KUWAIT COLLEGE OF SCIENCE & TECHNOLOGY
كلية الكويت للعلوم والتكنولوجيا

IEEE ICAIGE-S4IoT 2026
Final Conference Program & Book of Abstracts

October 13–15, 2026 | Kuwait College of Science & Technology (KCST) | Hybrid Format

PROGRAM SUMMARY & TRACK LEGEND

[IP] = In-Person Presentation | [OL] = Online Presentation

Track / Special Session	Accepted Papers
• Artificial Intelligence for Engineering Applications	20
• Renewable Energy, Distributed Generation, and Smart Grids	15
• IoT-Enabled Green Energy Systems and Smart Infrastructure	10
• Special Session 1: AI-Driven Smart Materials and Sensing Systems for Advanced Energy Technologies	8
• Special Session 2: Real-Time Embedded Artificial Intelligence Systems	8
• Artificial Intelligence for Automation of Energy Systems	5
• Power Converters and Energy Storage Systems	5
• Cloud-Edge Computing, IoT, and Renewable Energy	5

Visit our website

Website: icaige.com / s4iot.kcst.edu.kw **Email:** contact@icaige.com / S4iot@kcst.edu.kw **Mobile/WhatsApp:** +21623882186 / +9659973097

Tuesday, October 13, 2026

09:00 - 09:15	**Opening Address** *Hall C*
09:15 - 10:00	**Keynote Speaker I** **“AI-Driven Energy Transition: Intelligent, Flexible, and Resilient Integrated Energy Systems”** *Mousa Marzband, Professor at King Abdulaziz University, Saudi Arabia* *Hall C*
10:00 - 11:00	**Tutorial 1 (Part I): Artificial Intelligence for Energy Internet** **Intelligent Power Routing in Packetized Energy Networks** *Organized by: Hadi Y. Kanaan (Lebanon), Amani Fawaz (Lebanon)* *Hall C*
11:00 - 11:15	**Coffee & Networking Break**
11:15 - 12:30	**Tutorial 1 (Part II): Artificial Intelligence for Energy Internet** **Intelligent Power Routing in Packetized Energy Networks** *Organized by: Hadi Y. Kanaan (Lebanon), Amani Fawaz (Lebanon)* *Hall C*
12:30 - 13:15	**Lunch Break**

13:15 - 14:30 Hall C	**Artificial Intelligence for Engineering Applications (Part I)** *Session Chair: TBA*
ID022 [IP]	■ **A Danger-Aware Heuristic Pathfinding Algorithm for Safe Autonomous Navigation in 2D Engineering Environments.** — *Hajar Alsulaili, Muneerah Alshammeri, and Manar Almutairi*
ID032 [IP]	■ **Artificial Intelligence Enabled Curriculum Transformation Through an Intelligent Ecosystem for Future Workforce Readiness.** — *Ziad Salem, Hessa Alanjari, Frishtay Fatima, and Abdullatif Baba*
ID099 [IP]	■ **AI-Driven Crime Hotspot Prediction and Energy-Aware Patrol Allocation for Sustainable Smart Cities.** — *Mohamad Hoda*
ID086 [IP]	■ **Integrated Load-Flow and CCT-Based Transient-Stability Assessment of Photovoltaic Penetration in the IEEE 14-Bus System Under Variable Loading and Fault Conditions.** — *Mustafa Alhariri, Rashad M. Kamel, and Sultan S. Alanzi*
ID017 [OL]	■ **Uncertainty-Aware Physics-Informed Neural Networks for Degree of Polymerization Estimation in Oil-Immersed Power Transformers.** — *Alamera Nouran Alquennah, Mohammad Alshaikh Saleh, Ali Ghrayeb, Haitham Abu-Rub, Shady Khalil, Mohammed Al-Hajri, and Sunil Khatri*

13:15 - 14:30 TBD	**Renewable Energy, Distributed Generation, and Smart Grids (Part I)** *Session Chair: TBA*
ID087 [OL]	■ **Toward Smart Livestock Infrastructure: Sensor Benchmarking for Methane Monitoