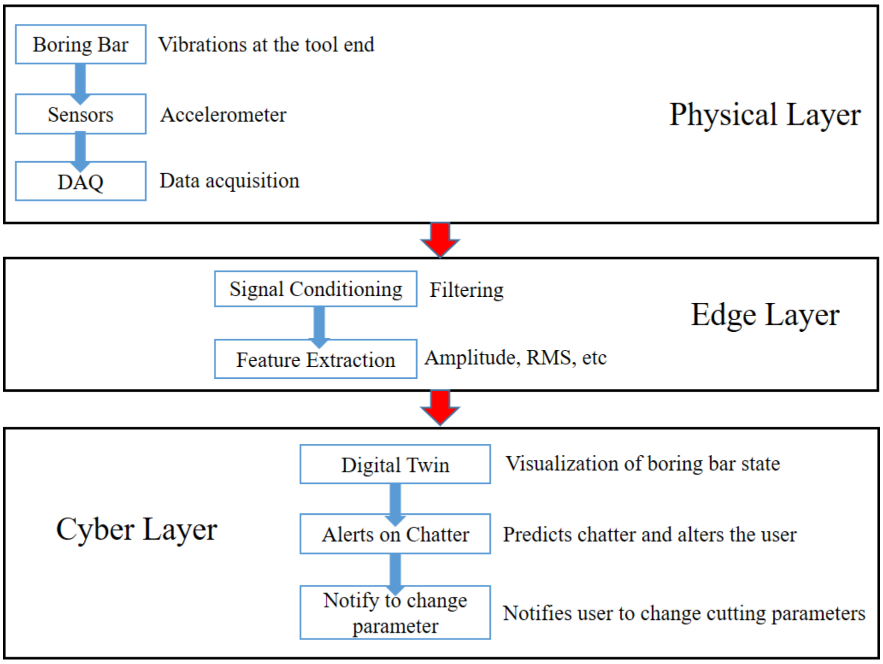

Figure 1. CPS architecture to predict and monitor chatter during machining.

Overall, this work contributes to the advancement of intelligent machining systems with real-time vibration monitoring and chatter prediction, paving the way toward more autonomous and resilient manufacturing environments.

Abstract ID: 3775

A Hybrid Contextual Reasoning Architecture for Fully Offline Conversational AI on Resource-Constrained Edge Devices

Karthik M[1]*, Kiran A V[1], Om M Doddamani[2], Kishor A V[1], Jeevan Nullipady Mallya[1], A B Rajendra[1]

[1]*Department of Computer Science & Engineering, Vidyavardhaka College of Engineering, Mysore 570017, India*

[2]*Department of Information Science & Engineering, Vidyavardhaka College of Engineering, Mysore 570017, India*

**vvce24cse0471@vvce.ac.in*

Abstract

Deploying capable conversational AI systems on edge devices is highly challenging due to severe limitations in memory, processing power, and energy. Most existing solutions depend on cloud connectivity, which compromises data privacy, increases latency, and reduces reliability in industrial and remote environments. This paper presents PRISM, a novel hybrid architecture designed to enable sophisticated, fully offline conversational AI on a commodity Raspberry Pi 4B with only 4 GB of RAM. The architecture consists of four tightly integrated components. The Contextual Reasoning Layer (CRL) performs deterministic intent classification and query planning using lightweight rule-based signals in under 3 ms. The Layered Memory System (LMS) maintains short-term conversational history in RAM and promotes important episodes to a long-term SQLite database based on a computed importance score. The Hybrid Semantic Retrieval (HSR) module combines BM25 term weighting with bigram co-occurrence scoring for effective local knowledge retrieval, significantly outperforming traditional TF-IDF. Finally, the Adaptive Shape Prompting (ASP) component dynamically adjusts response persona, verbosity, and hedging style. The system uses TinyLlama-1.1B through Ollama and incorporates real-time environmental sensing via a DHT22 sensor.

PRISM follows a deterministic four-layer pipeline in which every user query first passes through the Contextual Reasoning Layer (CRL) for intent classification into one of five categories: factual, conversational, analytical, environmental, and recall. Based on the identified intent, relevant context is retrieved from the Layered Memory System (LMS), which combines short-term conversational memory with long-term episodic storage, and the Hybrid Semantic Retrieval (HSR) module, which accesses the local knowledge base. The Adaptive Shape Prompting (ASP) component then constructs an optimized prompt before forwarding it to the TinyLlama-1.1B model running through Ollama. This hybrid design keeps computationally intensive operations lightweight until the final inference stage, enabling efficient execution on the Raspberry Pi 4B. Extensive evaluation was conducted using 600 diverse queries spanning all intent categories, with responses independently assessed by three human evaluators (inter-rater $\kappa = 0.81$). PRISM achieved a multi-turn contextual accuracy of 0.884 at a conversation depth of 12, representing a 22.6% improvement over a strong RAG-only baseline and a 46.4% improvement over vanilla TinyLlama-1.1B. The HSR module improved Precision@1 by 16.8 percentage points compared with TF-IDF while maintaining end-to-end response latency below 3 seconds for complex analytical queries. Peak RAM usage remained limited to 510 MB on the Raspberry Pi 4B, demonstrating the framework's suitability for resource-constrained hardware. Environmental sensing using the integrated DHT22 sensor also showed high accuracy, with mean absolute errors of 0.19°C for temperature and 0.53% for relative humidity. These results demonstrate that sophisticated reasoning, memory management, and retrieval capabilities can be achieved entirely offline, providing a practical foundation for privacy-preserving AI deployment in industrial and embedded environments.

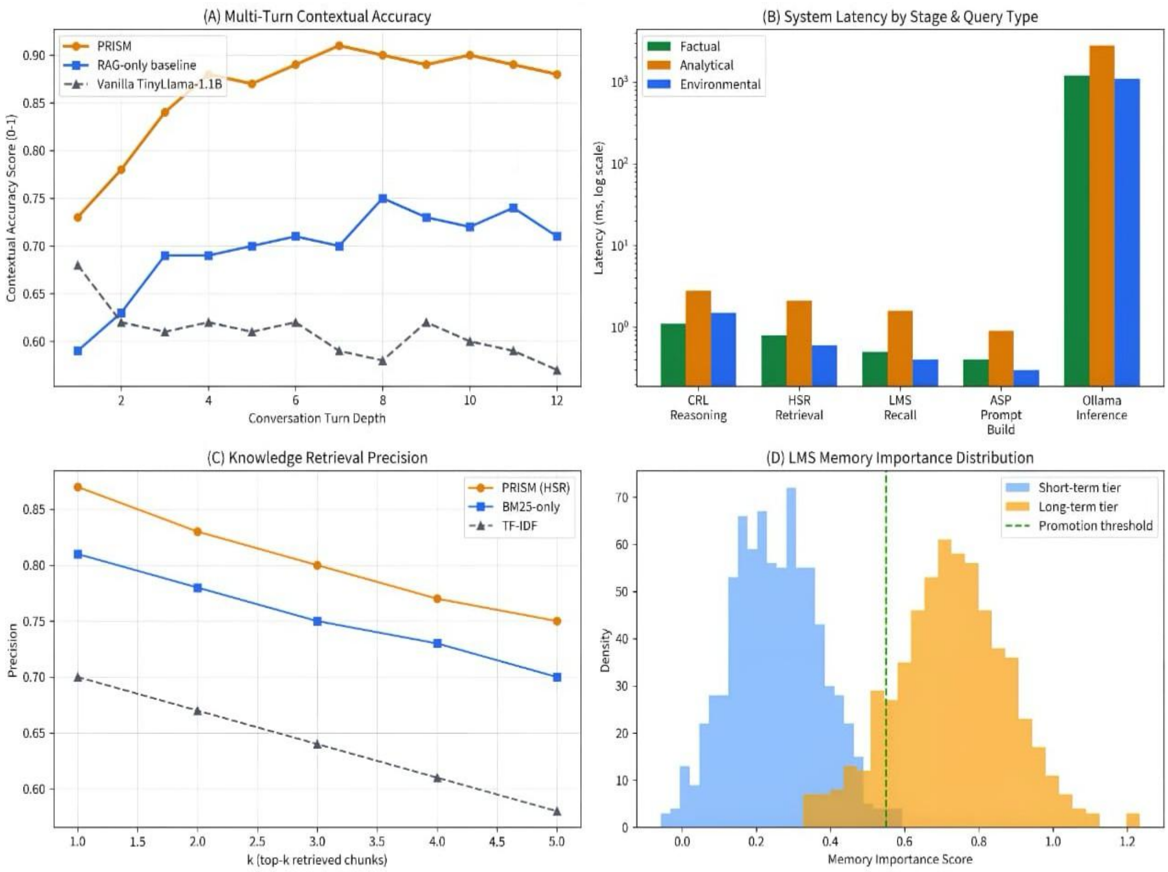

Figure 1. Multi-turn contextual accuracy (A), system latency by stage and query type (B), knowledge retrieval precision (C), and LMS memory importance distribution (D) for PRISM versus baseline systems.

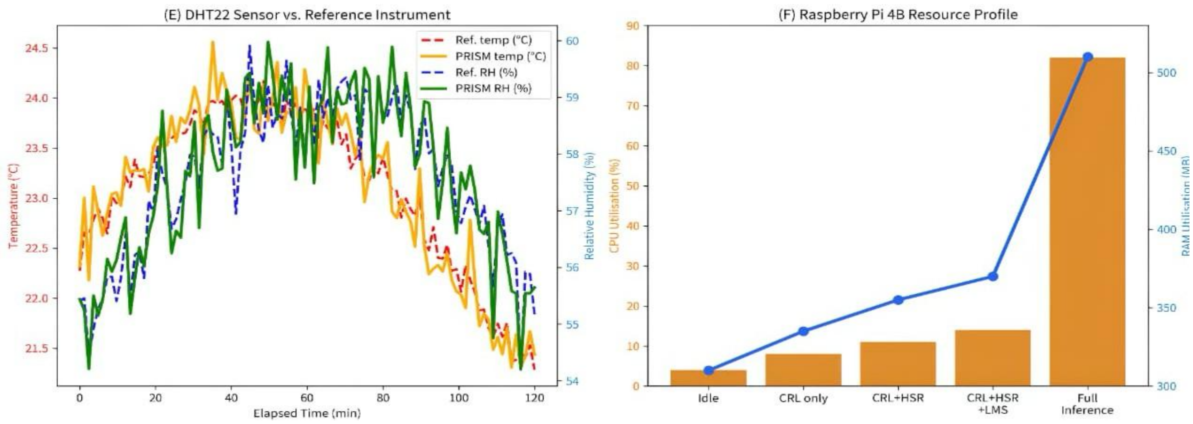

Figure 2. DHT22 sensor validation against reference instrument (E) and Raspberry Pi 4B resource utilization profile across pipeline stages (F).

Abstract ID: 2664

SMRITI: A Smart Manufacturing Reasoning Interface for Twin Integration using LLM-Enhanced Explainable AI

Rahul Kumar Ray*, Chiranjoy Chattopadhyay#

School of Computing and Data Sciences, FLAME University, Pune, Maharashtra, India

*rahul.ray@flame.edu.in

#chiranjoy.chattopadhyay@flame.edu.in

Abstract

Conventional Digital Twins (DTs) in smart manufacturing are primarily designed for passive real-time visualization and anomaly detection. However, most existing frameworks fail to explain the underlying causes of defects or provide actionable, on-the-fly remediation strategies, creating a significant "action gap." This paper proposes SMRITI (Smart Manufacturing Reasoning Interface for Twin Integration), a pragmatic framework that bridges the divide between perception and shop-floor intervention. Utilizing overhead conveyor camera feeds, SMRITI integrates: (1) real-time defect localization, (2) causal explanations of system failures, and (3) an integrated Large Language Model (LLM) that translates complex diagnostic data into simple, actionable instructions displayed on shop-floor dashboards. By shifting the focus from simple detection to transparent, real-time operational guidance, the SMRITI framework provides a highly deployable, human-centric solution for intelligent manufacturing and localized industrial upskilling.

Introduction

India's Industry 4.0 transformation requires manufacturing environments that not only detect errors but also actively assist the workforce in mitigating them. While Digital Twins (DTs) have been widely adopted for monitoring, they typically stop at the dashboard as passive observers. They highlight anomalies but do not explain their causes or guide remediation in real time. This limitation is acute in Indian MSMEs, where rapid decision-making on the factory floor is essential, yet expert diagnostic trainers are scarce. Recent vision-based methods achieve accurate defect detection on moving metallic surfaces, utilizing models such as NSLNet for environments with limited training datasets and FDRLNet for rapid defect localization. Simultaneously, Explainable AI (XAI) approaches like HG-XAI generate visual and textual justifications for machine wear. However, these systems often lack a user-friendly interface tailored to the fast-paced reality of a plant floor, where operators cannot rely on cumbersome XR headsets or complex data logs. To bridge this gap, we propose SMRITI. The core novelty of this interface is the integration of a Large Language Model (LLM) that acts as a real-time operational assistant. By processing live camera feeds and XAI outputs, the LLM generates instant, plain-language guidance on shop-floor monitors, transforming a passive monitoring system into an active daily operational tool.

Proposed SMRITI framework

The SMRITI framework operationalizes the transition from passive monitoring to active shop-floor assistance through three integrated layers:

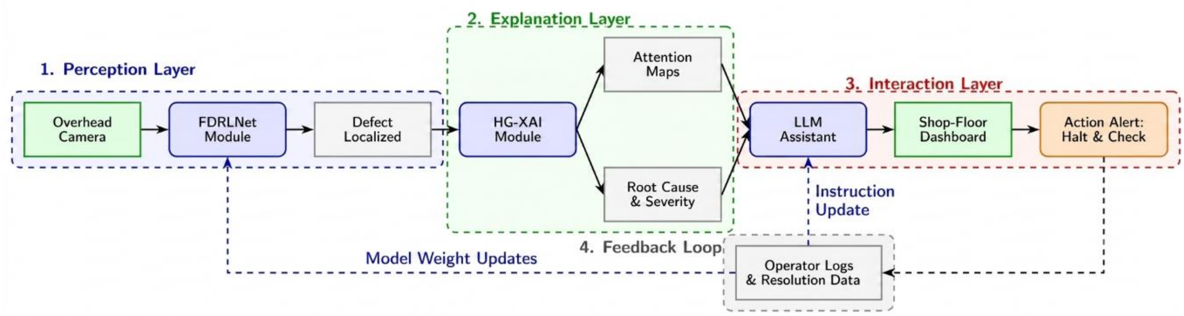

Figure 1. The SMRITI Framework Pipeline. Real-time defect detection and causal reasoning are processed by an LLM to generate instant, actionable remediation alerts for shop-floor operators.

Perception layer

Continuous Real-time Monitoring: This layer utilizes overhead cameras mounted above manufacturing conveyor belts for continuous visual monitoring. For steel surface defect localization, we adopt FDRLNet, which achieves fast recognition and bounding-box-level localization capable of keeping pace with line speeds.

Explanation layer

Causal XAI Generation: Raw detection outputs are passed to an Explainable AI module based on HG-XAI. This module generates visual attention maps highlighting the exact defective regions on the passing component, alongside structured technical descriptors detailing severity and probable root causes (e.g., roller misalignment or tool wear).

Interaction layer

LLM-Powered Dashboard Guidance: To make the XAI outputs immediately useful to a line operator, the framework utilizes an LLM-powered translation layer. (1) Operational Translation: The LLM ingests the highly technical HG-XAI descriptors and instantly translates them into concise, actionable alerts. (2) Shop-Floor Deployment: Instead of requiring wearable technology, these insights are pushed to ruggedized shop-floor displays or supervisor tablets. For example, rather than simply displaying an "Inclusion Type A" error, the dashboard provides a highlighted image feed alongside LLM-generated text: "Surface defect detected. High probability of debris on Roller 3. Recommend brief halt to clear roller before resuming."

Continuous feedback loop

As operators log the successful resolution of these LLM-guided prompts, the system records the data to refine both the confidence thresholds of the perception layer and the operational accuracy of future LLM instructions.

Proof-of-Concept: LLM-driven shop-floor dashboard

To validate the interaction layer, we designed a proof-of-concept dashboard interface tailored for ruggedized plant-floor tablets. The interface avoids complex data logs, instead utilizing a three-panel vertical stack: (1) Live Feed Simulator: Displays the continuous camera feed with FDRLNet bounding boxes highlighting detected anomalies (e.g., Surface Inclusions). (2) Causal Network View: Presents a simplified HG-XAI diagnostic tree linking the defect to machine parameters (e.g., Slag Entrapment), and (3) LLM Operator Guidance: The critical interaction zone where the LLM translates the above technical data into high-contrast, real-time actionable alerts. For example, rather than displaying "Inclusion Type A," the LLM generates: "CRITICAL ALERT: Halt line immediately. Purge casting tundish before resuming production," accompanied by a one-touch logging button to feed resolution data back into the system's learning loop.

Alignment with National Initiatives and Conclusion

The proposed SMRITI framework directly supports the Government of India's push for Industry 4.0 and Atmanirbhar Bharat by offering an indigenous, highly deployable diagnostic ecosystem. By avoiding the impracticalities of XR hardware on hazardous plant floors, this approach utilizes existing camera infrastructure combined with cutting-edge LLM and XAI integration. This ensures that Indian MSME workforces receive transparent, real-time guidance, reducing dependence on foreign diagnostic software and minimizing costly production downtimes. Future work will deploy this framework in a live steel manufacturing environment to evaluate the reduction in operational response times.

Abstract ID: 4338

Development of a Physics-Based Digital Twin Framework for the CNC End Milling Process

Akash K*, Tere Rajesh Bab, G.L. Samuel

Department of Mechanical Engineering, Indian Institute of Technology Madras, India
**akashkishore1905@gmail.com*

Abstract

Modern manufacturing systems have been greatly transformed by the rapid development of Industry 4.0 technologies. Digital Twins (DTs) play a vital role in achieving data-driven intelligence. In simple terms, a digital twin is a high-fidelity virtual representation of a physical system. It develops in parallel with the real system through continuous data exchange. The implementation of digital twins offers a promising platform for enhanced process understanding, predictive capability, and overall operational efficiency in machining processes such as CNC end milling, where multiple interacting physical phenomena occur simultaneously. Nevertheless, the optimization of productivity, tool life, and surface integrity remains a challenging task due to the highly nonlinear thermo-mechanical behaviour involved in metal cutting

In the current study, a physics-informed digital twin framework for the CNC end milling process is developed by integrating finite element simulations, machine learning, cloud computing, and real-time visualization techniques. Specifically, it aims to predict the key performance metrics of the machining process, namely tool wear and surface roughness, while enabling intelligent decision-making for process optimization. The proposed framework combines physics-based modelling with data-driven learning to achieve accurate and computationally efficient predictions without relying solely on extensive experimental investigations.

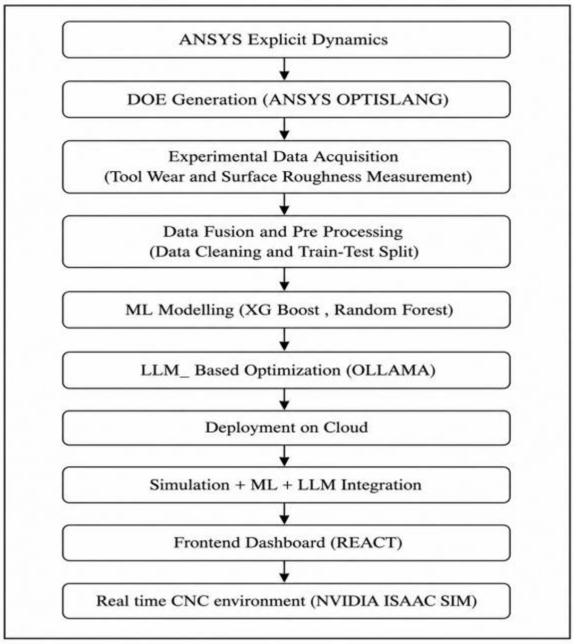

Figure 1. Digital twin framework for CNC end milling process.

To develop the physics-informed digital twin, a three-dimensional FE model of the end milling process was created using ANSYS Explicit Dynamics, which is well suited for modelling manufacturing processes involving large deformations, high strain rates, nonlinear contact interactions, and thermo-mechanical effects. The model captures realistic tool-workpiece interactions including contact mechanics, plastic deformation, frictional heat generation, and material removal. Spindle speed, feed rate, and depth of cut are varied over a predefined operating range to generate diverse machining conditions. Important process variables including cutting forces, temperature, contact pressure, equivalent plastic strain, and sliding velocity are extracted from each simulation to characterize the underlying machining physics. A Design of Experiments (DOE) methodology implemented through ANSYS optiSLang is employed to systematically sample the process parameter space. Furthermore, a Reduced Order Model (ROM) is developed using optiSLang to efficiently estimate machining responses across the design space, significantly reducing the computational cost associated with repeated finite element simulations.

To improve the robustness and practical applicability of the predictive framework, experimentally measured tool wear and surface roughness data are combined with the simulation-based dataset to form a hybrid dataset. This integration enables the model to benefit from both the physical understanding provided by simulations and the real-world variability captured through experiments. Based on this hybrid dataset, surrogate machine learning models are developed using XGBoost to predict tool wear and surface roughness from machining parameters. In addition, inverse prediction models are developed to recommend suitable machining parameters corresponding to desired quality requirements, thereby supporting process optimization and decision-making.

Scalability and real-time accessibility are achieved through a cloud-enabled architecture that integrates the physical CNC machine, edge computing, cloud services, predictive models, and visualization into a unified framework. Live machining data acquired from the CNC machine are transmitted through the cloud and streamed to a React-based user interface, enabling real-time telemetry monitoring and synchronized visualization of CNC motion using a URDF-based digital model. The interactive dashboard also provides predictive analytics, optimization recommendations, toolpath visualization from G-code, and an intelligent AI-assisted interface for interpreting machining predictions and offering process improvement suggestions. The framework further supports integration with NVIDIA Isaac Sim for immersive visualization of machining operations within a virtual environment. By combining physics-based modelling, hybrid surrogate machine learning, cloud-enabled digital twin architecture, and intelligent decision support, the proposed framework contributes towards the development of next-generation smart manufacturing systems capable of real-time monitoring, prediction, and process optimization.

Abstract ID: 3178

Physics-Based Digital Twin Development for Polymer Fused Filament Fabrication 3D Printing

Kunj Patel*, Tere Rajesh Babu, G.L. Samuel

Department of Mechanical Engineering, Indian Institute of Technology Madras, India
**kunjpatelhaker@email.com*

Abstract

One of the leading alternatives to traditional methods for creating parts and products is additive manufacturing (AM). Additive manufacturing, specifically fused filament fabrication (FFF), is one of the most widely used methods due to its ability to produce complex geometries and customized components. However, many of the technical challenges associated with manufacturing parts using FFF arise from the dependency of the process on variables, which limit part quality and dimensional accuracy and result in defects such as warping and residual stress. One method of overcoming some of these technical limits is to use digital twin technology by combining physical state representation, simulation, data acquisition, and real-time systems representation. This paper will present a physics-based digital twin framework for FFF to improve the accuracy of prediction and improve the understanding of the process by utilizing simulation, experimentation, and data-driven modelling together.

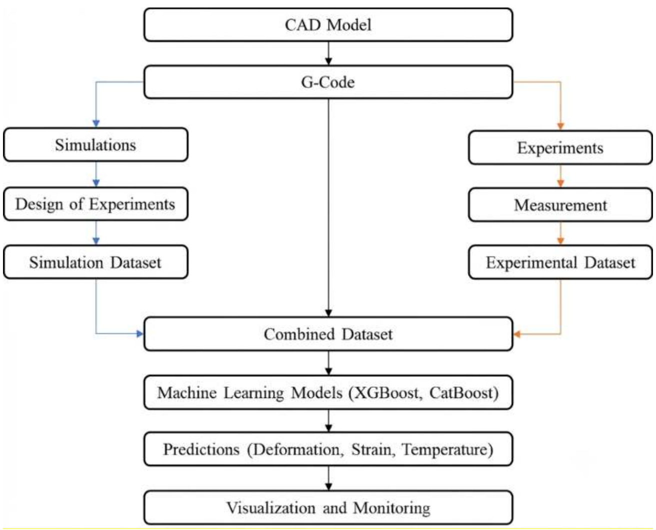

Figure 1. Digital twin framework for FFF integrating simulation, experimental validation, and machine learning.

The development of a digital twin framework requires the integration of physics-based simulation, experimental testing, and data-driven modeling to develop a digital twin of Fused Filament Fabrication (FFF) from a physics-based FFF simulation tool built in ANSYS, which uses transient thermal and structural analysis to simulate how time affects the temperature, deformation, and strain of the part throughout the print process, as well as how this will affect the part at various print speeds, layer thicknesses, and other process parameters. A Design of Experiments (DOE) methodology was utilized, which allows for systematic variations of important process parameters to result in a simulation dataset that can be used within the context of the digital twin. Simultaneously, experimental validation involves creating thin-walled geometries with the FFF process. Dimensional measurements obtained at many locations using a digital vernier caliper and a micrometer to determine deformation and strain by comparing the printed dimensions against their CAD model dimensions. Consistency between the simulations and experiments is ensured through the use of geometric measurements in multiple directions to determine equivalent strain and total deformation, which allows for direct comparison to simulation results. A combination of simulated and experimental data is used to create multiple machine learning models, such as Random Forest, XGBoost, and CatBoost trained from the combined dataset for multi-output prediction. XGBoost and CatBoost were the final models selected based on their superior predictive accuracy.

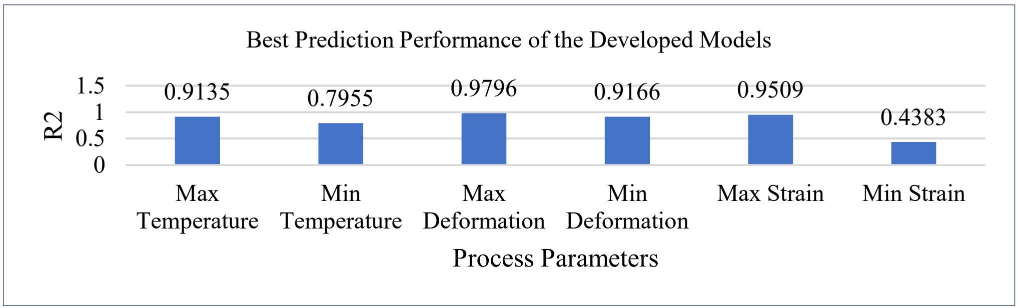

Figure 2. Best R^2 Performance (bar chart)

The backend framework develops the simulation and integrates a machine learning and data flow and the front end provides a dashboard to visualize results and provides a 3D printing environment in real-time.

The results show that all the trained models exhibit excellent accuracy across their outputs, with the highest accuracies achieved by the XGBoost and CatBoost algorithms. The developed model accurately predicts the maximum deformation, strain, and temperature of the process based on the process parameters. The comparison between actual and predicted maximum temperature responses shows a strong correlation, with a R^2 value of 0.9135. The strong correlation validates the predictive capability of the developed XGBoost model and confirms its suitability for integration into the proposed physics-based digital twin framework for reliable thermal response prediction.

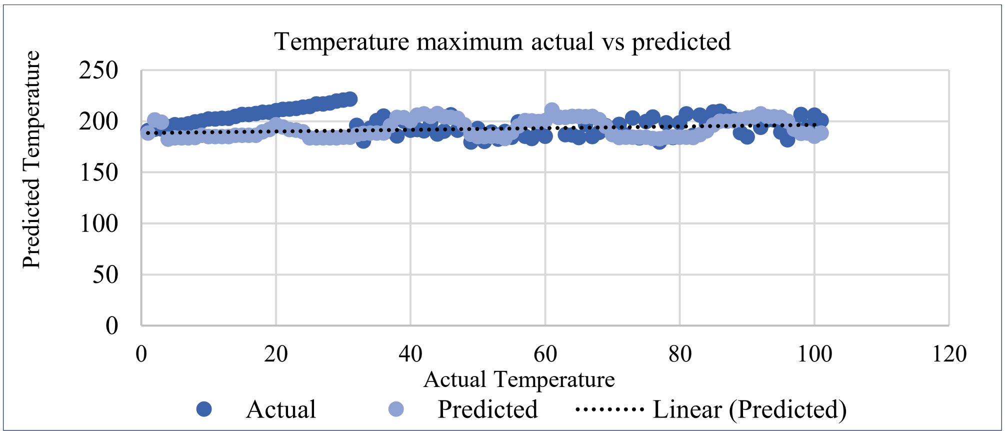

Figure 3. Actual vs. Predicted Maximum Temperature.

An effective approach that can be implemented to create a physics-driven digital twin of the fused filament fabrication process using simulation, experiments, and machine learning techniques. An effective method to enhance process comprehension and facilitate analysis of the residual stresses arising due to deformations, strains, and temperature outputs is introduced in this paper. However, there are limitations in the framework, mainly caused by a limited amount of data. The future research will focus on the extension of the database of both simulations and experiments.

Abstract ID: 3196

Adaptive Robotic GMAW Welding Simulation with DDPG Reinforcement Learning, Physics-Based Bead Geometry, and Gap-Adaptive Control for Marine Hull Fabrication

Dhanavath Gowthami*, Ganga Asok, Kondi Harshitha, Sneha Anand

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, 600036, India
**me23b165@smail.iitm.ac.i*

Abstract

Gas Metal Arc Welding (GMAW) is one of the most widely adopted fusion welding processes in heavy manufacturing and marine shipbuilding, owing to its high deposition rates, versatility across joint geometries, and compatibility with robotic automation. In hull fabrication, multi-pass welds bear the full structural load path of large steel panels subjected to cyclic wave loading, hydrostatic pressure, and corrosive seawater environments, mandating near-zero defect rates. Fixed-trajectory robotic GMAW systems fail under real manufacturing conditions because joint gap variations of 0–3 mm arise from thermal distortion during prior weld passes, camera visibility degrades due to arc glare, and process drift accumulates over extended runs. Weld rework on ship hulls accounts for up to 30% of total hull construction labour cost, underscoring the economic imperative for adaptive closed-loop control. This work presents a simulation-validated adaptive GMAW framework addressing these challenges through physics-based modelling and deep reinforcement learning.

The primary simulation environment is built in CoppeliaSim, with an ABB IRB4600 robot controlled via the ZeroMQ Remote API from Python. The weld seam is represented by four dummy waypoints; the robot's inverse kinematics (IK) solver tracks the path, while a Deep Deterministic Policy Gradient (DDPG) reinforcement learning agent applies real-time corrections. The system uses a dual-camera architecture: a navigation camera (96×96 px) scans the joint seam ahead of the torch and feeds a SeamCNN extracting a 48-D feature vector, while a pool camera mounted at the torch tip captures the melt pool and feeds a PoolCNN producing a 3-D feature vector. A column-brightness gap estimator converts pixel gap width to physical gap in millimetres and triggers a rule-based current boost when the detected gap exceeds a set threshold. Welding physics are modelled at every simulation step using semi-empirical correlations: bead width $W = kW \cdot I^{0.36} V^{0.12} v^{-0.27}$, penetration depth $P = kPI^{0.64} V^{0.13} v^{-0.35}$, and reinforcement height $H = kHI^{0.29} I^{0.09} I^{-0.21}$, where I, V, and v denote current, voltage, and travel speed, respectively. Heat input $Q = \eta \cdot I \cdot V / v$, peak HAZ temperature (Rosenthal moving-source approximation), cooling time $\Delta t_{8/5}$ (800°C–500°C), and HAZ width are computed as additional state variables. The DDPG agent operates on a 62-dimensional state vector comprising CNN features, gap measurements, heat variables, bead geometry (including aspect ratio W/H), path context, process parameters, and pass history. The 5-dimensional continuous action space commands lateral (±3.0 mm), standoff (±2.0 mm), and vertical (±2.0 mm) torch corrections together with current (±3.0 A) and travel speed (±0.5 mm/s) adjustments. Actor and critic networks each employ two hidden layers (256–128 neurons, ReLU activations), with Gaussian exploration noise (N (0, 0.02)) added to raw actions during training. A defect-aware reward function grants positive rewards for seam centring, target bead geometry (W: 6–9 mm, P: >2 mm, H: 2–3.5 mm), and heat input within 200–800 J/mm, while penalising six defect classes: undercut, lack-of-fusion, excess reinforcement, narrow bead, burn-through, and seam loss. A Lua script in CoppeliaSim synchronises visual bead rendering with the same physics equations used in Python, ensuring zero drift between state computation and visual

feedback. All 28 process signals — position, gap, current, voltage, speed, bead geometry, heat metrics, actions, rewards, and defect flags — are logged every simulation step as a CSV file for offline analysis.

A PyBullet prototype was developed in parallel to validate the core closed-loop seam-tracking and adaptive control pipeline on both flat-plate and curved ship-hull surface meshes, demonstrating the framework's applicability to realistic hull-profile geometries. Dashboard results from PyBullet runs confirm that lateral seam-tracking error remains below ±1 mm throughout the weld pass, peak HAZ temperature stabilises above the steel melt threshold, bead width tracks the 8 mm target under varying heat input, and thermodynamic bead width estimates converge with vision-based predictions after an initial transient. The controller includes a gap-adaptive current-boost rule, designed to scale current with detected gap width to prevent lack-of-fusion at wider joint openings without triggering burn-through; this behaviour was implemented but not exercised in the logged run, since the test scene-maintained a near-zero gap throughout. The CoppeliaSim implementation offers advantages in robot kinematic fidelity, sensor realism, and physics-consistent visual rendering, while PyBullet provides faster prototyping cycles and native support for arbitrary surface-mesh import. Together, the two environments span the full validation workflow from algorithm development to geometry-specific testing. Future work will replace the lumped-parameter heat model with a Goldak double-ellipsoid finite-element model for accurate residual stress prediction, substitute synthetic CNN inputs with networks pre-trained on real GMAW camera footage, conduct extended DDPG training episodes with curriculum learning (flat → curved → multi-pass), and ultimately deploy the CoppeliaSim model as a digital twin alongside physical welding hardware for hardware-in-the-loop validation on full 1 m × 1 m hull panel sections.

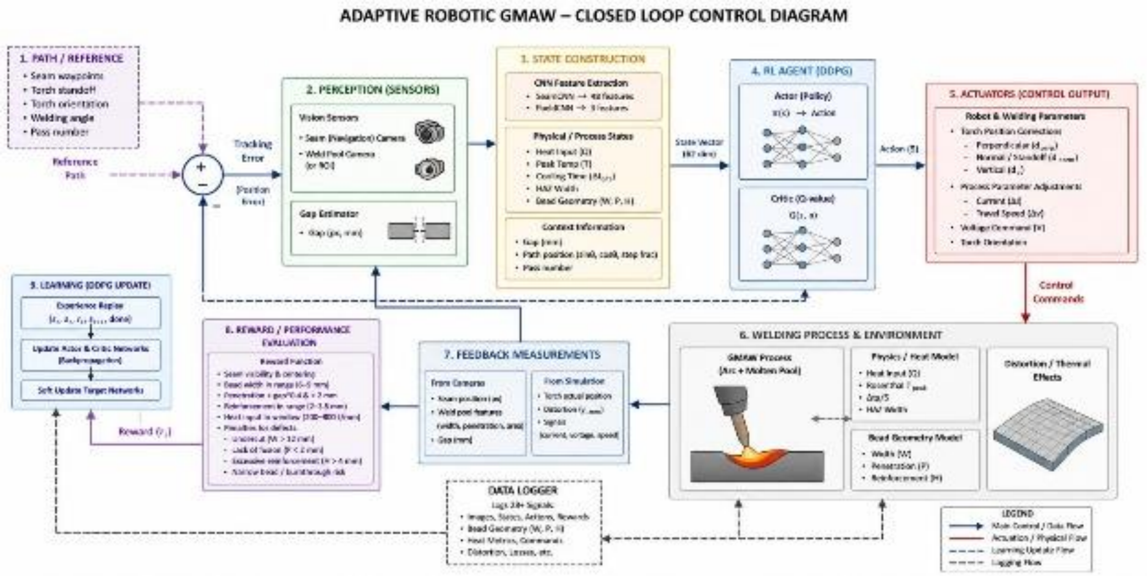

Figure 1. Closed-loop adaptive GMAW control architecture: seam waypoints and dual-camera perception feed a 62-D DDPG state vector; the actor outputs 5-D torch and process corrections; defect-aware rewards and 28-column CSV logging close the loop.

Abstract ID: 3120

Firmware-Native Industrial Digital Twin for Thermo-Kinetic Geometric Drift Estimation and Compensation in Smart Fused Deposition Modeling

Mourya Konathala[1, 2]*, Jayaganthan R[1, 2], Samuel G.L[1, 3]

[1]*Digital Twin for Autonomous Factory Lab, Indian Institute of Technology, Madras, Chennai 600036, India*
[2]*Department of Engineering Design, Indian Institute of Technology, Madras, Chennai 600036, India*
[3]*Department of Mechanical Engineering, Indian Institute of Technology, Madras, Chennai 600036, India*
**ed24s301@smail.iitm.ac.in*

Abstract

Industrial Digital Twins (IDT) under Industry 4.0 must transcend passive monitoring to become active cyber-physical agents. This work presents a complete three-layer firmware-native IDT for Fused Deposition Modeling (FDM) that closes the loop between virtual estimation and physical control. Unlike existing digital shadows, this framework implements a Closed-Loop Actuation Layer that injects corrective kinematic offsets (SET_GCODE_OFFSET) based on real-time thermo-kinematic drift analysis. The IDT exploits multi-layer firmware position data (stepper counts, kinematic transforms, commanded trajectory) to estimate geometric drift without external sensors. Experimental validation on a CoreXY system demonstrates a Z-axis thermal drift rate of -16.15 µm/hr strongly correlated with bed temperature ($R^2 = 0.87$). The closed-loop control reduces accumulated positional error from ±0.18 mm (open-loop) to ±0.05 mm (IDT-active), achieving a 72% improvement in dimensional accuracy. This confirms that firmware-native models can serve as true Industrial Digital Twins, enabling autonomous precision compensation without hardware modification.

Keywords: Industrial Digital Twin, Closed-Loop Control, Geometric Drift, Additive Manufacturing, Fused Deposition Modeling

Industrial Context and Problem Statement

Smart manufacturing demands cyber-physical systems where virtual models actively govern physical behavior. The Industrial Digital Twin (IDT) achieves this through three coupled layers: (i) Sensing/Data Layer ingesting physical state; (ii) Virtual Model Layer maintaining high-fidelity process representation; and (iii) Closed-Loop Actuation Layer issuing corrective commands. This third layer distinguishes an IDT from a passive digital shadow. In Fused Deposition Modeling (FDM), geometric drift systematic tool-end-effector deviation during extended cycles accumulates invisibly in open-loop stepper systems, degrading accuracy in precision applications.Primary mechanisms include: (1) thermal expansion/contraction of the frame; (2) stepper backlash/hysteresis; and (3) kinematic model mismatch. Existing works address post-mortem analysis or manual calibration; the characteristic write-back actuation channel required for true closed-loop DT behaviour remains absent. This paper presents a complete three-layer firmware-native IDT for real-time geometric drift estimation and compensation.

Digital Twin Architecture and Methodology

The IDT exploits multi-layer hierarchical position representations native to open-source firmware (Klipper, Marlin). Five firmware position layers are utilized: (i) raw stepper counts; (ii) stepper-to-millimeter transforms; (iii) forward kinematics; (iv) firmware-predicted toolhead position; (v) G-code commanded position. Inter-layer discrepancies constitute the state divergence signal,the virtual model's continuous quantification of physical process deviation from nominal. The Sensing/Data Layer operates in two modes: Offline (20–50 kHz kernel motion logs: trapq(motion queue) data which commands acceleration/velocity profiles per axis, stepq (stepper queue) data commands actual stepper execution positions, step_phase data which is sub-microstep resolution error signals. Online subscribes to live WebSocket firmware API streams (position, velocity, acceleration, temperature telemetry) for real-time state ingestion at 5 Hz.The Virtual Kinematic Model Layer processes ingested data to maintain a running consensus estimate of true tool head position. Instantaneous geometric drift is defined as:

$$\epsilon(t) = P_{live}(t) - P_{gcode}(t)$$

A rolling least-squares regression (200-sample window) estimates drift rate $d\epsilon/dt$ per axis. Layer-change events (Z step > 0.05 mm) are detected and timestamped to anchor layer-resolved error statistics. Multivariate thermal regression cross-correlates bed temperature (T_{bed}) and nozzle temperature (T_{nozzle}) against positional residuals to compute thermal coupling coefficients ($\partial\epsilon/\partial T$), enabling predictive capability for geometric state ahead of threshold exceedance. Closed-Loop Actuation Layer closes the cyber-physical loop between the virtual model and the physical FDM Printer. When estimated positional error $|\epsilon(t)|$ exceeds a threshold (default: 100 μm), the IDT autonomously issues a corrective SET_GCODE_OFFSET command to the firmware with magnitude $\epsilon(t)$, zeroing accumulated error without print interruption.

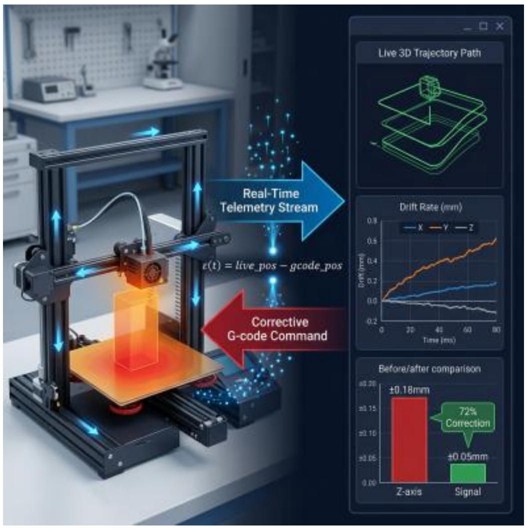

Figure 1. Firmware-Native IDT Architecture for FDM Geometric Drift Estimation and Compensation.

Results

The IDT virtual model identifies multi-axis drift rates from offline analysis as −7.3 μm/hr along the X-axis, −65.3 μm/hr along the Y-axis, and −16.15 μm/hr along the Z-axis, with the dominant Y-axis drift — approximately nine times that of the X-axis — indicating asymmetric Y-carriage compliance or kinematic backlash in the CoreXY architecture. The negative drift directions are consistent with thermal contraction behavior, where initial frame expansion during heatup is followed by net inward contraction as ambient cooling dominates the print cycle. Position residual standard deviation varies by motion type, with linear traversals showing σ ≈ 40 μm RMS and grid-fill infill patterns involving frequent direction reversals exhibiting σ ≈ 110 μm RMS, confirming backlash effects, while layer-transition spikes reach ±150 μm and decay within 0.5 seconds. Online API sensing exhibits ±20 μm quantization noise relative to offline analysis, with confirmed latency below 50 ms, and when the full three-layer IDT is activated, the actuation layer triggers 47 correction events over a 36-minute print (mean inter-event interval of approximately 46 seconds), with an average correction magnitude of 73 ± 31 μm per event (±42% relative variability, reflecting overshoot above the 50 μm trigger threshold from polling-interval delay).

Table 1. IDT closed-loop correction performance.

Performance Metric	Virtual Model Only	Full IDT Active	Improvement
Cumulative Z-axis positional error	±0.18 mm	±0.05 mm	72%
Per-layer dimensional tolerance	±0.12 mm	±0.03 mm	75%

Discussion And Conclusion

This work demonstrates a complete three-layer Industrial Digital Twin (IDT) for FDM geometric compensation, satisfying the core requirements of an updated virtual model, real-time coupling, and autonomous actuation, thereby advancing beyond conventional monitoring-only digital twin implementations toward true closed-loop control. A key industrial advantage lies in its firmware-native design, requiring no additional sensors and leveraging only standard commands such as SET_GCODE_OFFSET (Klipper) and M206 (Marlin), enabling immediate deployment across existing FDM systems without capital investment and providing a practical Industry 4.0 retrofit pathway. The tunable correction threshold allows adaptation to part-specific tolerances, while preliminary predictive control results (~85% error reduction) highlight a clear transition from reactive correction to predictive manufacturing intelligence. The system achieves a 72% reduction in cumulative Z-axis positional error (±0.18 mm to ±0.05 mm) and a 75% improvement in per-layer dimensional tolerance over 36-minute print cycles, validating its effectiveness as a self-correcting cyber-physical framework. Overall, this firmware-agnostic and industrially scalable IDT establishes that high-precision geometric compensation is achievable across smart FDM manufacturing environments without hardware modification.

Abstract ID: 2990

A Comprehensive Study on Streamlining Backward Facing Step Flow Simulation Using the PyAnsys Framework

Mohammed Abdul Razzaq[1]*, N. Sekarapandian[1], Abhishek Chitwar[2]

[1]Department of Thermal and Energy, School of Mechanical Engineering, Vellore Institute of Technology, Vellore, India

[2]Synopsys India, Pune, India

**razzaqmohammed22@gmail.com*

Abstract

The laminar backward-facing step (BFS) flow is a classical CFD benchmark valued for its geometric simplicity combined with rich flow physics – shear-layer separation, primary and secondary recirculation, and reattachment. Conventional Ansys Fluent workflows rely heavily on manual, GUI-driven geometry creation, meshing, and post-processing, which limits reproducibility and makes systematic parametric studies laborious. This work presents a fully scripted, Python-based CFD pipeline built on the PyAnsys ecosystem to address these limitations for the two-dimensional laminar BFS problem.

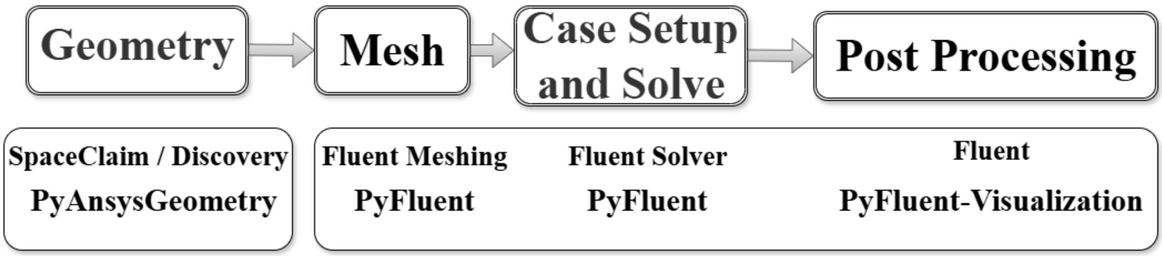

Figure 1. PyAnsys framework.

The pipeline couples four PyAnsys modules: PyAnsys-Geometry for parametric, sketch-based 2-D domain generation; PyFluent's Watertight Geometry meshing workflow for meshing; PyFluent solver mode for laminar case setup, in which the Reynolds number (100 to 800) is varied parametrically by modifying fluid viscosity while the inlet velocity and geometry are held fixed; and PyFluent-Visualization for streamline, wall- shear, and recirculation-length extraction. All four stages were executed and validated end-to-end for the single, simply-connected 2-D channel geometry considered here.

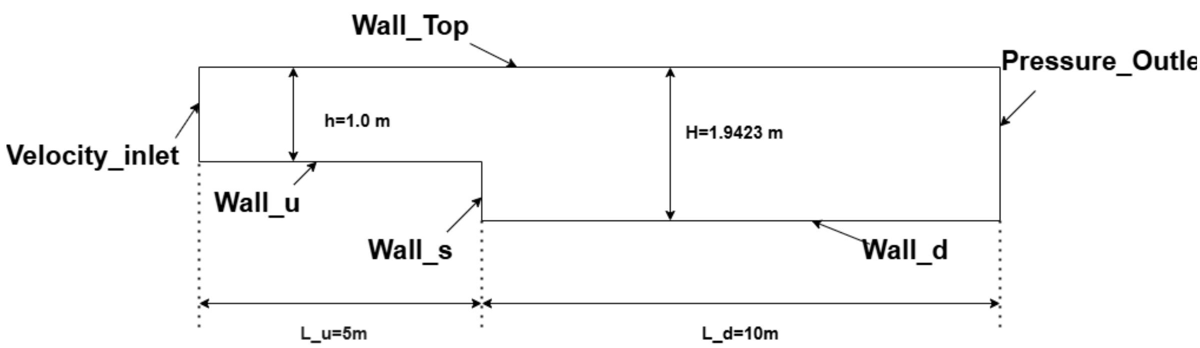

Figure 2. Schematic of the 2-D backward-facing step computational domain.

A six-level grid-refinement study at Re = 100 showed monotonic convergence of the primary reattachment length; a parallel study at Re = 800 confirmed convergence of all three reattachment/separation locations that appear at this higher Re. Primary recirculation lengths predicted across Re = 100–800 agree closely with the two-dimensional benchmark of Biswas et al. for Re ≤ 400, while both datasets show the expected progressive under-prediction at higher Re – a trend consistent with the onset of three-dimensional, sidewall-driven effects reported beyond Re ≈ 400.

Relative to established CFD-automation routes, the PyAnsys approach differs from legacy Fluent Scheme/TUI journal scripting and GUI macro-recording in exposing the full solver object model as native, introspectable Python, and from open-source Python bindings such as PyFoam for OpenFOAM in retaining direct, in-process access to a validated commercial solver rather than wrapping external text-based case files. A quantitative comparison of accuracy, run time, or scripting effort against these alternatives, or against SU2's Python-wrapped solvers, was outside the scope of the present validation and is identified below as future work.

Scope and limitations.

The present validation is intentionally narrow. The demonstrated geometry is a single, simply-connected 2-D rectangular expansion; the parametric control exercised was limited to viscosity-driven Reynolds-number sweeps at fixed geometry. Only the laminar viscous model was exercised – RANS/LES turbulence modelling, though accessible through the same PyFluent solver API, was neither invoked here. References in the wider PyAnsys ecosystem to machine-learning and generative-AI integration describe an architectural capability (solver objects are directly accessible to Python ML libraries); no ML model, optimizer, or generative-design loop was implemented in this study. Finally, no computational-performance benchmarking (wall-clock time, memory, or scripting overhead relative to manual GUI operation) was carried out; this quantification is planned for future work.

These results confirm that the scripted PyAnsys pipeline reproduces established 2-D laminar BFS benchmarks with high fidelity for Re ≤ 400 and captures the expected qualitative trends up to Re = 800. Future work will extend the framework to 3-D and turbulent regimes, quantify its performance against manual and alternative scripted workflows, and pursue ML-assisted parametric optimization.

Keywords: PyAnsys, PyFluent, backward-facing step, laminar flow, watertight geometry meshing, workflow automation

Abstract ID: 2890

Development of Digital Twin Framework for Compound Pendulum Experimentation empowered with IoT and Augmented Reality

Kiran Wakchaure[1]*, Pritish Vibhute[1], Manish Dhatrak[2]

[1]School of Engineering and Technology, Sanjivani University, Kopargaon, Maharashtra, 423603, India
[2]Department of Electronics and Computer Engg, Sanjivani College of Engineering, Kopargaon, Maharashtra, 423603, India
directorproject@sanjivani.edu.in

Abstract

The COVID-19 pandemic exposed critical vulnerabilities in laboratory-dependent engineering pedagogy, forcing students into remote learning environments entirely devoid of hands-on experimental access. In mechanical engineering curricula, the theory of the compound pendulum experiment—which requires physical manipulation of apparatus and real-time observation of dynamic phenomena—is foundational for developing intuitive understanding of rigid body dynamics, moment of inertia, and centre of gravity determination. The abrupt discontinuation of physical laboratory sessions revealed that existing virtual learning platforms, while adequate for theoretical instruction, offered no viable substitute for the empirical rigor intrinsic to experimental mechanics. This deficit motivated the present work, which proposes and validates a fully functional IoT and Augmented Reality (AR)-based digital twin framework specifically designed to replicate and in several measurable respects, surpass the conventional compound pendulum laboratory experience.

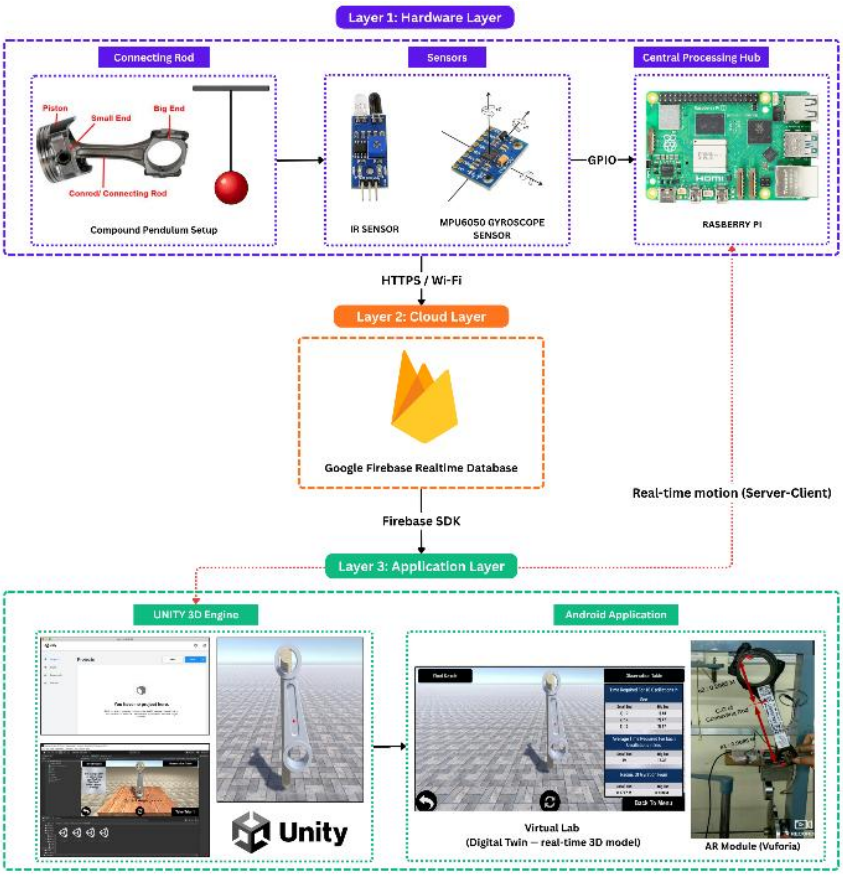

Figure 1. System Architecture — IoT and AR Based Compound Pendulum Digital Twin.

The proposed system integrates three principal technological subsystems into a cohesive experimental platform. The first subsystem comprises a physical IoT instrumentation layer built around a Raspberry Pi 4 (4 GB LPDDR4 RAM, dual-band 802.11ac wireless) serving as the central edge-computing node. Two sensors are interfaced via GPIO: an MPU-6050 six-axis MEMS inertial measurement unit providing simultaneous three-axis gyroscopic and accelerometric readings (configurable full-scale ranges: ±250 to ±2000 °/s gyroscope; ±2 to ±16 g accelerometer) for real-time angular displacement capture of the oscillating connecting rod, and an LM393-based infrared proximity sensor (detection range 2–30 cm) for high-accuracy oscillation counting and period measurement. The second subsystem is a cloud-based data pipeline wherein Python firmware on the Raspberry Pi streams sensor telemetry to Google Firebase Realtime Database over HTTPS, enabling persistent, globally accessible experimental records. The third subsystem is an Android mobile application developed in Unity 3D (version 2019.4 LTS), incorporating Vuforia-powered AR and a C# server–client communication layer for bidirectional data exchange with Firebase.

The experimental methodology replicates the classical compound pendulum protocol with sensor-augmented precision. The connecting rod is suspended at its small end and subsequently at its big end; at each configuration, the MPU-6050 precisely enforces a 4° initial displacement angle—a measurement traditionally prone to human error. The IR sensor autonomously records the time elapsed for ten complete oscillations across three consecutive trials, transmitting all observations to Firebase in real time. The Raspberry Pi executes the governing dynamics in Python: from the measured oscillation periods T_1 and T_2 corresponding to suspension at each end, the system solves the simultaneous frequency equations $f = (1/2\pi)\sqrt{[ga/(k^2 + a^2)]}$ to yield the radius of gyration K and the mass moment of inertia I_G about the centre of gravity. The gyroscope simultaneously streams live angular position data to the Unity application via the client–server architecture, animating a photorealistic 3D SolidWorks-derived model of the connecting rod in synchrony with its physical counterpart—realising a fully functional mechanical digital twin.

Experimental validation yielded the following quantitative outcomes. The IoT-instrumented system recorded average oscillation periods of 19 s (small end; three trials with zero inter-trial variance) and 17.33 s (big end; trials: 18, 17, 17 s), confirming high repeatability attributable to sensor-enforced initial conditions. Solving the governing equations produced a radius of gyration K = 0.04820 m, CG distances a_1 = 0.1747 m (from small end) and a_2 = 0.100 m (from big end), and mass moment of inertia I_G = 4.4611×10^{-3} kg·m^2. These results are simultaneously propagated to the Android application's observation table interface and confirmed against manual baseline measurements, demonstrating agreement within acceptable engineering tolerance. The total hardware cost is INR 4,199 (approximately USD 50), underscoring practical deployability in resource-constrained academic institutions. The augmented reality module constitutes the most pedagogically distinctive contribution of the system. Upon completion of all trial readings, the student activates the AR camera mode and directs the device towards a fiducial marker affixed to the physical connecting rod. The Vuforia engine recognises the target and overlays a spatially registered annotation directly onto the live camera feed, visually locating the calculated centre of gravity on the actual specimen with annotated offsets a_1 = 0.0685 m (small end to CG) and a_2 = 0.2065 m (big end to CG), rendered in the real-world coordinate frame. This AR visualisation closes the critical epistemic gap between abstract numerical results and physical spatial understanding—a gap that neither conventional laboratory instruction nor standard e-learning platforms adequately addresses. The virtual lab scene additionally provides a step-by-step guided walkthrough of the experiment within a navigable 3D environment, enabling students at any geographic location to develop procedural fluency independent of physical apparatus access. The framework offers several advantages over the conventional experimental modality. Sensor-enforced initial displacement eliminates the principal source of human error, yielding zero inter-trial variance in oscillation period at the small end suspension. The Firebase-mediated data pipeline provides tamper-evident, timestamped experimental records accessible to instructors for remote assessment—a feature absent from prior virtual laboratory implementations reviewed in the literature. - The digital twin paradigm realised here extends beyond passive simulation: real-time synchronisation of physical kinematics into the virtual model enables students to observe dynamic correspondence between physical and digital domains, reinforcing conceptual understanding of angular velocity, simple harmonic motion, and moment of inertia through embodied, spatially grounded interaction. The system's modular architecture supports straightforward extensibility; the Raspberry Pi can be substituted with an ESP32 microcontroller at approximately 80% cost reduction, and the Unity environment can accommodate additional mechanical experiments within the same IoT–AR pipeline.

In conclusion, the presented work constitutes a validated, low-cost, fully functional digital twin implementation for a canonical mechanical engineering experiment, combining real-time IoT sensor instrumentation, cloud-based data management, 3D virtual laboratory visualisation, and spatially registered augmented reality result presentation within a single unified Android application. The system achieved accurate CG determination ($K = 0.04820$ m, $I_G = 4.4611 \times 10^{-3}$ kg·m^2) with enhanced repeatability relative to manual experimentation, at a total hardware cost of INR 4,199. The work demonstrates that the confluence of IoT, digital twin technology, and augmented reality constitutes a practically viable and educationally superior paradigm for smart manufacturing-oriented engineering education, directly aligned with the objectives of Industry 4.0 and the thematic scope of ICIDTSM 2026.

Abstract ID: 3604

Adaptive Optimization of Surface Roughness in CNC Milling Using Experimental Machining Data

Mahesh Babu N*, Dr. Sivasrinivasu Devadula

Indian Institute of Technology Madras, Chennai, India
[*]me24d007@smail.iitm.ac.in

Abstract

Surface roughness is one of the most critical quality indicators in CNC milling, significantly influencing dimensional accuracy, fatigue strength, wear resistance, and the functional performance of machined components. Maintaining the desired surface finish is challenging due to the complex and nonlinear interactions among machining parameters such as cutting speed, feed rate, and depth of cut, as well as variations in machining conditions. Conventional machining practices generally rely on predetermined cutting parameters that are incapable of adapting to these variations, often resulting in inconsistent surface quality, reduced productivity, and increased tool wear. To address these limitations, this research proposes an Agentic Artificial Intelligence (Agentic AI)-based adaptive optimization framework for autonomous regulation of CNC milling process parameters to achieve the desired surface roughness. Experimental milling trials are conducted using a statistically designed Design of Experiments (DOE), where cutting speed, feed rate, and depth of cut are systematically varied to generate a comprehensive machining dataset. The corresponding surface roughness values are measured using a surface roughness tester and used to develop a robust machine learning prediction model based on the Extreme Gradient Boosting (XGBoost) Regressor. The trained model accurately predicts surface roughness from the machining parameters and is integrated into an Agentic AI architecture comprising perception, reasoning, planning, and action modules. The intelligent agent continuously evaluates the predicted machining performance, compares it with the desired surface roughness, and autonomously determines optimized combinations of cutting speed, feed rate, and depth of cut to improve machining quality. Unlike conventional offline optimization techniques, the proposed framework enables adaptive decision-making by iteratively refining machining parameters through a closed-loop optimization strategy based on predictive analytics. The effectiveness of the proposed methodology is experimentally validated by comparing the optimized machining parameters with conventional fixed-parameter machining in terms of surface roughness, machining time, dimensional accuracy, and machining efficiency. The proposed intelligent framework is expected to enhance process robustness, improve product quality, reduce operator intervention, and provide a practical foundation for the development of self-optimizing CNC milling systems aligned with the principles of Industry 4.0 and intelligent manufacturing.

Methodology

The proposed methodology begins with the development of an experimental dataset through a Design of Experiments (DOE), in which the primary machining parameters, cutting speed, feed rate, and depth of cut are systematically varied during CNC milling operations. For each experimental trial, the corresponding surface roughness (Ra) is measured using a surface roughness tester, generating a comprehensive dataset that captures the relationship between the machining parameters and surface quality. The collected experimental data are subsequently pre-processed through data cleaning, normalization, and feature analysis to improve model reliability and predictive performance. The processed dataset is then used to train an Extreme Gradient Boosting (XGBoost) Regressor, which learns the nonlinear relationship between the machining parameters and the resulting surface roughness.

The predictive model is evaluated using standard regression metrics, including the coefficient of determination (R^2), root mean square error (RMSE), and mean absolute error (MAE), to ensure prediction accuracy. The trained XGBoost model is integrated into an Agentic Artificial Intelligence (Agentic AI) framework, where the intelligent agent performs perception, reasoning, planning, and decision-making to predict surface roughness and determine the optimal combination of cutting speed, feed rate, and depth of cut required to achieve the desired surface finish. Finally, the optimized machining parameters are validated experimentally through additional CNC milling trials, and the proposed framework is evaluated by comparing its performance with conventional parameter selection methods in terms of surface roughness, prediction accuracy, and overall machining performance

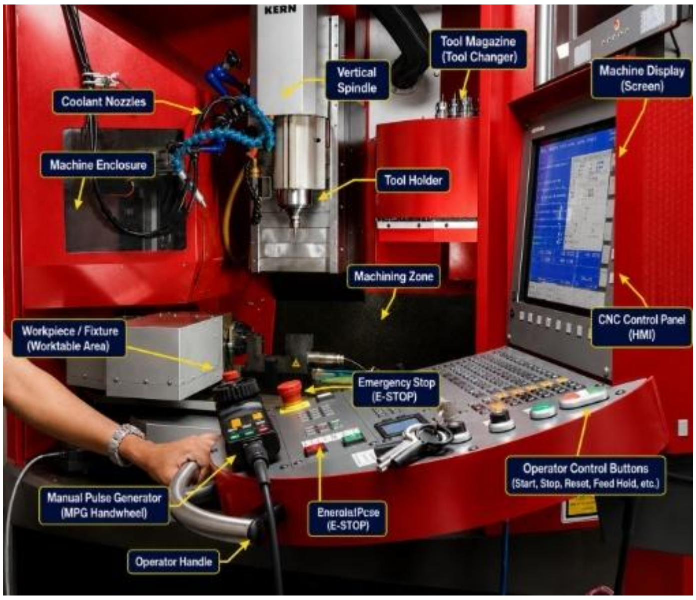

Figure 1. Experimental CNC Machining Setup.

Abstract ID: 4458

Development of a Real-Time Digital Twin Framework for Autonomous Monitoring and Control of a Motorized Spindle

Mohammad Taha Rizvi[1], Syed Ismail Sultan[2]*, Mohit Law[3#]

[1]*Interdisciplinary Centre for Artificial Intelligence, Aligarh Muslim University, Aligarh, India*
[2]*Department of Mechanical Engineering, Aligarh Muslim University, Aligarh, India*
[3]*Department of Mechanical Engineering, Indian Institute of Technology Kanpur, Kanpur, India*
**syedismailsultan10480@gmail.com*
[#]*mlaw@iitk.ac.in*

Abstract

Motorized spindles are among the most critical subsystems of modern CNC machine tools, directly influencing machining accuracy, productivity, and machine reliability. Continuous monitoring of spindle operating conditions is therefore essential for early fault detection, preventive maintenance, and improved operational safety. While existing standards provide guidelines for spindle condition monitoring, most implementations remain limited to data acquisition and visualization, with little emphasis on real-time bidirectional interaction and autonomous control. This work presents the development and implementation of a real-time digital twin (DT) framework for condition monitoring and supervisory control of a high-speed motorized spindle.

The proposed architecture consists of three integrated layers: (i) a physical spindle system, (ii) a data acquisition and processing layer, and (iii) a virtual representation component developed using the Unity Engine. The physical system comprises a naturally air-cooled RMCS 7008 motorized spindle (24,000 rpm maximum speed, 0.8 kW rated power) driven by a variable frequency drive (VFD). Spindle vibration is measured using an Endevco 44A16 piezoelectric accelerometer mounted near the front bearing housing, while spindle temperature is monitored using a PT100 sensor interfaced through a MAX31865 module. Vibration signals are acquired using a National Instruments NI9234 data acquisition card, and temperature measurements are collected through an Arduino-based acquisition system. Signal processing and feature extraction are performed in MATLAB.

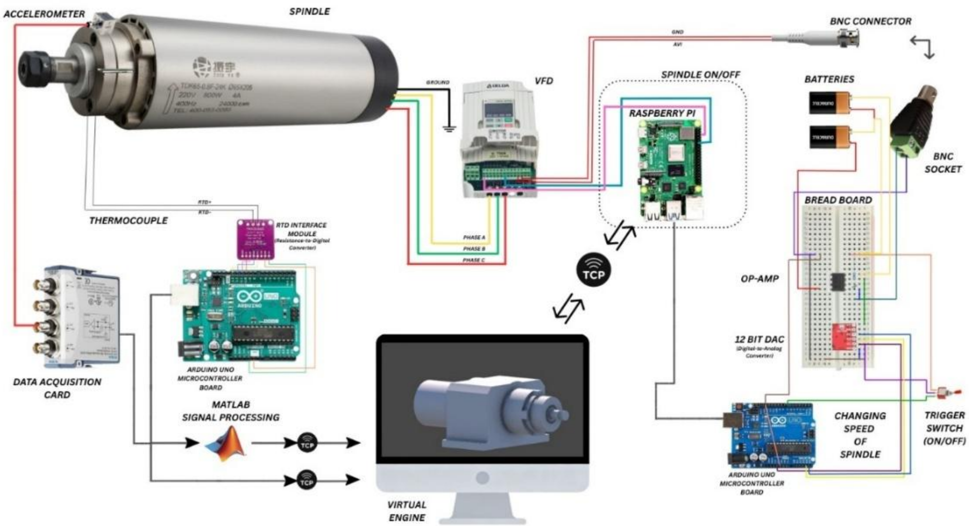

Figure 1. Real-time digital twin architecture for a motorized spindle, showing sensor-based acquisition of vibration and temperature data, TCP/IP communication with a Unity-based virtual representation, and bi-directional control of the physical spindle through microcontrollers and the VFD.

The processed vibration and temperature data are transmitted to Unity through TCP/IP communication, where a three-dimensional spindle assembly created in SolidWorks serves as the virtual representation. The digital twin reproduces the geometry, operating state, and rotational behavior of the physical spindle in real time. Live process variables, including spindle speed, vibration amplitudes, frequency spectra, and temperature, are continuously visualized within the virtual environment, enabling intuitive assessment of spindle health.

A key feature of the developed framework is bi-directional communication between the physical and virtual domains. Using a Raspberry Pi, Arduino Uno, MCP4725 digital-to-analog converter, and analog signal conditioning circuitry, commands generated within the Unity environment are transmitted to the VFD, allowing remote spindle start, stop, and speed control. Furthermore, threshold-based monitoring logic has been implemented, enabling automatic intervention when excessive vibration levels, abnormal spectral signatures, or critical temperatures are detected.

Experimental studies conducted across multiple spindle speeds demonstrate reliable real-time synchronization between the physical spindle and its digital counterpart with minimal communication latency. The virtual model accurately reflects the spindle's operating condition, while the control architecture successfully executes supervisory actions based on monitored health indicators. The proposed framework establishes a foundation for intelligent spindle monitoring and autonomous machine tool operation, providing a practical pathway toward cyber-physical manufacturing systems and next-generation smart machine tools.

Abstract ID: 3121

Condition Monitoring of Rotating Shafts Through Digital Twin and FFT-Based Vibration Analysis

Dhanalakshmi R*, Pavithra E, Arun Prakash C, Mani P

Department of Production Technology, Street Madras Institute of Technology, Anna University, Chromepet, Tamil Nadu 600044, India

abiiramamoorthy@gmail.com

Abstract

Condition monitoring of rotating machinery is essential for improving equipment reliability, reducing maintenance costs, and preventing unexpected failures. This work proposes a Digital Twin-based vibration monitoring system for rotating shafts using an MPU6050 accelerometer, ESP32 microcontroller, FFT-based signal analysis, and Unity 3D visualization Figure1. Real-time vibration and shaft speed (RPM) data are acquired and processed to identify mechanical looseness. A Digital Twin developed in SolidWorks and imported into Unity 3D is synchronized with the physical system through WebSocket communication. Machine conditions are classified into Healthy, Warning, and Fault states based on vibration thresholds. Experimental results demonstrate that the proposed system Figure2 provides low-cost, real-time condition monitoring, intuitive visualization, and supports predictive maintenance for Industry 4.0 applications.

Methodology

The proposed methodology integrates vibration sensing, shaft speed measurement, FFT-based signal processing, and Digital Twin visualization for rotating shaft condition monitoring.

The experimental setup consists of an induction motor coupled to a rotating shaft through a flexible jaw coupling and supported by pillow block bearings. Mechanical looseness is introduced by gradually loosening one of the bearing mounting nuts to simulate bearing support instability under different operating conditions. Figure 3 shows the bearing mounting before and after introducing controlled mechanical looseness.

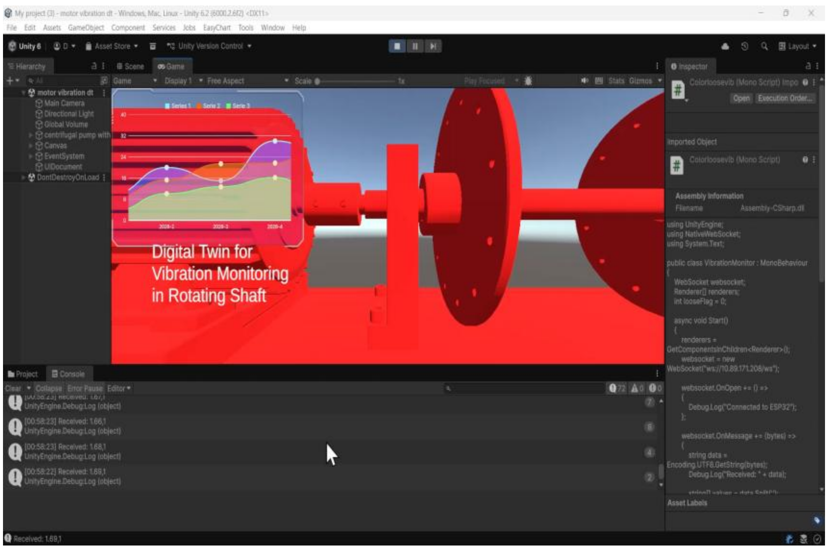

Figure 1. Digital Twin of vibration monitoring in induction Motor

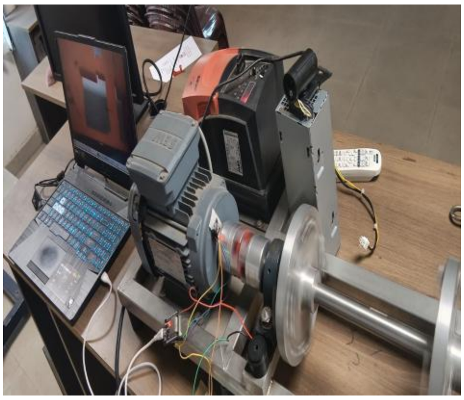

Figure 2. Digital Twin of Motor Shaft Assembly.

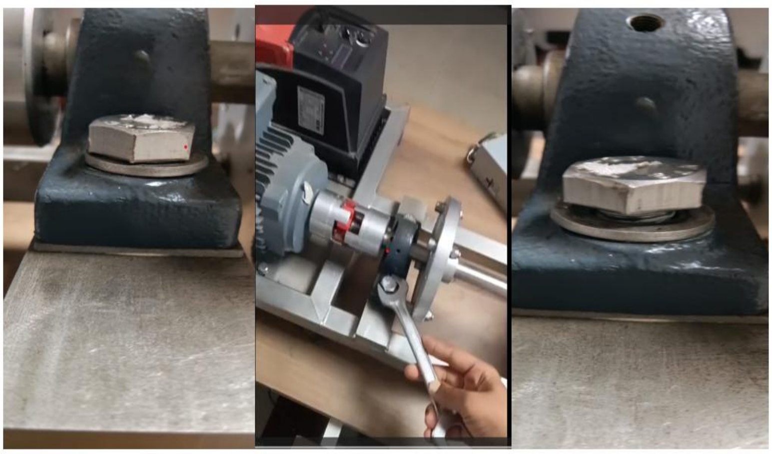

Figure 3. Bearing mounting before and after introducing controlled mechanical looseness in the rotating shaft system.

An MPU6050 three-axis accelerometer is mounted on the motor housing near the bearing to acquire vibration signals along the X, Y, and Z axes. The sensor is interfaced with an ESP32 microcontroller through the I²C communication protocol. In addition to vibration measurement, the shaft rotational speed is measured using an infrared (IR) proximity sensor, which detects each shaft revolution and calculates the shaft speed in revolutions per minute (RPM). The ESP32 simultaneously acquires both vibration and rotational speed data and continuously transmits them to a host computer for further processing. Figure 4 shows the real-time shaft rotational speed (critical speed monitoring) obtained during the experiment. The measured RPM is used to synchronize the rotational motion of the Digital Twin with the physical rotating shaft.

The acquired vibration signals are analysed using the Fast Fourier Transform (FFT) to convert the time-domain signals into the frequency domain. The dominant frequency components are used to identify abnormal vibration caused by mechanical looseness. Figure 5 shows the FFT spectrum of the acquired vibration signal.

Results and Discussion

The proposed Digital Twin-based monitoring system was experimentally validated using a rotating shaft test rig under healthy and looseness conditions. The MPU6050 successfully acquired real-time vibration data, while the IR sensor measured the shaft rotational speed. FFT analysis identified dominant frequency components corresponding to shaft rotation and revealed additional frequency peaks under looseness conditions, indicating mechanical instability. The threshold-based algorithm accurately classified the machine condition into Healthy, Warning, and Fault states. The Unity-based Digital Twin synchronized successfully with the physical system through WebSocket communication, updating the virtual shaft motion and machine condition in real time. Colour-based visualization and buzzer alerts improved operator awareness and enabled early fault detection. Overall, the proposed system demonstrates an effective, low-cost solution for intelligent rotating machinery condition monitoring and predictive maintenance. The Digital Twin enables continuous real-time monitoring.

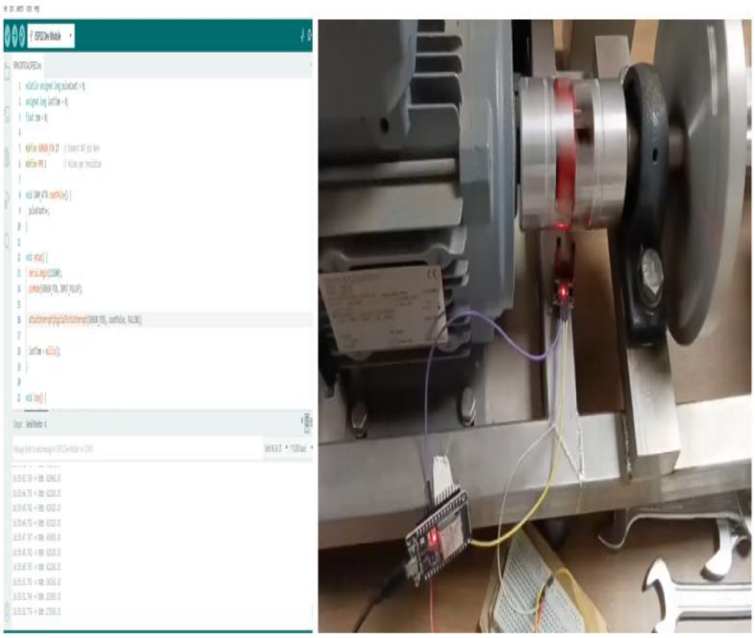

Figure 4. Real-time shaft rotational speed (RPM) measurement using an infrared proximity sensor and ESP32 for monitoring the operating speed of the rotating shaft.

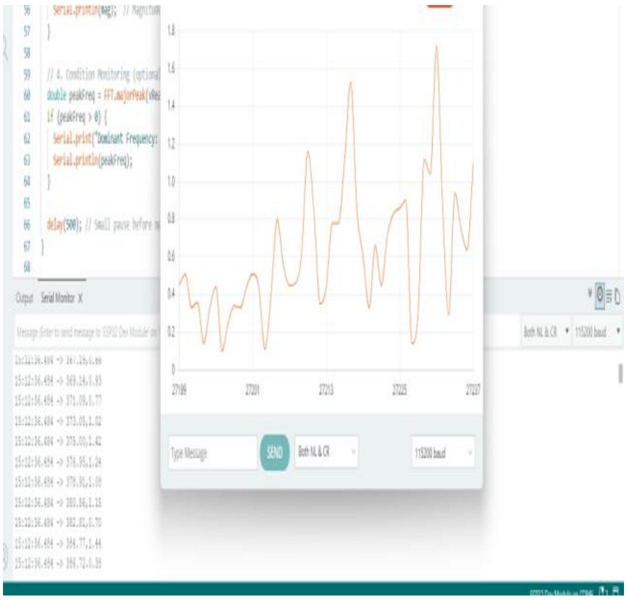

Figure 5. FFT spectrum of the vibration signal showing the dominant frequency components used for rotating shaft fault diagnosis.

Abstract ID: 3156

Smart Factory Predictive Maintenance Hub: An Edge-to-Cloud Architecture for RUL Prediction and Fault Classification in Rotating Industrial Machinery

Jaikrishna R*, Yogeswaran I, Sandhya C, Krithn Aadityan G, Kishore V

Meenakshi Sundararajan Engineering College, Anna University, Kodambakkam, Chennai, India
*jaikrishna5125@gmail.com

Abstract

The evolution of Industry 4.0 has led to the emergence of the need for intelligent and real-time systems for managing asset health, which operate at the edge but perform computation in the cloud. In this work, the authors introduce the Smart Factory Predictive Maintenance (PdM) Hub, which is an end-to-end architecture for the edge-to-cloud continuous monitoring of the rotating assets, including motors, pumps, and bearings (Figure 1). The edge module could be implemented as a Tuya T3 microcontroller with a specially designed C/C++ firmware on the tuyaOpen platform collecting 6-DOF vibration data using the MPU6050 and temperature data using the DS18B20 sensors that can be further encrypted using AES and then sent using MQTT over TLS to the cloud ingestion pipeline.

The ML intelligence component carries out

(i) RUL regression for the NASA C-MAPSS FD004 dataset, and

(ii) four-class fault classification (healthy, inner race, outer race, rolling element) for the CWRU Bearing Dataset.

For RUL regression, models were trained on raw sensor and operating setting features of C-MAPSS after performing feature selection using a variance-based technique, applying a piecewise RUL constraint, and carrying out engine-wise 10-fold cross-validation. In the case of fault classification, CWRU vibration data were divided into windows containing 1024 samples each, and eight time-domain statistics were derived and validated through 10-fold stratified cross-validation. Different machine learning approaches, such as CNN, LSTM, and their hybrid versions, have been compared with Random Forest and XGBoost algorithms.

The highest accuracy for CWRU fault classification, according to pure CNN models (99.63% total), was closely followed in second place by hybrid CNN-LSTM (99.10%) and third place LSTM (98.37%) splits, as shown in Figure 2. XM/XG Boosted models attained similar mean cross-validated scores of 92.47% (± 1.36%) vs. 92.03% (± 1.15%) while both being significantly less than deep learning model methodologies used above. The best results of RUL prediction for the C-MAPSS data set were again produced from LSTMs (RMS qual:29.344 and PHM qual:7098.7) followed by 1D CNNs (RMS qual = 34.086, PHM qual:12934.38), and hybrids (RMS qual:35.044, PHM qual:22026.9). Thus, demonstrating that the hybrid architectural combinations of models were not guarantee successful results with regards to RUL regression predictions as depicted in Figure 3. Both Random Forest and XGBoost provided a high level of predictive accuracy with RMSE and PHM scores that were competitive to the deep learning models (RF: RMSE 29.97, PHM 7.7561; XGBoost: RMSE 30.11, PHM 8.1724), while also providing approximately 10x faster inference times, making them more appropriate for latency-sensitive edge deployment.

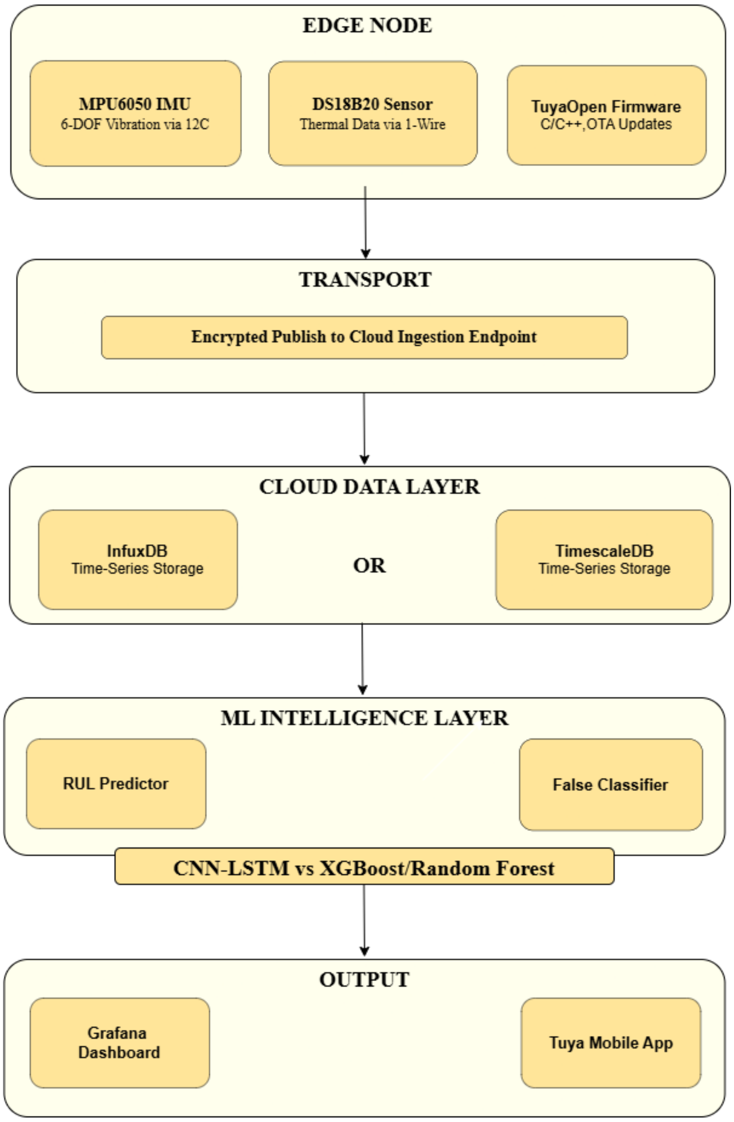

Figure 1. Proposed smart factory predictive maintenance hub architecture.

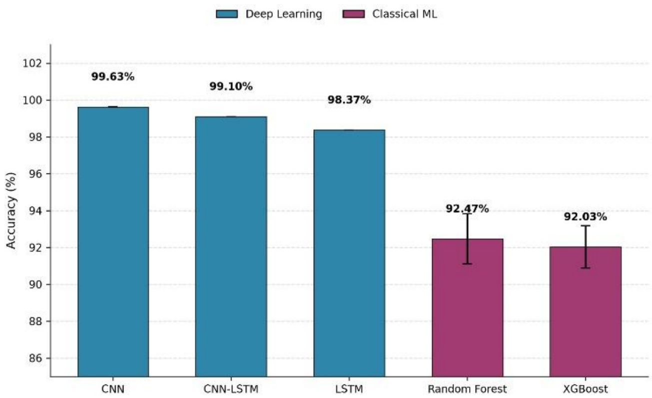

Figure 2. CWRU Bearing Fault Classification - Model Accuracy Comparison.

As part of the new edge node definition, this device is designed using a custom KiCad 7 PCB board that utilizes hardware-based crypto acceleration and supports over-the-air (OTA) updating using the tuyaOpen SDK. The edge node will utilize device authentication, encrypted data communication, and firmware management from remote locations to support industrial deployment. At the edge, data is protected with AES encryption, and all devices will communicate with one another using TLS-secured

MQTT transport. Also, using Docker-based containerization can provide horizontal scalability and the ability to move to a cloud-based platform while preserving overall system architecture.

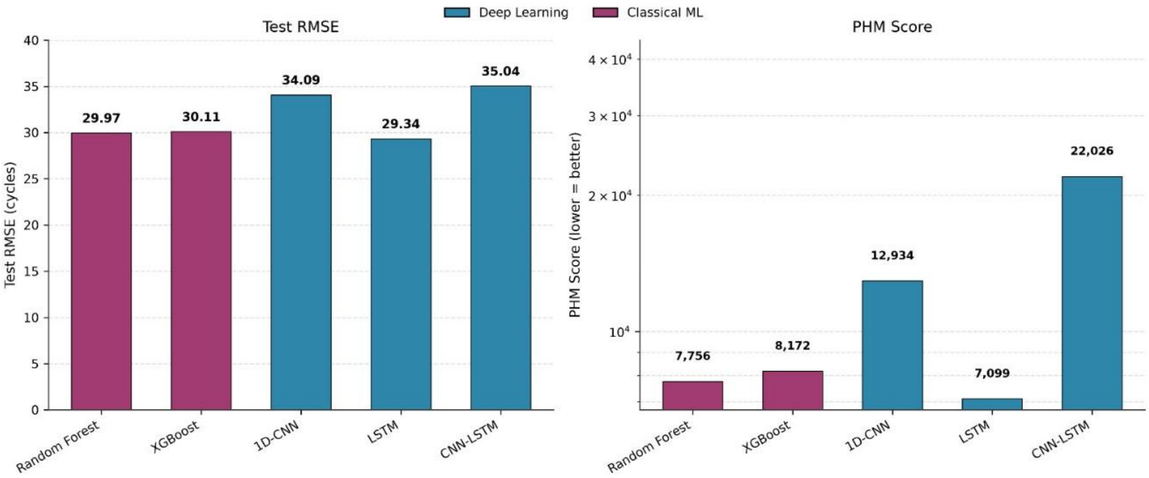

Figure 3. C-MAPSS FD004 - RUL Prediction Performance Comparison

Abstract ID: 2560

Physics-Informed Digital Twin for Bearing RUL Prediction with Bi-LSTM

Abdelhalim M.H.A.Eldosh[1]*,Girish Kumar[2]

[1]*Department of Production and Industrial Engineering, Delhi Technological University, Delhi, India*
[2]*Department of Mechanical Engineering, Delhi Technological University, Delhi, India*
**abdoeldosh840@gmail.com*

Abstract

Remaining Useful Life (RUL) estimation of rolling element bearings is critical for smart manufacturing based on predictive maintenance, but it is still a challenge due to the lack of run-to-failure data, the nonlinearity of the degradation process, and the limited interpretability of deep learning models. Bearings are ubiquitous and account for the majority of industrial failures, and thus robust prognostics are critical for reducing the associated downtime and maintenance costs.

To overcome these challenges, a physics-informed hybrid framework is proposed that combines a Digital Twin-based simulation model with a Bidirectional Long Short-Term Memory (Bi-LSTM) network. The Bi-LSTM model is selected because it captures temporal dependencies in both forward and backward directions, enabling more effective learning of nonlinear bearing degradation patterns and improving RUL prediction accuracy. The Digital Twin is implemented as a physics-informed virtual model for synthetic data generation, not as a real-time system fully synchronized with the physical system. It is realized with bearing fault physics in the form of Ball Pass Frequency Outer Race (BPFO) equations, which enables the generation of realistic synthetic degradation signals, which is used for pre-training the model. The model is then fine-tuned on real vibration data from the NASA IMS dataset (Run 2, Bearing 1), due to its well-defined run-to-failure behavior and clear fault progression. The study focuses on outer race faults due to their clear representation in the impulsive nature of vibration-based degradation signals. Three time-domain statistical features, namely Root Mean Square (RMS), Kurtosis, and Crest Factor, are extracted from the vibration signals to characterize bearing degradation and serve as inputs to the Bi-LSTM model. Finally, a monotonic constraint is introduced to enforce the physical behavior of the RUL evolution. The proposed hybrid modeling approach leverages domain knowledge with data-driven learning capability for better generalization across different operating conditions. This approach provides improved RUL predictions than existing approaches that consider physics-based modeling and deep learning independently. The physics-informed Digital Twin is integrated with bidirectional temporal learning and explainable AI into an end-to-end pipeline for RUL prediction.

The proposed model obtains an RMSE of 0.102 and 0.057 for MAE and R^2 of 0.75, which indicates a strong predictive ability in nonlinear degradation. In comparison to the Bi-LSTM model without Digital Twin pre-training, the RMSE is about 40% better, which proves that physics-informed data augmentation yields a significant improvement. The prediction results follow the actual degradation trend, especially during the rapid failure period. Moreover, the higher R^2 value indicates a stronger correlation between degradation trends of the predicted value and actual value, which further shows the effectiveness of the proposed methodology.

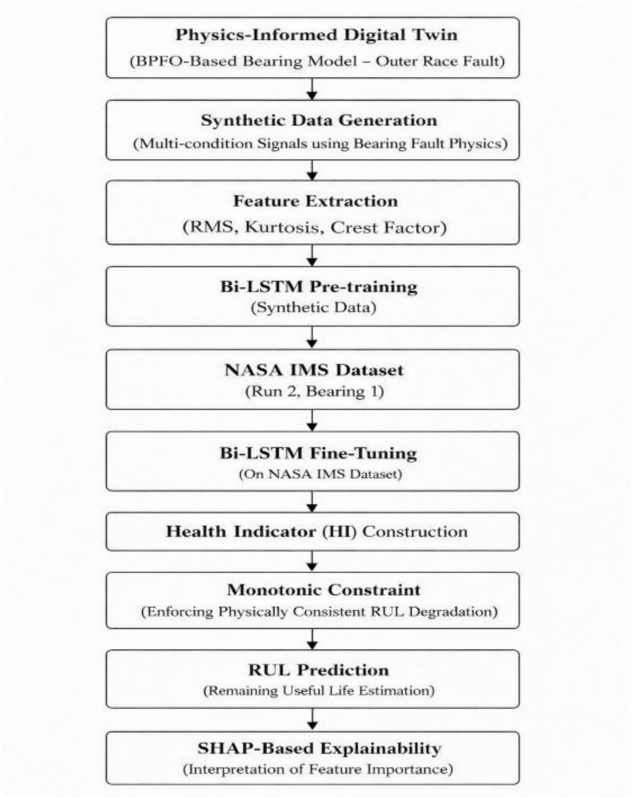

Figure 1. Proposed physics-informed hybrid framework for bearing RUL prediction.

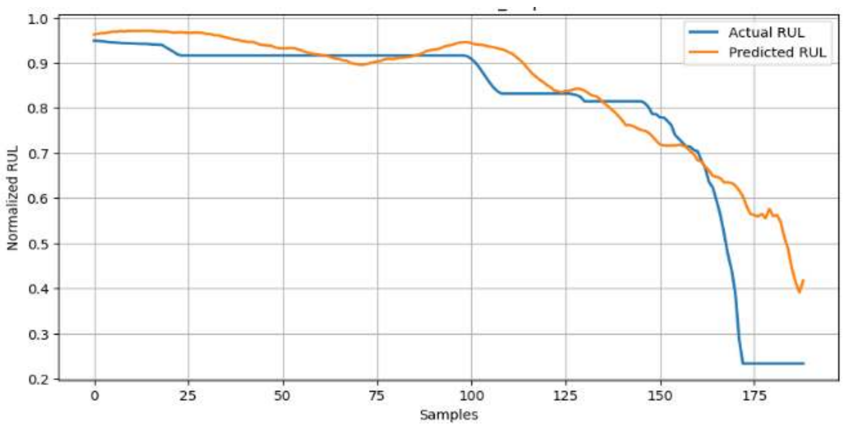

Figure 2. Actual vs Predicted RUL.

Furthermore, explainability based on SHAP is used to interpret model predictions. Analysis indicates that RMS is the most important feature in later stages of degradation, which is consistent with the physical behaviour of outer race faults. This suggests that the model has learnt meaningful mechanical patterns instead of incidental correlations. Bidirectional LSTM further improves temporal learning by considering both past and future context in degradation modelling, which is essential in the industrial context where maintenance decisions require physically explainable rationales.

From an industrial perspective, the overall framework provides reliable and interpretable maintenance decision-making through failure mitigation and maintenance schedule optimization. The Digital Twin-based simulation, together with deep learning, offers a reliable and scalable solution for predictive maintenance in data-limited industrial settings.

Abstract ID: 2698

Dual-Fault Diagnosis in Single-Stage Spur Gear Systems

Sudha Radhika, Anbumani P, Charanya G, Dominic Sanju James R*

Department of Electrical and Electronics Engineering, BITS-Pilani Hyderabad Campus, India
**dominicsanjujamesrex0501@gmail.com*

Abstract

Tooth root crack initiation and progressive surface wear are two degradation modes in a single-stage spur gear drive that are addressed by the fault-tolerant control architecture presented in this study. In order to track both failure processes, the underlying gear dynamics model marries Paris law crack formation with Archard wear accumulation and accounts for time-varying mesh stiffness, rotational inertia, and backlash nonlinearity. A weighted objective that balances torque tracking error against crack severity is minimized by a Model Predictive Controller working across a finite receding horizon. Sequential quadratic programming is used to solve the optimization at each time step. Additionally, as wear depth approaches a critical limit, a Lyapunov-based controller gradually derates torque and speed using a composite energy function to enforce closed-loop stability. MPC significantly reduces the peak crack health index compared to uncontrolled operation, according to MATLAB/Simulink closed-loop data, and the Lyapunov controller quantifiably slows wear accumulation before maintenance is required.

Keywords: Lyapunov control stability, Paris crack law, Archard wear law, Extension of fatigue life, gear prognostics, and structural and integral health monitoring.

Introduction

Gear failures are one of the primary causes of mechanical drivetrain interruptions in automobiles, aircraft, and industrial environments. Two frequent failures in gears are pitch-line surface wear from persistent sliding contact under load and tooth root fatigue cracking from repeated bending stress focused at the tooth fillet. Both of these concerns worsen with continued operation, and they are not self-healing. The Paris-Erdoğan power law describes crack propagation, and the wear rate given by Archard's law is influenced by contact pressure and speed.

Mathematical model

Gear Dynamics

The simulation of gear pair dynamics employs two rotational degrees of freedom. Mesh contact is made based on relative angular displacement inside the plane of action. Time-dependent mesh stiffness accounts for both cyclical fluctuations in tooth contact and gradual stiffness reduction due to cracking, demonstrating a clear relationship between structural integrity and dynamic responsiveness. The mesh torque, which includes elastic and damping components, drives the equations of motion for each shaft, while viscous damping and external load torque serve as opposing forces.

Tooth Root Crack: Paris Law

Crack progression is observed by monitoring the diameters of the crack (a_1 and a_2) caused at the pinion and gear roots. Each tooth root's bending stress is

$$\sigma_i = \frac{|\,T_m\,|}{K_{s_i}} \tag{1}$$

The above expression (1) defines the cyclic stress intensity factor,

$$\Delta k_i = Y_i \sigma_i \sqrt{\pi a_i} \tag{2}$$

the crack propagation in the gear is given by,

$$\dot{a}_k = C_k(\Delta k_k)^{m_k} \tag{3}$$

here C_k is the Paris constant and the health index of the crack is given by,

$$HI = \max\left(\frac{a_1}{a_{Crit1}}, \frac{a_2}{a_{Crit2}}\right) \in [0, 1] \tag{4}$$

when there is a fracture in the gear, HI is 1.

Surface Wear: Archard Law

Wear depth changes in accordance with

$$\dot{h}_i = K_{w,i}\, p(t)\, v_{slide,i} \tag{5}$$

here the sliding velocity at the contact of the gears is $v_{slide,i}$ and contact pressure is p(t). For control, the wear error state is

$$e_{wear} = h_1(t) - h_{1,safe}. \tag{6}$$

when the e_{wear} becomes positive; this implies the wear in the gear exceeded the safe limit of operation.

Control strategies

Model Predictive Control for Crack Mitigation

MPC chooses the torque sequence over N = finite steps at each time step that minimizes,

$$J(U) = \sum_{i=1}^{N}\left[w_1 HI_i^2 + w_2(T_i - T_{ref})^2 + w_3(\Delta T_i)^2\right] + \sum_{i:HI_i>0.9}(HI_i - 0.9)^2 \tag{7}$$

A large penalty is imposed prior to fracture based on the health index value HI = 0.9 so that the soft constraint is activated and the cost weights balance crack severity, efficiency of the torque, and smoothness in the actuation, respectively. The Paris rule is applied to the candidate torque to propagate fracture sizes forward at each horizon step, ensuring that the optimization is both computationally tractable and physically grounded.

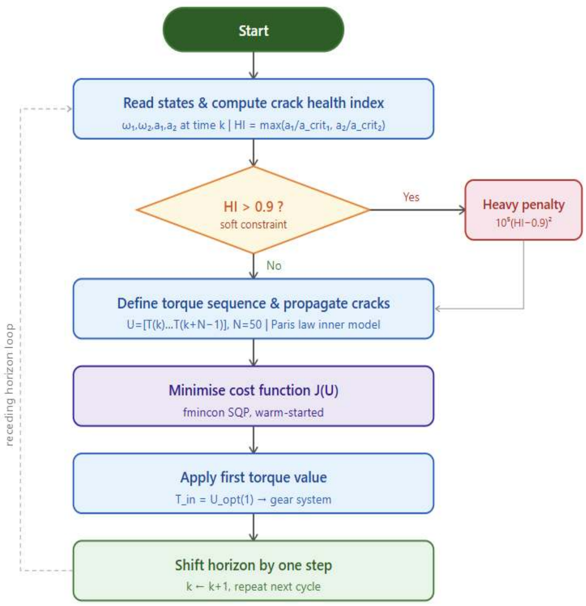

Figure 1. Proposed MPC-Crack Mitigation.

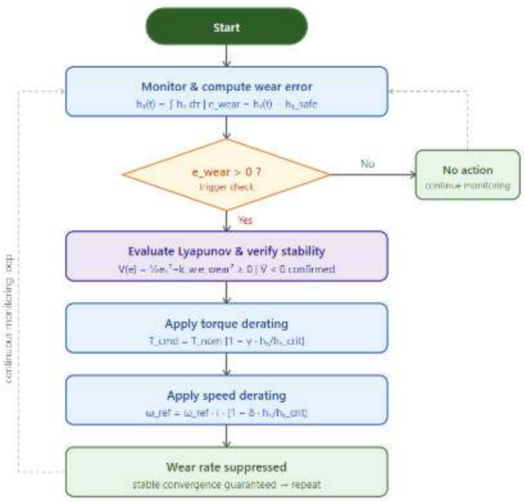

Figure 2. Proposed Lyapunov – wear suppression.

Lyapunov-Based Control for Wear Suppression

The composite candidate function is used by the Lyapunov controller.

$$V(e) = \frac{1}{2}(\omega_2 - \omega_{2,ref})^2 + \frac{10^6}{2}(h_1 - h_{1,crit})^2 + \frac{10}{2}\int \mid Tm \mid dt^2 \geq 0 \quad (8)$$

Only at the healthy operating point is this function zero. The controller initiates proportionate derating when accumulated wear exceeds the safe threshold ($e\_wear > 0$).

$$T_{in,cmd} = T_{nom}\left[1 - \gamma \cdot \frac{h_1}{h_{1,crit}}\right] \quad (9)$$

$$\omega_{1,ref} = \omega_{2,ref} \cdot i \cdot \left[1 - \delta \cdot \frac{h_1}{h_{1,crit}}\right] \quad (10)$$

Taking the time derivative of V confirms,

$$\dot{V} = -K \mid e \mid < 0 \quad (11)$$

This guarantees that the regulated states align with the safe operating envelope.

Simulation setup

Simulations performed in MATLAB and the metrics reported include time to failure, crack and wear trajectories, input torque profiles, and composite health indices.

Conclusion

This research described a dual-fault prognostic control architecture for a single-stage spur gear drive. A modelling environment was created to include Paris law fracture propagation, Archard surface wear, crack-induced stiffness loss, and backlash all in a single nonlinear model. An MPC method with a limited step receding horizon actively reduces the crack health index by optimizing input torque, whilst a Lyapunov-based controller reliably suppresses wear rate by gradual torque and speed derating once wear reaches the safe limit. Both strategies work together to extend gear life considerably beyond what uncontrolled operation allows, resulting in a practical and theoretically viable approach to active fault management.

Building on these simulation results, future studies will focus on physical validation with a digital twin system.

Abstract ID: 3164

Robustness Analysis of Federated Learning Defenses Against Local Model Poisoning Attacks

Palak Aggarwal[1]*, John Augustine[2], Rashami Sarode[3], Preeti Mishra[1], P. Vinod[4]

[1]*Department of Technology, Doon University, Dehradun, Uttarakhand, India*
[2]*Department of Computer Science, Indian Institute of Technology, Chennai, India*
[3]*Centre for Cybersecurity. Trust and Reliability (CyStar), Indian Institute of Technology, Chennai, India*
[4]*Department of Computer Applications, Cochin University of Science and Technology, India*
**pagg123456@gmail.com*

Abstract

The rapid integration of emerging technologies such as artificial intelligence, machine learning and cloud-based decision systems has fundamentally transformed modern industries. These advancements have redefined the scope of connected machinery and autonomous robotic systems, enabling more adaptive, data-driven, and intelligent operations.

Consequently, the industrial and manufacturing landscape has undergone a significant shift, with the urgent need for scalable and integrated solutions for managing these cyber-physical systems. However, this transformation is accompanied by several challenges. The reliance of smart industries on internet-enabled infrastructures and distributed architectures has increased their exposure to cybersecurity threats, including data breaches, privacy breaches, malicious intrusions, and systemic vulnerabilities.

As a result, cybersecurity is of vital importance in the development and deployment of smart manufacturing ecosystems. Ensuring robust protection against cyber threats is essential for maintaining system integrity, safeguarding sensitive data, and enabling the secure evolution of autonomous smart industries.

Federated Learning (FL) is a technology that has gained significant traction in smart manufacturing environments. This approach introduces an additional layer of privacy, as sensitive raw data remains localised on individual clients rather than being transmitted to a central server or aggregator. At the same time, each participating client benefits from periodic updates of the aggregated global model, which captures collective intelligence across the entire network.

Despite its advantages, FL remains susceptible to several security vulnerabilities. One such threat is local model poisoning, where malicious clients deliberately manipulate their locally trained model updates before submitting them to the central aggregator. These adversarial updates can corrupt the aggregation process, thereby compromising the integrity and reliability of the global model.

This paper investigates the impact of local model poisoning attacks on a range of aggregation-based defense mechanisms like Krum, Multi-Krum, Trimmed Mean, Median, Bulyan, FoolsGold, and FLAME to evaluate their robustness and performance under adversarial conditions. Unlike prior studies that predominantly rely on simplified benchmark datasets such as MNIST and CIFAR-10, this work employs a realistic network traffic dataset (CICIDS2017) and the Gotham 2025 dataset distributed across multiple clients. Here, we implement the Fang Attack for full-knowledge and partial-knowledge scenarios and also extend it to IID and non-IID scenarios.

Existing literature mostly considers this attack for Krum, Trimmed Mean, and Bulyan and uses the extended version of Bulyan on other aggregation-based defenses. Here, we try to extend this attack to all the above-mentioned defenses. The parameters taken into consideration for this study are Attack Success Rate (ASR), Accuracy, Precision, F1-Score, Recall, and Validation Loss. The paper provides a

detailed robustness analysis of these defenses mechanisms and suggest the mechanisms that are likely to withstand local model poisoning attacks.

This enables the development of federated learning–based intrusion detection systems (IDS) that more accurately reflect real-world deployment scenarios. Consequently, it provides a more reliable basis for assessing the effectiveness and resilience of robust aggregation defenses against malicious participants in practical environments.

Abstract ID: 3003

A Digital Twin-Based Cybersecurity Framework for G-code Anomaly Detection in Additive Manufacturing Systems

S Harish Mrithyunjayan[1]*, M Selvakumar[2], G L Samuel[2], P B Pankajavalli[1], Heramb Tansen Chaudhari[3]

[1]*Department of Computer Science, Bharathiar University, Coimbatore, Tamil Nadu – 641046, India*

[2]*Department of Mechanical Engineering, Indian Institute of Technology, Madras, Chennai - 600036, Tamil Nadu, India*

[3]Department of Computer Engineering, Vishwakarma Institute of Information Technology, Pune – 411048, India

**harish.mrithyunjayan@gmail.com*

Abstract

This work is about a Digital Twin-enabled cybersecurity framework. It is designed to protect printers from bad people manipulating G-code instructions. Traditional ways of keeping cyber- systems safe have not worked well for 3D printers. They use general detection strategies and signature-based methods. These methods often fail to catch behavior that is specific to the system. They also do not provide real-time information. The proposed framework is different. It focuses on the security challenges of G-code-driven manufacturing processes. G-code is like a language that tells the printer what to do.

It controls the movement, extrusion and other parameters of the printer. If someone manipulates the G-code it can cause problems. The printer might make products or move in a way that is not safe. The system uses a Digital Twin to monitor the printing process in time. The Digital Twin is like a copy of the printer. It shows what the printer is doing and helps catch any behavior.

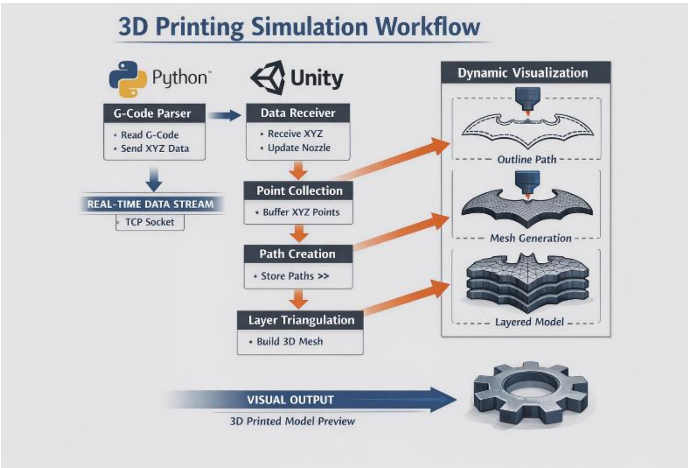

Figure 1. 3D printing simulation workflow.

The framework looks for four types of threats: Toolpath manipulation, Extrusion tampering, Sudden motion anomalies, and command injection. It covers both the safety of the products and the operation of the system. The methodology starts with getting the G-code. It can come from a file or a real-time stream. Each instruction is parsed to get the information. This includes coordinates, extrusion values and control parameters. The system uses a state- tracking mechanism to keep track of what is happening. It computes features like displacement, velocity and extrusion-to-distance ratios. These features show how the printing process is behaving. They are used to detect any anomalies.

The anomaly detection pipeline uses a framework that combines machine learning techniques. A statistical layer uses Mahalanobis Distance to quantify deviations from behavior which detects spikes and outliers in motion and extrusion. A learning-based model uses XGBoost to capture relationships and identify subtle anomalies. A lightweight rule-based validation layer ensures safety constraints are met. It checks for things like temperature commands or invalid parameter ranges. The system was tested with normal manipulated G-code datasets which gave the results showing that the hybrid detection framework is much better than statistical approaches. It can detect low-intensity manipulations like toolpath deviations and extrusion inconsistencies.

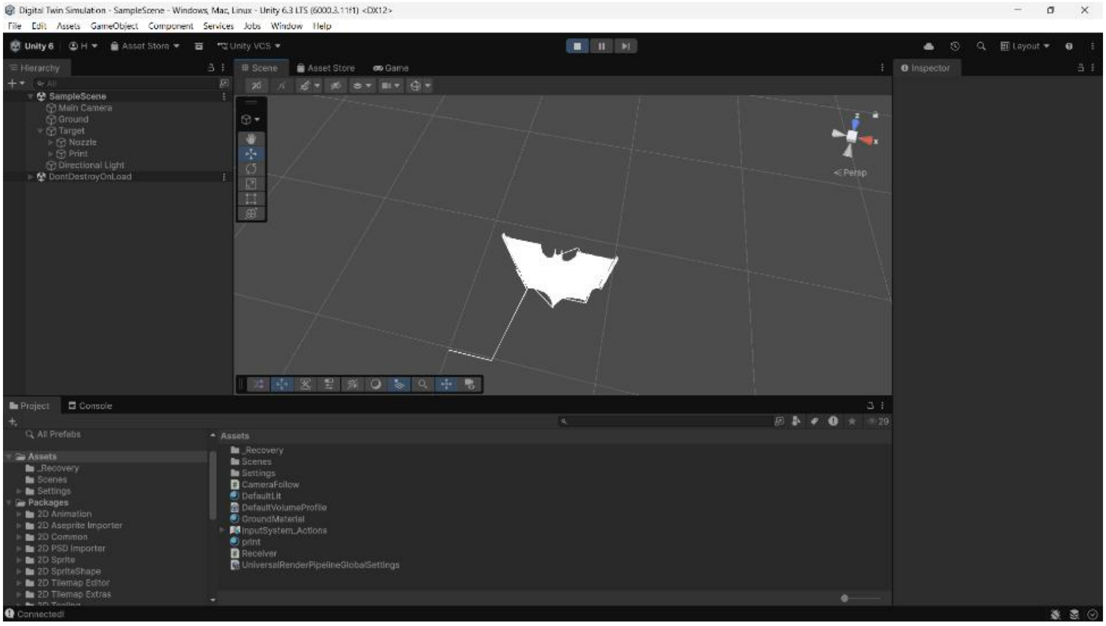

Figure 2. Digital Twin simulation.

The Digital Twin makes it easier to understand what is happening. It shows anomalies alongside their effects. The system can detect problems in time making it suitable for practical use in additive manufacturing workflows. The findings show that monitoring G-code in time enables active cyber-physical protection. The proposed framework provides a scalable and system-specific solution for securing 3D printing systems. It improves detection accuracy. Enhances transparency and situational awareness.

This makes additive manufacturing environments safer and more resilient. The Digital Twin-enabled cybersecurity framework is a step forward in protecting 3D printers from bad people. It uses a combination of real-time simulation and hybrid anomaly detection to keep the printers safe. This approach is better, than ways of keeping cyber-physical systems safe. It is an effective way to protect 3D printing systems from multiple classes of attacks.

Abstract ID: 3015

AI-Based Cybersecurity Framework for Anomaly Classification in an Autonomous Factory Using a Confidence Weighted Voting Ensemble of SALSTM, TCN, and DBN

Siddesh R[1*], M Selvakumar[2], Niketh S[2], D Ramyachitra[1], G L Samuel[2]

[1]*Department of Computer Science, Bharathiar University, Coimbatore-641046, Tamil Nadu, India*

[2]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai-600036, Tamil Nadu, India*

**siddeshramesh2004@gmail.com*

Abstract

Autonomous manufacturing environments are becoming more reliant on fleets of interconnected machines, which has made the cyberattack surface much larger and less protected. This paper describes a two-stage intrusion detection system (IDS) that uses an Isolation Forest (IF) pre-filter to send each incoming window along either an attack confirmation path or a stealthy-attack inspection path. The final decisions are made by a confidence-weighted voting ensemble that combines a Self-Attention LSTM (SALSTM), a Temporal Convolutional Network (TCN), and a Deep Belief Network (DBN). The framework was initially built on a physics-constrained synthetic dataset and verified on the generalisability of four public benchmark datasets. The framework gets an F1 score of 0.9926 and an AUC-ROC score of 0.9985 on the synthetic dataset. It also has perfect per-class precision of 1.0000 for all six attack types. This is backed up by SHAP-based explainability and a SHA-256 forensic audit trail.

Traditional intrusion detection tools, be they rule-based or single-model classifiers, do not adequately capture the temporal depth and multi-domain complexity inherent in actual factory network behaviour. Low-amplitude attacks like Gradual Drift are still mostly invisible to threshold-based detectors. On the other hand, high-intensity events like DoS/Dropout need quick, high-confidence identification to keep production from being disrupted. The proposed architecture tackles both ends by using signal-processing-based feature enrichment, an IF routing mechanism, and a heterogeneous ensemble classifier tested on a synthetic but operationally relevant dataset that includes six types of attacks: Command Injection, Coordinated Attack, DoS/Dropout, Gradual Drift, Replay Attack, and Sensor Spoofing. The same pipeline was also tested on UNSW-NB15, TON_IoT, CICIDS2017 and Edge-IIoTset with F1-scores of 0.9932, 0.9807, 0.9934 and 0.9994 respectively.

A simulation dataset with 100,800 time-stamped records that covered 82 features across 45 network-layer metrics and 37 physical-layer sensor channels was split into three parts based on time (70/10/20) and then into sequences of length 30. CUSUM (Cumulative Sum Control Chart) ($k = 0.5$, $h = 5.0$) and Daubechies-4 wavelet decomposition were used to expand each window from 82 to 258 dimensions. These methods were only used on normal traffic to keep the data secure. The IF pre-filter then sends windows that are above the anomaly threshold to an attack confirmation path and those that are within the normal range to a stealthy-attack inspection path. This makes sure that Gradual Drift is not quietly skipped. Each of the three independent classifiers, SALSTM, TCN, and DBN, gives an attack probability. When the ensemble score goes above 0.74, a weighted vote based on validation F1 (0.3352, 0.3348, 0.3300) sounds an alarm. SMOTE and inverse-frequency weighting help with class imbalance, and the SHAP Gradient Explainer shows that NET_new_connection_flag, NET_error_rate, NET_plc_command_count, A_path_deviation, and NET_response_time_ms are the five most important features. This shows that physical-layer channels work well with network telemetry. Every

inference event is logged in a SHA-256-chained log that is checked against 2,000 entries to make sure it follows IEC 62443.

The ensemble gets an accuracy of 0.9968, a precision of 0.9951, a recall of 0.9902, an F1-score of 0.9926, and an AUC-ROC of 0.9985. A precision of 1.0000 across all six attack types shows that there are very few false positives when compared to normal traffic. DoS/Dropout and Sensor Spoofing both get perfect F1 scores of 1.0000. Gradual Drift, the hardest class, gets a recall of 0.9732 and an F1 score of 0.9864. The ensemble correctly identifies over 97% of these events. Real-time streaming validation of 500 consecutive samples maintained a classification accuracy of 100%, with ensemble probabilities consistently exceeding 0.90 for all verified attack windows. Edge-IIoTset outperformed the synthetic baseline, achieving the best benchmark result (accuracy 0.9991. The conservative threshold of 0.74 makes sure that alarms only go off when at least two models agree with high confidence. This cuts down on unnecessary production stops. The framework is a credible, IEC 62443-aligned foundation for industrial IDS deployment because it has few false positives, SHAP attribution that is specific to the type of attack, and a verified audit trail. The future extension of this work is to validate the framework with real-time data. The proposed framework is designed for use in real-world systems like smart manufacturing, autonomous warehouses, and cyber-physical production systems that follow Industry 4.0 standards. In these systems, fleets of machines are always at risk from new cyber threats. People also think that the method will work for other important areas of infrastructure, like autonomous power grids, smart transportation systems, and healthcare automation networks, where operational safety standards call for clear decision-making and tamper-evited anomaly detection.

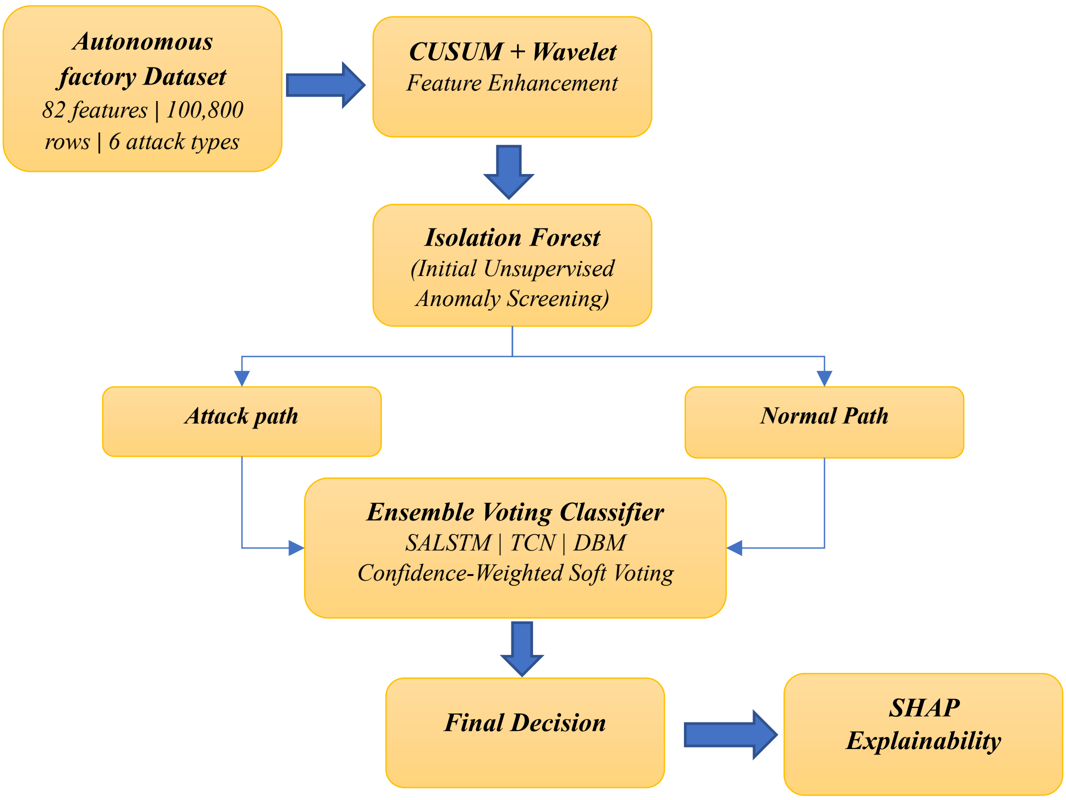

Figure 1. Methodology

Abstract ID: 2893

Multi-Sensor Orchestrated Cyber Physical System Framework: Interactive Explainable Digital Twin for secure vehicular network

Sivakumarn B[1], Deepa Jose[2]*

[1]*Computer Science and Engineering, KCG College of Technology, Karapakkam, Chennai, India*
[2] *Electronics and Communication Engineering, KCG College of Technology, Karapakkam, Chennai, India*
**deepa.ece@kcgcollege.com*

Abstract

Automobile cyber-physical systems (CPS) demand comprehensive security solutions that are real-time but also provide human-readable explanations. However, traditional Artificial Intelligence (AI) models provide good but limited detection of anomalies and are often seen as black boxes, leading to cognitive overload for operators caused by having too many sensor alerts at the same time. This study introduces the Interactive Explainable Digital Twin (I-XDT) framework, which ensures the security of key parts of a vehicle, such as its speedometer, indicator, and coolant systems, which are operated autonomously. The two major innovations in the framework are Contextual Feature Filtering to reduce the mental effort of the operator and an Active Learning Loop that improves the AI engine using feedback from the operator in both directions. Built with C# and Python, the prototype can synchronize data from multiple sensors within 3-6ms and can have an F1 score of 0.916 for detecting anomalies while correctly rejecting out-of-range error data with 100% accuracy.

Introduction

The automated cyber-physical systems are very complex and heavily rely on complicated cyber and physical components, making them particularly vulnerable to sophisticated cyber-attacks. A powerful application for providing real-time virtual monitoring has been found in the use of digital twins with sensor streams. To explain the decisions that are made regarding anomalies, explainable AI (XAI) tools such as SHAPley Additive exPlanations (SHAP) have been introduced to identify the sensors that contribute the most to the anomaly detection decisions made by the AI. Current state-of-the-art frameworks, like ENIGMA, are unidirectional and can only try to understand the data properties understood by the operator, without adapting them with respect to the human domain experience. In addition, when an emergency is in a critical phase, giving all sensors' readings at once will result in high cognitive noise and delay in emergency mitigation. The right approach to fill these gaps is to dynamically prioritize telemetry according to the state of operation of the vehicle and to introduce an active feedback loop with human verification. The proper way to address these gaps is to dynamically prioritize telemetry according to the operational state of the vehicle and to introduce an active feedback loop that uses human verification.

Methodology

The framework I-XDT's operational architecture works through 3 key phases in order to ensure secure monitoring of the speedometer, indicator, and coolant sensors.

The generation of a digital twin for data collection

The real-time sensor measurements (speed, temperature) are combined and submitted to a cloud service (such as Azure IoT) to create a unified software simulation of the vehicle. This is where an XGBoost classifier is being used to identify a data anomaly in the incoming data stream.

Contextual Feature Filtering

Standard SHAP values [4] are improved by a score $R_i(t)$ representing the context-relevance, defined as:

$$R_i(t) = |\phi_i(t)| \times \omega_j(S) \quad (1)$$

Where $\phi\_i\ (t)$ is the SHAP value of the sensor i at time t.

The contextual weight is denoted by $\omega\_j\ (S)$, which is related to the operational state S of the vehicle (e.g., the brake operational state, the turning operational state, and the cruising operational state).

A certain threshold τ is set, and operators are only given the most critical variables; those that fall below τ are removed from the mix.

Human-in-the-Loop Active Learning

Evaluated alerts are directed to a diagnostic dashboard. As operators enter the system to mark or confirm false positives, their input is used to modify the diagnostic rules in a two-way flow, enabling the system to evolve in real time to deal with new cyberthreat patterns.

Conclusion

The proposed I-XDT framework presents a highly secure, context-aware, and human-assisted platform for vehicular cyber-physical networks. The system reaches an effective level of data noise reduction of 75%, record-low latency of 3–6 ms, and enables a bidirectional active learning loop, eliminating the risk of cyber-physical attacks and reducing operator cognitive load. This architecture is very similar to the human-centric and resilient principles promoted by Industry 5.0.

Abstract ID: 3146

Digital Twin–Integrated Attack Detection Framework for Adaptive Cruise Control in Cyber-Physical Vehicle Networks

Thulasi M Santhi*, K.Srinivasan

National Institute of Technology Tiruchirappalli, Trichy, 620015, Tamil Nadu, India.
*thulasimsanthi93@gmail.com

Abstract

Intelligent transportation systems built based on cyber-physical concepts are transforming urban mobility by improving traffic flow, safety, energy efficiency, and reducing pollution. A key technology enabling this transformation is Adaptive Cruise Control (ACC), which maintains safe inter-vehicle distances and enhances road capacity. Despite these benefits, connected vehicle systems remain highly vulnerable to cyber threats due to their reliance on communication networks and distributed control. Past research has largely focused on single-attack scenarios such as denial of service (DoS), false data injection, replay, or jamming, but real-world environments often expose vehicle platoons to multiple simultaneous attacks. Coordinated threats, including camouflage, false data injection attacks, DoS attacks, actuator measurement manipulations, and sensor faults, pose a particularly severe risk by misleading vehicles, disrupting communication, and destabilizing platoon dynamics. Distinguishing between the sensor faults and attacks is challenging because both show similar characteristics. Existing resilient control strategies frequently assume leader vehicle information is always available, yet in practice this assumption may fail due to communication limits or sensor faults. This undesired system behaviour engenders physical operational security and economic loss.

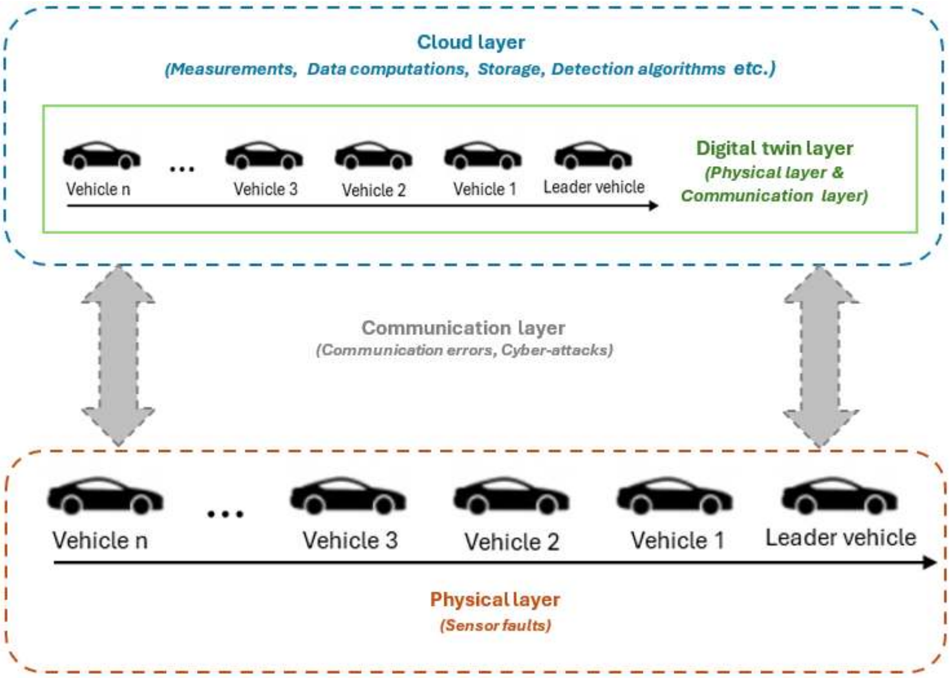

Figure 1. Digital twin based cyber-attack detection.

To address such challenges, a dual-layer attack detection framework shown in Figure 1, that enhances resilience in connected vehicle systems is introduced in this work. The architecture integrates a digital twin, implemented in the cloud as a virtual representation of the physical layer, which mirrors the communication topology and interacts in real time with the physical system. Complementing the digital twin, a machine learning–based detection algorithm is employed using variational autoencoders to identify anomalies and classify both attacks and sensor faults, thereby strengthening stability, safety, and efficiency in intelligent transportation systems operating under adversarial conditions. The vehicle platooning positions without overlapping of vehicle is given in Figure 2.

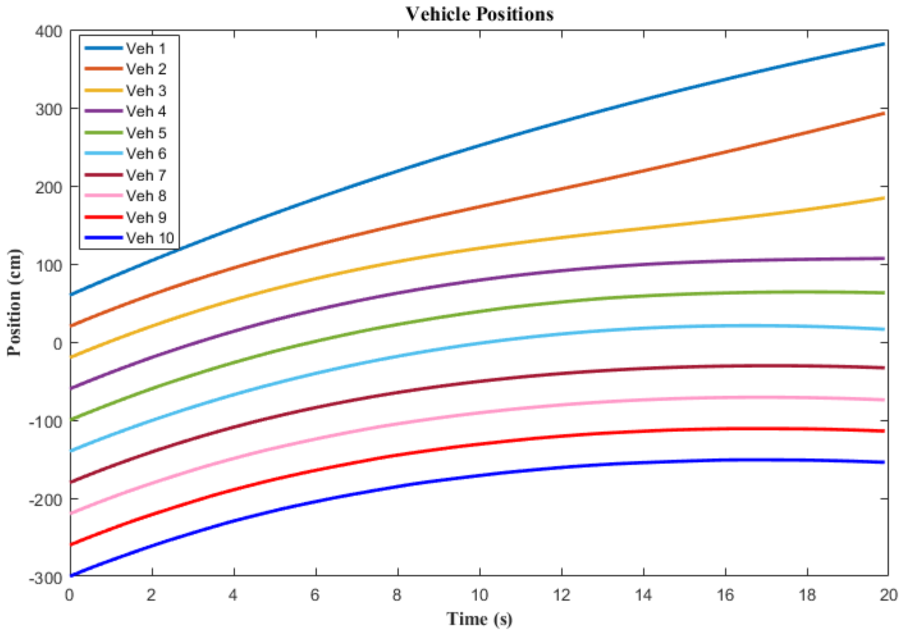

Figure 2. Vehicle positions in adaptive cruise control.

Connected vehicle systems can be modelled as comprising one leader vehicle and multiple follower vehicles, where the dynamics of each follower are expressed as $\dot{x}_i = Ax_i + Bu_i$ with x_i representing the state and u_i the control input, where $i = 1,2, \dots . n$. The leader vehicle follows the dynamics $\dot{\varphi} = A\varphi$, and the sensing model of each follower is defined as $x_i = C\varphi$. Communication among follower vehicles is bidirectional, meaning each vehicle exchanges information with both its preceding and following neighbors. Importantly, follower vehicles do not directly receive information from the leader; instead, they must predict the leader's state, ensuring coordinated control and stability within the platoon despite limited direct communication. This research considers four major types of cyberattacks and sensor bias faults that can compromise the performance of connected vehicle systems.

Camouflage attacks: mislead the platoon by disguising a malicious vehicle as the leader

False data injection (FDI) attacks distort the information exchanged among vehicles, altering the transmitted state measurements. $y_{aj} = y_j + a_j$, where y_j denotes the true state of neighbour vehicle j, a_j represents an attacking signal, and y_{aj} denotes the manipulated measurement transmitted to the vehicle i.

Actuator attacks: manipulate the control inputs of follower vehicles, disrupting their command signals and destabilizing dynamics. $y_{aj} = y_j + a_j$, where a_j denotes the actuator attacker signal and y_{aj} represents the compromised control input.

Denial-of-service (DoS) attacks interrupt communication links, with the most destructive form being zero-topology attacks - blocking all connections during active periods.

Together, these attack vectors pose significant risks to system stability and safety, requiring resilient detection and mitigation strategies. Along with this, a bias fault is also introduced in the sensor to analyse and detect the coordinated threat – attack and fault. $y_{fj} = y_j + f_j$, where y_j denotes the true state of neighbour vehicle j, f_j represents a fault signal, and y_{fj} denotes the faulty measurement.

The digital twin is designed in the proposed framework to detect most forms of attacks, except for denial-of-service (DoS) disruptions. In addition to the aforementioned method, a machine learning strategy using a variational autoencoder is developed to ensure adaptive cruise control (ACC) under simultaneous denial-of-service (DoS) and actuator attacks. By defining the tracking error between the actual and predicted states, the controller employs adaptive parameters and time-varying gains to suppress the influence of actuator disturbances. The framework also addresses practical issues such as packet loss and communication latency, showing that the proposed scheme can tolerate intermittent communication failures and benefit from high-speed wireless technologies or data compression to reduce latency. Overall, the proposed approach strengthens the robustness, safety, and efficiency of intelligent transportation systems, providing a scalable solution for secure platoon coordination in adversarial environments.

Abstract ID: 3163

To design and implement a stable and reliable defense mechanism for detecting label-flipping attacks in federated learning

Deepti Joshi[1]*, Prof. John Augustine[2], Dr Rashmi P Sarode[2], Dr Preeti Mishra[1], Dr Vinod P[3]

[1]*Department of Computer Sciences, Doon University, Dehradun, India*
[2]*Department of Computer Sciences, IIT Madras, Chennai, India*
[3]*Cochin University of Science and Technology, India*
**deeptijoshi1532@gmail.com*

Abstract

In federated learning (FL), several clients cooperate to train ML models while maintaining their data privacy because of the local data storage principle. However, this feature leads to the vulnerability of FL algorithms to various types of attacks, including label flipping attacks (LFA). Current Federated Learning driven security architecture for Network Intrusion Detection System (NIDS) are prone to noisy dataset where an adversary may flip the label of dataset during training time. However, the robustness evaluation of such architectures in both symmetric and asymmetric label flipping attack scenario is less explored. "LFAs" refers to the action performed by malicious clients aimed at corrupting the training labels to lower the quality of the trained ML model. There are two kinds of label flipping attacks: Symmetric and Asymmetric LFA. In the symmetric label flipping case, all labels noise is done randomly, while in the asymmetric label flipping case, the labels are systematically misclassified/mapped to target labels, making detection more challenging. Symmetric label flipping weakens the global model as it contaminates all classes and asymmetric is known for its stealthy contamination in targeted classes (e.g, Dos --> Portscan, etc). These vulnerabilities are particularly critical in smart manufacturing environments, where distributed edge devices such as sensors and industrial control systems collaboratively train models for tasks like intrusion detection, predictive maintenance, and anomaly detection. Compromised local updates in such systems can lead to unreliable decision-making and potential system failures.

Prior to exploring defenses for such attacks, we explored existing approaches to mitigate LFAs, especially clustering-based techniques, while working with them we got several limitations, as they imply assumptions regarding the data distribution among the clients. Therefore, this work is a comprehensive empirical study of the impact of client data distribution on clustering-based defenses. As a result, it is discovered that these techniques demonstrated significant instability in their results across the training iterations, leading to unreliable detection of adversarial activity.

We have proposed an impact-based defense that works in four steps. First, the server records the baseline validation loss of the current global model. Second, for each client, it temporarily loads that client's weights, measures validation loss, then restores the original weights. Third, it computes the loss difference for each client— positive means the client hurts the model, negative means it helps. Fourth — and this is the key — instead of flagging everyone with positive loss diff, which would flag benign clients too, we use a statistical threshold of mean plus one standard deviation. Only clients whose loss increase is significantly above average are flagged as malicious. The empirical analysis has shown that the impact of the client updates on the global model, measured through validation loss, is a more reliable criterion for determining whether the client is malicious or not because malicious updates are likely to degrade model performance, while the honest updates improve it.

Based on this observation, a simple and effective impact-based defense mechanism is developed. The approach evaluates each client update by measuring the change in validation loss before and monitors

client-wise training loss deviation and identifies potentially malicious clients that are contaminating the datasets. For both the attack scenarios, we are using text-based data (CICIDS2017), although many works in this domain are explored on the image-based data- (e.g., MNIST, CIFAR-10). In our work for symmetric attacks, a signature-based filtering strategy is used, while for asymmetric attacks, an adaptive threshold based on the statistical distribution of loss differences is employed to detect more subtle adversarial behavior.

Experiment results demonstrate that the proposed approach achieves stable and consistent detection across training rounds and hence eliminates the instability observed in clustering-based methods. The method effectively reduces the attack success rate (ASR) and maintains model performance under realistic adversarial settings up to 70% malicious clients in the IID data distribution approach. Overall, the proposed impact-based evaluation framework provides a simple, interpretable, and computationally efficient solution for defending against label flipping attacks in federated learning, making it well-suited for secure and reliable deployment in smart manufacturing systems.

Abstract ID: 3162

Secure and Robust Federated Learning Framework for Detecting and Mitigating Anonymous Free-Rider Attacks

Ayush Maindola[1]*, John Augustine[2], Rashmi P. Sarode[2], Preeti Mishra[1], P. Vinod[3]

[1]Department of Computer Science and Engineering, Doon University, Dehradun, India

2Department of Computer Science and Engineering, IIT Madras, Chennai, India

[3]Department of Computer Applications, Cochin University of Science and Technology, India

*ayushmaindola2004@gmail.com

Abstract

Federated Learning (FL) is an emerging machine learning technique in which multiple independent machines, called clients, collaborate to train a global model in a privacy-preserving manner, where these clients never share any of their private datasets with anyone; they only share the updated model weights with the global aggregator. In such a setting, a free-rider attack is one of the most common and critical attacks, where free-rider clients contribute little to no valuable updates to the global model, but still benefit from it. These free-rider attacks damage model performance, the fairness of the federated system as well as increases the number of rounds and time taken by the model to converge.

In this paper, a reputation mechanism is developed that effectively handles anonymous free-rider attacks. The mechanism identifies, flags, and discards these malicious clients completely from the training phase. The system maintains reputations for each client based on the quality of updates they contribute. Any client whose reputation is constantly below the cutoff score is first flagged and kept in a probation period for a few rounds. If their reputation still does not improve, they are discarded from the training. This work is also relevant to smart manufacturing environments in Industry 4.0 and Industry 5.0, where multiple connected machines, sensors, and cyber-physical systems collaboratively train models using federated learning while preserving data privacy. By mitigating anonymous free-rider attacks, the proposed mechanism can reduce the convergence time, and the number of rounds increased by the anonymous free-riders.

A reputation system was developed. In the developed reputation system, the global model only shares the initial global weights, and afterwards, the global average weights are not shared with the clients for the initial 25 rounds; from the 26th round, the global model average weights start being shared with each client. The working of the developed reputation system is fairly simple; it works on two separate checks :

Reputation system based on giving the clients a reputation score based on the metrics, i.e., Accuracy, Precision, Recall, F1 score, and loss.

A score is given to the clients by calculating the Euclidean distance between the initialized global model weights and the weights updated by the clients.

For the first check, at the very first round each client is given zero (0) as the reputation score and then the training starts normally, for the initial 15 round the reputation system is not active, from the 16th round the reputation system initiates and flags the clients as free-rider whose reputation score is less than the Cutoff, for the developed reputation system the cutoff is calculated as :

The threshold controls the sensitivity of the reputation mechanism in identifying suspicious clients. In this work, the threshold is set to 1, meaning a client is flagged if its score falls more than one standard deviation below the mean of all client scores. This choice provides a balance between detection sensitivity and fairness. A lower threshold increases sensitivity but may lead to false positives by

flagging honest clients, while a higher threshold reduces false detections but risks allowing malicious clients to remain undetected. Preliminary experiments with threshold values in the range [0.5, 1.5] showed that a value of 1 offered the most stable performance in terms of convergence time and false detection rate. Therefore, the threshold acts as a balancing factor between accurately detecting free-riders and preserving honest client participation. As the reputation system is not active for the initial 15 rounds, the metrics of each client are averaged out for the initial 15 rounds, and then the reputation scores are calculated as:

$$Reputation\,Score = \frac{Accuracy\ + Precision\ +\ Recall\ +\ F1\,Score\ +\ Invereted\,normalized\,loss}{5}$$

Here, all the other metrics are "higher the better," but the loss is "lower is better" and the composite needs "higher is better" for everything. So loss is min-max normalized and flipped:

$$Inverted\,normalized\,loss\ =\ 1 - \left[\frac{Average\,loss\ -\ Min\,loss}{Max\,loss\ -\ Min\,loss}\right]$$

After the reputation scores for all the clients are calculated, they are compared with the cutoff score, and clients with reputation scores lower than the cutoff are saved as a set of possible free-rider clients.

For the second check, at the very first round, a copy of the initial global model weights is saved before sending to the clients to compare them afterwards by calculating the Euclidean distance between the initial global weights and the updated weights by each client; for an honest client, this is a meaningful non-zero number. For an anonymous free-rider, returned = sent, so update norm ≈ 0. By taking the union of both sets, a client is flagged if caught by either method. After the clients are flagged as free-riders, they are put in a probation period for 5 more rounds. Afterwards, the average of the client's last 5 reputation scores from probation is calculated. If their score improves to 85% of the honest clients' average, they are reinstated. If the client's reputation score does not increase in the probation period, they are completely removed from the training. The probation period is kept in the reputation system so that the clients who were falsely flagged as Free-riders in the warm-up rounds due to their poor quality of data or any other reasons can get another chance to get back into the training and can save themselves from getting removed from the training altogether. The reputation system identifies, flags, and completely removes anonymous free-rider clients from the FL training.

Experimental Results: It was seen in the baseline FL model that in the absence of any attack implementation, the model achieved an accuracy of 99.71, Precision 99.01, Recall 76.88, and loss 0.112, with the number of rounds taken for the global model to converge being 106 and the convergence time 130 min (without defence). On introducing attacks with 20%, 40%, 60%, and 80% of anonymous free-rider clients, the metrics did not significantly vary. However, the number of rounds taken for the global model to converge (without defence) was 141, 146, 183, and 263, and the convergence time was 155 min, 119 min, 99 min, and 80 min, respectively.

Once the defense was added, the number of rounds taken for the model to converge with 20%, 40%, 60%, and 80% of free-rider clients became 118, 144, 123, and 110 rounds, respectively, and the time taken was 134 min, 117 min, 71 min, and 32 min, respectively.

At 20% and 40% free-rider clients, the defence gave only minor improvements in rounds and time, +23 rounds and +21 min, and +2 rounds and +2 min, respectively, whereas at 60% and 80%, +60 rounds and +28 min, and +153 rounds and +48 min savings were realized, respectively. Thus, this defence is highly effective in detecting free-rider clients and excluding them in FL; the benefits were more prominent with increased free-rider fraction.

Abstract ID: 2518

A DPDP-Compliant Edge-AI Human Digital Twin Using Physics-Driven ODE and Adaptive Reinforcement Learning

Nagharajhan K P K*

Independent Researcher, Chennai, Tamil Nadu, India
**naghuirtt@yahoo.co.in*

Abstract

Industrial digital twins have matured through physics-driven modelling for asset state estimation, prognostic forecasting, and adaptive maintenance. This work investigates whether these architectural principles transfer to modelling a human being on a commodity smartphone, with/without wearable sensors or cloud infrastructure. We present a complete on-device Human Digital Twin that couples a 7-coefficient ordinary differential equation (ODE) for continuous biological-state estimation with epsilon-greedy reinforcement learning (RL) for user-specific adaptation. The twin additionally issues short-horizon predictions that it later verifies against realised outcomes and self-corrects, tightening the RL reward and yielding calibrated, honestly communicated confidence. In a 21-day pilot with 23 participants stratified across eight occupational profiles (agricultural labourers to software developers), the system achieved fatigue-trend prediction accuracy of 79.3% (±5.8%), sleep-phase classification agreement of 82.6% against Health Connect ground truth, and user-feedback concordance of 75.1%—all executing on-device at 0.33 Hz with sub-200 ms sensor-to-prediction latency. The current implementation makes no network calls of any kind—no cloud, telemetry, analytics, or third-party SDKs—and provides on-device data export and erasure, thereby satisfying the data-minimisation, purpose-limitation, and data-principal-rights provisions of India's Digital Personal Data Protection (DPDP) Act, 2023 by architecture rather than by post-hoc data governance.

Introduction and Motivation

While digital twins for industrial assets combine physics-based state models with sensor fusion and adaptive updating, their application to human health remains largely confined to cloud-dependent platforms that require dedicated wearable hardware. This limits reach: fewer than 2% of the Indian population owns a health-grade wearable, yet over 750 million carry smartphones equipped with accelerometers, gyroscopes, and screen-interaction sensors. We therefore ask whether a physics-driven digital-twin architecture—structurally comparable to those used for turbine or bearing prognostics—can be implemented entirely on a mobile edge device using only built-in phone sensors, while achieving meaningful personalisation through reinforcement learning. The principal contributions are: (i) a continuously integrated biological ODE executing at 0.33 Hz on ARM processors; (ii) an adaptive Q-learning framework that converges to user-specific health policies; (iii) a self-verifying prediction loop in which each forecast is checked against its realised outcome and used to correct the model; (iv) occupation-aware prognostic modelling validated across eight workforce strata; and (v) a multilingual natural-language interface for human–twin interaction.

Methodology

The system follows a five-layer Sense→Infer→Model→Predict→Communicate pipeline (Figure 1). Sensing: the accelerometer is polled at 200 ms with low-pass gravity separation (α=0.85); frequency-domain analysis classifies gait (walking 1.8–2.2 Hz, running 2.5–3.5 Hz), complemented by screen-unlock frequency, battery-drain rate, and Health Connect signals. Inference: a context engine fuses 14 sensor channels into activity states, including occupation-specific patterns (tractor whole-body vibration for farmers, loom-beat rhythm for weavers, delivery-rider stop–start cycles). Feature engines:

four parallel engines (Steps, Sleep, Heart, Posture) with independent Q-adaptive thresholds. ODE core: the biological state is H = [fatigue, stress, recovery, focus, energy], evolving as $dH/dt = f(H, S, \Theta, C)$, where S is the sensor vector, Θ is a 7-coefficient personal fingerprint learned from data, and C is temporal context (circadian phase and cultural calendar, including agricultural seasons and festival schedules). The ODE integrates Bannister's two-component model (CTL/ATL, τ=42 d/7 d), Borbély's two-process sleep model, Van Dongen's sleep-debt dynamics, Malik's RMSSD-based autonomic assessment, and EPOC-gated recovery. Adaptation: Θ adapts via epsilon-greedy Q-learning ($\varepsilon=0.18$, $\alpha=0.12$ decaying as $1/(1+n/100)$); each prediction is an action and each user tap a reward (±1.0), with Bayesian confidence reported via a Beta(α,β) posterior. A persistent background service continues state capture while the app is closed, maintaining an uninterrupted longitudinal record. Communication: Layer 5 renders ODE output as natural-language messages in 13 Indian languages with occupation-aware vocabulary. All computation—state estimation, ODE integration, Q-updates, and language generation—executes on-device, so no raw personal data leaves the phone.

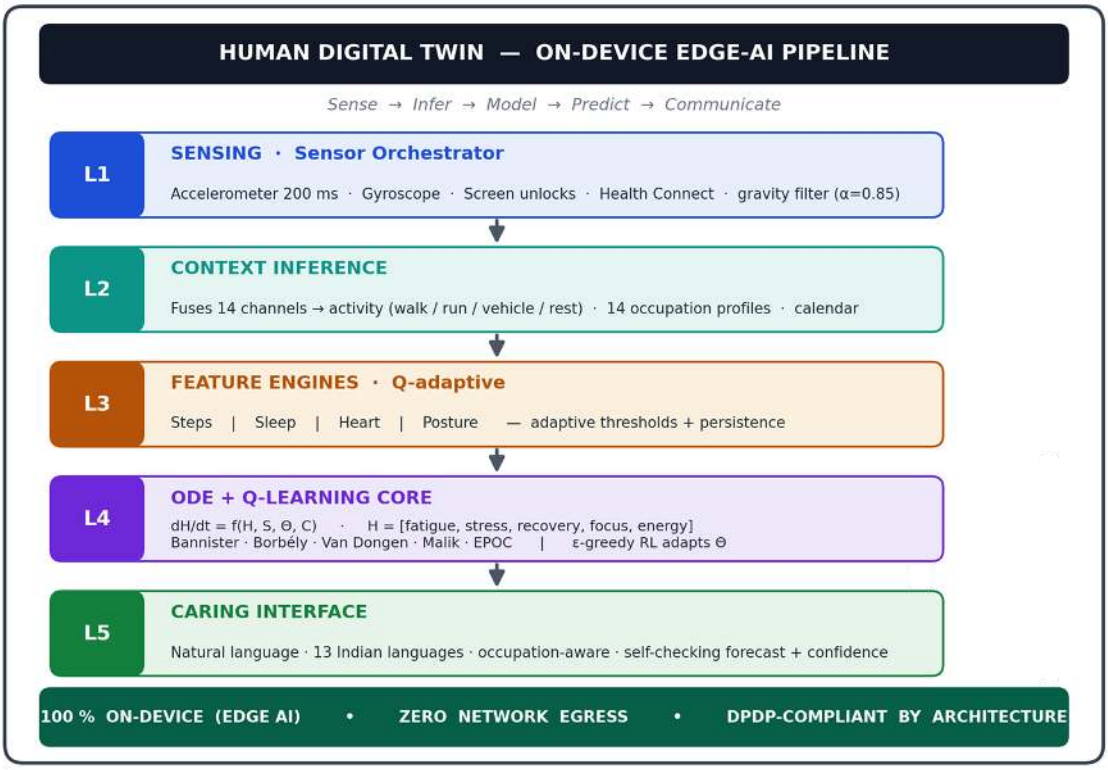

Figure 1. Five-layer on-device pipeline of the Human Digital Twin. The ODE + Q-learning core adapts a 7-coefficient personal fingerprint Θ from user feedback; all processing runs on-device with zero network egress (DPDP-compliant by architecture)

Results and Validation

The system was piloted with a Flutter-based, user-in-the-loop architecture optimised for ARM processors; the Android version is in pilot and the iOS version is in development. The 23-participant cohort was stratified across eight profiles (e.g., agricultural labourer, weaver, delivery rider; n=2-3 per stratum), assessing three metrics. (a) Fatigue-trend prediction: the ODE's 6-hour energy/fatigue forecast, compared against self-reported energy (5-point Likert, 3×/day), reached 79.3% (±5.8%) directional agreement by day 14, up from 54.1% on day 1, confirming Q-learning convergence. (b) Sleep-phase classification: accelerometer-based four-phase detection (awake/light/deep/REM), validated against Samsung Health via Health Connect (n=8 with wearables), reached 82.6% epoch-level agreement (30-min windows), with deep sleep at 75.4% and wake detection at 92.1%. (c)

Recommendation concordance: binary user feedback on oracle predictions improved action-selection accuracy from 51.8% (near-random, days 1–3) to 75.1% (days 15–21), with Bayesian posterior confidence exceeding 0.74. Occupation-specific thresholds showed statistically significant variance ($p<0.05$) between sedentary and physically active profiles, validating the need for occupation-aware modelling.

Discussion

The results indicate that the Sense→Model→Predict→Act architecture established for industrial cyber-physical systems transfers to human-health prognostics when the physics model is grounded in exercise physiology, giving predictions consistent with commercial wearables. On-device Q-learning together with the self-verifying prediction loop provides both personalisation and calibrated trust, while the zero-egress design ensures data sovereignty under the DPDP Act and enables cost-effective, smartphone-based human-centric digital twins across occupational sectors. A Phase-II study with 200+ participants and clinical-grade actigraphy is planned to extend these pilot findings. This work contributes to the ICIDTSM themes of Digital Twins and Cyber-Physical Systems, Physics-Driven Modelling and Simulation, AI and Machine Learning, Prognostics and Health Management, Sensing Technologies, and Human–Machine Collaboration, operationalising Industry 5.0's human-centricity by treating the worker not merely as the operator of a digital twin but as its subject.

Keywords: Human Digital Twin; Physics-Driven ODE; Adaptive Q-Learning; On-Device Health Prognostics; Cyber-Physical Systems; Human–Machine Collaboration; Edge AI; Industry 5.0

Abstract ID: 3472

Dynamic Beam Shaping using Coherent Beam Combining

Swathy S[1]*, Satyajit Maji[2], Balaji Srinivasa[2]

[1]International School of Photonics, CUSAT, India
[2]Indian Institute of Technology Madras, India
**swathys1210@ug.cusat.ac.in*

Abstract

This work investigates beam shaping in tiled-aperture Coherent Beam Combining (CBC) systems for the generation of near-uniform flat-top intensity profiles in the far field. The optimized dimension of the target square-shaped flat-top region is concluded from numerical simulation of such a CBC system by the use of phase optimization using a suitable composite metric function.

Introduction

High-power laser sources are integral to laser-based additive manufacturing, where beam quality directly governs process efficiency and part integrity. A flat-top (top-hat) beam profile, with uniform energy distribution across the interaction zone, has been shown to improve process stability, reproducibility, and material utilisation. Traditional beam shaping approaches, including refractive optics, diffractive optical elements, and spatial light modulators (SLMs), are inherently static, wavelength-sensitive, or limited by low optical damage thresholds. Coherent Beam Combining (CBC) overcomes these barriers: a multi-channel fibre laser array driven by a common seed source can produce programmable far-field patterns through phase-only control.

Methodology

The CBC system is modelled as a tiled-aperture hexagonal array of Gaussian beamlets (waist w_0 = 5 mm, λ = 1064 nm, spacing d = 12 mm) whose field is coherently superimposed with a controllable phase ϕ_i. A thin-lens phase transformation followed by the band-limited angular spectrum method (BL-ASM) provides accurate, alias-free far-field intensity computation. A composite merit function combining the Power-in-Target-Region (PITR) and intensity uniformity across the target is defined as:

$$J(\phi) = \mathrm{PITR} \cdot U \qquad (1)$$

where PITR measures power contained within a square target region of specified size, and U quantifies the flatness of the intensity distribution. The Adagrad adaptive gradient algorithm is employed to optimise the phase of each channel. An optimal mask size range is obtained based on the maximization of the merit function J.

Results

A 37-channel and 61-channel beam combining system in a hexagonal configuration is considered, and the corresponding source and target planes are shown in Figure 1 for the ideal phase-synchronized case with a circular diffraction-limited spot and when optimized for flat-top square intensity. For both configurations, the optimised far-field intensity distributions evolve through three distinct regimes as the target mask size (L) is increased, as shown in Figure

Though PITR increases with mask size, the uniformity shows initial enhancement, an intermediate plateau region, and afterward a slow degradation with mask size normalized with respect to the diffraction-limited central lobe size R_{PIB}. The results reveal a trade-off between power concentration and uniformity as the target mask size increases, and an optimal mask size ($L_{opt}/R_{PIB} \approx 4$) is found suitable for both 37-channel and 61-channel beam combining when equal weightage is given to the contained power and the intensity uniformity, i.e., w1 = w2 = 0.5.

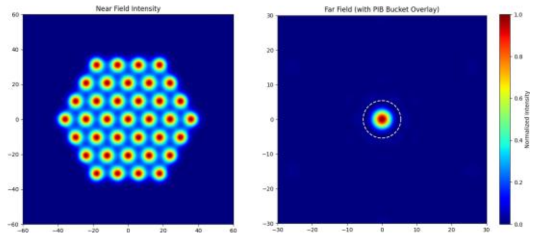

(a) 37-channel PITR

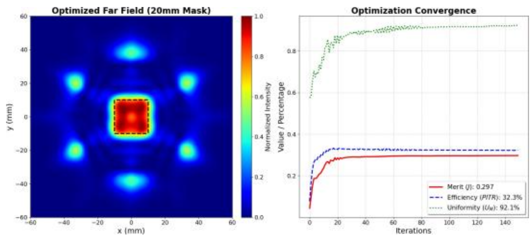

(b) 37-channel Uniformity

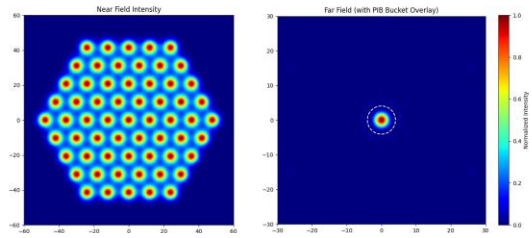

(c) 61-channel PITR

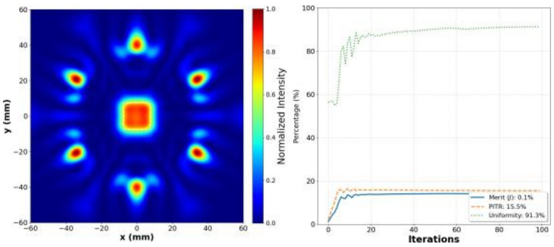

(d) 61-channel Uniformity

Figure 1. Square flat-top beam shaping performance for 37- and 61-channel CBC systems with active phase control.

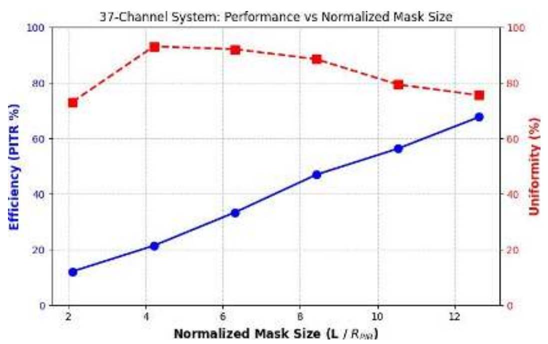

(a) 37-channel results

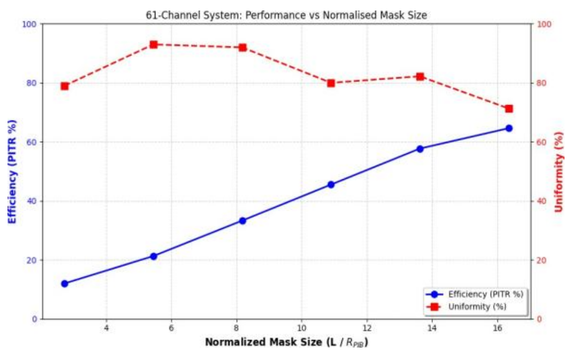

(b) 61-channel results

Figure 2. Variation of PITR and Uniformity with respect to mask size.

Furthermore, a preliminary study incorporating experimental phase noise data of high-power fiber amplifier reveals that the optimized system maintains beam quality under realistic conditions, with only moderate degradation in performance which will be presented at the conference.

Conclusion

This work demonstrates the feasibility of adaptive phase optimisation for scalable, programmable beam shaping in multi-channel CBC systems tailored to laser additive manufacturing. Identifying a universal normalized operating point ($L_{opt}/R_{PIB} \approx 4$) provides a physically grounded guideline for target mask selection, transferable between array configurations. This work demonstrates that adaptive phase optimization using CBC provides a scalable and programmable approach for dynamic beam shaping, with strong potential for high-power laser applications. Future work will address alternative mask geometries, systematic noise robustness characterisation, and experimental validation of the simulation framework.

Abstract ID: 3246

Capsule-Enhanced Generative Adversarial Networks for Financial Time Series: A Comparative Study of RNN Architectures and Mamba

Ankush Goyal*, Richa Golash, Ajay Kumar

SCSET, Bennett University, Tech Zone II, Greater Noida, Uttar Pradesh, India
**ankushgoyalmlk@email.com*

Abstract

Financial time series generation remains challenging due to inherent non-linearity, stochasticity, and complex temporal dependencies in market data. This paper introduces a systematic investigation of Capsule neural network (Caps-Net) enhanced generative adversarial networks (GAN) for multi-feature financial time series synthesis, particularly focusing on the dynamic routing characteristic of Caps-Net and preservation of hierarchical relationships. We evaluate through comprehensive ablation studies of four architectural variants that systematically decompose the contributions of temporal processing modules (Mamba state-space models vs. LSTM) and feature aggregation mechanisms (Capsule Networks vs. standard attention), implemented within a Wasserstein GAN framework with a gradient penalty. We demonstrate that capsule-based routing mechanisms significantly enhance distributional fidelity and cross-asset correlation modeling. The LSTM-Capsule variant achieves superior performance with the lowest Wasserstein-1 distance (WD) of 0.4049 (9.6% improvement over LSTM-Attention) and optimal correlation matrix preservation of 0.1045 (5.5% improvement), while the Mamba-Capsule architecture attains a Wasserstein distance (WD) of 0.430 (11.9% improvement over Mamba-Attention). This research is especially helpful in analysing three key aspects: (i) cross-asset dependency modelling, (ii) preservation measurement of multi-modal feature hierarchy, and (iii) evaluation of distributional fidelity across assets. In the future, we can extend the framework to multi-asset class synthesis with macroeconomic conditioning to enable scenario-based risk management and counterfactual policy analysis for institutional applications.

Abstract ID: 3119

Impact of Multimodal LLMs and Agentic AI in Smart Manufacturing – a systematic review

Devasena G*, Senthil Sivakumar M, Velmurugan C, Kamakshi K

Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli, Tamil Nadu, India
*devasenag@iiitt.ac.in

Abstract

The shift towards smart manufacturing in Industry 4.0 and Industry 5.0 is fuelled by the combination of cyber-physical systems, the Industrial Internet of Things (IIoT), and artificial intelligence. Two emerging AI approaches, Multimodal Large Language Models (MLLMs) and Agentic AI, are gaining attention for their potential to create intelligent, autonomous, and adaptive manufacturing systems. These technologies address the weaknesses of traditional AI, which is often limited to specific tasks and struggles with context in changing industrial settings. MLLMs improve on standard language models by combining different data sources, such as text logs, sensor data, machine vision images, audio signals, and time-series process data. This blending of data helps to provide a complete picture in manufacturing systems. Recent studies show that combining multimodal data greatly improves predictive maintenance, process optimization, and quality assurance in smart manufacturing systems.

A review by Leng et al. emphasizes the importance of multimodal deep learning in smart manufacturing. This approach combines sensor data and visual inspection to improve decision accuracy and robustness. The study shows that the move from unimodal to multimodal systems is recognized as a significant technological change that supports next-generation intelligent systems. In, Zhao et al. discuss the concept of Industrial Foundation Models (IFMs), where LLMs serve as a core component for intelligent manufacturing systems, enabling semantic reasoning and decision support. These models facilitate improved human-machine collaboration by allowing operators to interact with systems through natural language. A comprehensive review in Artificial Intelligence Review shows that LLMs can act as autonomous agents that interact with tools and external systems to perform complex tasks.

Agentic AI represents the next step in AI systems, emphasizing autonomy, goal-directed thinking, planning, and teamwork among different agents. Agentic systems can break down complex manufacturing tasks into smaller parts, use tools, and coordinate actions across various systems. In smart manufacturing, Agentic AI allows for smart agents that can oversee production processes, improve scheduling, and respond to immediate issues. The combination of MLLMs and Agentic AI creates a strong decision-making structure for smart factories. MLLMs provide context and understanding of multimodal data, while Agentic AI manages execution and allows for independent decision-making and workflow management. This collaboration supports intelligent automation throughout the entire manufacturing process, including design, production, inspection, maintenance, and supply chain management.

A recent study in Manufacturing Letters proposes an Agentic AI framework for smart manufacturing that integrates LLM-based agents with a Data–Model–Knowledge (DMK) layer. The framework enables perception, reasoning, planning, and optimization through coordinated multi-agent systems, as shown in Figure 1. Further, Ren et al. provide a systematic review of LLM-based agents and Agentic AI, highlighting their role in enabling autonomous decision-making and intelligent process optimization in manufacturing environments. The study identifies Agentic AI as a key driver for achieving adaptability and resilience in Industry 4.0 systems. The combination of Multimodal LLMs and Agentic AI is a new area of research that merges perception and decision-making into one framework.

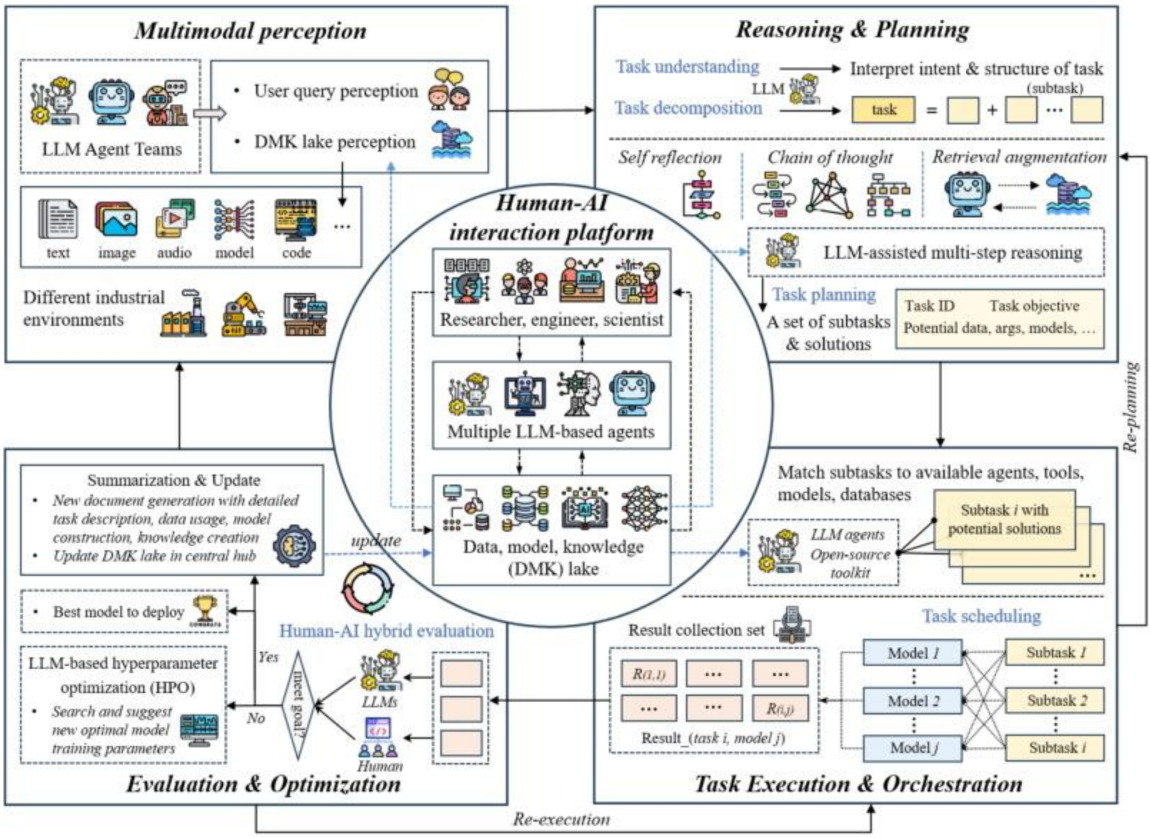

Figure 1. A novel agentic AI framework in smart manufacturing.

Recent studies show that MLLMs greatly improve agent abilities by offering better contextual understanding and enhanced reasoning across various data types. These systems support applications like real-time monitoring, anomaly detection, and adaptive control. A 2026 review on Industry 5.0 points out that the joining of generative AI and agentic systems will be crucial for human-robot collaboration and smart manufacturing environments. Agentic AI enables systems to independently plan and execute manufacturing tasks, reducing human intervention while maintaining adaptability. Continuous feedback between agents and physical systems enables real-time optimization and self-improving manufacturing processes. Natural language interfaces powered by MLLMs enhance usability, allowing engineers to interact with complex systems intuitively. Multi-agent architectures support distributed manufacturing environments and dynamic reconfiguration, thus increasing the efficiency of the production process.

Despite promising capabilities, several challenges remain: data privacy and industrial security need to be addressed without compromising the quality of the manufacturing system. Instead of developing the new system, the advanced technology should be integrated with legacy manufacturing systems. As it is necessary to process multimodal data in a smart manufacturing system, the computational cost of implementing Agentic AI models would be a great challenge. The explainability and trust need to be enforced since the agents are making autonomous decisions with less human intervention. Moreover, in an Agentic AI system, small agents are communicating with each other. Hence, protocols should be refined for efficient communication. The role of MLLM and Agentic AI has a lot of scope in a smart manufacturing system. Thus, the need for MLLM and the incorporation of Agentic AI in smart manufacturing are systematically reviewed, and their impacts are analysed.

Abstract ID: 3123

Developing a Generative DTDL Framework for Smart Manufacturing

Brijith Jacob, Santhosh Kumar G*

Department of Computer Science, Cochin University of Science and Technology, Kochi, India
**brijithjacob@cusat.ac.in*

Abstract

Smart manufacturing, driven by Industry 4.0 technologies, demands agile and adaptable systems. This work proposes the development of a novel framework that leverages Generative AI techniques and the Digital Twins Definition Language (DTDL) to revolutionize the modeling and simulation of Cyber-Physical Systems (CPS) in smart manufacturing environments. Digital Twins (DTs) have been an active area of research since its introduction in 2002. The emergence of DTs as pivotal enablers in CPS has catalyzed a growing interest in their systematic modeling and development. The project critically examines the current landscape of DT modeling, with a concentrated focus on the modeling methodologies, standards, architectural paradigms, and established frameworks that aim to facilitate generalizable and interoperable DT implementations. Ultimately, this study serves as a foundation for future research into the development of an integrated development environment (IDE) tailored for the design, modeling, and deployment of DTs aiming to streamline the creation process and ensure alignment with international standards and best practices. This proposal primarily focuses on the efforts to ease modeling of DTs by using DTDL, a widely accepted digital twin prototype specification. As an approach we propose a framework which will automatically generate domain specific DTDLs while text description about the domain components is given as input.

To automate the generation of DTDL, the framework uses a RAG (Retrieval Augmented Generation) based context aware LLM (Large Language Model) architecture. The framework utilizes a knowledge corpus implemented using PostgreSQL utilizing the pgvector extension. DTDL models from GitHub repositories are used as source documents after validation. These DTDLs are transformed into semantic vector embeddings using an embedding model and indexed within this database. This enables high-precision retrieval of relevant structural patterns based on user intent. When the user gives a text description of a component as input, the vector corresponding to the input text is created and using this a set of vectors is retrieved from the corpus. The retrieval is based on the similarity between the input vector and the document vectors in the corpus. Most similar two documents will be retrieved and passed as context to the LLM along with the input. This is the basic architecture of the proposed system. To specialise this system to a specific domain, say smart manufacturing, it requires enough source documents in this domain. For experimenting we have identified a manufacturing ontology based repository. To ensure that the automatically generated DTDLs follow not only semantic rules but also syntactical rules, one more source for giving syntactic context is also added. This source will have DTDL samples collected from Microsoft Azure Digital twin documentation. These samples follow different valid syntactic variations. Thus while prompting the LLM to generate the DTDL as per user input, the system gives two types of context. One context gives domain specific semantics based on the domain ontology and the second one corresponding to the syntax rules for DTDL. By maintaining the ontology based semantics which is domain specific, in a dedicated corpus stored offline, the system need not compromise the privacy and security of sharing proprietary/ confidential details related to the product. This is very important in industrial deployments. The architecture is compliant with both offline and online LLM architectures. Thus if performance takes priority over security and privacy an online LLM engine can be opted and in cases where security and Privacy demands high priority an

offline LLM engine can be used. Following gives an illustration of the architecture of the proposed system.

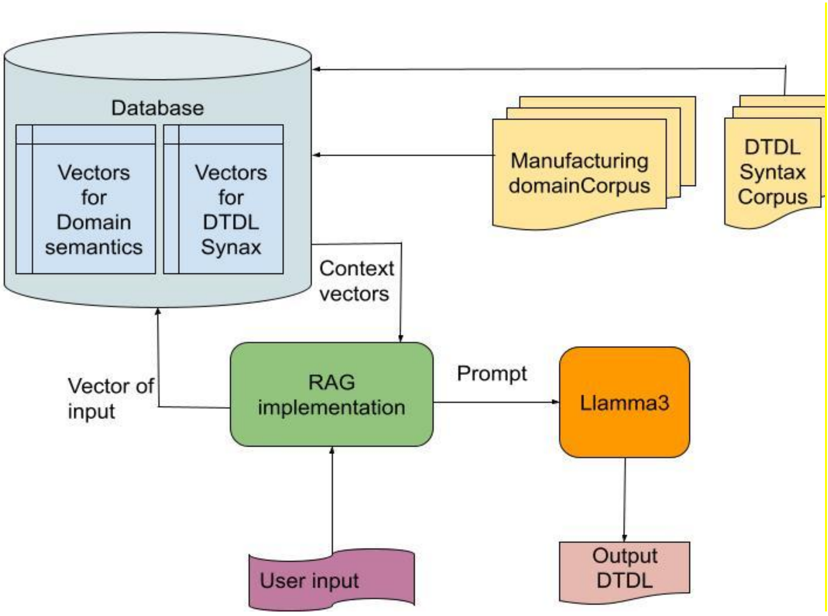

Figure 1. System architecture overview.

To ensure industry-grade structural integrity, all generated outputs are passed through an open-source DTDL validation engine. DTDL documents passing this validation are added to the pgvector store, facilitating a self-improving, automated learning loop. We plan to evaluate the framework's efficacy by comparing the automatically generated DTDL documents with the actual DTDL documents from the repositories. For this twenty percent of DTDL documents will be reserved and will not be stored as corpus. These documents will be used for comparison and the similarity between these documents will be used as the evaluation metric. We have also planned to experiment by adding a third context which is also semantic context but created using wot-json-manufacturing ontology documents (not in DTDL format). The resulting DTDLs generated as output will be compared to the results of the previous set up. The results will be compared for any improvement in the quality of DTDLs generated.

An experimental evaluation for multiple open-source LLMs through a manufacturing case study, quantifying performance via the percentage of syntactically valid DTDL models generated, is also planned as future work. Manual DTDL authoring is time-intensive and error-prone. An automated approach like the one proposed will reduce the overall design and development efforts to a large extent. Especially fine granular level component description in a standardised format will improve reusability in a larger context.

Abstract ID: 3149

Tiny ML-Compatible Lightweight Physics-Informed Digital Twin for Edge Systems: A Simulation Study with Exploratory Federated Unlearning

Avantika V*, Kavya R, Sandhya Chengaiyan, Jailani Beevi M

Department of Electronics and Communication Engineering, Meenakshi Sundararajan Engineering College, Chennai, India

* *parthaven17@gmail.com*

Abstract

The evolution of industrial digital twins highlights the need for systems that are lightweight, adaptive, and capable of operating in edge networks that are limited in resources. Physics-Informed Neural Networks (PINNs) provide predictions that follow the laws of physics and remain consistent, which is preferable for a digital twin that represents a real-time system. Another recent development is Federated Unlearning, which performs the removal of corrupted or unreliable data in distributed or decentralized systems. Federated Unlearning promotes post-hoc corrections, which enable the model to be updated even after the introduction of corrupted data. Recent studies have explored PINN-based learning, federated learning, and unlearning, contributing significantly to the development of next-generation digital twins. However, the integration has been focused more on complex infrastructure digital twins. This work aims to integrate Physics-Informed Learning, Federated Learning, and Unlearning for resource-limited devices that run using Tiny-ML.

This work proposes a feasibility-focused simulation pipeline that is compatible with Tiny-ML and a lightweight Physics-Informed Digital Twin. Performance was evaluated under simulated sensor variability due to noise and contextual disturbances. A lightweight version of PINN was implemented using physics-based regularization, along with smoothness and bounded behaviour constraints, to ensure that the outputs are physically consistent, while maintaining low computational complexity for resource-constrained Tiny ML devices.

For the digital twin, a reduced-order behavioural model of a CNC spindle system was implemented in Simulink. Healthy and faulty simulated sensor data, namely rotational speed (RPM), motor current, and feed rate, were generated to produce a healthy baseline snapshot (57 samples) and a faulty dataset (3009 samples). Both datasets were normalized before developing the model. The healthy baseline was used as the reference. A lightweight version of PINN, consisting of a three-input, two-hidden-layer multi-layer perceptron of 64 neurons per hidden layer, resulting in 4,611 trainable parameters, was implemented. The training phase integrated supervised reconstruction loss, a physics-informed residual, smoothness, and bounded-behaviour regularization for physically consistent predictions. The model was trained using the Adam optimizer with gradient clipping, validation-based early stopping, best-model restoration, and a convergence-based stopping criterion to prevent update saturation. The fault dataset was divided into three virtual federated clients that refined local models before global aggregation using Federated Averaging (FedAvg).

Physics residuals were then evaluated as a physics-guided anomaly indicator for corrupted client behaviour, after which the client with the largest residual was excluded using Exact Re-Aggregation. Finally, three lightweight post-unlearning stabilization rounds were performed before evaluation through mean squared error (MSE), Recovery Rate (ρ), Model Drift (Δ), and host-side computational benchmarking.

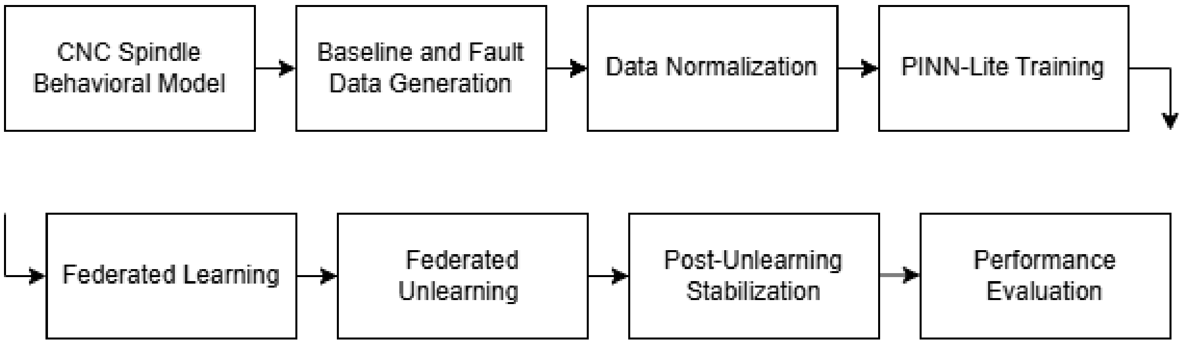

Figure 1. Workflow of the Proposed Tiny ML-Compatible Physics-Informed Digital Twin with Exploratory Federated Unlearning

The results show that the framework was successful in reconstructing the normal operating behaviour of the CNC spindle and maintaining consistency through physics-informed regularization. Training convergence was stable through an early stopping mechanism, with a best validation loss of approximately 0.0724; a low physics residual was maintained, which indicates physically consistent spindle behaviour. The healthy baseline MSE was 0.067427, which increased to 0.076689 due to the corrupted client data. After Federated Unlearning using Exact Re-Aggregation and lightweight post-unlearning stabilization, the reconstruction error decreased to 0.071536, with Recovery Rate (ρ) of 0.5563, indicating partial recovery of model performance and corruption removal. Physics residual analysis identified Client 3 as the anomalous participant with the highest residual value (0.233518), which served as the criterion for initiating the unlearning process. The recovered model produced a Model Drift (Δ) of 0.116982 and a relative parameter drift of 1.6116%, which indicated that stability improved after removing corrupted updates.

The lightweight PINN architecture needed only 4,611 trainable parameters, approximately 18.01 KB of parameter storage, a memory footprint of approximately 18.52 KB, and achieved an average host-side inference latency of 0.0805 ms, demonstrating low host-side inference latency suitable for future evaluation on resource-constrained digital twin platforms.

The proposed framework demonstrates the feasibility of integrating a reduced-order behavioral CNC spindle model with lightweight physics-informed learning and Federated Unlearning for resource-constrained industrial applications. The physics-informed residual keeps network predictions consistent while the additional smoothness and bounded-behaviour constraints improve prediction stability.

In the federated learning process, physics residuals serve as a physics-guided anomaly indicator to identify corrupted client updates before reconstructing the global model. Exact Re-Aggregation and lightweight post-unlearning stabilization were able to restore part of the degraded model performance, with relatively low model drift, which implies stable recovery after corruption removal. Furthermore, the implemented convergence-control mechanisms that included validation-based early stopping and update-saturation detection, prevented unnecessary optimization after convergence had been achieved.

However, the study is limited by the use of simulated federated clients, a relatively small healthy baseline dataset, host-side computational bench-marking rather than embedded hardware deployment, and fixed physical coefficients. Therefore, concepts including long-term autonomous adaptation and self-referential drift are discussed as future research directions. Future work will focus on embedded Tiny ML deployment, adaptive estimation of physical parameters, larger distributed federated networks, and real-time autonomous digital twin operation.

Abstract ID: 3136

Designing a Distributed Ledger-Enabled Trust-Aware Federated Learning Framework for Auditable Backdoor Detection in Collaborative Intrusion Detection Systems

Anurag Bansode[1]*, John Augustine[2], Rashmi P. Sarode[2], N Renugadevi[1]

[1]*Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli, India*
[2]*Indian Institute of Technology Madras, Chennai, India*
**242101@iiitt.ac.in*

Abstract

Federated Learning systems enable training of models among varied clients without the need of revealing its private data. The decentralized nature of preserving confidential or private data is also a source of hostile manipulation, which is heavily exploited with a backdoor attack. Recent studies in this field have highlighted that the attackers introduce malicious parameters into global models without degrading the overall performance. Techniques such as FedESP and Distributed Backdoor Attacks have shown that attackers, with the use of adversarial training and trigger optimization, can hide the triggers that bypass existing defensive mechanisms. Malicious attackers can rejoin the system in the absence of robust identity accountability with the creation of a different identity, referred to as whitewashing.

To address these challenges, we propose a trust-aware federated learning system that assumes that no single rule aggregation can fully address the modern attacks. A layered strategy is adopted for the defence, which integrates identity anchoring that prevents whitewashing, malicious update scoring, which is an anomaly-based scoring to detect and penalize malicious model updates, and auditability through distributed ledger technology, making it more resilient for secure federated learning.

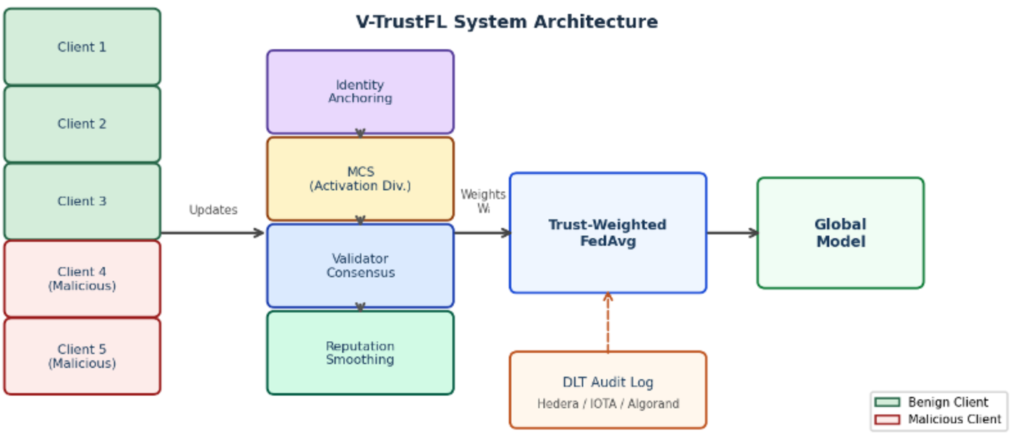

Figure 1. V-TrustFL Architecture Diagram

Methodology

Experiments were conducted on the CICIDS2017 dataset, where the class was recategorized to 6-class intrusion detection, SECOM (a semiconductor manufacturing dataset), and CMAPSS (a 3-class equipment health dataset). The system trains a tabular Neural Network under non-IID federated data partitioning and uses 10 clients for 100 rounds. Dirichlet-based partitioning is done to emulate the heterogeneous client data. The framework records global accuracy, weighted F1, Attack Success Rate (ASR), False Positive Rate (FPR), False Negative Rate (FNR), and whitewashing metrics.

A set of defensive signals is incorporated to evaluate the client updates as such:

Identity Anchoring and Reputation Management: Each client is assigned a persistent identity with the help of a SHA-256 hash of its public key. This enables the system to track behaviour across rounds, accumulate exclusion counts and implement banning policies, which mitigate whitewashing attempts.

Malicious Confidence Scoring (MCS): Based on the principle mentioned in the framework, utilized the differential testing and neuron activation fingerprinting, which monitors the scenario in which a particular set of neurons is activated based on the deviation from the majority. The model with high deviation is then identified as an anomaly and is excluded.

Validator-Based Behavioural Scoring: A pool of validators is created based on which checks for the model performance on a private set of data of the selected clients, which ensures the global model has good convergence. Divergence between the client output and the group consensus is measured using robust median-based consensus to minimize the influence of any client

Trust-weighted Aggregation: Based on the MCS and reputation score, the system calculates the effectiveness of the client's contribution to the global model. The system penalizes the client that it finds to be suspicious.

Forensic accountability with Distributed Ledger: A hash-chained local audit log is maintained, which logs the decision made by the framework. These can also be anchored to Distributed Ledgers like hedera, IOTA, or Algorand, which ensure that the decisions taken are tamper-evident and verifiable by external parties.

Results and Discussion

The result showed that the baseline had a mean global accuracy of 88.22%, weighted F1 of 82.74%, mean loss of 57.976, and ASR of 92.53%. In the case when defence is applied, the accuracy and weighted F1 are similar for CICIDS2017. However, there were behavioural differences where the mean FPR rises to 22.86%, which indicates that the defence ignores some benign clients. A mean FNR of 33% shows that the malicious partition is not suppressed and escapes a ban. This means that the influential clients escape successfully.

For SECOM, the baseline had an accuracy of 92.42% and a weighted F1 of 89.69%. ASR was low in the dataset, which was 20.35%. CMAPSS was the most difficult to attack as the accuracy was as low as 0.23, while the ASR was extremely high; the trust-aware solution did increase the accuracy by 1%. The false positive behaviour varied by dataset. CMAPSS had 18% compared to SECOM, which was 5%.

These findings suggest that the framework's contribution of operationalizing identity continuity, per-round anomaly detection validator consensus, and DLT auditability demonstrates that multiple trust mechanisms can be integrated into a single FL system. Future work will focus on mitigation of backdoor attacks and better calibration between MCS and validators, further complemented by an investigation into DLT's role from passive auditability to active security enforcement.

Abstract ID: 3132

Autoencoder-Driven Prediction-Error Triggered Digital Twin Synchronization on TinyML Edge Nodes for Bandwidth-Efficient IIoT Predictive Maintenance

Abirami J*, Deeptipriya D, Sandhya Chengaiyan, Gaurav R V, Anoab Balan T

Department of Electronics and Communication Engineering, Meenakshi Sundarajan Engineering College, Chennai – 600 024, Tamil Nadu, India

**abiramijeyaseelan05@gmail.com*

Every Industrial Digital Twin deployed today makes the same silent assumption: that a physical machine must continuously stream its sensor data to remain faithfully mirrored in its virtual replica. This assumption wastes information. The acoustic signature of a machine operates with predictable patterns, which result in almost no new information across all sensor windows during normal operation. The Shannon framework establishes that expected events carry no information. A factory with 500 machines streaming raw acoustic features at 31.25 windows per second generates over 2.5 million feature vectors per hour, the overwhelming majority redundant. This paper demonstrates that Digital Twin synchronization should match the frequency of machine anomalies instead of following the rate of sensor data acquisition from machines.

Methodology:

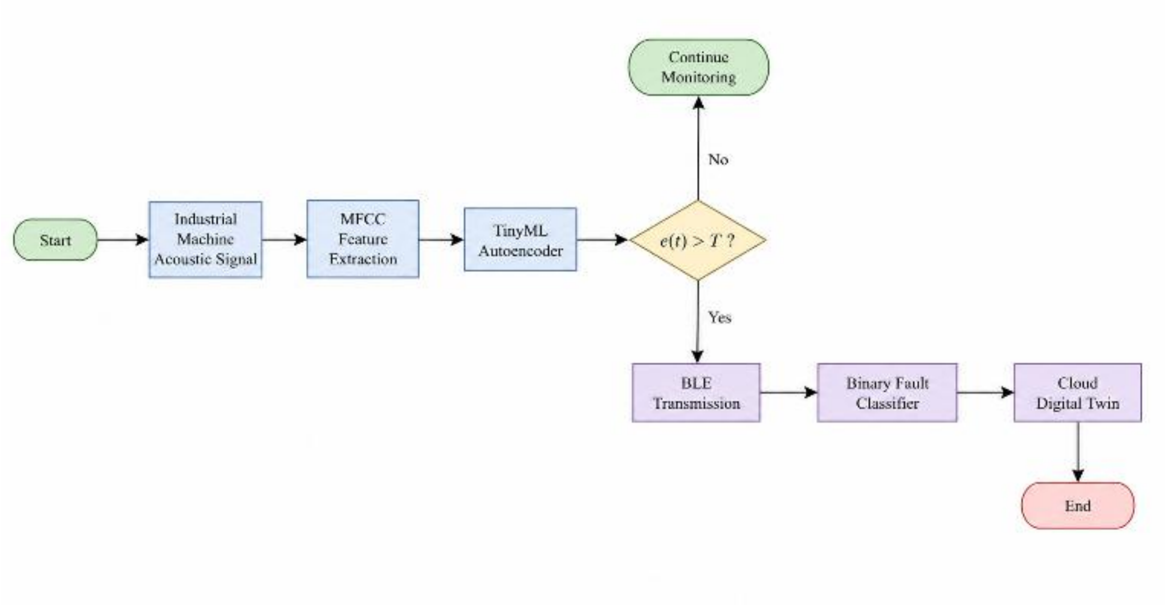

Figure 1. Methodology.

We present the Prediction-Error-driven Digital Twin (PEDT) framework, implemented as a verified dual-branch TinyML architecture on an ESP32-S3 microcontroller. The system operates as follows. Audio is captured at 16 kHz and transformed into 68-dimensional feature vectors comprising 7 time-domain statistics, 9 spectral features, and 52 MFCC plus delta coefficients per 500 ms window with 75% overlap. These features feed two INT8-quantized models deployed simultaneously via TensorFlow Lite for Microcontrollers.

The first branch is an autoencoder (Dense: 68→48→24→8→24→48→68 with batch normalization, 19.1 KB INT8) trained exclusively on normal-operation acoustic data using the unsupervised protocol on MIMII Fan id_00 at 0 dB SNR — 11,468 normal frames with no anomaly labels used during training.

The autoencoder computes per-window reconstruction error $e(t) = \|x - \hat{x}\|^2$. The Digital Twin synchronization gate uses this error threshold: if $e(t) < \theta$ (threshold = 0.6333, set at the 95th percentile of training error), the edge node transmits nothing over BLE. If $e(t) \geq \theta$, an 8-byte BLE packet is transmitted containing the error magnitude and timestamp, triggering the second branch.

The second branch is a supervised binary classifier (Dense: 68→64→32→16→2 with batch normalization and dropout, 13.4 KB INT8) that identifies the fault state at 85.3% accuracy. The classifier activates only upon autoencoder-triggered BLE events, providing confirmed fault identification without continuous transmission. The Digital Twin state is formally defined as:

$$\text{Twin}(t) = \text{Prior}(t) + \text{ErrorCorrection}(t)$$

where the cloud-side prior model runs continuously and applies error corrections only when the edge triggers transmission.

Both models are converted to INT8 TFLite format and verified end-to-end: the autoencoder correctly scores normal samples below threshold (MSE: 0.507) and anomalous samples above threshold (MSE: 0.681); the classifier correctly identifies healthy samples at 88.3% confidence and anomalous samples at 99.6% confidence. Total combined flash footprint is 32.5 KB — 0.4% of ESP32-S3 available flash

Results:

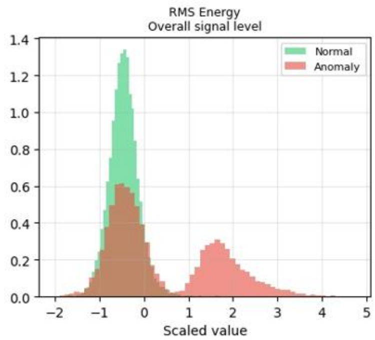

Figure 2. RMS energy distribution showing higher energy in fault states, enabling anomaly-triggered transmission.

Evaluated on 4,762 held-out test frames (2,381 normal, 2,381 anomalous) from MIMII Fan id_00 at 0 dB SNR, the classifier achieves 85.3% overall accuracy with 88% precision on anomaly detection and 82% recall. The autoencoder registers a false alarm rate of only 5.6% — meaning 94.4% of normal operation windows generate zero BLE transmission, yielding a theoretical bandwidth reduction of 91.2% compared to continuous 68-feature streaming (equivalent to 1,088 bytes/sec vs. 12,240 bytes/sec at 31.25 Hz). Feature importance analysis identifies RMS energy (impact: 0.245) and spectral entropy (impact: 0.117) as the dominant discriminative features, with MFCC coefficients 0, 4, and 8 collectively contributing 0.328 impact — consistent with known acoustic fault signatures in rotating machinery. The dual-branch architecture resolves the fundamental tradeoff in single-model approaches: the autoencoder provides conservative, low-false-alarm gating (5.6% FPR) while the classifier provides high-confidence fault confirmation (85.3% accuracy) only when the gate triggers.

Discussion and Conclusion:

The PEDT framework's central contribution is the architectural unification of anomaly detection and Digital Twin synchronization triggering into a single edge operation. The reconstruction error is not merely a fault indicator — it is the synchronization intelligence of the twin, determining when the virtual replica needs to know something has changed. The system eliminates separate detection, scheduling, and transmission pipeline stages while its synchronization bandwidth scales according to machine acoustic anomalies.

The unsupervised training protocol — requiring only normal-operation recordings collected during machine commissioning — is a decisive practical advantage. No historical fault data is needed. The autoencoder detects any deviation from learned normality regardless of failure mode, making the PEDT framework immediately deployable on new machines without labeled datasets. The complete dual-branch system operates within microcontroller-class hardware constraints, enabling deployment at machine level without gateway infrastructure.

Future work will extend the PEDT framework to multi-machine federated architectures and evaluate synchronization bandwidth reduction across all four MIMII machine categories under variable SNR conditions.

Abstract ID: 3066

Low-Power AI Inference Accelerator: MAC Unit Optimization via Operand Isolation and Switching Activity Reduction

Ali Mohamed Ali Elsharef[1]*, Abhishek Kumar[1], Prof. S.C. Sharma[2]

[1]*Department of Applied Physics, Delhi Technological University,*

[2]*Department of Engineering Physics, Delhi Technological University*

**alimovlly@gmail.com*

Abstract

With the rapid increase of edge computing and the proliferation of Artificial Intelligence (AI) applications on IoT and embedded systems platforms, there is an increasing need for high-inference throughput hardware accelerators operating within tight energy constraints. General purpose processors such as CPUs and GPUs are not a good fit for such applications because of the high dynamic power consumption required while inferring deep neural networks. The Multiply Accumulate (MAC) operator that forms the basis for computations within Deep Neural Networks (DNNs) has become the main focus in achieving lower power consumption. This research proposes a low power design for an AI inference accelerator using a FIR approach, implemented in MATLAB.

This methodology utilizes a two-stage process. In the first stage, a baseline FIR filter implementation is performed with a conventional MAC design that uses 10 symmetric taps for 200 random Gaussian inputs. Each input invokes one MAC operation, giving a total of 1,910 MAC operations. In the second stage, a modified approach is used whereby a threshold-based technique is applied to identify operands close to zero. If the magnitude of an input is less than 0.2, it is categorized as a near-zero operand, and the corresponding MAC operation is ignored.

Switching activities are calculated based on Switching MAC Operation Count, wherein MAC operations that do not involve processing non-zero state-changing data are excluded. The power consumption is estimated according to $P = \alpha C V^2 f$, with α representing the normalized switching activity, C equal to 1 pF, V equal to 1 V, at frequencies from 0.5–10 MHz.

Experimental results from seven separate trials show that there is always reduced power. The baseline MAC count is 1,910, whereas the optimized MAC count lies between 1,544 and 1,665. The average switching activity is 0.852, resulting in an average dynamic power saving of 15.4% and average EDP improvement of 28.4%.

Table 1. MAC operation count and power reduction results (Baseline MAC = 1910).

Trial	Optimized MAC	Switching Ratio (α)	Power Reduction (%)	EDP Improvement (%)
1	1656	0.867	13.3	24.8
2	1589	0.832	16.8	30.8
3	1650	0.864	13.6	25.4
4	1623	0.850	15.0	27.8
5	1544	0.808	19.2	34.7

Trial	Optimized MAC	Switching Ratio (α)	Power Reduction (%)	EDP Improvement (%)
6	1665	0.872	12.8	23.9
7	1585	0.830	17.0	31.1
Mean	1616	0.846	15.4	28.4

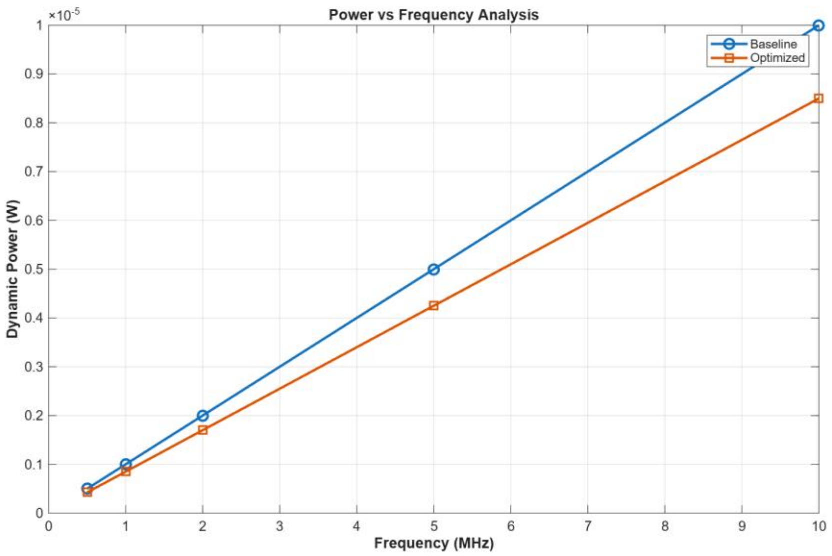

Figure 1. Power vs. Frequency analysis: Baseline vs. Optimized MAC (C = 1 pF, V = 1 V, 0.5–10 MHz). Optimized design consistently reduces dynamic power across all frequencies.

Figure 1 verifies that the optimized design continues to exhibit superior performance over the baseline in terms of power at all frequencies within the operating range. Dynamic power for the baseline is observed to be around1.0×10^{-5} W and 0.85×10^{-5} W for the optimized design at 10 MHz Both graphs show linear growth with frequency according to the $P = \alpha C V^2 f$ formula.

The results obtained from this study corroborate those documented in the existing literature related to RTL-based power minimization strategies. Gundrapally et al.3 reported a maximum of 29.4% reduction in power consumption using Enhanced Clock Gating (ECG) and Local Explicit Clock Enable (LECE) for the ResNet-20 algorithm implemented in an FPGA. We have been able to achieve the same level of optimization, on average, at 15.4% through operand separation in the current case before actual hardware synthesis. Similarly, Liu et al.6 recorded a maximum switching activity reduction of 65% using zero-skipping (LazyGPU) under high sparsity conditions. Our current design employs a threshold value of 0.2 to skip approximately 15.8% of the multiply-accumulate operations in a Gaussian distribution. In summary, operand isolation based on threshold values at the MAC operation level achieves consistent reductions in dynamic power consumption in the range of 12.8%-19.2%. The Switching MAC Operation Count approach represents a strong and data-independent methodology for estimating power in high-level simulations. These findings lay down the foundation for an implementation in full RTL form where additional benefits can be obtained using clock gating and zero skipping techniques. Future work will implement this design in Verilog/VHDL and synthesize on FPGA to obtain physical power measurements via Xilinx Power Estimator for direct comparison with the simulation benchmarks presented here.

Abstract ID: 3180

A Real-Time Virtual Reality-Based Digital Twin Framework for Thermal Error Prediction of Machine Tool Spindles

Anirban Tudu*, Adithya Anand, Shashi Bhushan Gunjan, D.S. Srinivasu

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai-600036, India
*anirbanjumech13@gmail.com

Abstract

The manufacturing industries struggle to balance high production rates and precision in parts machined by high-speed motorized spindles in the era of Industry 4.0. The spindle's thermal stability - driven by its complex geometry, transient heat flows, and dynamic loads - directly causes structural deformations that compromise machining accuracy. Thermal errors represent a significant challenge in high-precision machining, often accounting for a large percentage of total spindle displacement. Spindle thermal error can introduce up to 75 % of total error in machined component. Traditional error prediction methods typically lack real-time interactivity. Recent digital twin-based solutions address this by enabling real-time predictions with minimal delay. The present digital-twin (DT) based thermal error predictions in literature lacks intuitive visualization.

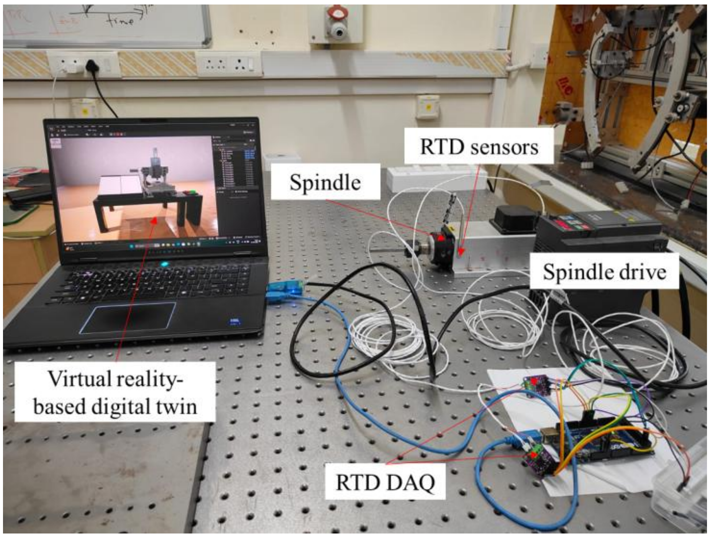

Figure 1. SpindleTwin: Digital twin for machine tool spindle.

This paper proposes a digital twin framework that leverages co-simulation between Simulink and Unreal Engine to create a predictive environment with virtual-reality (VR) based intuitive visualization. The framework is designed to handle complex thermal error prediction models, providing users with a photorealistic "SpindleTwin" that reflects actual hardware behavior while predicting future thermal states (Figure 1). This study introduces a multi-domain digital twin framework for the real-time prediction, visualization of thermal errors in high-precision machine tool spindles, addressing a critical bottleneck in advanced manufacturing. By seamlessly integrating Unreal Engine 5 (UE5) for photorealistic, immersive VR visualization, Simulink for coupled multi-physics modeling of thermal, structural, and dynamic behaviors, and a Hardware-in-the-Loop (HIL) architecture, the framework bridges virtual simulations with physical hardware to deliver high fidelity and interactivity. Central to the system is a robust communication gateway via USB-to-RS485 adapter, enabling bidirectional data exchange with a Delta MS300 Variable Frequency Drive (VFD) and distributed physical sensors (e.g., RTDs, and displacement probes). The closed-loop framework integrates real-time sensor data with predictive machine-learning models (multi-physics simulations augmented by physics-informed neural networks in this study) on an Intel® Core™ Ultra 9 185H processor with 32 GB RAM computational facility, producing thermal error forecasts with sub-micron accuracy. Unlike conventional DTs limited to offline analysis, this framework can be scaled to live HIL tuning of cooling or error compensation strategies, enhancing process stability and productivity with its seamless two-way communication. The proposed approach not only advances DT scalability across multi-domain interactions but also sets a benchmark for AI-augmented manufacturing, with extensible applications to predictive maintenance and precision machining processes.

The digital twin framework is divided into four primary layers (Figure 2):

- *Interaction and Visualization Layer (Unreal Engine 5):* Acts as the front-end interface. Users can operate the spindle through AR/VR headsets (e.g., Meta Quest), adjusting parameters like virtual speed or toggling different thermal error scenarios.
- *Computational and Control Layer (Simulink):* Serves as the "brain," housing the acquisition of data and running machine learning models for thermal error prediction.
- *Communication Layer (Modbus):* Provides the bridge between software and hardware. A USB-to-RS485 adapter is used for low-latency data exchange between Simulink and the controller, while Modbus RTU facilitates communication with the VFD.
- *Physical Layer:* Comprises the Delta MS300 VFD, the spindle motor, and the thermal sensors.

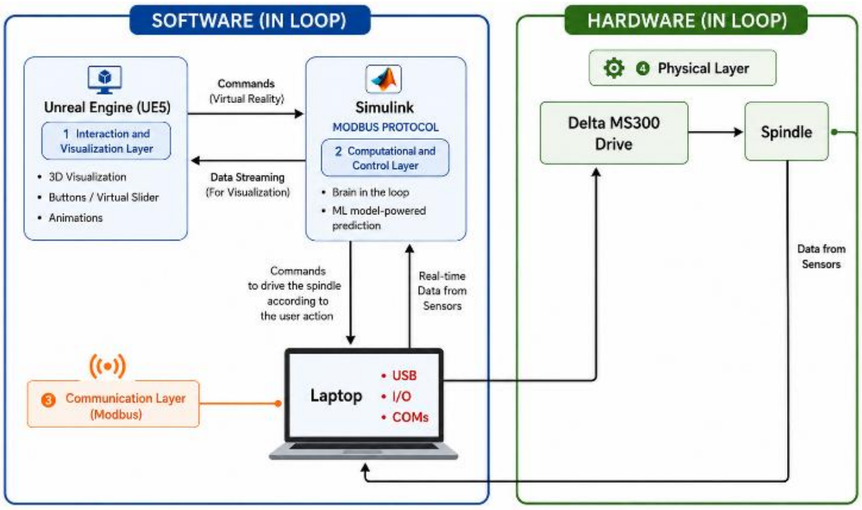

Figure 2. Schematic diagram of the "four primary layers".

The framework employs a USB-to-RS485 adapter for its industrial communication capabilities. The controller is wired to the Delta MS300 via RS485 for Modbus RTU communication. The thermo-mechanical behavior of the spindle is modeled in Simulink. This model receives the desired setpoint from the VR interface and actual temperature feedback from the RTD sensors. The controller computes the necessary commands to drive the Delta MS300 while simultaneously simulating the thermal expansion based on the current operational state. The data from Simulink is fed back into UE5 to drive the spindle animation. Beyond simple rotation, the framework can visualize different thermo-mechanical changes, such as heat-affected spindle dynamics. Predicted thermal errors can be overlaid on the 3D model, allowing researchers to inspect deformation and chatter within a virtual workspace.

The execution sequence follows a strict loop to ensure synchronicity:

1. *VR Interaction:* A user in a headset sets a target speed via a UI widget or controller input.

2. *Setpoint Transmission:* UE5 sends the setpoint to Simulink via a UDP link or native co-simulation blocks.
3. *Physical Actuation:* Simulink writes the command to the USB-to-RS485 adapter Modbus node, which translates it into a Modbus write to the Delta MS300.
4. *Sensor Feedback:* The physical spindle rotates; the speed sensor and temperature sensor feed pulses back to the Simulink, which computes the thermal error.
5. *Twin Update:* The measured temperature and error are sent back to Simulink and Unreal Engine, ensuring the virtual spindle reflects the predicted thermal displacements of the physical asset.

The framework achieved average MAEs of 1.15°C at the front bearing, 0.95°C at the rear bearing, and0 2.5°C at the motor. Additionally, the method yielded an average RMSE of 2.42μm for axial tool deformation. Therefore, the proposed framework offers a robust, low-cost solution for developing digital twins focused on machine tool spindle research. By combining the high-fidelity visualization of Unreal Engine with the deterministic control and modeling in Simulink, the present work can effectively predict and visualize thermal errors in an immersive environment.

Abstract ID: 3118

Integration of 3D Material Printing, AI, and IoT for Smart Manufacturing: A Review on Industry 4.0 Automation

Senthil Sivakumar M*, Velmurugan C, Devasena G, Jagadeeshwar Pashikanti

Indian Institute of Information Technology Tiruchirappalli, Trichy, Tamil Nadu, India
**msenthilsivakumar@iiitt.ac.in*

Abstract

The recent development of Industry 4.0 has transformed traditional manufacturing standards by integrating digital technologies, cyber-physical systems, and intelligent automation. Additive manufacturing, popularly known as 3D material printing, is an emerging technology used as a foundation for smart manufacturing. The conventional 3D printing process has difficulty in achieving flexibility in printing complex geometries, rapid prototyping, and customised production. To overcome these challenges, the computer vision-based AI-driven techniques are increasingly used in the 3D printing process. This paper presents a comprehensive overview of the 3D printing process and the integration of digital technologies towards industrial automation, as shown in Figure1(a). The main objectives of this paper are to highlight the benefits and challenges of integrating digital technologies such as Artificial Intelligence (AI) and Internet of Things (IoT), and the benefits of AI-driven 3D printing in process monitoring, defect detection, optimal process control, and predictive maintenance. A convolutional neural network (CNN) is a popular deep learning approach used in 3D manufacturing for real-time process monitoring, defect detection, and surface quality assessment, thereby supporting process optimization and predictive maintenance, and subsequently improving the quality of 3D-printed parts. In addition, it enhances the reliability and efficiency of industrial 3D material printing. The review highlights the use of popular machine learning and deep learning algorithms to improve the design and manufacturing workflow in 3D material printing, especially in the development process of Industry 4.0. Incorporation of Sensors and IoT devices into the 3D material printer autonomously captures high-quality images at regular intervals, enabling remote monitoring of the printing process and supporting self-diagnosis of printing defects, thereby creating real-time feedback loops for smart parametric adjustments to enable self-correction. Remote monitoring through sensors and intelligent decision-making using AI models delivers built-in diagnosis, adjustments, and corrections to enable data-driven optimization. A review of the insights on popular AI models and IoT modules integrated with 3D material printers is presented, focusing on self-optimized intelligent manufacturing, the capability of performing adaptive process control, automatic defect detection, and self-correction mechanisms to achieve enhanced accuracy in manufacturing. In addition, it highlights the advantages of integrating AI and IoT, including print quality improvement, manufacturing sustainability, scalability, and waste reduction for industrial deployment.

However, the collection of large datasets, training and handling AI models, and system integration cost are the major challenges of integrating AI and IoT with the 3D printing process. In addition, the system setup cost, security risk of data management, procedural complexities, limited available materials design options, and the ability to monitor and control mechanisms lead to uncertainty in the manufacturing process. The detailed parametric, integration, and implementation challenges are discussed in this review to provide potential solutions and future research directions. The comparative study of traditional manufacturing methods and AI and IoT-technology-enabled additive manufacturing methods presented in the review demonstrates the technological benefits that can be achieved in terms of performance improvement and improved material utilization with economic impacts. In addition to technological printing advancements, the future directions are also presented, focusing on the

development of technology-aided 3D material printing processes for smart manufacturing and Industry 4.0. The employability of fully integrated digital twins and AI-based defect detection and optimization procedures project the focus of next-generation smart and autonomous manufacturing systems. presented a CNN-based deep learning approach to detect warpage autonomously in 3D-printed parts by gathering high-quality image datasets using a Raspberry Pi, as shown in Figure1(b). By the focus of industry 4.0, have developed a smart shipyard in China using AI-driven IoT and sensor-based control environment to save the front-line workers from noise, paint mist pollution, and highly dangerous health problems with 70% increased production efficiency. This complete transformation allows developing completely self-controlled production systems for smart factories with the support of self-optimization methods.

This review explores how 3D printing, artificial intelligence, and Internet of Things technologies interact inside smart factories. This detailed study and analysis help researchers and industrial experts to understand and shape them for future intelligent automated production. Although focused on integration, the study reveals pathways toward self-guided fabrication methods through digital connectivity. Efficiency, long-term viability, and expansion capacity emerge when machine learning guides layered manufacturing tasks. Rather than standing alone, these systems gain strength by linking physical output with data-driven decisions.

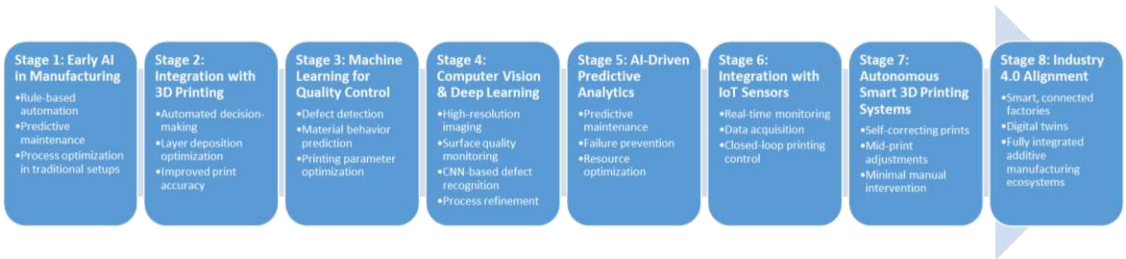

(a)

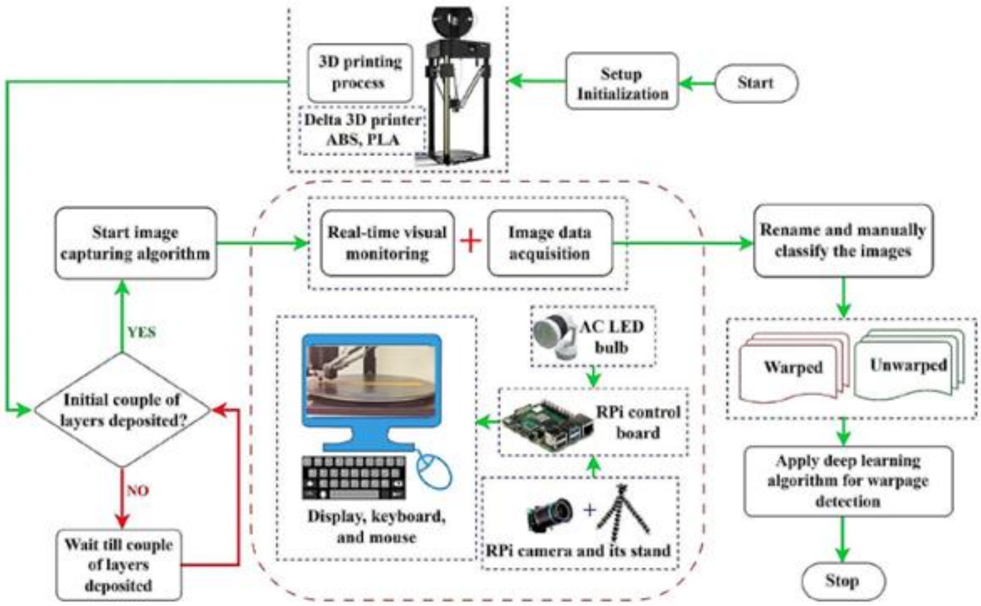

(b)

Figure 1. (a) Integration of AI, IoT sensors, and computer vision into manufacturing workflows, (b) Experimental setup using Machine learning and IoT.

Abstract ID: 2701

Orchestrating Industrial Complexity: A Proposed Human-Swarm Interaction Framework for Agentic Business Process Digital Twins

Vijayakumar Kempuraj[1]*, Dr. C. Lakshmi[2#]

[1] Department of Computer Science, SRM Institute of Science & Technology, Chennai, India
[2] Computational Intelligence Department; Associate Chairperson, School of Computing, SRM Institute of Science & Technology, Chennai, India
**vijimechanical@gmail.com*
#lakshmic@srmist.edu.in

Abstract

Current industrial landscapes, particularly in the automotive and energy sectors, are characterised by highly siloed departments and fragmented business processes. From demand forecasting and capacity planning to execution and quality control, the lack of horizontal integration often leads to buildability mismatches, pricing errors, and delayed responses to macroeconomic shocks. To address these inefficiencies, this research proposes the development of a digital twin for business processes that integrates Agentic AI through a Human-Swarm Interaction (HSI) framework. Unlike traditional Human-Machine Collaboration, which relies on a 1:1 ratio, this proposed approach aims to enable a single human operator to manage a decentralised swarm of autonomous agents representing different functional silos. By designing "swarm-level" command languages, the framework will be developed to allow for collective intent and rapid re-alignment of the entire value chain in response to market dynamics, ensuring that human judgment remains in the loop without the burden of micro-management. The conceptual architecture of the proposed framework is illustrated in Figure. 1.

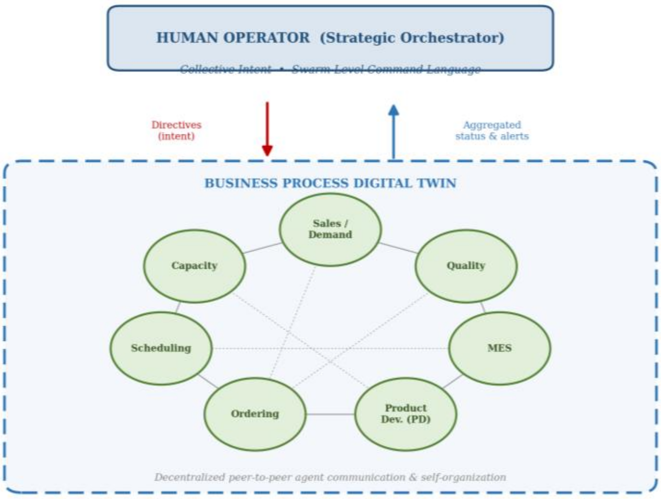

Figure 1. Human-Swarm Interaction framework for the Autonomous Twin.

Introduction

In modern industrial companies, business processes are heavily fragmented across a massive human hierarchy. A typical product lifecycle involves a sales team conducting free-demand forecasting based on market trends, followed by a separate capacity-planning team that studies factory and supply-chain constraints. Allocations are then made based on high-margin rules, production is scheduled, and execution is managed by coordinating with parts suppliers. Simultaneously, an ordering team takes customer configurations, which must be validated by the Product Development (PD) group and sent to the Manufacturing Execution System (MES). This siloed approach creates significant friction.

Customer configurations are frequently priced incorrectly or cannot be built due to engineering constraints. Furthermore, when market dynamics shift abruptly—such as a geopolitical conflict causing fuel prices to rise, GDP to drop, and vehicle sales to plummet—the organisation struggles to adjust free demand, realign allocations, and modify production schedules fast enough to avoid severe over-inventory issues. To address this, we propose the conceptualisation and design of a digital twin of the business processes themselves, powered by Agentic AI and governed by Human-Swarm Interaction (HSI).

Proposed Methodology

Agentic AI and Human-Swarm Interaction: Traditional Human-Machine Collaboration (HMC) generally focuses on a 1:1 interaction paradigm. However, managing an enterprise-wide business process requires coordinating dozens of moving parts. We propose the design and implementation of an HSI framework in which a single human operator collaborates with a "swarm" of autonomous agents. The overall workflow is summarised in Figure 2, and the methodology rests on three core pillars:

- **Agentic Representation:** Each functional silo (Sales, Capacity, Scheduling, Ordering, PD, MES, and Quality) will be represented by an autonomous AI agent within the digital-twin environment.
- **Swarm-Level Command Language:** Rather than micro-managing each department's agent, we will develop an interface in which the human operator inputs "collective intent." For example, the operator could set a swarm-level directive to "optimise for margin while reducing overall production volume by 15 percent due to market contraction."
- **Decentralised Communication Protocols:** The agents will be programmed to communicate with one another peer-to-peer. We will model scenarios in which, if the Quality Agent detects an incoming supplier defect, it instantly notifies the Ordering and Scheduling Agents to halt the affected configurations and re-route supply chains autonomously.

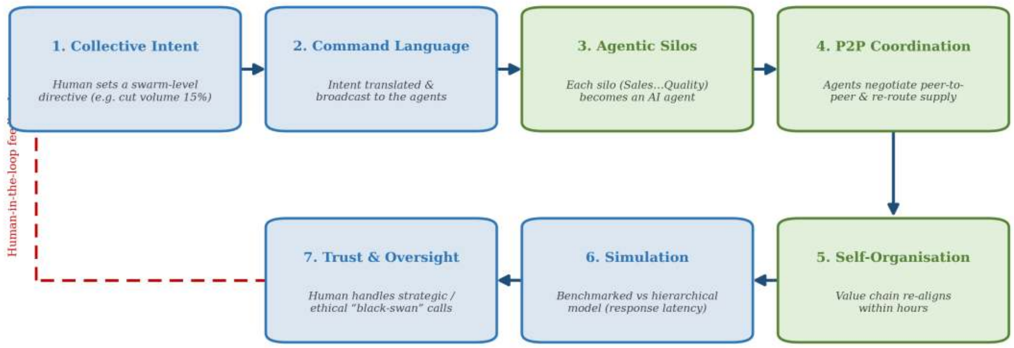

Figure 2. Methodology flowchart of the proposed HSI-enabled business process Digital Twin

Expected Results and Discussion

Through planned simulations of this framework using historical supply-chain and manufacturing data, we anticipate observing a dramatic reduction in response latency to external shocks. We plan to benchmark the proposed HSI-enabled digital twin against traditional hierarchical models. In a simulated scenario involving a sudden drop in market demand, the traditional model typically takes several weeks to propagate adjustments across all suppliers, engineering teams, and plant teams. It is hypothesised that the HSI-enabled digital twin will allow the swarm to self-organise within hours: the Demand Agent would update the forecast, the Capacity Agent would re-allocate resources to high-margin vehicles, and the MES Agent would adjust the factory-floor schedule accordingly. A key area of discussion and future evaluation will be “Trust Calibration.” By moving away from micro-management, we will investigate how the human operator can effectively focus on high-level ethical and strategic decisions, while the agentic swarm handles the multi-variable optimisation. We expect this to ensure that, even in “black-swan” events, the business process remains resilient and buildable.

Conclusion

To truly leverage the concept of an Industrial Digital Twin, organisations must look beyond physical factory assets and digitise their business processes. The proposed integration of Agentic AI and Human-Swarm Interaction holds the potential to break down functional silos, ensuring that what the customer orders is exactly what can be engineered, priced, and built. Ultimately, this research aims to demonstrate how this framework can create a more resilient supply chain and elevate the human worker to a strategic orchestrator of autonomous systems.

Abstract ID: 2979

Phased Integration of 5G Technology in a Digital Twin-Based Autonomous Manufacturing Factory

Sharon Blessy J V[1]*, G. L. Samuel[2], N. Renugadevi[3]

[1]*Department of ECE, IIIT Tiruchirappalli, India*
[2]*Department of Mechanical Engineering, IIT Madras, India*
[3]*Department of CSE, IIIT Tiruchirappalli, India*
**sharonblessyjv@gmail.com*

Abstract

Digital twin-based autonomous manufacturing systems require reliable, scalable and low - latency communication between physical machines and their virtual counterparts. In most factory environments, such communication is typically achieved using wired Ethernet and Wi-Fi networks. While this setup ensures deterministic local control, it introduces limitations in terms of scalability, flexibility, and dependence on fixed broadband infrastructure. Furthermore, mobile components such as Autonomous Mobile Robots (AMRs) experience connectivity constraints due to interference-prone Wi-Fi environments and limited coverage. While prior works have demonstrated fully 5G-enabled digital twin implementations and wireless control architectures, such approaches often require significant redesign of existing industrial communication systems and introduce challenges in protocol adaptation and control-layer reliability.

This work investigates the feasibility of a phased migration from the traditional Ethernet and Wi-Fi communication to 5G Technology in such factories. It involves practical integration of 5G as a backhaul communication mechanism for a digital twin autonomous factory comprising CNC machines, robotic arms, and AMRs. Instead of redesigning the existing control network, a phased integration approach is adopted where 5G is introduced at the network backhaul level, preserving the stability of the internal LAN used for machine-level communication. This approach aligns with practical industrial migration strategies, where critical control loops remain wired while higher-level communication is progressively transitioned to wireless technologies. The architecture is further contextualized within broader 5G digital twin ecosystems and incorporates edge-assisted communication principles to support latency-sensitive interactions.

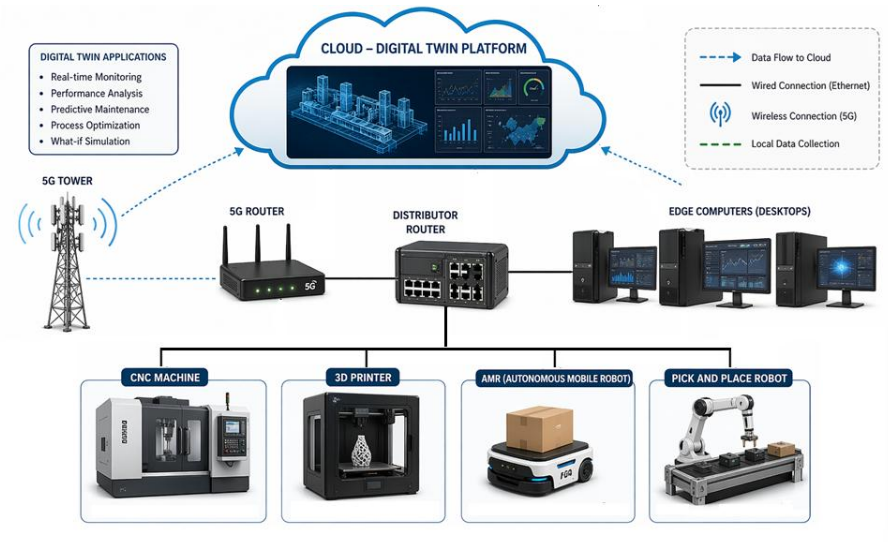

Figure 1. Architecture of the autonomous factory with 5G backhaul connection.

This integration is achieved by replacing the existing broadband connection with a commercial 5G backhaul using a 5G Industrial Router or Customer Premises Equipment (CPE), as shown in Figure 1. The proposed architecture is validated through practical deployment in an autonomous manufacturing laboratory to evaluate the feasibility of introducing 5G without modifying the existing machine-level communication network. The implementation focuses on verifying internet backhaul connectivity between the factory network and cloud-hosted digital twin services while observing network stability, deployment flexibility, and the practical challenges encountered during migration.

Initial deployment demonstrates that digital twin communication can be successfully established through the commercial 5G backhaul while preserving the existing Ethernet-based control network. Experimental observations indicate that the overall system functionality is maintained; however, network performance is influenced by the quality of the available cellular signal within the factory environment. The study highlights that commercial 5G backhaul provides a practical alternative to fixed broadband infrastructure while emphasizing that reliable cellular coverage is an important requirement for consistent communication performance, consistent with observations reported in previous studies on 5G-enabled digital twin systems and industrial wireless communication architectures.

The study establishes a practical migration pathway toward 5G-enabled autonomous factories and discusses key deployment considerations, including network coverage, infrastructure flexibility, and incremental integration strategies. Future work will extend this implementation through comprehensive performance evaluation under optimized 5G network conditions, investigate latency-sensitive industrial applications, evaluate the proposed architecture over a private 5G network, and explore fully wireless machine connectivity and remote operation supported by 5G-enabled digital twins.

Abstract ID: 3166

Real-Time Chatter Detection and Adaptive Control in High-Speed Micromilling Using Normalized Spectral Entropy and Spindle Speed Regulation

Suraj Kumar[1]*, Rinku K. Mittal[2], Ramesh K. Singh[1]

[1]*Department of Mechanical Engineering, Indian Institute of Technology Bombay, Powai, Mumbai, India*
[2]*Department of Mechanical Engineering, Indian Institute of Technology Guwahati, Guwahati, Assam, India*
**suraj_me@iitb.ac.in*

Abstract

High-speed micromilling is widely used for fabricating micro-scale features; however, the low flexural stiffness of micro end mills makes the process highly susceptible to regenerative chatter, resulting in poor surface finish and premature tool failure. Conventional stability lobe-based approaches are often limited in real-time applications due to process variability, uncertainty in dynamic parameters, and modeling complexity. Entropy-based methods have shown promising potential for chatter onset detection. To address these challenges, this work presents a unified framework for real-time chatter detection and control using normalized spectral entropy (NSE) combined with adaptive spindle speed regulation. The objective is to develop a robust, computationally efficient, and parameter-independent methodology capable of both early chatter detection and its suppression during continuous machining.

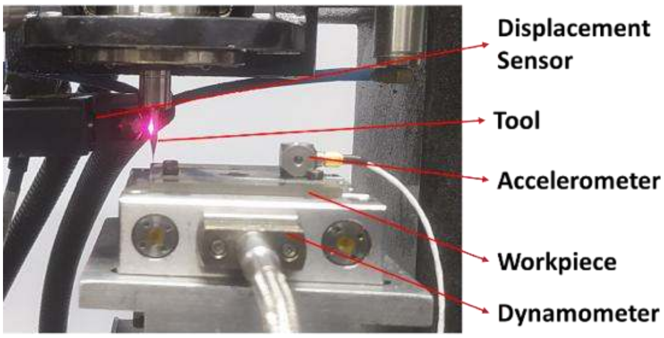

Figure 1. Experimental Set-up.

The chatter detection methodology builds upon our previous work, where NSE is used as a customized feature for identifying chatter onset. Acceleration signals acquired during machining are pre-processed using a high-pass Butterworth filter to remove spindle rotation and tooth passing frequency components below 2500 Hz, thereby enhancing sensitivity to chatter-induced dynamics. Spectral entropy, computed from the normalized power spectrum, quantifies the distribution of energy across frequency components. Stable cutting exhibits broadband characteristics (high entropy), whereas chatter leads to dominant frequency components (reduced entropy). The NSE, defined as the ratio of spectral entropy during machining to its initial stable value, serves as a normalized indicator independent of cutting parameters. A threshold value of approximately 0.95 is used for chatter detection. Experiments are conducted on a high-speed micromachining center (140,000 rpm spindle) shown in Figure 1, where cutting forces and vibrations re measured using a dynamometer and a triaxial accelerometer (0–12 kHz bandwidth) with a sampling rate of 20 kHz. Slot milling of Ti6Al4V is performed using a 500 µm diameter two-flute end mill under varying spindle speeds (20,000, 40,000, and 60,000 rpm) and feed per tooth values (1, 3, and 5 µm/tooth), while the depth of cut is gradually increased from 0 to 200 µm to induce chatter.

The complete methodology is shown in Figure 2. It is divided into two parts. The first part presents real-time chatter onset detection using NSE. After data acquisition, a moving-window FFT is performed periodically, followed by the computation of spectral entropy and subsequently NSE. The second part focuses on real-time chatter control using spindle speed regulation. In this stage, the dominant chatter frequency is identified from the spectrum by selecting non-harmonic frequency components with amplitudes higher than the tooth-passing frequency. Based on this, a modified spindle speed, close to the current operating speed, is calculated using Equation 1 to maintain a 90° phase difference ($\varepsilon = 0.25$) between the current and previous regenerative tool paths.

$$N = \frac{1}{Z}\frac{f}{n + 0.25} \tag{1}$$

The results demonstrated in Figure 3 are about the evolution of signal characteristics with increasing depth of cut for chatter onset detection. Initially, the signal exhibits low amplitude (or moving energy), which progressively increases with depth of cut. At 20,000 rpm, chatter onset introduces fluctuations in moving energy, although an overall increasing trend is observed. Spectral entropy shows significant variation in its initial value across different cutting conditions. To eliminate this dependency, normalization is performed using the initial stable entropy, resulting in NSE, which consistently starts at unity. Across all cutting conditions, NSE values below 0.95 reliably indicate chatter onset. This is further validated through corresponding surface topography and synchronized temporal markers. Thus, NSE is established as a robust and reliable feature for chatter detection under varying machining conditions.

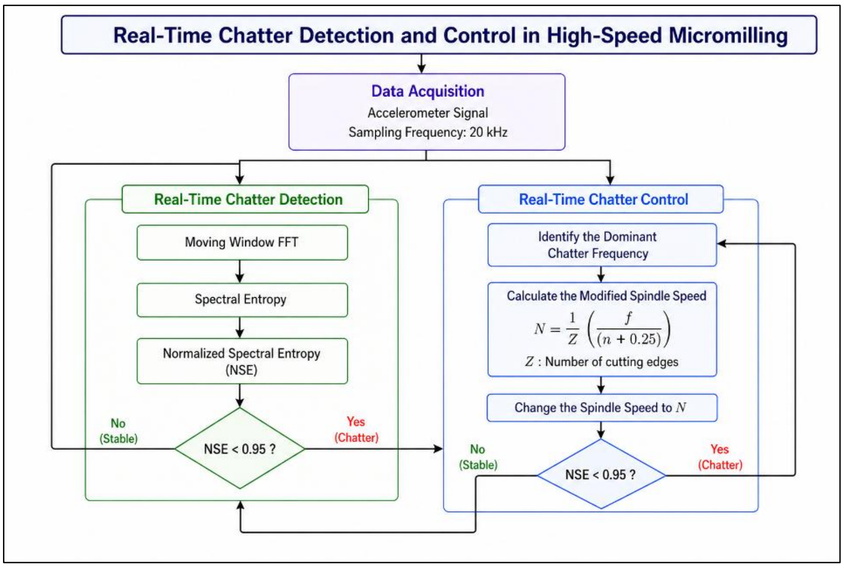

Figure 2. Methodology for real-time chatter detection and control.

For real-time chatter control, an adaptive spindle speed regulation strategy based on Liao et al. is implemented. The method maintains a phase difference of 90° ($\varepsilon = 0.25$) between successive tool engagements to suppress regenerative chatter. Upon detection of chatter (NSE < 0.95), the dominant chatter frequency is extracted from the power spectrum, and the spindle speed is adjusted accordingly to satisfy the phase condition. The modified spindle speed is selected to be as close as possible to the current speed and is updated at a controlled rate of 500 rpm/s to ensure smooth operation without interrupting the machining process.

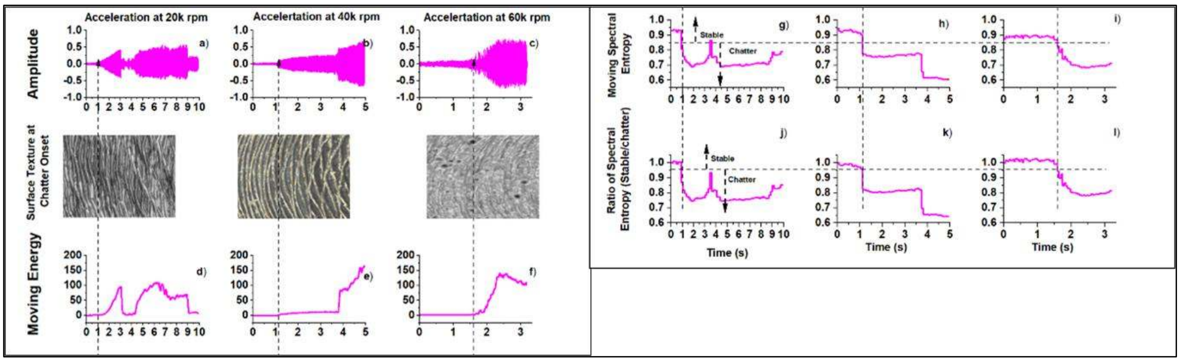

Figure 3. Evolution of features with chatter onset. (a) acceleration data at 20000 rpm, (b) acceleration data at 40000 rpm, (c) acceleration data at 60000 rpm, (d) moving energy at 20000 rpm, (e) moving energy at 40000 rpm, (f) moving energy at 60000 rpm, (g) moving spectral entropy at 20000 rpm, (h) moving spectral entropy at 40000 rpm, (i) moving spectral entropy at 60000 rpm, (j) Normalized spectral entropy at 20000 rpm, (k) Normalized spectral entropy at 40000 rpm, and (l) Normalized spectral entropy at 60000 rpm

The proposed framework is computationally efficient, does not require prior system identification, and is robust to varying cutting conditions. This makes it highly suitable for integration into intelligent machining systems and digital twin-based manufacturing environments. Future work will focus on the complete implementation and experimental validation of the real-time chatter control strategy in high-speed micromilling.

Abstract ID: 3115

From P&IDs to Knowledge Graphs: Enabling Process Intelligence through Digitization

Prateek Gaurav[1], Ashish Mehta[2], Subhadra Devi Saripalli[2], N. Renugadevi[1], Rajagopalan Srinivasan[2, 3]*

[1]*Department of Computer Science and Engineering, Indian Institute of Information Technology Tiruchirappalli, Tiruchirappalli, India*

[2]*Department of Chemical Engineering, Indian Institute of Technology Madras, Chennai, India*

[3]*American Express Lab for Data Analytics, Risk and Technology, Indian Institute of Technology Madras, Chennai, India*

**raj@iitm.ac.in*

Abstract

Piping and Instrumentation Diagrams (P&IDs) serve as the foundational "master data" for chemical and process plants, encoding critical information about equipment, pipelines, instrumentation, control logic, and their connections. Despite their central role in plant design, operation, and safety analysis, P&IDs remain largely confined to unstructured formats such as raster images and PDF documents, especially the legacy P&IDs. This disconnect prevents seamless integration with modern digital engineering systems and limits their utility in automated workflows, leaving critical engineering knowledge inaccessible in a structured, machine-readable form. This limitation introduces a significant operational bottleneck. Engineers must manually interpret and analyze P&IDs for applications such as process simulation, compliance validation, procurement planning, and hazard and operability studies (HAZOP). This Manual interpretation of P&IDs is time-intensive and error-prone, particularly in large-scale industrial systems. Consequently, there is a growing need for automated systems that can bridge the gap between visual diagrams and structured engineering knowledge. Recent advancements indicate a shift from traditional detection pipelines toward context-aware semantic understanding. Shteriyanov et al. address the "semantic gap" in P&IDs by proposing a hybrid AI system capable of format-aware interpretation of pipeline metadata, achieving 91.1% end-to-end accuracy on industrial datasets. Similarly, Partanen demonstrates the effectiveness of modern OCR systems such as PaddleOCR (PP-OCRv5), achieving 89% word-level accuracy on engineering drawings through customized preprocessing. Parallel developments in document intelligence, such as the integration of YOLOv12 with PaddleOCR, highlight improvements in structural extraction from dense layouts.

Earlier works in engineering diagram understanding have explored deep learning-based symbol recognition and pipeline extraction –, as well as OCR-based approaches. Additional advancements in document and diagram understanding have been reported in, while few-shot and symbol recognition challenges have been studied in. Despite these advancements, existing approaches exhibit key limitations: (i) reliance on rigid OCR pipelines with limited semantic interpretation, (ii) sensitivity to symbol variability and class imbalance, (iii) fragmented text detection under dense layouts, and (iv) lack of unified frameworks for integrating text, symbols, and topology into a coherent representation. To address these challenges, we propose an end-to-end computational pipeline for structured digitization of P&IDs, integrating computer vision, document understanding, and graph-based modeling. The system consists of four modules: text analysis, symbol detection and classification, pipeline topology reconstruction, and relational grounding, resulting in a structured knowledge graph. This architecture delivers several key advantages over conventional approaches. By unifying text recognition, symbol understanding, and topological reasoning within a single coherent framework, the system eliminates the fragmentation inherent in earlier multi-stage pipelines. The integration of semantic validation and spatial grounding ensures that extracted entities are not detected in isolation

but are grounded within the engineering context. Using a knowledge graph as the output representation enables structured, machine-readable encoding of P&ID content that is directly queryable and interoperable with downstream digital engineering tools. Additionally, the framework's ability to generalize under data scarcity through augmentation and few-shot classification makes it practically deployable in industrial environments where annotated data is limited. Experimental evaluation on synthetic and real-world datasets demonstrates strong performance, achieving 97.5% accuracy in text detection and over 90% in symbol detection (excluding major equipment classes). The digitized knowledge graph supports a broad range of high-value industrial applications. The key applications include process simulation, compliance verification, maintenance planning, and safety analysis (e.g., HAZOP), by providing a structured, queryable representation of plant information that can be validated against codified domain knowledge and engineering rules. Overall, the proposed framework bridges the gap between unstructured engineering diagrams and intelligent, data-driven process systems.

Abstract ID: 3158

Metaverse-Based Digital Twin of a Smart e-Yantra Laboratory for Industrial IoT and Smart Manufacturing Applications

N V Jatin Pochiraju*, Palash Rai

Department of Computing Technologies, SRM Institute of Science and Technology, Kattankulathur, Tamil Nadu, India

**nvjatinpochiraju@gmail.com*

Abstract

Industrial digital twins are a fundamental component of Industry 4.0, facilitating real-time synchronization of physical assets and their virtual counterparts for monitoring, control, and simulation. This work presents a metaverse-based digital twin of a smart e-Yantra laboratory, developed as a prototype to demonstrate the integration of industrial IoT, cyber-physical systems and immersive visualization in smart manufacturing environments. The system simulates the TP1508 e-Yantra laboratory using a Unity-based 3D environment, thus maintaining a one-to-one mapping of the physical layout and infrastructure. The authors have developed a small physical setup that is IoT-enabled to simulate the industrial subsystems. The physical setup consists of LED indicators for visualizing the system state, DC fans to simulate the ventilation mechanism, a 12V solenoid lock for access control, and a DHT11 sensor to measure the temperature and humidity. These components form important parts of smart industrial systems such as environmental sensing, safety control, and device actuation.

The methodology is based on a three-layered digital twin architecture. The Physical Layer is represented by IoT devices interfaced by microcontrollers. The Communication Layer allows the digital twin and the physical setup to exchange data in real-time, in both directions. The interactive metaverse interface for visualization and control is provided by the Digital Twin Layer, implemented in Unity. The data from sensors embedded in the physical system are transferred to the virtual environment in real-time. Actions performed by users in the digital twin can cause actions in the physical model, thus making it a closed loop.

Figure 1. Unity-based digital twin representation of the e-Yantra laboratory environment.

The system achieves real-time synchronization with minimal latency. The sensor data is accurately represented in the digital twin, and the actuation actions like controlling the LEDs, fans, and door locks work effectively in response to commands issued from the virtual end. The system performs consistently even when operated continuously, hence proving its appropriateness for real-time applications.

Results show that the framework proposed improves situational awareness and remote interaction through a metaverse interface that is intuitive. The users will be able to observe their environment, manage infrastructure, and even simulate situations without being physically present. In comparison to traditional monitoring systems, the application of immersive visualization helps to understand system behaviour more effectively.

In conclusion, the developed system is an actual demonstration of the utilization of a digital twin in the metaverse, as guided by the principles of smart manufacturing. The system can be scaled up for use in other large-scale industrial settings such as factory floors. In the future, we intend to incorporate artificial intelligence and data integration techniques.

Abstract ID: 3100

ABSA-GraphRec: Perception-Aware Constraint Learning via LLM-Enhanced Aspect Sentiment for Graph-Based Tourism Recommendation

Dhivya K*, Suganya R

Vellore Institute of Technology, School of Computer Science and Engineering, Chennai, India.
**dhivya.k2024@vitstudent.ac.in*

Abstract

Tourism recommender systems face a fundamental tension: real-world feasibility constraints – distance, budget, accessibility, and crowd conditions – are inherently subjective, yet existing systems model them as static, objective filters. Post-hoc constraint enforcement treats all users identically, ignoring that a "crowded" venue may be acceptable to a casual visitor but unacceptable to a senior traveller. This gap between objective constraint modelling and subjective user perception remains largely unaddressed in the literature. We propose **ABSA-GraphRec**, a novel framework that (i) extracts aspect-level sentiment signals from user reviews via a hybrid VADER–TextBlob pipeline with domain-adapted tourism keywords, (ii) integrates these signals into a Graph Attention Network (GAT) as adaptive constraint-penalty weights, and (iii) jointly optimises recommendation accuracy and sentiment alignment via a sentiment-regularised ranking loss — embedding perception inside the graph-learning process rather than as a post-processing filter. An ablation study on a Yelp review subset (N=3,000) confirms that the sentiment-regularisation term yields substantial gains in both ranking accuracy and recommendation quality over an identical GAT architecture trained without it.

Keywords: Tourism Recommendation, Graph Attention Network, Aspect-Based Sentiment Analysis (ABSA), Large Language Model (LLM), Sentiment Regularisation.

Methodology

The proposed framework comprises three interconnected modules. (i) **LLM-Enhanced ABSA Module**: A hybrid VADER–TextBlob sentiment pipeline, augmented with domain-adapted keyword lexicons for four travel-specific aspects — price, access, crowd, and distance — extracts fine-grained, aspect-level sentiment from raw review text. Sentence-level matches against each aspect's keyword set are scored by averaging VADER's compound score and TextBlob's polarity score, then aggregated per Point-of-Interest (POI) / business to yield a structured perception matrix. (ii) **Perception-Aware Graph Attention Layer:** A bipartite user–POI interaction graph is constructed (users and items as nodes, ratings as edges), over which a two-layer Graph Attention Network (GAT) learns latent embeddings. Aggregated aspect sentiment is attached to each edge as a perception signal that modulates the sentiment-regularisation term during training, allowing the model to suppress items that are objectively feasible but subjectively undesirable to the user community. (iii) **Sentiment-Regularised Ranking Optimisation:** Recommendation is framed as implicit-feedback ranking, trained with a Bayesian Personalised Ranking (BPR) loss jointly with a sentiment-regularisation penalty that discourages high predicted scores for negatively perceived items. The combined objective is:

$$L_{BPR} = -\sum \ln \sigma(s_{ui} - s_{uj}) \tag{1}$$

$$L_{sentiment} = [\ \mathrm{ReLU}(-w_{sent}) \cdot \mathrm{ReLU}(s_{ui})\ ] \tag{2}$$

$$L = L_{BPR} + \lambda \cdot L_{sentiment} \tag{3}$$

where s_{ui} denotes the predicted affinity between the user u and item i, w_{sent} is the aggregated aspect-sentiment weight on the edge (u, i), and λ (set to 0.3 in our experiments) controls the strength of sentiment regularisation. $ReLU(-w_{sent})$ activates the penalty only for negatively perceived items, so the model is pushed away from recommending POIs the user community has described unfavourably.

Experimental Results and Discussion

Two datasets serve complementary roles. The curated TN-Tourism dataset (10 manually collected Google Maps reviews across 5 South Indian POIs – Meenakshi Temple, Brihadeeswarar Temple, Marina Beach, Government Museum Chennai, and Guindy National Park) is used qualitatively to illustrate the ABSA pipeline's ability to extract culturally specific, fine-grained aspect sentiment; at this sample size, it is not used for ranking-metric evaluation. The Yelp Open Dataset subset (N=3,000 reviews) provides the quantitative evaluation: an 80/20 train–test split on user–item interactions, with Recall@10 and NDCG@10 computed on held-out pairs. Table 1 reports a controlled ablation isolating the contribution of sentiment regularisation: identical GAT architecture and training schedule, with λ = 0.3 (ABSA-GraphRec) versus λ = 0 (GAT-only baseline, i.e., no sentiment signal).

Table 1. Sentiment-regularisation ablation on the Yelp subset (N=3,000).

Model	λ (sentiment)	Recall@10	NDCG@10
GAT-only baseline	0	0.398	0.196
ABSA-GraphRec (ours)	0.3	0.592	0.349

Adding the sentiment-regularisation term (λ = 0.3) lifts Recall@10 from 0.398 to **0.592** (+48.7%) and NDCG@10 from 0.196 to **0.349** (+78.1%) relative to the identical GAT architecture trained without any sentiment signal, isolating the contribution of perception-aware regularisation from architectural confounds. As visualised in Figure 1, the TN-Tourism case study shows Government Museum Chennai recording a comparatively muted crowd sentiment, consistent with review mentions of queues at the ticket counter, while Brihadeeswarar Temple records the highest crowd sentiment (+0.51), consistent with descriptions of its peaceful, uncrowded atmosphere. Marina Beach shows a distinctive split — strongly positive on price and access, but negative on crowd and distance — a pattern invisible to any single aggregate rating, and exactly the kind of per-aspect, per-POI nuance that static constraint filters cannot represent. This is the core novelty of embedding ABSA directly inside the graph-learning objective rather than applying it as a post-hoc filter.

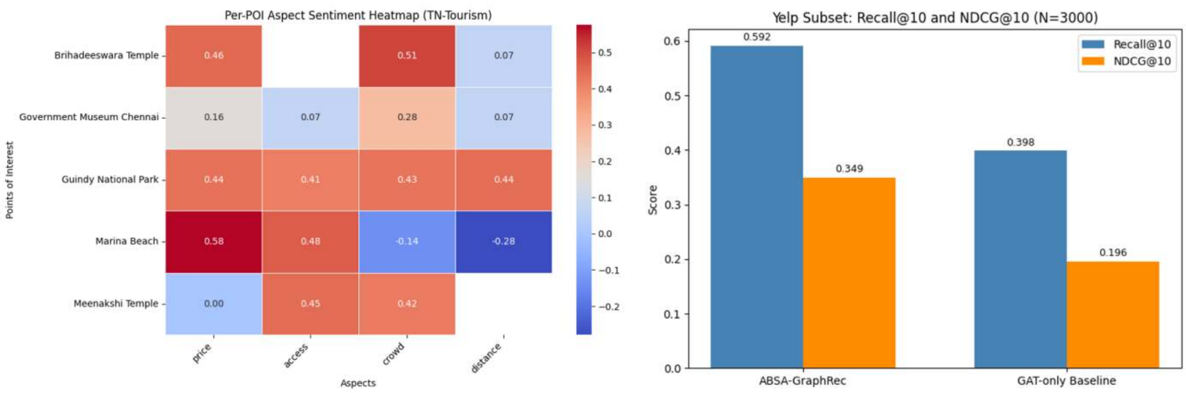

Figure 1. POI aspect-sentiment heatmap on the TN-Tourism case study (left) and Recall@10/NDCG@10 sentiment-regularisation ablation on the Yelp subset, N=3,000 (right).

Conclusion

ABSA-GraphRec advances tourism recommendation by treating feasibility as a dynamic, perception-aware construct embedded directly within graph-attention learning rather than enforced as a static post-hoc filter. A controlled ablation on the Yelp subset confirms that sentiment-regularised training yields substantial, isolated gains in both Recall@10 and NDCG@10 over an architecturally identical baseline, while the TN-Tourism case study demonstrates the pipeline's ability to surface culturally-specific, aspect-level perception signals — including subtle crowd and accessibility nuances relevant to the South Indian tourism context — that conventional rating aggregates obscure. Future work will scale the quantitative evaluation to graph-based baselines such as LightGCN and KLLMs4Rec under identical splits, extend the curated TN-Tourism dataset, and explore multi-modal review inputs and real-time sentiment streaming for cross-city generalisation.

Abstract ID: 2709

Digital Twin of Mechanical Machines using Unity

Samuel Mathias Menezes*, Shantaram D.S. Borker, Arantes Sebastiao Caiado, Rayan Vaz, Dr. Saeesh Verenkar

Padre Conceição College of Engineering Mechanical Department, Goa University, Verna, Goa, India
**samzes182003@gmail.com*

Abstract

With the advancements in Technology, the concept of Digital Twin has emerged as a powerful tool for bridging the existing gap between physical systems and their virtual counterparts. In the field of mechanical engineering, maintaining systems and their reliability is a critical challenge. The systems (machines) have to be continuously monitored for breakdowns and failures, and a lack of predictive maintenance and monitoring systems leads to increased downtime and frequent losses for the industries. However, large-scale implementation can be costly; small-scale implementation can prove to be very effective, cost-efficient, and helpful in reducing losses. This paper focuses on evaluating the continuous efficiency of the machines and anomaly detection in a machine. This article gives a comprehensive framework for developing a Digital Twin using Unity and ROS as the digitalization and communication layers. A prototype model of a packaging line was designed, simulated, tested, and fabricated for experimentation.

Methodology

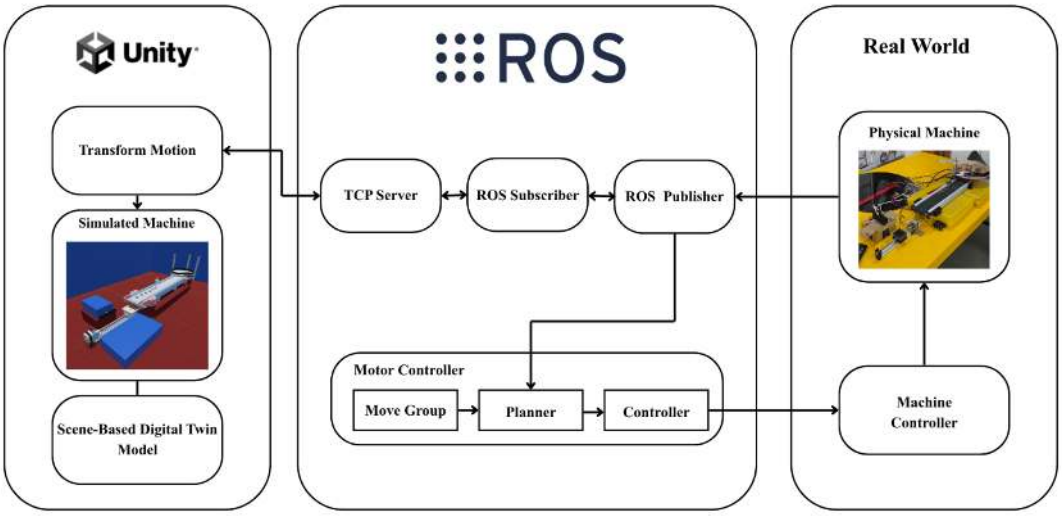

Figure 1. Architecture of the implemented Digital Twin Model.

Presented in Figure 1 is the methodology used in the project, which presents the main system architecture and the tools used, along with the flow of information for the system.

Tools used in the prototype are Unity 2022.3.62f1 for the visual representation of the physical system, ROS2 Jazzy Jalisco, along with IoT sensors, were used to continuously monitor the machine parameters to detect any changes in the system that can help identify the cause and time of failure and can prove to be useful for predictive maintenance.

The virtual replica of the packaging line consisting of a turntable, conveyor, linear guide, and ejection piston was designed using Fusion 360. These models were imported to Unity. Motion scripts were coded in C# and were applied to each machine. These machines were run to check the sequencing and alignment, after which the fabrication was started. The machines were fabricated as designed, and the motors and IOT sensors were installed. The motors and sensors were controlled using an Arduino Uno microcontroller. ROS was used to code the motor sequence controller, sensor interface, and ROS-TCP-Endpoint to communicate between ROS and Unity. A machine learning model was coded in ROS for detecting any anomaly during the operation of the machines.

Results

The developed Digital Twin successfully replicated the operations of the packaging line consisting of turntable, conveyor belt, linear guide, and ejection piston in unity. The Digital Twin was tested independently for optimizations during the design phase. After fabrication, the Digital Twin was linked with the physical model. The synchronization between the physical model and the digital twin was evaluated based on object movement in both the virtual and physical environments. Based on the comparison of cycle times between the physical packaging line and the virtual model, the synchronization accuracy of the Digital Twin was calculated to be approximately 90%. Real-time sensor data, including motor temperatures, ambient temperature, and object count, were successfully transmitted to Unity for accurate visualization of machine states and process conditions.

Discussion

The results obtained from the Digital Twin demonstrate the improvements in design optimization and process planning. The use of the Digital Twin during the design stage enabled testing and optimization of machine sequencing, alignment, and process flow before fabrication, thus reducing the need for repeated physical modifications and reducing developmental costs. The real-time sensor data such as motor temperatures and ambient temperature, helped in evaluating the thermal behaviour of the motors and estimating the temperature difference between the motors and the surrounding environment. The implementation of machine learning-based anomaly detection provides the foundation for future predictive maintenance applications by enabling early identification of abnormal operating behaviour.

The results demonstrate the feasibility of developing cost-effective Digital Twin solutions for small and medium-scale manufacturing industries, where large-scale industrial Digital Twin implementations may not be economically viable.

Abstract ID: 2663

Synergizing Multi-Source Energy Management: An Intelligent, Unified Architecture for Next-Generation Electric Vehicles

Satvika Gobi*, Annie Alex#

School of Electrical Engineering, Vellore Institute of Technology, Chennai, Tamil Nadu, India

**satvika.gobi2025@vitstudent.ac.in*

#annie.alex2025@vitstudent.ac.in

Abstract

The rapid advancement of electric vehicle (EV) technologies demands intelligent energy management systems capable of overcoming major limitations such as limited driving range, battery degradation, inefficient energy utilization, and poor coordination among multiple energy sources. Conventional EVs primarily rely on a lithium-ion battery, while auxiliary technologies such as hydrogen fuel cells, regenerative braking, and wireless charging are generally implemented as independent subsystems. This fragmented approach results in inefficient power distribution, increased conversion losses, reduced system efficiency, and greater dependence on the battery.

To address these challenges, this paper proposes a unified multi-source intelligent EV energy management architecture integrating a lithium-ion battery, hydrogen fuel cell, regenerative braking system, and dynamic wireless charging under a centralized intelligent control strategy. A High-Voltage Direct Current (HVDC) bus serves as the common power distribution platform, enabling seamless integration of all energy sources while reducing multiple AC–DC and DC–DC conversion stages, switching losses, and hardware complexity.

The battery and hydrogen fuel cell directly supply DC power to the HVDC bus, while wireless charging energy is converted from AC to DC using a dedicated rectifier before integration. During regenerative braking, the Interior Permanent Magnet Synchronous Motor (IPMSM) operates in generator mode, converting kinetic energy into AC power that is rectified to DC and returned to the battery through the HVDC bus. Power conditioning is performed using a Z-source converter, which provides single-stage buck–boost DC–AC conversion for the traction system, while an Integrated Gate-Commutated Thyristor (IGCT)-based switching network dynamically routes power according to load demand, source availability, and operating conditions.

The overall energy management strategy is governed by a spintronic-based intelligent controller that continuously monitors battery state of charge (SoC), vehicle speed, power demand, thermal conditions, and source availability. Predictive algorithms estimate future energy requirements and optimize power distribution by prioritizing battery power during acceleration, fuel cell operation during cruising, regenerative braking during deceleration, and wireless charging whenever infrastructure is available. The IPMSM is controlled using Direct Torque Control (DTC), providing rapid torque response, reduced torque ripple, and high operating efficiency. To improve safety and reliability, Polymeric Positive Temperature Coefficient (PPTC) devices provide self-resetting overcurrent protection for the battery subsystem, while a SCADA-based platform enables real-time monitoring, diagnostics, and fault detection.

Performance analysis of the proposed architecture indicates that the HVDC network reduces conversion losses and improves voltage stability by minimizing unnecessary power conversion stages. The Z-source converter and IGCT switching network ensure efficient power conditioning, seamless source transition, and uninterrupted energy delivery under varying operating conditions. Intelligent energy scheduling balances power among the battery, fuel cell, regenerative braking system, and wireless

charging, thereby reducing battery stress, extending battery life, and maximizing overall energy utilization. Efficient recovery of braking energy and integration of wireless charging further improve vehicle efficiency and reduce dependence on onboard energy storage. In addition, the DTC-controlled IPMSM delivers smooth speed regulation and high dynamic performance, while the PPTC protection and SCADA monitoring framework enhance operational safety, fault tolerance, reliability, and maintenance capability.

Overall, the proposed multi-source architecture results in a higher energy efficiency, extended driving range, improved regenerative energy recovery, reduced battery degradation, superior motor performance, stable power management, and enhanced system reliability. Its scalable and modular design also supports future integration of renewable energy sources, advanced charging technologies, and intelligent transportation infrastructure, making it a practical and future-ready solution for next-generation electric mobility.

Abstract ID: 2659

Digital Twin and AI-Driven Test point Prediction for Calibration Process Optimization in Energy Meter

Richie Garg[1]*, Silvio Rizzuto[2], Leema Christy Devasagayam[1]

[1]*Schneider Electric, Bengaluru, 560067, India*
[2]*Schneider Electric, Eybens, 38320, France*
**richiegarg14@gmail.com*

Abstract

The calibration of a precision energy meter is a time-intensive process in manufacturing environments. Calibration workflows typically involve three sequential phases: Capture, Calibration, and Verification, with total durations approaching 18 minutes per unit. In an era of increasing demand for smart energy metering solutions, reducing calibration cycle time without compromising measurement accuracy is a critical operational challenge. The problem statement is to predict a part of the surface map of calibration coefficients (Figure 1a). It is complex since it involves many variables like frequency, 3-phase voltage, current, and phase shift, which have non-linear behaviour. Presently, 16 test points composed of voltage, current, phase shift, and frequency values are physically injected into the energy meter, and the challenge is to predict half of them (Figure 1b).

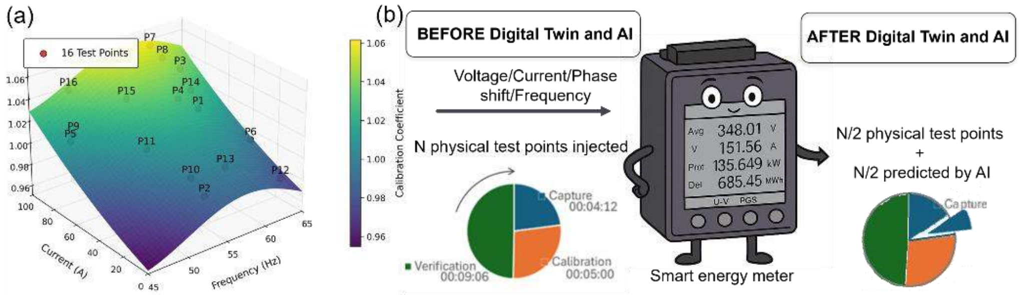

Figure 1. (a) Calibration coefficient surface map with 16 test points, (b) Reduction of physically injected test points by half using AI and digital twin.

This work presents two approaches that leverage digital twin technology and machine learning models, including deep learning. The digital twin approach incorporates the data from the electronic schema of voltage sensors, current transformers and an analogue-to-digital conversion simulation of the device's expected electrical behaviour in Simscape. The AI model approach employs partial least squares regression and a neural network in MATLAB that are trained on a two-year manufacturing data history from the production line. The model could learn hierarchical patterns in calibration outcomes across 53000 units under test to predict the optimal subset of test points.

The system operates in two steps. In the Step 1 Capture phase, the AI model analyses incoming component parameters and predicts half of the test points. In step 2, the Correction phase, the product itself applies corrective factors to the firmware, enabling adaptive and targeted calibration. In the future, for real-time deployment for a digital twin, these models could learn patterns in energy meter sensor (already existing) measurements and component tolerances from traceability reports. The verification phase remains unchanged to ensure compliance with metrological standards. AI model results demonstrate a measurable reduction in the number of active calibration test points during the Capture phase, accelerating the overall calibration process by 11%. The Verification phase duration of 9 minutes remains constant as it is dictated by regulatory requirements, while the Capture (4 minutes) and Correction (5 minutes) phases benefit from optimized test point coverage. The AI and Digital Twin approaches demonstrate significant potential for scaling across the product family and other smart energy meters. Future work will focus on further training on verification data augmentation, cross-product generalisation of the AI model, and integration with real-time quality feedback loops on the production line. The advantages are faster calibration and verification cycle time, greater flexibility, lower manufacturing cost by 13%, reduced operator time, and a reduction in six test benches operating simultaneously in three shifts in a day. The benefit could be cascaded to more smart metering device plants.

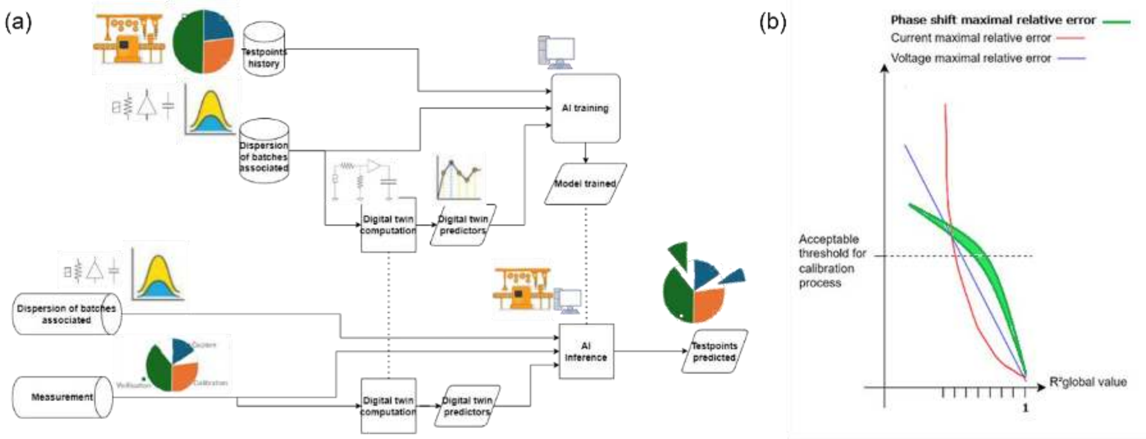

Figure 2. (a) AI approach based on test point data history and digital twin based on manufacturing batch dispersion, (b) R^2 coefficient as a key indicator of prediction quality.

This digital twin is a physics-based virtual representation of the energy meter, continuously updated using real dispersion data of electronic and electromagnetic components coming from manufacturing batches (Figure 2a). Its main value is to move from nominal design assumptions to realistic production-aware predictions. By injecting real component dispersions, the digital twin defines true prediction limits, confidence intervals, and sensitivity to variability, instead of single ideal values. The model computes relevant physical descriptors, such as internal electrical and electromagnetic states, which are not directly measurable on the production line.

These descriptors are physics-informed, ensuring that predictions remain explainable, robust, and causally linked to real phenomena. Because it is physics-based, the digital twin avoids black-box behaviour, limits extrapolation risks, and ensures relevant observation of performance drivers, especially under varying conditions. The digital twin is updated at each change of production lot, focusing on the components that have the strongest impact on performance and metrology accuracy. New dispersion data are injected to refresh prediction boundaries and maintain alignment with the real manufacturing context. Integration into the existing process is enabled by Model-Based Design. The approach is first validated through a virtual demonstrator using historical test and production data, proving accuracy and added value. Once validated, automatic code generation is used to deploy the digital twin directly into the production test bench. This allows real-time inference, prediction of test points, and decision support on the manufacturing line. We rely on the maximum error encountered, statistically estimated from the lower bound of the error distribution measured by our demonstrator on the available test dataset (Figure 2b). Alongside this, we use the R^2 coefficients a key indicator of prediction quality and model efficiency. These metrics together provide a robust and representative assessment of predictive performance. Accurate prediction is critical, as it directly drives the quality of the energy meter calibration. Therefore, improving prediction reliability is essential to ensure calibration accuracy and overall system performance. In summary, this digital twin transforms component variability into a controlled input, delivers trustworthy physics-based predictions, and integrates seamlessly into industrial production through model-based design and code generation.

Abstract ID: 2656

A Digital Twin Framework for Vibration-Based Condition Monitoring of a Single-Stage Spur Gearbox

Harish Muthu S[1], C Krithik Kesav[1], Sailesh Patra[1], Rachana Krishna[1], Sabareesh G R[1]*, Sudha Radhika[2]

[1]*Mechanical Department, BITS Pilani Hyderabad Campus, Hyderabad, India*

[2]*EEE Department, BITS Pilani Hyderabad Campus, Hyderabad, India*

**sabareesh@hyderabad.bits-pilani.ac.in*

Introduction and Background

Rotating machinery, such as gearboxes and bearings, is essential across many industries for load transmission and motion conversion. If small defects in these parts go unnoticed, they can lead to major equipment failures and costly downtime. Condition monitoring (CM) helps prevent this by tracking machine health and scheduling maintenance only when needed, which is a key part of predictive maintenance. Traditional vibration-based diagnosis captures the distinct dynamic behaviour of gears and bearing defects, providing a reliable foundation for health assessment. However, modern predictive maintenance is advancing through digital twin frameworks, which synchronize real-time physical data with high-fidelity virtual models. To efficiently process the complex vibration data produced by these systems, data-driven machine learning techniques are employed to automate feature extraction and fault classification. Traditional ML lacks physical interpretability, while pure mathematical models lack real-time adaptability. Therefore, integrating physics-based mathematical modelling with advanced machine learning creates a powerful digital twin that offers a robust, real-time solution for accurate gearbox diagnostics.

Mathematical Modelling

The mathematical model simulates the gearbox's vibration signatures using fundamental kinematic parameters, such as shaft speeds and gear tooth counts. Specifically, it first converts the operating speed of the motor into a baseline input shaft frequency. The output shaft frequency is then derived by scaling this input based on the gear ratio. These rotational speeds dictate the gear mesh frequency (GMF), representing the fundamental rate at which the teeth engage. The amplitude of vibration for a healthy condition of the gearbox is given by:

$$y(t) = A \, sin\,(2\pi . GMF . t) \tag{1}$$

Fault in the pinion

When a crack develops on a pinion tooth, it creates a slight physical impact every time that tooth enters the mesh. Figure 1(a) visually captures this phenomenon. The healthy gear profile displays a single, clean amplitude spike at the GMF. In contrast, the faulty pinion profile illustrates the modulation effect. Vibration amplitude in this case is given by the following expression, where m is the modulation index and f_1 is the frequency of the pinion.

$$\begin{aligned} y_{fp}(t) &= A \, sin\,(2\pi . GMF . t) + 0.5 . A . m \, sin\,(2\pi .(GMF + f_1). t) \\ &\quad + 0.5 . A . m \, sin\,(2\pi .(GMF - f_1). t) y(t) \\ &= A \, sin\,(2\pi . GMF . t) \end{aligned} \tag{2}$$

Fault in the outer race of the bearing

When a defect occurs on the stationary outer race, the specific rate of transient impulse generated every time a rolling element passes over the crack is mathematically calculated as the Ball Pass Frequency

Outer (BPFO). Figure 1(b) shows a bearing with an outer race fault. In the given expression, number of balls (n), ball diameter (d), Average bearing diameter (D), Rotational frequency of shaft (fr), Contact angle (α).

$$BPFO = n.\, f_r \,/\, 2.(1 - d \,/\, D \cos\alpha) \tag{3}$$

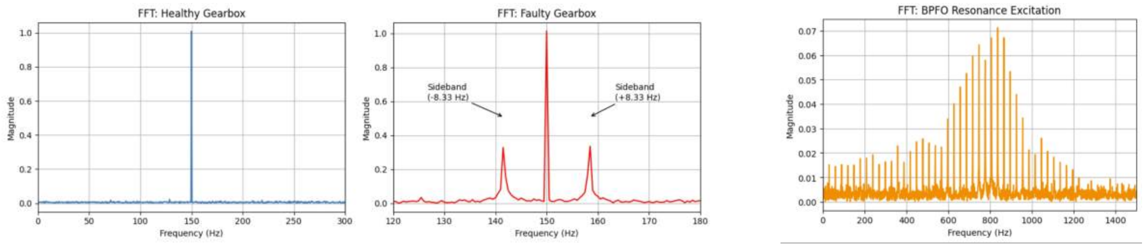

(a) (b)

Figure 1. (a) Fast Fourier Transform (FFT) spectra comparing a healthy gear mesh (left) to a faulty pinion exhibiting characteristic sidebands (right), (b) FFT spectra depicting an outer race bearing defect displaying BPFO harmonics.

Condition Monitoring and Fault Classification

Vibration signals from the gearbox are captured simultaneously by two accelerometers mounted on the bearing. The raw data is fed into a Simulink model that replicates the gearbox dynamics and reconstructs virtual sensor outputs for further processing. Key diagnostic features, including RMS, kurtosis, GMF sidebands, and bearing defect frequency envelope spectra (BPFO), are extracted by applying FFT. The processed sensor data is then trained using random forest framework trained to recognize four distinct states, including a healthy baseline, pinion damage, gear tooth faults, and bearing defect. The classification result is displayed instantly on the GUI, confirming a healthy state or raising a fault alert that identifies the specific defect class, enabling maintenance scheduling well ahead. Figure (2) shows the signal flow of the CM process.

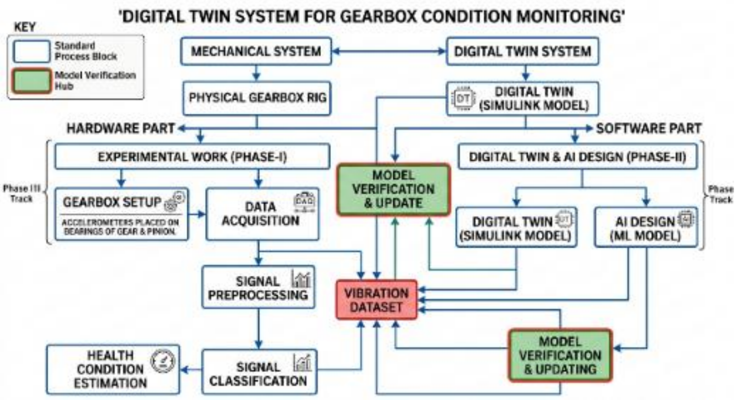

Figure 2. Condition monitoring pipeline

Conclusion

The current work presents a vibration-based CM framework for a single-stage gearbox that combines mathematical modelling with ML classification. Closed-form expressions derived for gear mesh dynamics and bearing impulse responses demonstrate how characteristic fault signatures manifest as sidebands and periodic transients in the frequency domain. Dual-accelerometer data acquisition feeds a feature-extraction pipeline, and an ensemble ML classifier distinguishes among four machine states with the diagnosis reported through a dedicated GUI.

Abstract ID: 2606

A Distributed Sensor Network-Based Digital Twin for Real-Time Environmental Monitoring and Intelligent Decision Support in Indoor Workspaces

Shriiraam Satish*, Vikash Vardhan V S

Bharathidasan Institute of Management, Trichy, India
**shriiraam8112002@gmail.com*

Abstract

Modern shared indoor workspaces operate under dynamically varying environmental conditions influenced by occupancy patterns, localized external factors, and centralized control systems. Parameters such as temperature, air quality, humidity, and lighting are typically regulated using static configurations that do not adapt to real-time usage or spatial variability. These environmental factors, although often overlooked, can significantly influence occupant comfort, cognitive performance, and productivity when suboptimal. Existing monitoring approaches, including recent digital twin-based IEQ monitoring frameworks, are largely fragmented, rely on limited or single-point sensing rather than distributed spatial sensing, and primarily support passive visualization rather than informed decision-making.

This work proposes a distributed sensor network-based digital twin framework for real-time environmental monitoring and analysis in indoor workspaces. The architecture consists of distributed sensing nodes positioned across spatially distinct zones to capture localized variation in indoor environmental conditions. These sensing nodes capture key parameters such as temperature, air quality, humidity, lighting, and occupancy-related signals. Data from the sensing nodes are aggregated at a local edge gateway/server, with node-to-gateway communication enabled over Wi-Fi using an MQTT-based data exchange mechanism. The framework is designed to be extensible to multiple sensing units within a scalable distributed sensing architecture.

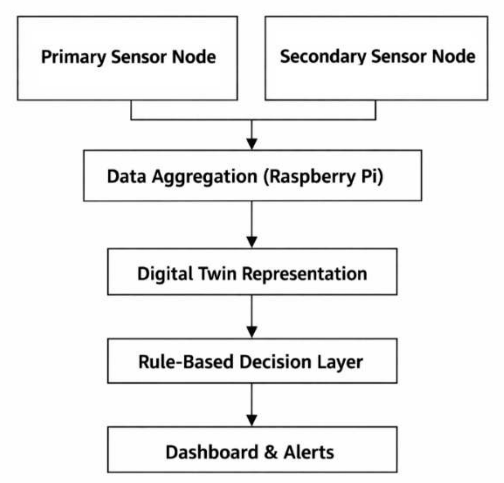

Figure 1. Distributed digital twin system architecture for indoor environmental monitoring and decision support.

The system maintains a synchronized virtual representation of environmental states across zones, enabling continuous tracking and evaluation of indoor conditions, similar in concept to prior IEQ-focused digital twin implementations. The present prototype maintains a synchronized, real-time virtual representation for environmental monitoring, early warning, and decision support, with closed-loop interaction with building control systems retained as a future extension. In addition to real-time state synchronization, the prototype is designed to support lightweight short-horizon estimation of selected indoor environmental states, particularly CO_2 concentration. This estimation is performed using a lightweight rule-based environmental state estimation model based on current environmental conditions, occupancy information, and predefined ventilation assumptions.

A rule-based decision layer using multi-sensor fusion is incorporated to map environmental states, occupancy conditions, and predefined thresholds to context-informed insights. Intelligent decision support in this context is achieved through structured interpretation of combined environmental and occupancy data, enabling identification of condition-specific inefficiencies. For example, elevated CO_2 levels under active occupancy conditions can indicate inadequate ventilation, while inconsistencies between occupancy patterns and thermal or lighting conditions can indicate potential inefficiencies. This rule-based fusion approach is inspired by prior corroborative sensing work, where multiple signals are considered jointly to improve inference reliability without reliance on complex models.

Indoor environmental conditions often exhibit spatial variation across zones that cannot be captured through single-point sensing. The proposed framework is designed to capture such cross-zone variation and occupancy-driven environmental dynamics, enabling more representative assessment of indoor conditions. By combining distributed sensing with occupancy-aware interpretation and short-horizon state estimation, the system is intended to support identification of conditions such as air quality degradation, ventilation inefficiencies, and localized comfort variation. Prior studies have established that elevated CO_2 concentrations and poor indoor air quality negatively impact cognitive performance and occupant comfort, and the proposed framework enables these conditions to be evaluated within a real-time environmental state representation.

Unlike a conventional IoT dashboard that primarily displays individual sensor readings, the proposed framework maintains a synchronized zone-wise virtual representation that integrates distributed sensing, occupancy-aware inference, and predictive early warning. Together, these capabilities form a unified decision-support architecture for actionable indoor workspace monitoring. This work contributes a low-cost distributed digital twin framework for actionable indoor workspace monitoring, with relevance to smart buildings and industrial workspaces where environmental conditions affect operational efficiency, worker well-being, and sustainability. While the current implementation focuses on sensing, state representation, and decision support, the architecture is designed to support future integration with building control systems. The proposed approach highlights the potential of combining IoT-based sensing, digital twin modeling, and multi-sensor fusion to enable practical, context-aware environmental management aligned with Industry 4.0 and Industry 5.0 paradigms.

Abstract ID: 2605

Digital Twin-Enabled Smart Grid for Hybrid Renewable Energy Integration with Self-Powered Sensing and Advanced Storage

Chakhyusa Saisrita Mittra, Pranav Jaganathan*

School of Electrical Engineering, Vellore Institute of Technology, Chennai, Tamil Nadu, India
**chakhyusa.saisrita2025@vitstudent.ac.in*

Abstract

The transition towards cleaner, more sustainable energy sources presents new challenges. The challenge is no longer just to create new renewable energy sources (which renewable technologies such as Solar Photovoltaic Systems, Wind Power Generation, and Biomass Energy successfully do), but to create stable, reliable, and usable energy from those sources in real-world situations. Each of these renewable energy sources can provide some challenge to stability because of their variable nature; renewable energies will be dependent upon and influenced by available environmental resources. For instance, Solar energy will only be produced when there is direct sunlight; wind energy will be further impacted by the weather; and Biomass energy relies on a constant management approach for continual energy production. Given the above, simply increasing the use of renewable energy sources will not necessarily provide a stable electrical system. Therefore, more must be done to create a power system that continuously monitors itself, understands how it is functioning, and intelligently responds to any changes. This paper provides an overview of a digital twin-enabled smart grid framework designed to address the above challenges through a more integrated, adaptive approach. The hybrid renewable energy system will use a combination of solar, wind, and biomass to generate green energy from one clean energy source called a coordinated grid.

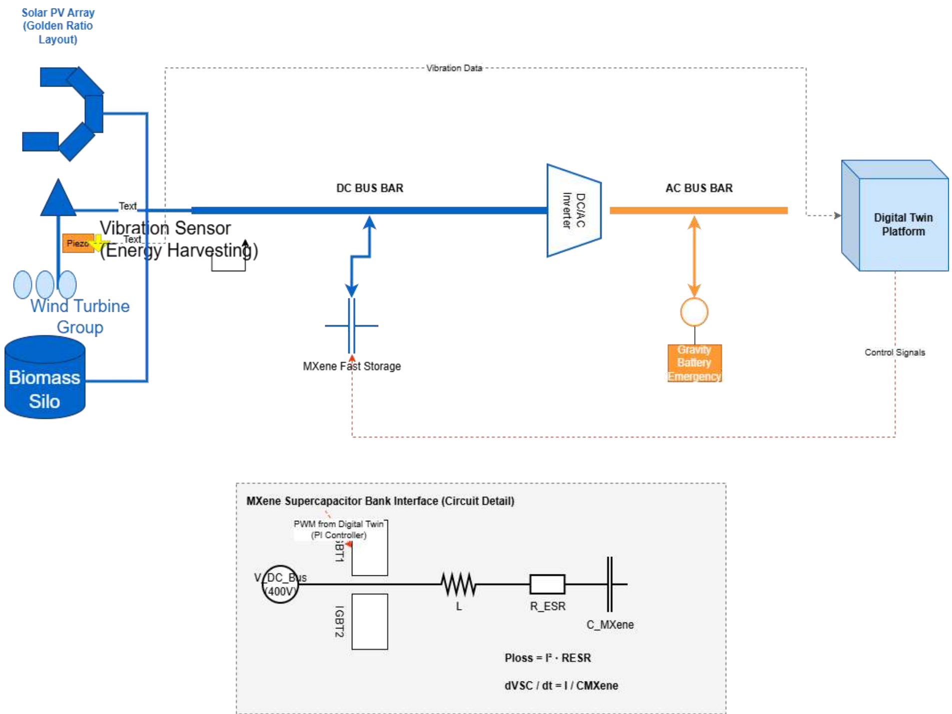

Figure 1. Schematic diagram of the Digital Twin – Smart Grid

In the solar area, silicon-based photovoltaic (PV) panels and advanced tandem solar cells will be used to produce reliable and efficient solar power. For both the solar and wind systems, the use of design patterns based on the Golden Ratio will help improve the collection of sunlight by minimising shading losses and optimising the use of available space. Finding many new ways to improve performance systems through simple design principles is an area of innovation for designers today. Wind safety is also part of this system. Wind safety will include both large traditional wind turbines (for high wind areas) and small windmills mimicking the shape of leaves (for lower wind speeds). Using a combination of these will ensure that energy production will continue even when the average wind speed is below ideal levels to maximise production. The biomass energy component of this hybrid system, it provides a steady and controllable form of power. Since biomass is less subject to changes in the environment than solar or wind power, it can be used as an important stabiliser within this hybrid energy system. A critical part of the proposed system is its ability to perform self-monitoring in a long-term sustainable way and through the efficient use of power generated by the very energy that would otherwise be lost to the environment. The system will generate its own electrical power using mechanical vibration and structural motion converted to electricity by Triboelectric Nanogenerator and lead zirconate titanate-based piezoelectric materials, and when to expect certain types of failures, thus changing the role of the grid from a static passive delivery of energy source to an active and self-aware energy delivery source that can respond to real-time changes. It is essential to ensure that there is coordination among different distributed devices to guarantee the efficient functioning of the system, since differences in time will result in an inconsistent process and delayed response. To solve the problem, the new framework suggests using a synchronization technique based on the idea of discrete time crystal systems, where the system exhibits a stable periodic behaviour despite being influenced by external forces. This paper focuses on incorporating hybrid energy sources and advanced energy storage systems into the system to manage variations in generation and load demand. MXenes have been considered high-performing materials for fast charging and discharging cycles and are therefore useful in dealing with short-term changes in generation. On the other hand, a gravity battery system is included in the system for the sake of emergency use during significant instability periods. The sensors that utilize power from solar panels and wind turbines to monitor the parameters, including temperature, vibration, and structural strain will transmit real-time information about all changes to a central database and process them using digital twin technology that creates a virtual copy of the physical grid. The use of comparison between the current state and the simulation in the digital copy will allow for anomaly detection, recognition of inefficiencies, and prediction of trends. Instead of regarding the generation of energy, monitoring of its performance, storing the information, and synchronizing all operations in an energy network separately, this method unites them in one system. It leads to the creation of a robust system, allowing adaptation and improvement, thanks to constant monitoring of all changes. In conclusion, the paper proves that the application of technologies in combination allows creating smarter energy networks.

Abstract ID: 3733

Deep Learning Enabled Microstructure–Property Correlation in A356 Aluminium Alloy: Focus on Silicon Refinement and Grain Globularity by GIPUC Method

Hamsashree S[1,2]*, Nilesh Kumar[1], S. K. Panigrahi[1], Vani V[2]

[1]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai 600036, India*

[2]*Department of Computer Science and Engineering, National Institute of Technology, Puducherry 609609, India*

**iithamsashree@gmail.com*

Abstract

Machine learning-driven computational frameworks based on convolutional neural networks (CNN), region-based convolutional neural networks (RCNN), segmentation network (SegNet), and boundary-constrained module- segmentation network (BCM-SegNet) architectures. Along with artificial neural networks (ANN), they were developed to systematically analyze, segment, quantify, and correlate the critical microstructural features governing the mechanical performance of A356 aluminium alloy. The primary objective of this study is to establish a robust integrated experimental and computational framework capable of accurately extracting microstructural descriptors from scanning electron microscopy (SEM) micrographs and predicting the resulting mechanical properties under varying casting conditions. The proposed framework focuses on essential microstructural parameters including eutectic silicon length refinement, globularity index of eutectic silicon particles, spatial distribution of silicon phases, dendrite fragmentation behaviour, and the globularity of α-Al grains formed during solidification. These microstructural attributes play a significant role in determining the strength and ductility relationship, hardness, and overall mechanical response of A356 aluminium alloy.

The CNN architecture was employed for automated feature extraction and classification of microstructural characteristics from SEM images. To further enhance segmentation accuracy and boundary localization of silicon particles and α-Al grains, RCNN models were integrated for region-based feature identification and object-level refinement. In addition, SegNet and BCM-SegNet architectures were implemented for semantic segmentation and pixel-level microstructure mapping, enabling highly accurate identification of eutectic silicon networks and globular grain morphologies. The incorporation of BCM-SegNet improved segmentation robustness through boundary correction mechanisms and enhanced contextual feature learning. Comparative analysis among CNN, RCNN, SegNet, and BCM-SegNet models was carried out to optimize the prediction and segmentation performance while evaluating the corresponding error percentages and classification losses associated with each architecture. The optimized framework demonstrated improved feature extraction capability, reduced prediction error, and enhanced consistency in microstructural quantification.

The extracted microstructural parameters were subsequently utilized as inputs to an ANN- based predictive model to establish correlations between processing conditions, microstructural evolution, and the resulting mechanical properties of the alloy. The dataset used for training, validation, and testing was generated through a scientifically engineered gas-infused pulsated ultrasonic casting (GIPUC) process. This advanced casting technique integrates ultrasonic streaming, controlled gas infusion, and vertical pulsation to manipulate melt hydrodynamics and solidification kinetics in a controlled manner. The synergistic interaction of these mechanisms significantly promotes dendrite fragmentation, enhances heterogeneous nucleation, refines eutectic silicon morphology, improves silicon globularity

index, and facilitates a more homogeneous spatial distribution of silicon particles throughout the aluminium matrix. Simultaneously, the process transforms coarse dendritic α-Al grains into near-globular morphologies, thereby improving the structural uniformity of the alloy.

The developed machine learning framework effectively captures the complex relationships between processing parameters, microstructural transformations, and mechanical performance metrics such as tensile strength, elongation, and hardness. Experimental observations combined with deep learning-based segmentation reveal that reduced silicon length, increased silicon globularity, uniform particle dispersion, and enhanced α-Al grain globularity contribute significantly toward improving the mechanical behaviour of the alloy. The comparative study of CNN, RCNN, SegNet, and BCM-SegNet models also demonstrates that advanced segmentation architectures substantially improve the precision of microstructural quantification and reduce overall error percentages in feature detection and property prediction. Through this integrated experimental and artificial intelligence-assisted methodology, the study establishes a reliable framework for intelligent microstructure characterization, predictive materials engineering, and process optimization of A356 aluminium alloy, thereby highlighting the critical influence of silicon morphology and grain globularity on the strength–ductility response and overall metallurgical performance of cast aluminium alloys developed by the GIPUC casting method.

Keywords: CNN, RCNN, ANN, SegNet, BCM-SegNet, GIPUC, A356

Abstract ID: 3458

Leveraging the potential of an innovative solid-state manufacturing route with an integrated Machine learning model for the recycling of Aluminum alloy.

Clement Rufus Raj A[1]*, D. K. Patel[2], Aveek Mohanty[2], P. Thiyagarajan[1], S. K. Panigrahi[2]

[1]*Department of Computer Science, Central University of Tamil Nadu, Neelakudy, Thiruvarur, India*
[2]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Madras, India*
**clementrufuscutn@gmail.com*

Abstract

Aluminium alloys are widely used as structural materials in key strategic sectors such as aerospace, marine, and automotive, owing to their high strength-to-weight ratio, low density, excellent corrosion resistance, and exceptional recyclability. Additionally, the demand for lightweight materials has driven increased consumption of aluminium and its alloys, leading to greater aluminium scrap. Thus, to minimise environmental degradation while reducing carbon footprints and greenhouse gas emissions, the recycling of aluminium and its alloys becomes critical. However, conventional aluminium recycling technologies predominantly rely on remelting and are often associated with (i) high energy consumption, (ii) non-uniform material properties, and (iii) the generation of hazardous dross and flux emissions, posing significant environmental and health risks. To overcome these drawbacks, the present study explores a innovate solid-state manufacturing technology for recycling aluminium alloy chips into a useful product. The proposed recycling route leads to (i) reduced processing time, (ii) a uniform microstructure, and (iii) enhanced diffusion kinetics compared to conventional methods. Furthermore, a machine learning model is integrated to efficiently predict mechanical properties, namely density and micro-hardness, as a function of four key input process parameters: sintering temperature, power rating, holding time, and heating rate. Machine learning models, specifically Random Forest and XGBoost regressors, were trained on the experimental dataset to predict mechanical properties, and feature importance analyses using the Pearson coefficient and SHAP were performed to identify the dominant process parameters. The integrated approach offers an environmentally friendly, data-driven predictive model for recycling aluminium alloys. It aims to minimise experimental trials, optimise the process, improve material performance, and support sustainability and circular economy principles.

Abstract ID: 3229

Cloud-Enabled Digital Twin Framework of an Industrial Robot for Seamless Integration and Task Automation in Autonomous Factories

Deep Singh[1], Shiwam Shorya Sharma[2]*, N. Arunachalam[1], Pramod Kumar Shrivastava[3], Varun Balachandra[3]

[1]Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, Tamil Nadu, PIN – 600036, India

[2]IITM Pravartak Technologies Foundation, B5-01, IITM Research Park, Tharamani, Chennai, Tamil Nadu, PIN - 600113, India

[3]Accenture Global Solutions, RMZ Ecospace, Ecospace Business Park Accenture Innovation Hub, Bellandur, Bengaluru, Karnataka, PIN - 560103, India

*shiwamshoryasharma@gmail.com

Abstract

The current technological advancements and adoption of digital technologies by industries necessitate the implementation of autonomous and digital twin systems to enhance production in terms of quality and quantity. The industrial robot is one of the key components required in order to automate the shop floors. The use case of industrial robots is not limited to material handling and machine tending but is also being implemented for the assembly of components, welding, quality inspection with an integrated vision system, etc. Apart from these, various research demonstrates the application of industrial robots in the field of manufacturing, such as drilling, milling, grinding, robo-forming, etc. With the expanding use cases of industrial robots, the current era necessitates the development of a digital twin of industrial robots, which not only monitors the state of health but also allows autonomous operation within the shop floor environment.

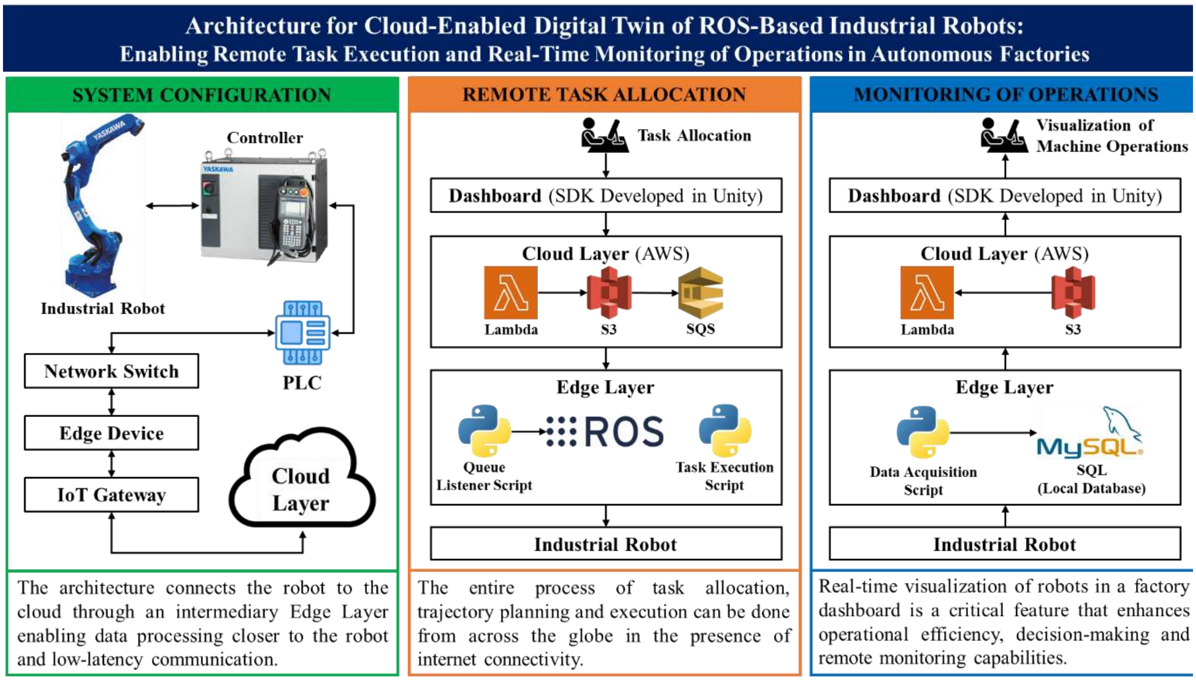

Figure 1. Comprehensive cloud-enabled Digital Twin architecture of ROS-based Industrial Robots.

The current research focuses on the development of a cloud-enabled digital twin for industrial robots, designed to automate tasks assigned from any location worldwide. The notion of DT in Industry 4.0 is the creation of a digital representation of a physical system that is constantly updated with real-time data from the actual world. This data interchange is bidirectional. By exploiting internet connectivity, the system allows continuous, near real-time communication between the physical robot and its digital counterpart. It also enables interaction with other physical systems in an autonomous factory to enable efficient task scheduling, monitoring, and execution. While industrial robots are commonly used today for repetitive tasks, with integration into factory systems being relatively straightforward through Programmable Logic Controllers (PLC) and similar automation systems. New challenges arise when implementing these robots in autonomous factories to work in collaboration with other physical systems. In such environments, tasks can be assigned remotely via cloud systems, requiring the robot to not only perform its assigned work but also collaborate with other physical systems within the factory. This level of inter-system cooperation and flexibility is crucial for achieving the full potential of autonomous or smart factories. Moreover, integrating diverse robotic systems with cloud-based services is complex.

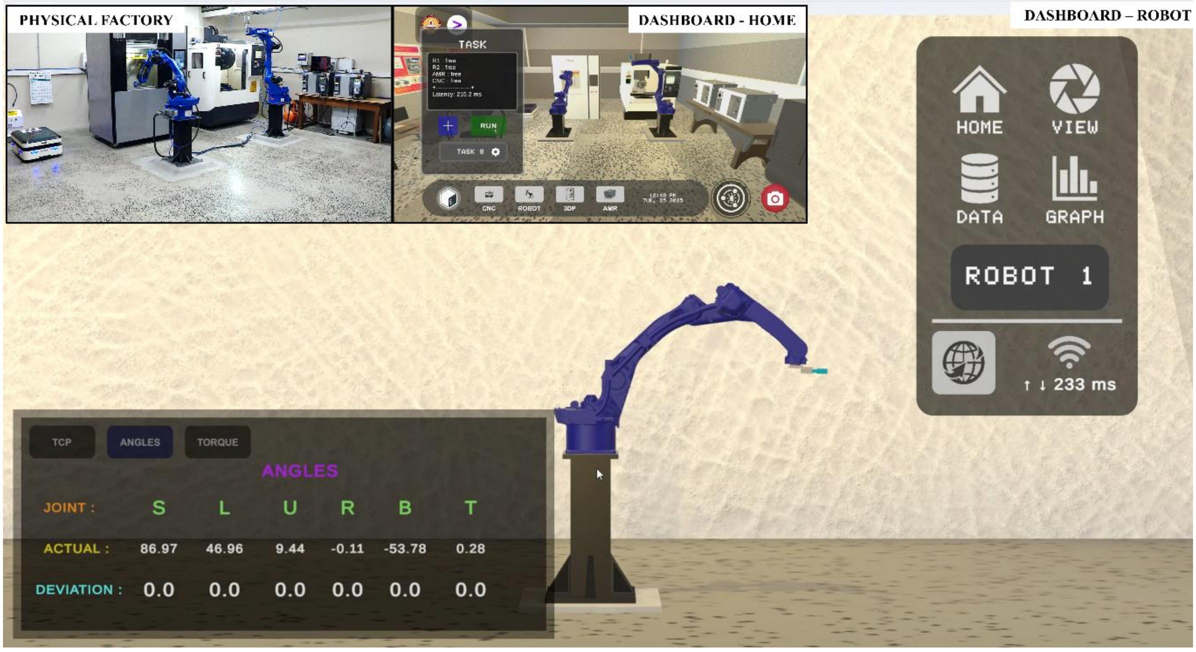

Figure 2. Digital dashboard of an autonomous factory developed in Unity Game Engine for remote task allocation, execution and monitoring.

A digital twin framework (as presented in Figure 1) has been developed to mimic the robot's physical and operational characteristics, enabling task allocation and performance evaluation. Global connectivity is established using cloud-based technology to allow users to assign tasks using the user-friendly dashboard (Figure 2), developed using the Unity Game Engine, from across the globe ensuring smooth and near-real-time communication between the physical system and its digital counterpart. The integration of the robot with other systems, coupled with the ability to dynamically alter itself in conjunction with numerous machines, is a primary focus of this research. With an emphasis on task execution, operational autonomy, and inter-system collaboration, it examines the framework and application of the cloud-based digital twin architecture. The developed framework showcases the potential of cloud-enabled digital twins to revolutionize the manufacturing industry through the implementation of remote, autonomous operations for collaborative, cloud-based operations in autonomous factories but not limited to it.

Abstract ID: 3228

Machine Learning-based Fault Diagnostics for Industrial Robots using Accelerometer Signals

Manikala Prem Sainadh[1]*, Deep Singh[2], Shiwam Shorya Sharma[3], N. Arunachalam[2]

[1]*Department of Computer Science, National Institute of Technology Puducherry, Puducherry, India*
[2]*Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India*
[3]*IITM Pravartak Technologies Foundation, Tharamani, Chennai, India*
**premsainadh99@gmail.com*

Abstract

With the advancements of Industry 4.0 technologies, manufacturing is transitioning from traditional mass production shop floors to autonomous smart/dark factories operating round the clock with minimal human intervention. Industrial robots (IRs) play a crucial role in enabling autonomous operation on manufacturing shop floors due to their varied applications, including material handling, machine tending, welding, parts assembly, machining, polishing, etc. As with any electromechanical system, the IRs and their subsystems are prone to degradation and failure. Hence, continuous monitoring is of utmost importance when such electro-mechanical systems are deployed within unmanned manufacturing shop floors. IRs consist of a series of links and joints, of which the links are rigid, and the joints comprise actuators, gears and bearings, which are prone to degradation and wear. Although IRs are intended for high-precision applications, gradual degradation at the joints can lead to failure and unpredictable shutdowns of the equipment and the entire production line. If monitored continuously, various faults and their severity can be detected well in advance, enabling condition-based maintenance or planned maintenance with limited downtime.

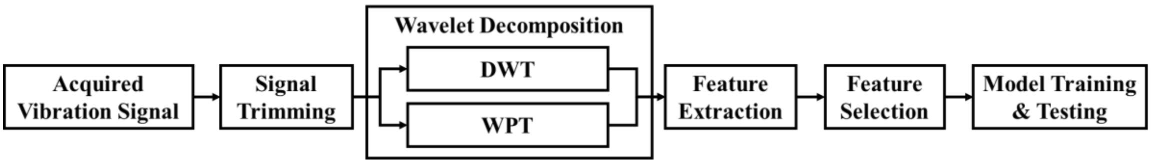

Figure 1. Proposed methodology for gear fault severity estimation using DWT and WPT.

(a) Healthy flexspline (b) L_1 Severity (c) L_2 Severity (d) L_3 Severity (e) L_4 Severity

Figure 2. Faults induced in the flexspline of the harmonic reducer.

Table 1. Severity index of flank wear fault of harmonic reducer.

Fault Class	Severity Index (Level)				
F2	0 (F0)	0.38 (L1)	0.46 (L2)	0.65 (L3)	1 (L4)

The current study demonstrates a data-driven architecture (as shown in Figure 1) for fault diagnostics of an IR. For this study, flank wear fault has been seeded on one of the flexspline teeth of the harmonic gear reducer at joint 4 of the Yaskawa GP12 IR. The flank wear fault has been seeded at four levels as presented in Table 1, and the corresponding flexspline wear image has been presented in Figure 2. A tri-axial accelerometer has been mounted on joint 4 of the robot to capture the vibration signal of the robot during its operation at an 8 kHz sampling frequency. The joint 4 of the IR is operated at five different angular velocities (48.7°, 96.1°, 141.7°, 186.8°, and 229.3° s^{-1}) covering both low-speed and high-speed regimes for 50 iterations under healthy conditions of the reducer and also at each level of severity. Acquired signals are then trimmed, reflecting the motion of the robot joint from -165° to +165°, which corresponds to 56 meshing cycles per tooth of the flexspline with the circular spline. The trimmed dataset is then decomposed using Discrete Wavelet Transform (DWT) and Wavelet Packet Transform (WPT) because the vibration signals are non-stationary and possess transient fault signatures. DWT implements a hierarchical decomposition in which only approximation (low-frequency) components are further split at each level, whereas WPT decomposes both approximation and detail (high-frequency) components at every level.

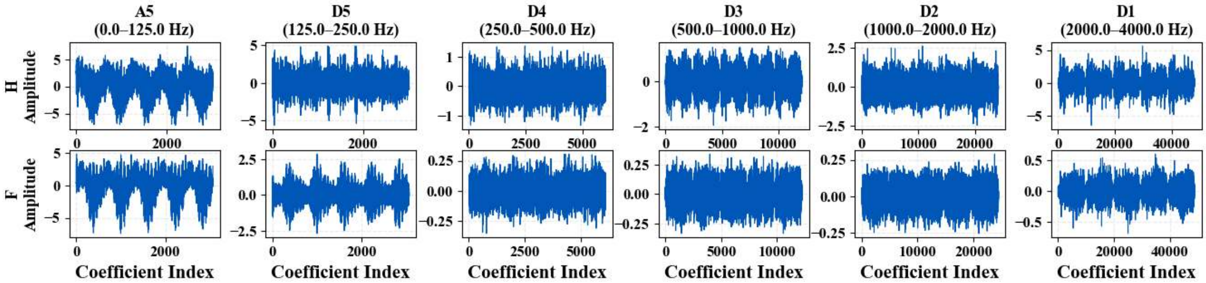

Figure 3. Five-level DWT decomposition of healthy (H) and faulty (L 4) vibration signals of the reducer.

From each wavelet band, a comprehensive set of statistical features, including root mean square (RMS), energy, variance, kurtosis, crest factor, and entropy, is extracted to construct the feature dataset. Feature importance is then evaluated using a Random Forest algorithm for both Discrete Wavelet Transform (DWT) and Wavelet Packet Transform (WPT) representations. The top 10 most significant features obtained from the DWT-based analysis are illustrated in Figure 4. Compared to conventional filter-based methods, which neglect feature interactions, Random Forest is employed due to its ability to capture complex non-linear relationships and interdependencies among features. Based on the ranked feature set, a regression model is trained to estimate the fault severity. The performance of the regression model is assessed by comparing the predicted severity values with the ground truth, as shown in Figure 5 for the DWT-based features. The DWT-based model achieves a Mean Absolute Error (MAE) of 0.1055, Mean Squared Error (MSE) of 0.0360, and an R^2score of 0.9830. In comparison, the WPT-based model yields an MAE of 0.0742, an MSE of 0.0646, and an R^2score of 0.9694.

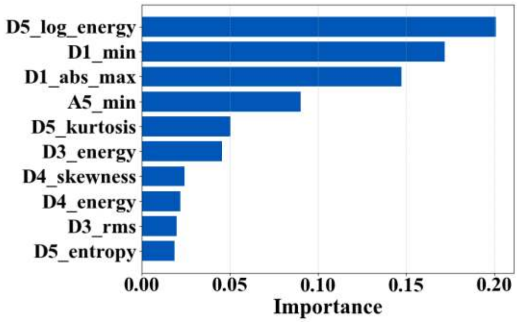

Figure 4. Feature importance ranking.

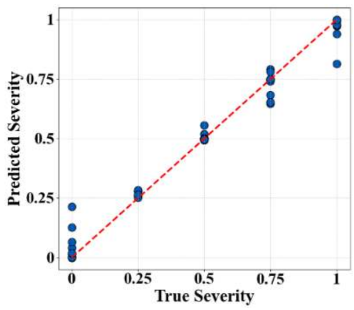

Figure 5. True vs predicted fault severity.

Although the WPT-based model exhibits a lower MAE, the DWT-based model demonstrates superior overall performance in terms of variance explanation (R^2) and consistency with true severity trends. This indicates that the DWT-based feature representation more effectively captures fault characteristics, making it a reliable approach for fault-severity estimation.

Abstract ID: 3204

Automated Painting of Aerospace Components: An End-to-End Pipeline from Geometrical Pre-processing to Robot Execution

Lavanya Rajagopalan Iyer*, Shreya Tejal, Aruna V S, Mythrayee K L

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai, India
**me23b178@smail.iitm.ac.in*

Abstract

This project presents an end-to-end automated spray-painting pipeline for axisymmetric aerospace components. Automated painting offers higher consistency, lower paint wastage, and lower long-term operational cost than manual application—advantages that matter most in aerospace manufacturing, where coating uniformity directly affects surface finish, weight, and aerodynamic performance. The pipeline centres on a shared geometric pre-processor feeding two independent trajectory-generation methods: a helical sweep validated through PyBullet physics simulation, and a coordinated-motion approach implemented as executable ABB RAPID code and verified in RobotStudio.

Methodology

Geometric pre-processor

Several restraints that were placed upon the workpiece geometry must be acknowledged. First, the workpiece is reducible to a single profile revolving about a central axis. This accounts for the fact that a significant number of common aerospace components are axisymmetric and reduces the computational power required for trajectory generation and simulation. Second, the meridional profile must be continuously differentiable. This ensures that every point on the surface has a well-defined, unique tangent plane, enabling the trajectory generator to accurately compute normals at every point along the path. Finally, the local radius must be sufficiently large that the spray cone's projection onto the surface remains approximately circular rather than elongating into an ellipse, which doubles deposition density at the minor axis and produces a stripe of excess film thickness.

The geometric pre-processor was built using Python's Open3D library. It accepts the 3D model of the component in STL file format and transforms it into a mesh, whose topology is then validated to ensure watertightness and orientability. Once validated, the pre-processor then extracts the triangles relevant to the coverage area, resulting in a clean, open-shell mesh used as the common input to both trajectory-generation methods that follow. Violation of the smooth-profile constraint produces abrupt normal divergence between adjacent waypoints, generating position discontinuities that cannot be tracked at constant TCP velocity — the dominant factor for uniform deposition. Violation of the minimum-radius constraint causes the footprint to elongate into an ellipse, doubling deposition density at the minor axis. Both are therefore enforced as entry conditions rather than soft guidelines.

Method 1 - Helical Sweep with Stationary Workpiece

The first trajectory generator converts a point cloud sampled from the pre-processed mesh into a helical path traced by a stationary fixed-base manipulator. The conversion module extracts per-point surface normals, recovers the workpiece symmetry axis via PCA, enforces a consistent outward-normal convention, and offsets each point by 250 mm along the surface normal to produce end-effector spray poses. Points are re-ordered along the helical sweep parameter $s = z_a + \left(\frac{p}{2\pi}\right)\theta$, with pitch p set to deliver 50% footprint overlap — the consensus optimum for Gaussian deposition profiles. Spray-dynamics parameters (cone half-angle 22.5°, footprint radius $R \approx 103.5$ mm, $\sigma = R/3$, traverse speed

300 mm/s) are drawn from peer-reviewed literature. The IK solver was configured such that any pose not reproducible within 0.1 mm is flagged rather than silently approximated. Orientation tolerance is bounded by tilt within ±15° from the surface normal; zero nominal tilt is enforced by construction at every waypoint. Validation is simulation-based; no physical pilot trials were conducted.

Method 2 - Coordinated Motion with Rotating Positioner

The second trajectory generator extracts the meridional profile of the workpiece, along which the robot makes passes while the workpiece continuously rotates on a positioner. It sequentially sections the workpiece along perpendicular planes to obtain a planar polyline representation of the meridional profile, which is offset by the requisite standoff distance. However, as the workpiece curvature changes, standoff distances must change accordingly to ensure uniform coverage and thickness, which are critical aspects in high-performance aerospace applications. Therefore, by assuming a fixed angular overlap fraction K_o and that the local standoff distance d scales with the local radius R, we obtain the following relationship:

$$d = R\sin\left(\frac{\Delta\theta}{2(1-K_o)}\right)\cot\left(\frac{\alpha}{2}\right) \tag{1}$$

where α is the spray cone full-angle and Δθ is the angular sector subtended at the symmetry axis. The RobotStudio simulation uses an IRB 580 paint robot with an IRP-A rotary positioner under RobotWare 6.16.03. Every waypoint is encoded as a unit quaternion, eliminating gimbal lock and enabling smooth SLERP interpolation. The wobjdata frame attached to positioner unit STN1 applies $P_{world} = T(\theta)P_{local}$ in real time, absorbing encoder error rather than accumulating it in the robot joints, and the z2 zone enforces constant TCP velocity throughout. No physical pilot trials were conducted; validation is RobotStudio simulation-based.

Results

Helical Sweep with Stationary Workpiece

The trajectory was validated through five progressive PyBullet simulation iterations. Four geometric invariants were verified: standoff-radius error of 0.000 mm against the expected 300 mm; 0/480 misaligned spray-axis orientations; quaternion–matrix consistency of 2.29×10^{-16}; and a sweep_index–axial-position Pearson correlation of 0.972. In the final PyBullet run over 11,600 simulation steps, 634 frames recorded arm-to-part contact (5.5% of the run), comprising 121 distinct touch–leave events and a deepest penetration of −16.99 mm. The estimated real-world cycle time at 300 mm/s traverse over the 100 mm axial length of this component is approximately 21 s per circumferential pass.

Coordinated Motion with Rotating Positioner

The profile extraction yielded 25 robtarget waypoints (pPaint001 to pPaint025) covering the full axial extent of the nose-cone profile. A scaled 1:2 variant of the trajectory was generated and executed successfully in RobotStudio, confirming pipeline invariance to workpiece size — the curvature-adjusted pitch formula automatically rescales d with R. The painting cycle executed without controller errors: approach and retreat moves with ±100 mm offsets bracketed the path cleanly, WaitTime delays of 0.1 s on valve activation and 0.05 s on deactivation absorbed transient flow, and no TRAP fault interrupts were triggered. Complete outer-surface coverage was confirmed visually in the RobotStudio simulation environment.

Discussion

The collision events observed in the helical sweep are an expected consequence of using a fixed-base manipulator for full-circumferential coverage — the arm must physically reach across the workpiece for far-side waypoints, which PyBullet's collision-blind IK cannot prevent. The coordinated-motion method sidesteps this constraint structurally by having the positioner present each surface sector to the gun rather than requiring the arm to reach it, and correspondingly achieves complete coverage without collision in RobotStudio. Together, the two methods demonstrate that the same geometric pipeline can serve both simulation-first and deployment-first workflows, with the choice between them governed primarily by available hardware. Extending this work to non-axisymmetric components would require patch decomposition and a full collision-aware planner — directions the current pipeline is structured to support.

Abstract ID: 3198

Bubble Detection and Parameter Optimisation for Sterile Vial Filling

Aditi Vaidya*, Arshiya Singh, Hemanshini Desai, Kanishka Pareta

Department of Mechanical Engineering, Indian Institute of Technology, Madras
**me23b248@smail.iitm.ac.in*

Abstract

The pharmaceutical industry offers a variety of problems that can be solved via AI integration. One such class of problems is combating the issue of air bubbles and foams in vials. Air bubbles and foams not only compromise dose accuracy but may also lead to injection discomfort, product rejection, and serious regulatory issues. Air bubbles and foams are primarily caused by a high flow rate, leading to turbulent flow and incorrect nozzle height. Air bubbles and foams take different forms depending on the cause. We aim to train a model that learns optimal parameters for vial filling and is capable of mapping cause and effect. Further, we would like to integrate model suggestions as feedback into the manufacturing workflow.

Our Framework: In this work, we propose a unified framework that combines computer vision-based detection with safe, data-driven optimization to address bubble formation in sterile vial filling. Our approach is structured in two key stages: (i) real-time detection and characterization of air inclusions, and (ii) safe online adaptation of CPPs based on observed process deviations.

The first component focuses on real-time bubble detection using a convolutional neural network based on the YOLO v11 architecture for high-speed object detection. It solves a logistic regression problem of classifying the images based on the presence of bubbles and foam and counting the bubbles if they are present. There is a strong dependence of the formation of foam on the sudden changes in volume flow rate, as illustrated in this paper. Foaming in the liquid occurs when the liquid surface is disturbed. This may be caused by a high-volume flow rate coming into the vial, a sudden change in flow rate, or a sudden influx of fluid, which causes a splash at the start. At higher flow rates, the liquid exits the nozzle with greater momentum, increasing the impact velocity upon contact with the liquid surface in the vial. This enhances inertial forces relative to surface tension, effectively increasing the Weber number of the system. This results in air entrapment and the formation of foam. As a first step, we start with the observation that the initial fill should be slowed to eliminate the splash, while the later fill appears stable and thus it can be sped up without generating significant foam. Excessive foam generation can lead to overflow, compromising fill volume accuracy and increasing contamination risk. The dependence of bubbles is related to fluid flow rate and nozzle height as explained here. When the nozzle is positioned higher above the liquid, the falling jet creates strong splashing and traps air, leading to increased bubble formation. In contrast, lowering or submerging the nozzle significantly reduces air entrainment and thus minimizes bubbles.

Directly training a model to modify CPPs using bubble formation as a feedback signal presents significant practical challenges. With a goal to train a model to change CPPs based on bubble formation as a feedback mechanism may cause the process parameters to change drastically, which is infeasible to implement. While large datasets of static vial images are readily available for training detection models, these images typically lack corresponding process metadata, making it difficult to establish a reliable mapping between observed air inclusions and the CPPs that produced them. Consequently, learning a direct relationship between process parameters and bubble formation would require extensive controlled experimentation, which is both time-consuming and resource-intensive. To circumvent this

problem, we adopt an online learning approach wherein the system iteratively improves based on real-time observations.

It is important to note that in most industrial settings, approximately optimal CPPs are already identified through prior experimentation and process design. An online algorithm, as we propose, may be used to further fine-tune the process. We plan to evaluate the use of the SAFEOPT approach to change the control process parameters in a feasible manner. The requirement of sterile vial filling aligns with Bayesian optimization. The objective is to minimise the number of bubbles formed, which is an output of the CNN model, while maintaining speeds suitable for production. The SAFEOPT algorithm is a safe Bayesian optimization framework designed to optimize an unknown objective function under safety constraints. It models both the objective and constraint functions using Gaussian Processes, enabling probabilistic estimates of performance and safety across the parameter space. At each iteration, SAFEOPT maintains a set of parameters that are known to be safe with high confidence and selects new evaluation points by balancing two objectives: expanding the safe region and improving performance within it. This is achieved by leveraging confidence bounds derived from the Gaussian Process, ensuring that any newly selected parameter configuration satisfies the safety constraints with high probability. Starting from an initial set of known safe parameters, the algorithm incrementally explores neighbouring regions, exploiting the assumed smoothness (Lipschitz continuity) of the underlying functions to infer safety in unexplored areas.

In the context of sterile vial filling, this framework admits a one-to-one correspondence. The parameter domain corresponds to the space of critical process parameters (CPPs), including fill speed, nozzle height, and pressure. The objective function is defined in terms of minimizing bubble and foam formation, as quantified by the outputs of the vision-based detection model. Safety constraints map to the maximum number of admissible bubbles. The initial safe set is naturally provided by nominal CPPs obtained through prior process design and validation. Within this formulation, SAFEOPT enables iterative, data-driven adjustment of CPPs that remains within safe operating bounds while progressively improving process performance, making it particularly well-suited for real-time, PAT-aligned control of sterile filling systems.

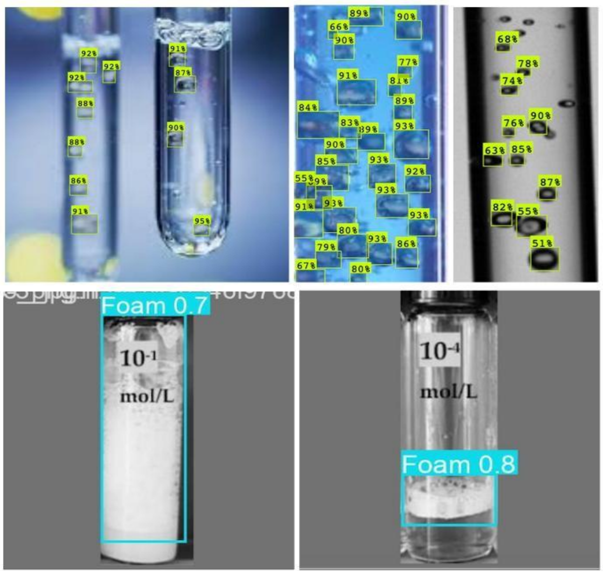

Figure 1. Detection of bubbles and foam by the CNN model.

Note: All authors are equal contributors, listed alphabetically

Abstract ID: 3159

Tribological Behavior of Wire Arc Additive Manufactured Super Duplex Stainless Steel using Laser Shock Peening

R. Kalaiyarasan*, J. Jerald, M. Sivakumar

Department of Production Engineering, National Institute of Technology Tiruchirappalli, Tiruchirappalli, Tamil Nadu. India-620 015.

**rajakalaiyarasa@gmail.com*

Abstract

Super Duplex Stainless Steel (SDSS) components fabricated using Wire Arc Additive Manufacturing (WAAM) offer distinct advantages for high-performance applications in corrosive and wear-intensive environments such as offshore, marine, and chemical processing industries. In this study, SDSS specimens were produced using the Cold Metal Transfer (CMT) WAAM process with optimized parameters: deposition current of 122 A, voltage of 11 V, wire feed rate of 2.75 m/min, welding speed of 220 mm/min, and shielding gas flow of 15 L/min argon, resulting in an average wall thickness of 7 mm and layer height of 2.5–2.6 mm. Dry sliding wear tests were conducted under two loading conditions: 10 N load, 15 mm stroke, 5 Hz frequency, and 20 N load, 10 mm stroke, 10 Hz frequency. Results showed low weight loss (0.00447–0.0101 g), wear rates ranging from 10^{-6} to 10^{-7} cc/s, and coefficients of friction (COF) between 0.53 and 0.60, with surface roughness (Ra) values ranging from 0.591 to 0.765 μm. To further improve surface integrity and tribological performance, Laser Shock Peening (LSP) was applied at energies of 2 J and 3 J under the more severe 20 N / 10 Hz condition. Post-LSP, wear resistance showed significant improvements, demonstrating the effectiveness of the hybrid CMT-WAAM and LSP approach. This combination offers enhanced surface properties while retaining the geometric flexibility and cost-effectiveness of WAAM, making it a promising method for producing high-performance SDSS components suitable for demanding industrial environments.

Keywords: Additive Manufacturing, Wear, Friction, Laser Shock Peening, Cold Metal Transfer.

Introduction

Super duplex stainless steel fabricated by the Cold Metal Transfer-based Wire Arc Additive Manufacturing process is gaining attention because of its high strength, corrosion resistance, and suitability for large engineering components. However, layer-by-layer deposition produces directional microstructure, repeated thermal cycles, and surface heterogeneity, which can affect wear performance under sliding contact. The present work investigates the wear behaviour of CMT-WAAM-fabricated SDSS before and after laser shock peening.

Methodology

CMT-WAAM fabricated SDSS samples were prepared and divided into untreated and laser shock peened conditions. Reciprocating sliding wear tests were performed using different load, stroke length, and frequency conditions. The untreated samples were tested at 10 N, 5 mm stroke length, 5 Hz frequency and 20 N, 10 mm stroke length, 10 Hz frequency. The LSP-treated samples were tested at 20 N, 10 mm stroke length, and 10 Hz frequency with laser energies of 2 J and 3 J. Wear rate, coefficient of friction, and surface roughness were measured before and after testing. The worn surfaces were compared to identify abrasive wear, adhesive wear, plastic deformation, and surface damage. ASTM G133 is more suitable for this type of test because it covers linearly reciprocating ball-on-flat sliding wear of metals and other wear-resistant materials.

Results and Discussion

The untreated CMT-WAAM SDSS samples showed a relatively stable coefficient of friction in the range of 0.53–0.60 after the initial running-in period. A higher load of 20 N combined with a frequency of 10 Hz resulted in greater wear than the lower-load condition. Laser shock peening markedly improved the wear resistance, reducing the wear rate from the order of 10^{-6} to 10^{-7} cc/s. Among the treated samples, the 3 J condition produced the lowest wear rate, with a slight improvement over the 2 J condition. The three-dimensional surface profiles also showed that LSP altered the original surface morphology through localized plastic deformation.

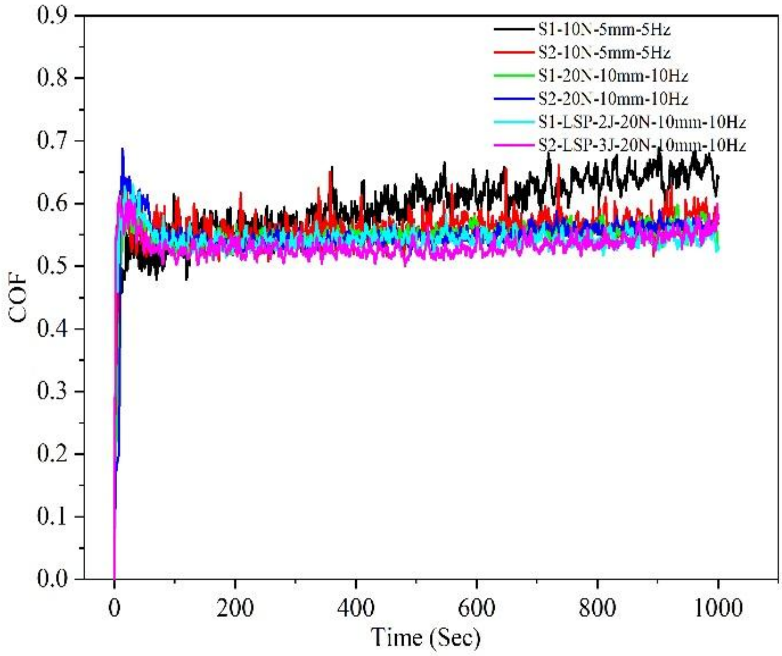

Figure 1. Variation in the coefficient of friction with sliding time for the investigated samples.

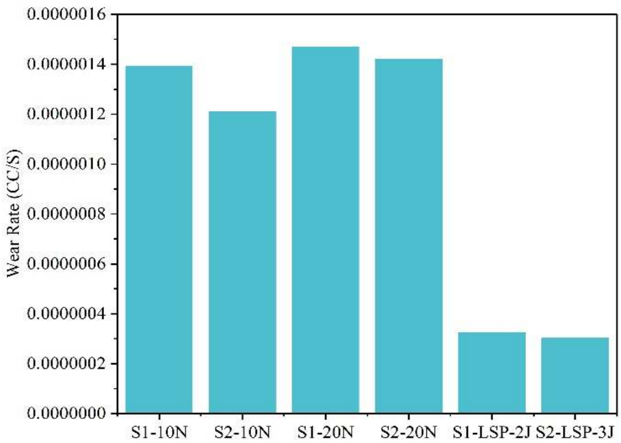

Figure 2. Comparison of the wear rates of the investigated samples.

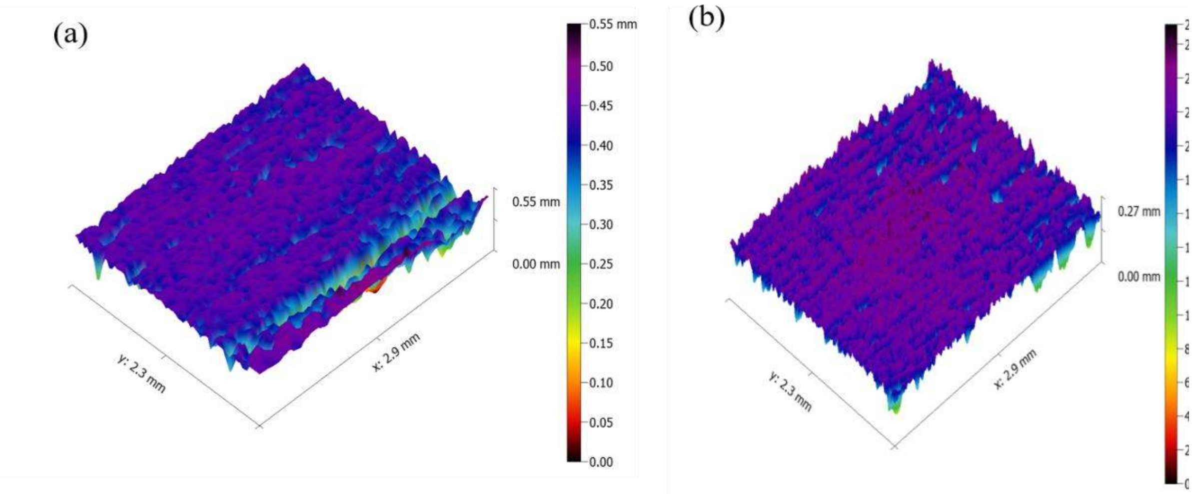

Figure 2. Three-dimensional surface topography of the investigated samples: (a) as-built condition and (b) laser shock peening-treated condition.

Abstract ID: 3151

AI-Based Predictive Maintenance Framework for Wind Turbine Using Real-Time SCADA Data

Srihari Prasad*, Lavlin Jaison, Parth K Makati, Sreerag P N

Department of Mechanical Engineering, Indian Institute of Technology Madras, Chennai 600036, India
**me23b249@smail.iitm.ac.in*

Abstract

Wind turbines have critical components whose failures result in significant operational downtime and maintenance costs. Many wind farms are in remote areas and exposed to harsh natural environments such as strong winds and high temperatures, which can lead to variable conditions. This means the wind turbines often endure continuously varying loads, which can result in failures in transmission chain components. Conventional SCADA-based monitoring systems rely on threshold-based alarms, which detect faults only after critical temperature limits are exceeded. This work proposes an AI-driven predictive maintenance framework for proactive thermal health monitoring of wind turbine gearboxes using real-time SCADA data.

Figure 1. A high-fidelity digital twin of the Vestas V82 Wind Turbine (modelled in Fusion 360)

Using real-time SCADA data for two weeks from a Vestas V82 wind turbine, our work integrates multi-target temperature forecasting - predicting gear oil, front bearing, and rear bearing temperatures approximately 20 minutes ahead – using gradient-boosted regression models trained on lag, rolling, and other operational and thermal features, including wind speed, generator rpm, ambient temperature, and Nacelle internal temperature. Based on these predictions, our end goal is to develop a predictive-maintenance layer using residual analysis, thermal imbalance detection, and threshold-based risk estimation, enabling early detection of abnormal heating behaviour, bearing asymmetry, and estimation of time-to-threshold conditions. As summarized in Table 1, all models achieved strong predictive accuracy with R^2 values exceeding 0.89 and MAE remaining below 1.3 °C across all targets.

Table 1. Performance metrics of XGBoost regression models for 20-minute ahead forecasting of gearbox oil and bearing temperatures on the Vestas V 82 SCADA dataset.

Target	MAE (°C)	RMSE (°C)	R^2
Gear Oil (GrO)	0.476	0.659	0.942
Front Bearing (GrBF)	1.288	1.679	0.895
Rear Bearing (GrBR)	0.932	1.300	0.893

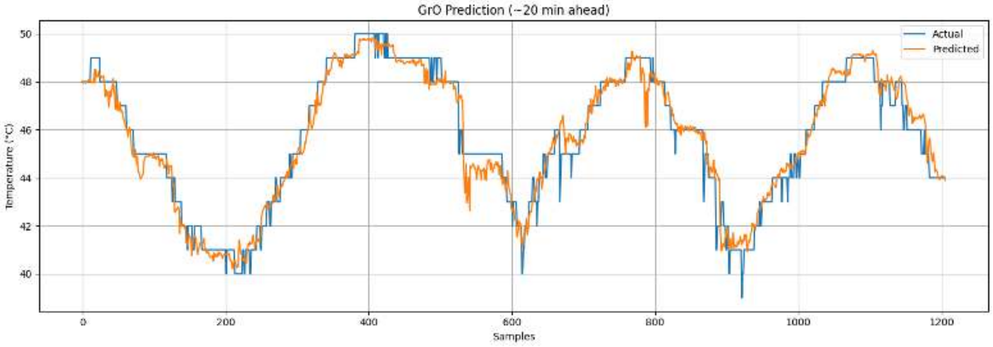

(a)

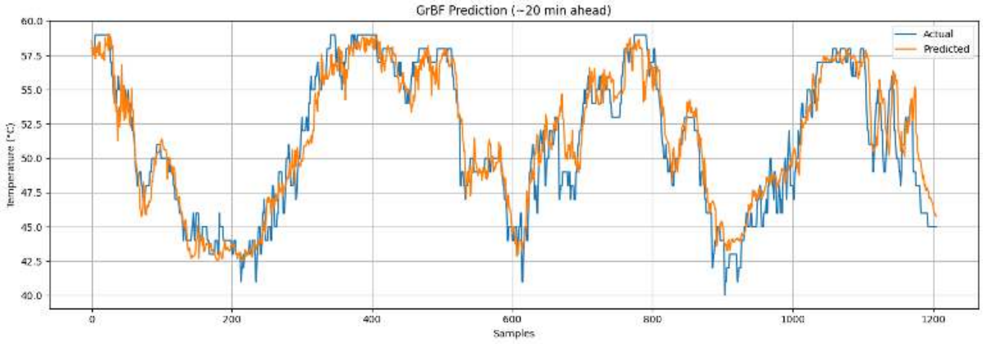

(b)

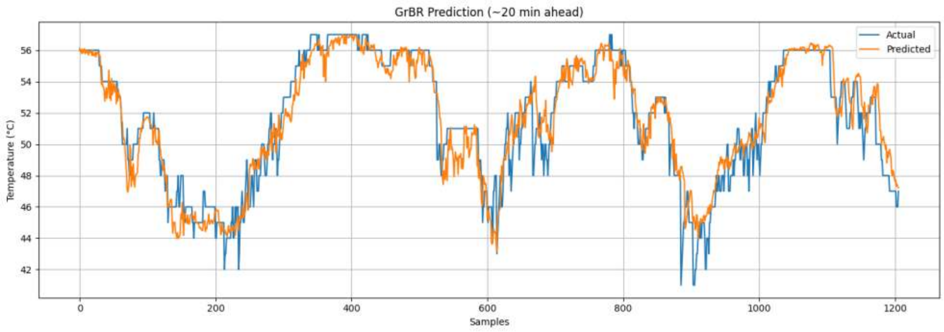

(c)

Figure 2. Actual versus predicted temperature plots (GrO) over the test set, demonstrating 20-minute ahead forecasting using a gradient-boosted regression model trained on Vestas V82 SCADA data, (a): Gear oil temperature (R^2 = 0.942, MAE = 0.476°C), (b): Gear front bearing temperature (R^2 = 0.985, MAE = 1.288°C), (c): Gear rear bearing temperature (R^2 = 0.893, MAE = 0.932°C).

A framework is proposed, to go beyond thermal forecasting to deliver a comprehensive gearbox and turbine health monitoring system, validated on six months of Vestas V82 SCADA data. Our XGBoost-based regression models were trained on engineered lag, rolling, and thermodynamic features to forecast gear oil, front bearing, and rear bearing temperatures approximately 20 minutes ahead, with results presented in Figures 2(a)–(c). Complementing the predictive models, a site-specific power curve was developed accounting for local elevation and air density, enabling benchmarking of actual power production against OEM standards. A high-fidelity digital twin modelled in Fusion 360 was further integrated with real-time SCADA streams, enabling dynamic visualisation of blade pitch, rotor RPM, and nacelle orientation — bridging predictive analytics with physical asset behaviour. Future work will focus on implementing LSTM-based time-series models and comparing their predictive accuracy and computational runtime against the current framework to assess suitability for real-time deployment.

Abstract ID: 3107

Development and Implementation of ROS 2 Architecture for Differential-Drive Indoor AMRs

N. Prabhu Ram*, Prof. G.L. Samuel

Indian Institute of Technology Madras, Department of Mechanical Engineering, Chennai, Tamil Nadu, PIN-600036

**prabhujeofficial@gmail.com*

Abstract

Autonomous Mobile Robots (AMRs) are essential for modern, flexible automation. However, making them navigate reliably in the real world—bridging the "**Sim-to-Real**" gap—remains a major challenge. This paper presents the complete development of a differential-drive AMR using **ROS 2 Jazzy** to address this exact problem. It covers the entire process from Gazebo simulation and **SLAM**-based mapping to physical hardware integration using a Raspberry **Pi 4** and **custom sensor circuitry**. Crucially, also integrated a priority-**multiplexed joystick** as a human-in-the-loop fail-safe, creating a robust, safe, and scalable foundation for dynamic **Industry 4.0** environments.

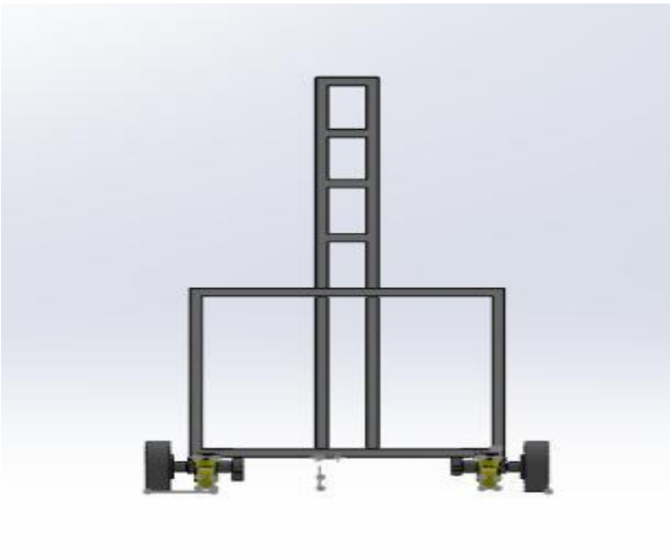

Figure 1. CAD design of the Robot.

This project focuses on the electrical integration and software development of an Autonomous Mobile Robot (AMR) utilizing a differential drive mechanism. The primary objective was to build a robust, scalable architecture for autonomous indoor navigation and state estimation. The hardware phase involved the complete physical implementation of the robot, including custom circuit design to interface the sensor suite (LiDAR and IMU) and motor drivers with an embedded Raspberry Pi 4 running an Ubuntu Server environment. To enable autonomous capabilities, the system was developed using the ROS 2 Jazzy middleware.

The complete navigation stack was implemented from the ground up, encompassing closed-loop wheel odometry, SLAM-based localization, and dynamic path planning. Before physical deployment, the robot's kinematics and navigation algorithms were rigorously modelled and validated within a Gazebo simulation environment, ensuring a reliable “Sim-to-Real” transfer. Furthermore, to guarantee human-in-the-loop safety during physical testing, a priority-multiplexed Bluetooth joystick was integrated as a deterministic, sub-millisecond teleoperation fallback. Ultimately, this project demonstrates a comprehensive, end-to-end integration of embedded electronics, sensor-to-actuator control, and modern autonomous navigation frameworks on a physical robotic platform.

Problem statement

As industries transition to smart automation, the demand for adaptable, open-architecture mobile robots is rapidly growing. However, deploying these autonomous systems in unpredictable physical environments often exposes a "Sim-to-Real" gap, leading to localisation drift and unreliable actuator responses. Furthermore, trusting purely autonomous planners without an immediate manual override poses severe safety risks during unmapped edge cases. This project addresses these challenges by engineering a custom differential-drive AMR with ROS 2 Jazzy and a Raspberry Pi 4 edge compute node. By combining Gazebo simulation with physically embedded circuitry and integrating a priority-multiplexed joystick as a fail-safe override, this work ensures a robust, human-in-the-loop architecture for real-world deployment.

Methodology

Began by defining the AMR through a complete virtual URDF model to establish baseline kinematics and test basic odometry control in a simulated environment. Next, integrated 2D LiDAR and IMU sensors to feed perception data into the SLAM Toolbox and AMCL particle filters for accurate mapping and localization. This foundation allowed us to deploy the core autonomous navigation stack, using multi-layer costmaps, A*/Dijkstra global planners, and Behavior Trees to govern robot logic. Finally, the project moved to real hardware on a Raspberry Pi 4 main brain, implementing a crucial twist_mux priority multiplexer to manage seamless, safe switching between full autonomy and a manual Bluetooth joystick fallback.

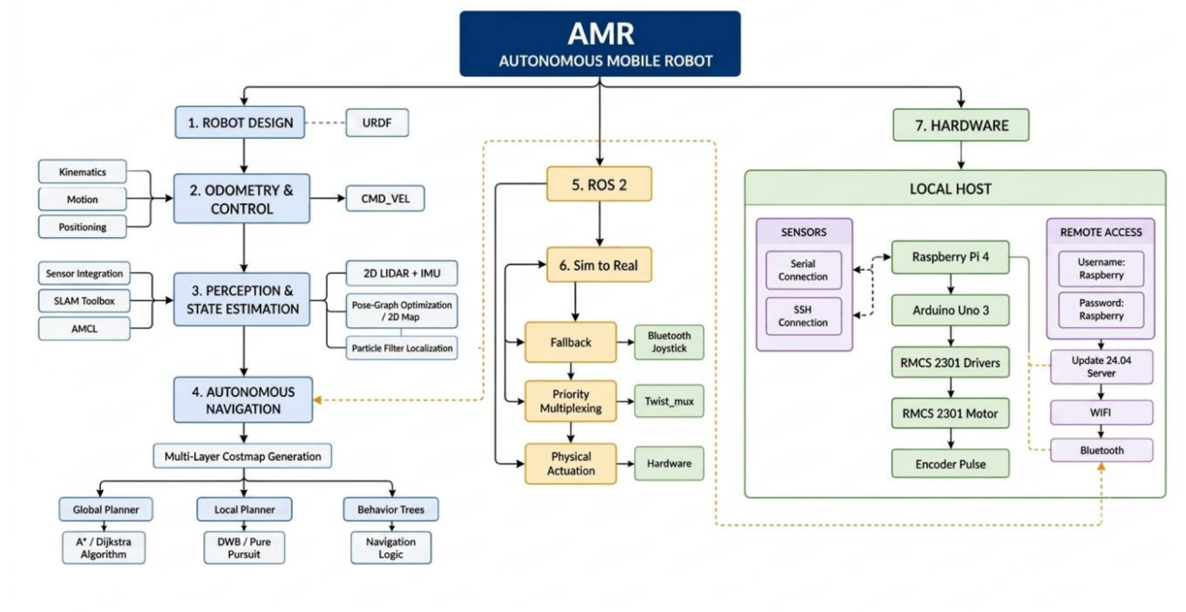

Figure 2. Methodology flow chart.

Abstract ID: 2674

Design and Fabrication of LPBF-Based H13 Helical Mandrel with Conformal Cooling for Enhanced Thermal Management in Automotive Brake Caliper Casting

Porrkodi Mohan[1]*, Balakrishnan Subi[1], Savyasachi nellikode[2], Murugaiyan Amirthalingam[2], Saravana Kumar[2], Vivek Singh[2]

[1] *Department of Chemistry, Central University of Tamil Nadu, Thiruvarur, Tamil Nadu, India*

[2] *Department of Metallurgical and Materials Engineering, Indian Institute of Technology Madras, Tamil Nadu, India.*

**porrkodimohan@gmail.com*

Abstract

Automotive brake calipers require high dimensional accuracy, structural integrity, and thermal stability to ensure reliable performance under operating conditions. Although conventional casting remains the preferred manufacturing method due to its cost-effectiveness and suitability for mass production, it often suffers from low material yield. A significant portion of molten metal is lost through gating systems, sprues, risers, and machining allowances, resulting in poor casting efficiency. Conventional riser design approaches, including empirical methods such as Caine's curve, primarily focus on optimizing feeding requirements to improve casting yield. Moreover, the complex geometry of brake calipers leads to non-uniform cooling during solidification, causing defects such as shrinkage cavities, porosity, residual stresses, and dimensional distortion. Consequently, thermal gradients remain a major cause of defects, limiting the quality and reliability of cast components. To overcome these challenges, conformal cooling has emerged as an effective thermal management strategy. Unlike conventional straight cooling channels, conformal cooling channels follow the geometry of the mould cavity, enabling more uniform heat extraction and controlled solidification. In this work, helical mandrels are proposed to create continuous conformal cooling pathways within the mould. The helical geometry increases the heat transfer surface area, promotes uniform temperature distribution, and improves the mechanical integrity and dimensional accuracy of the final casting. However, the intricate internal geometry of helical mandrels cannot be fabricated using conventional manufacturing techniques. Laser Powder Bed Fusion (LPBF), an advanced metal additive manufacturing process, offers a practical solution by enabling the fabrication of complex internal geometries with high precision. Through layer-by-layer selective melting of metal powder, LPBF facilitates the direct manufacture of helical conformal cooling channels, thereby enhancing thermal management while minimizing secondary machining operations. H13 tool steel is selected as the mandrel material because of its excellent hardness, wear resistance, thermal stability, and resistance to thermal fatigue, making it highly suitable for tooling applications subjected to repeated thermal cycling. Although rapid solidification during LPBF may introduce residual stresses and cracking, these challenges can be effectively mitigated through appropriate optimization of process parameters, enabling the production of dense, defect-free, and durable components. This work highlights the integration of LPBF-fabricated H13 helical mandrels with conformal cooling as a promising approach to improve casting yield, enhance thermal management, reduce casting defects, and achieve high-quality automotive brake caliper production.

Keywords: Caine’s curve, Conformal cooling, Mandrel, LPBF, H13 tool steel.

Abstract ID: 3079

Multimodal Tutor Kit for Blind Children

Jeevika Kannan, Lavanya M, Sandhya Chengaiyan, Dhanushree J, Dhatchayeni M*

Electronics and Communication Engineering, Meenakshi Sundararajan Engineering College, 363, Arcot Road, Kodambakkam, Chennai, India

[*]jeevikakannan164@gmail.com

Abstract

An array of lingering barriers is preventing visually impaired children from accessing education in the early years of their schooling, primarily due to the lack of interactive, accessible, and self-directed learning tools. A traditional education relies heavily on visual methods and has created barriers for blind students to fully interact and participate in the educational process and acquire important skills to learn independently. Although many children can develop their initial literacy skills using a tactile braille system, this system is limited because it is not interactive in real-time; therefore, no haptic or auditory feedback is received, and there is no opportunity to be actively engaged in the learning process. To address this gap, the Multimodal Tutor Kit (MTK) for Blind Children was developed as an alternative assistive educational technology that combines auditory instructions, tactile braille input, and natural language speech synthesis to help facilitate a high-engagement, inclusive learning experience for children in Primary Stage Education.

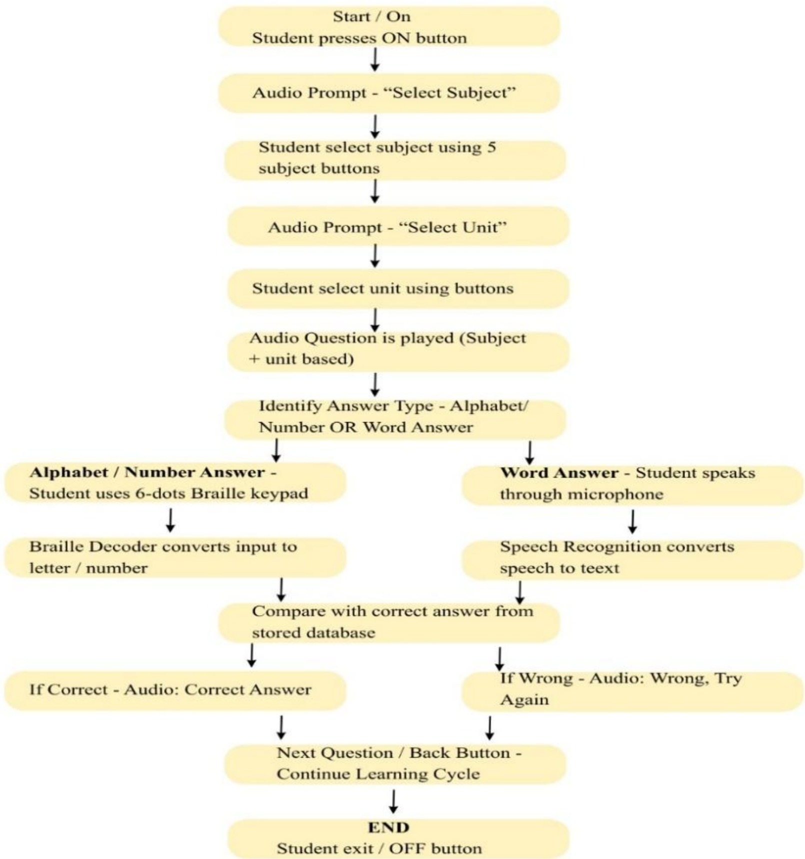

Figure 1. Algorithm for multimodal tutor kit for blind children.

Offering a standardized and accessible design for blind or visually impaired individuals through clear, automated audio instructions upon Turn-On of the system allows for an intuitive means of operating.

After subject selection has been made using distinct tactile inputs, there are four primary learning domains – Language, Maths, Science and Social Studies (as illustrated by Figure 1)- that students can work through. Once a subject area of study has been selected, an instructional unit is prompted for by way of indicated syllabus units to provide for organized and structured learning. All instructional materials and questions will be delivered as auditory output without any reliance on visual support. Users will provide input using a multiple-digit keypad ("Braille Keypad") when providing numeric objectives and/or using an integrated microphone when providing descriptive/word type answers to questions. Concurrent processing is performed on the user input relative to the preconfigured dataset of possible responses, providing instantaneous, real-time auditory feedback regarding user input accuracy when the user input accuracy when the user connected with the system via either the Braille Keypad or microphone input method. User will continue to interact with the system and to learn through the responsive design; accordingly, an integrated navigation control ("Back" button) enables users to move easily between various stages of operation.

The system will be built around an embedded architecture that uses a microcontroller. The integration of both hardware and software elements into one comprehensive system is the main objective. The input system will use tactile push buttons to navigate through the subject and unit menus. A Braille keypad will allow the user to enter text characters in six-dot Braille encoding. The output system will use an audio module with a built-in speaker to provide synthesized or pre-recorded instructions, questions, and real-time auditory feedback to the user. When the user provides an answer that requires a complex description, a speech recognition module will convert the spoken words to text characters. Input processing, algorithmic decision making, and output generation will be coordinated through a central processing unit that can be either an Arduino (or similar) or an ESP32 (or similar) or a combination of both. Persistence of data will be maintained in a structured database containing pre-defined subjects, units, questions, and answer keys. The evaluation engine will decode the tactile Braille input into a text character, compare the alphanumeric and/or verbal response against the database entries, and use pre-defined matching algorithms to validate that the response is both accurate and complete.

This new system will enable more people to have a chance to learn, since it will help reduce the barriers between learning with Braille and using modern technology as educational tools. For example, there are many different ways a student can interact with material through our device, such as tactile, auditory, and vocal. Using this type of interaction will encourage children who are blind to engage their cognitive abilities and to retain information; therefore, allowing all students to be able to actively participate in their learning. This method gives students the ability to learn on their own without having to depend on their teacher or caregivers, and also provides students with an organized curriculum that is specific to the grades that are being taught. In addition, the system design will have a great amount of capacity for future expansion through things like (not limited to) multilingual support, cloud-based storage of/notices of assignments, localized tracking of performance, as well as Artificial Intelligence (AI) driven Adaptive Learning Models. These systems are designed to be scalable so that they may be developed in ways that will help increase the use of technology to empower children with disabilities through education and create equal access to learning.

Abstract ID: 2557

Smart Prescription-Based Automated Medicine Dispenser

Anitha G*, Navami Krishna R, Pooja S, Ranjhana S, Samprithi S, Sanjana A

Department of Electronics and Communication Engineering, RMD Engineering College, RSM Nagar
**ga.ece@rmd.ac.in*

Abstract

This work tackles the issue of errors and inefficiencies in traditional medicine dispensing systems. In these systems, pharmacists manually read prescriptions and provide medicines. This process takes a lot of time and can lead to mistakes, especially when workloads are heavy. Current solutions, like barcode or QR-based systems, need digital input first. IoT-based dispensers only offer reminders without smart prescription analysis. Therefore, an automated system is needed to directly interpret physical prescriptions and ensure accurate and secure dispensing.

To address these limitations, we propose an automated medicine dispensing system that uses OCR, RFID, embedded systems, and IoT technologies. The system incorporates a camera module to capture the doctor's prescription. Python-based image processing techniques improve the image and extract key details like medicine name, dosage, and quantity using Optical Character Recognition (OCR).

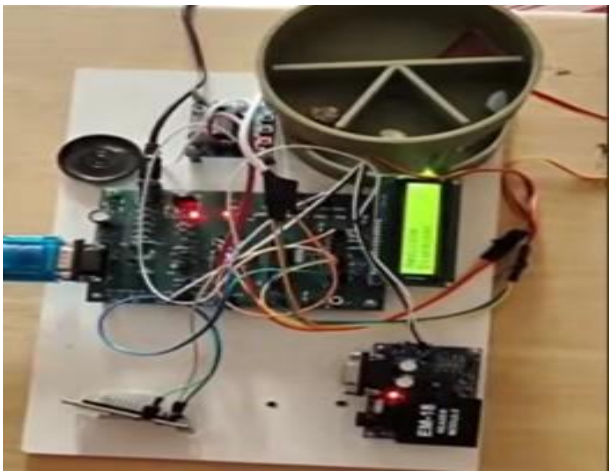

Figure 1. Medical Dispenser Prototype.

For security, we implement RFID-based authentication to ensure that only authorized users can access the system. The Raspberry Pi serves as the main controller, processing and validating the prescription data. Once verified, the data is sent to a PIC microcontroller, which operates a motor-driven dispensing mechanism to release the exact number of required tablets. Additionally, IoT integration allows for real-time monitoring of system activity, including medicine inventory and user transactions, through a cloud database. The system functions through stages like initialization, authentication, image capture, OCR processing, data transmission, and automated dispensing.

Overall, this proposed system reduces human errors, improves efficiency, ensures secure access, and enhances patient safety. It is well-suited for modern healthcare and smart pharmacy applications. Future improvements will include incorporating AI for more precise prescription identification as well as extending the application to cover additional pharmaceuticals and patients. Furthermore, mobile and cloud-based monitoring alongside security measures such as biometric authentication make efficient large-scale healthcare delivery possible.

Abstract ID: 3101

Extended Abstract for ICIDTSM 2026 Conference Agentic AI – Cement plant operations.

Velayutham Loganathan

TTANI.AI, Chowdry Nagar Main Road, Valasaravakkam, Chennai-600087. INDIA
ceo@cementplantengg.com

Agentic AI orchestration

Agentic AI models designed with neural network patterns are deployed in parallel and in series to govern the process of cement making. From the mining of Limestone up to the grinding of cement, Agentic AI supervises, plans, and executes the operations at each stage. The output parameters of one Agentic model are the input model for another Agentic model. AI Agents deployed in series/parallel work together in the right architecture in harmonious activity. Any variation in Raw materials composition activates the neural network function and adjusts the operational parameters such as the amount of additive materials, Temperature control, and pressure variations. The cement manufacturing process is a multi–input process and multi-output process, with operation parameters influencing others in a non-linear manner.

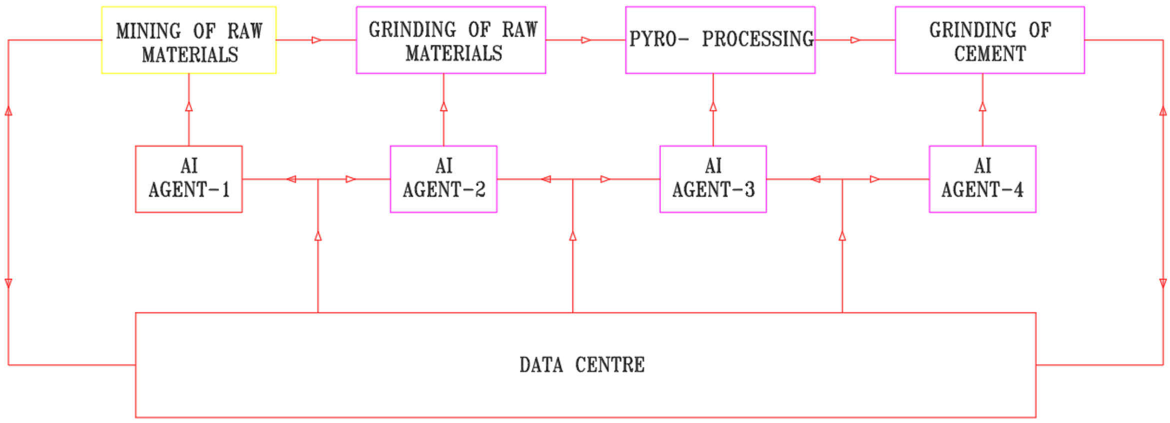

Figure 1 Agentic AI ORCHESTRATION

The derivation of the activation function of a neural network involves deep scientific research. The regular activation function of Foundational models is not aligning with the equations of chemical changes in the Raw material/limestone composition during the pyro-process. The process of cement making is diversified in nature and varies with the types of cement produced. Based on the type of cement, Raw additive materials also vary, and hence the neural network model varies. The model is to be trained in line with the chemical reaction of materials. Any deviation in the Raw materials composition is noticed at First Agentic AI, will subsequently alter the equations of the regression pattern in the downstream.

The Standard patterns of regression models are not in line with the equations of chemical changes in the limestone composition during pyro process.

The activation functions of neural network is not standard functions as in Foundational models. The Clinkerisation process of Raw materials involves physical and chemical changes. The influence of input parameters is plotted in non-linear regression model. The Agentic AI models are trained based on this activation functions. The activation functions are derived from the domain expertise and are based on data collection.

Agentic AI network, supervise, plan and execute the processes of Grinding of raw materials and Pyro processing.

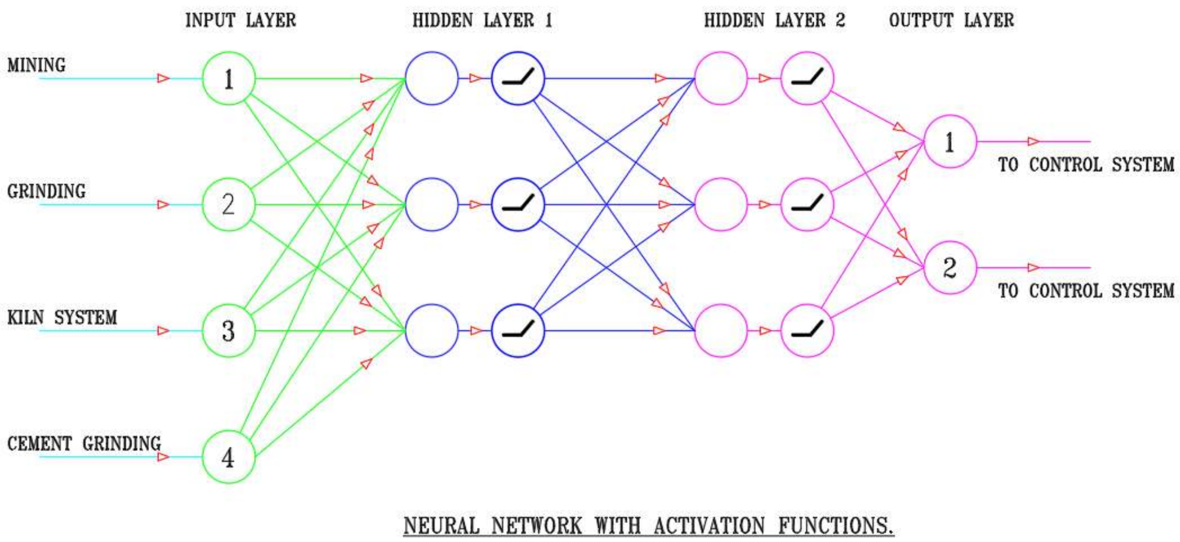

Figure 2. Neural network.

AI Brain network:

The Agentic AI orchestration of cement plant operation is coordinating the tasks in the same concept like how neurons act in the human brain. The left brain is composed of data collection, analytics, and acts as a data center for Agentic models. The Right brain in the upstream tracks the data and controls the operations of the plant. The first part of the neural design is part of the left brain, and the second part of the neural network is the right brain.

Brief Integration of Agentic AI with Thermo Chemical reactions in Plants

The mixture of Raw materials gets heated, de-moisturized, and de-carbonated. There is a sequence of phase changes. The agentic AI model controls the phase changes to get the desired strength of cement at the process completion. The required properties of cement having consistent quality and compressive strength are the outcome of the raw materials composition, Fineness of particles ground. The networking of Agentic AI controls, alters, and simulates the operational parameters from Mining, then Grinding, and Kiln systems.

Abstract ID: 2498

Multi-Domain Digital Twin Framework for Real-Time Monitoring and Object Identification Across Air, Sea, and Ground Domains

Ajith Kumar A[1]*, Muthukumar K A[2]

[1]*Shinkai Labs Private Limited, RMZ Millenia Business Park, Chennai, Tamil Nadu, India*

[2]*Department of Computer Science and Engineering, SSN College of Engineering, Chennai, Tamil Nadu, India*

**ajithkumar@shinkai-labs.com*

Abstract

Digital twin (DT) technology has emerged as a transformative paradigm in Industry 4.0, enabling synchronized virtual replicas of physical systems for real-time monitoring, predictive analytics, and intelligent decision support. While DT frameworks for single-domain assets are well established, a critical gap persists in jointly modeling, monitoring, and identifying objects across heterogeneous operational environments spanning aerial, maritime, and terrestrial domains simultaneously. This paper presents the Multi-Domain Combat Agent (MDCA) Digital Twin Framework a unified three-tier architecture comprising a Digital Twin Engine (DTE) for physics-driven synthetic data generation, a Multi-Domain Surveillance System (MDSS) for real-time cross-domain sensor fusion via Optimal Transport (OT), and an Advanced Tactical Simulator (ATS) for Live-Virtual-Constructive (LVC) validation with multi-agent reinforcement learning (MARL) driven object identification. Validated over 500+ simulated hours, the framework achieves 94.7% object classification accuracy, 47 ms end-to-end latency, and 6.1% false alarm rate, representing improvements of 13.4 pp, 3×, and 67%, respectively, over single-domain baselines.

Industry 4.0 and emerging Industry 5.0 paradigms have positioned digital twin technology as a foundational enabler for intelligent cyber-physical systems (CPS). A digital twin creates a continuously updated virtual replica of a physical system, integrating real-time sensor streams, physics-based models, and AI-driven analytics to enable monitoring, prognosis, and autonomous control. While significant progress has been achieved for single-asset DTs, the challenge of constructing cohesive multi-domain digital twins where heterogeneous systems across aerial, maritime, subsurface, and terrestrial environments must be jointly modeled, monitored, and identified remains largely unaddressed. Multi-domain environments introduce compounding challenges: (i) sensor modality heterogeneity across radar, EO/IR, sonar, and kinematic sensors; (ii) dynamic, contested topologies; and (iii) severe ground-truth data scarcity. The proposed MDCA framework, illustrated in Figure 1, directly addresses these gaps through a purpose-built three-tier digital twin architecture.

As shown in Figure 1, the MDCA framework comprises three hierarchical tiers. The **Digital Twin Engine (DTE)** constructs physics-based virtual replicas of air (fixed-wing, rotary-wing, UAV), naval (surface and subsurface), and ground (armored, wheeled) platforms, producing multi-modal synthetic sensor streams radar cross-section, EO/IR signatures, sonar returns, and 6-DoF kinematics at 100 Hz update rates and 12,000 labeled samples per minute via GAN-based augmentation, without reliance on classified sources. The **Multi-Domain Surveillance System (MDSS)** aggregates heterogeneous sensor streams via Earth Mover's Distance (EMD)-based Optimal Transport, producing a unified domain-agnostic situational picture with MARL-based object identification using Graph Attention Networks (GAT) and PPO adversarial self-play. The Advanced Tacti**cal Simulator (ATS)** integrates LVC exercises with 10,000+ scenarios per hour and a three-tier edge–fog–cloud inference pipeline compressed via INT8 quantization and structured pruning to achieve sub-12 ms platform-level latency.

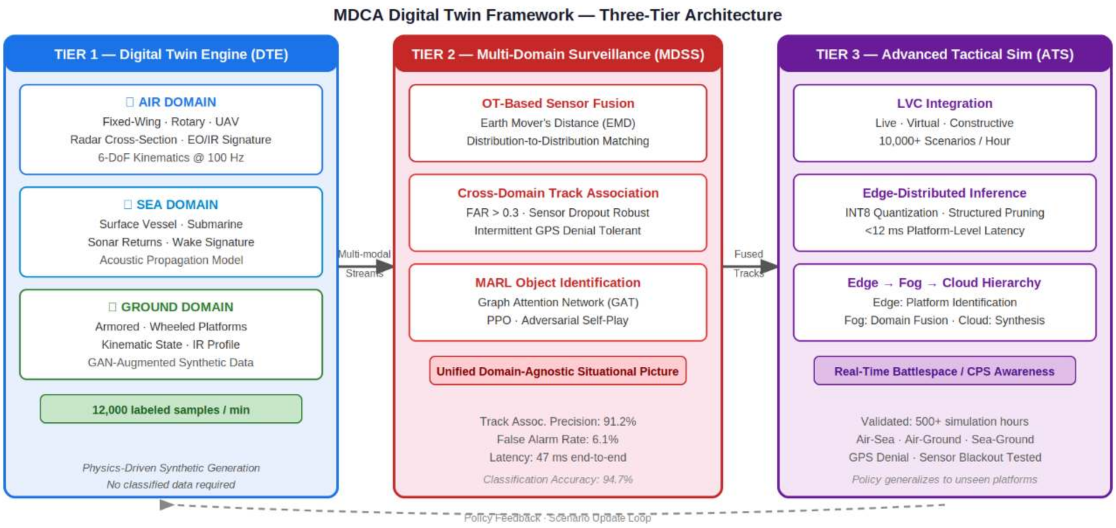

Figure 1. MDCA Three-Tier Digital Twin Architecture: DTE (physics-driven data generation) → MDSS (OTfusion + MARL identification) → ATS (LVC simulation + edge inference).

Synthetic Data Generation. Domain-specific physics models parameterized by radar cross-section databases, aerodynamic transfer functions, and acoustic propagation equations generate temporally coherent multi-modal data. A GAN-based augmentation module introduces realistic sensor noise, clutter, and environmental perturbations, eliminating reliance on classified data sources. **OT-Based Sensor Fusion.** EMD-based OT associates multi-source tracks by mapping state-space probability distributions, achieving robust cross-domain data association under clutter densities exceeding 0.3 and intermittent sensor dropouts, without requiring hard point-to-point correspondences. MARL Object Identi**fication**. Cooperative identification agents share partial observations via a Graph Attention Network and co-evolve robust classification policies through adversarial self-play using PPO, yielding policies that generalize to unseen platform configurations without retraining.

Results and Discussion

Simulation across 500+ hours of air–sea, air–ground, and sea–ground scenarios demonstrates: object classification accuracy 94.7% (vs. 81.3% single-domain baseline, +13.4 pp); cross-domain track association precision 91.2% (vs. 74.6%); end-to-end latency 47 ms (vs. 143 ms, 3× improvement); false alarm rate 6.1% (vs. 18.4%, 67% reduction). OT-based fusion maintained track continuity through simulated GPS-denial and sensor blackout events. MARL self-play policies generalized to previously unseen platform configurations without retraining, and the synthetic data pipeline sustained 12,000 labeled samples per minute throughout, confirming data pipeline scalability for large-scale CPS monitoring.

Conclusion

The MDCA Digital Twin Framework delivers a unified three-tier architecture for real-time multi-domain monitoring and object identification. Physics-driven synthetic data generation eliminates data scarcity, OT-based fusion ensures cross-domain track continuity under adverse sensing conditions, and MARL produces generalizable identification policies deployable within edge hardware constraints. Future work will pursue hardware-in-the-loop integration with physical sensor arrays and transfer learning for rapid adaptation to new platform classes.

Keywords: Digital Twin; Multi-Domain Monitoring; Optimal Transport; Sensor Fusion; Multi-Agent Reinforcement Learning; Synthetic Data Generation; Object Identification.

Abstract ID: 3152

Development of a VR-Based Digital Twin for Drone Threat Simulation and Counter-UAV Response Training

A Nageswara Rao[1]*, Agnishwar J[2], Rithika Mohan[1, 2], Karthikeyan[1, 2], Kavya Santhoshi B[1], Kummathi Chenna Reddy[1], Ramanna Madhusudhana[1], Sreenivasulu Putta[1], V· M· S· N·Pavan Kumar Ch[1], Nitalaksheswara Rao K[1], Thakur Singh[1], P S Londhe[1]

[1]AICTE Industry Fellow, Garuda Aerospace, Chennai, Tamil Nadu, India

[2]Garuda Aerospace Chennai, Tamil Nadu, India

**nagesh803@gmail.com*

Abstract

The widespread use of unmanned aerial vehicles (UAVs)/Drones, especially inexpensive nano and micro drones, has created significant threats to the security of military facilities and infrastructure. UAVs are now used for scouting, reconnaissance, and attack missions, posing a threat to security. The Russian invasion of Ukraine has shown the effectiveness of drone warfare, which includes swarm drone attacks and fast-moving first-person-view (FPV) strike drones. It's crucial to train personnel in identifying and countering UAV threats. But conventional training approaches face challenges in terms of cost, safety, and the inability to effectively practice in realistic threat environments. Virtual Reality (VR) technologies hold the promise of immersive and affordable training. Implementing Digital Twin concepts enables real-time virtual models of physical systems and processes for monitoring, simulation, and decision-making. In this study, the concepts of VR and Digital Twin are combined to create a simulation framework for Counter-UAV (C-UAV) training through scenarios. Our approach uses a Digital Twin model of a military base, along with simulations of drone breaches and mitigation strategies. The aim is to create a scalable and adaptive system for training and testing C-UAV strategies.

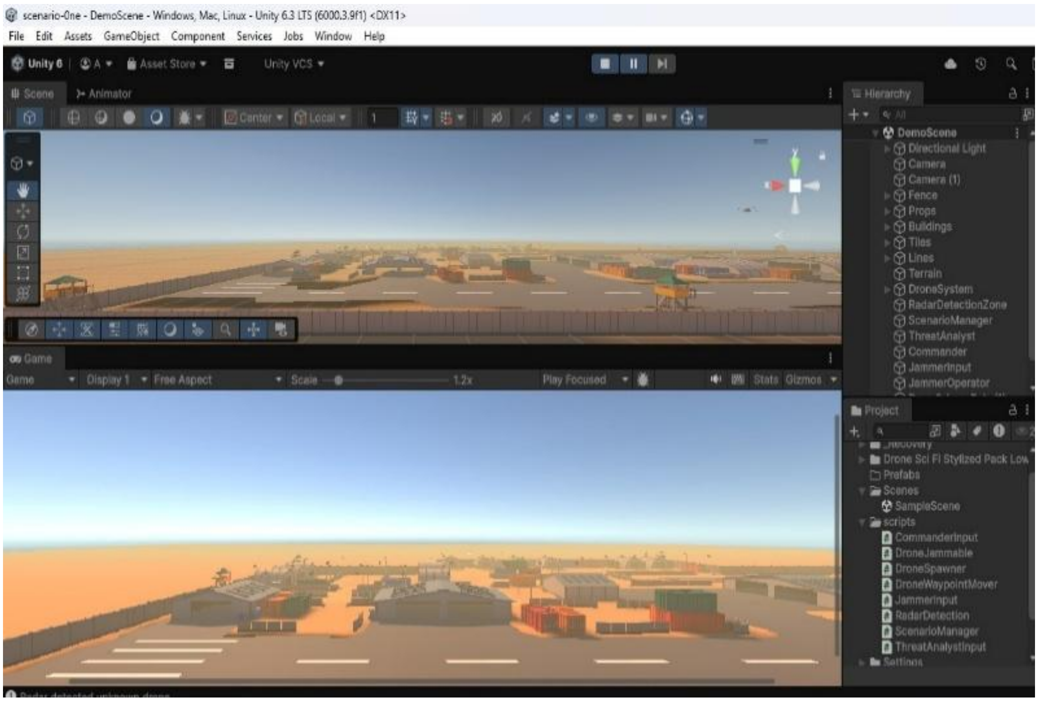

Figure 1. Scenario for Army basin enter a UAB.

Digital Twin Approach to C-UAV Simulation

Digital Twin systems involve the interplay of a real system and a virtual model with data exchange and decision-making functions. In this paper, the Digital Twin system comprises three main layers:

Physical Layer: Includes physical assets (military base) and threats (UAV). While presently it doesn't integrate with sensors in real-time, it is capable of integration with radar and RF sensing systems in the future.

Digital Twin Layer: The virtual environment is implemented in Unity Engine, and the military base and the UAV are simulated. The Digital Twin models the layout, UAV behavior, and responses.

Decision and Control Layer: It uses the Digital Twin for information and selects the countermeasures. It has rules for threat recognition and mitigation, allowing for operational decision-making.

This interaction allows a closed-loop simulation system, where responses and performance metrics are generated based on events in the virtual environment. This approach is consistent with cyber-physical system models prevalent in smart manufacturing and intelligent customer response management.

Scenario Design and Implementation

A key scenario for drone infiltration is defined in the proposed framework. A UAV intrudes into a military area, and a radar system is used for detection, based on proximity detection methodology. The UAV is identified as an unrecognised or hostile target. The decision module assesses the threat and allows RF jamming. Activating it will cause the drone to lose contact, leading to the display of hovering and landing behaviours, and eventually "bring it down". The system operates in an event-driven manner, including detection, classification, decision, activation of countermeasures, and assessment of outcome. The system uses Unity and C# programming. Modules include drone navigation, detection, decision-making, countermeasure reaction, and event history. The system is designed to be easily extendable to new scenarios and features. To further test the effectiveness of the simulation, a scenario of a faster-moving drone is included. This involves the UAV travelling faster, creating more time pressure and complexity for the system. This scenario showcases the flexibility of the Digital Twin model in various threat scenarios.

Results and Discussion

The system produces the event logs that record the simulation. These logs record the detection of a drone, threat assessment, activation of countermeasures, jammer activation, loss of communication and complete neutralization. The event-driven nature of the simulation allows for its traceability and understanding of the system's behaviour.

Key observations include:

- The system is able to simulate the entire counter-UAV workflow.
- The event-based approach guarantees the proper order of actions.
- The fast drone scenario increases the challenge, showcasing the system's capacity to simulate different threat scenarios.

In terms of Digital Twin, the system exhibits features including virtualisation, simulation, and monitoring. While not currently synchronised with the physical system in real time, the system lays the groundwork for such connections.

The system also showcases the versatility of Digital Twins beyond their typical use in manufacturing. By simulating defence scenarios, the system demonstrates the potential of Digital Twins for other real-time simulation and decision-making applications, including smart cities and critical infrastructure.

Conclusion and Future Work

In this paper, we reported on the design of a VR Digital Twin tool for drone threats simulation and anti-drone training. The framework combines threat detection, decision-making and threat countermeasure simulation in a virtual environment. This project shows the potential of Digital Twin systems for dynamic threat and training simulations. Efforts will be made to build upon the system by integrating real-time data from physical monitoring sensors, threat identification using AI, and simulation of drone swarms. Furthermore, the incorporation of immersive VR equipment and collaboration among users will enhance user training. The framework offers a flexible and modular approach to Digital Twin simulation systems, which will help to develop smart and resilient systems

ICIDTSM
2026

SPONSORS

JULY 20–22, 2026 · ICSR HALL, IIT MADRAS

Department of Mechanical Engineering, Indian Institute of Technology Madras
Chennai, Tamil Nadu – 600036

icidtsm2026@gmail.com | ge.iitm.ac.in/icidtsm-2026

Scan to visit website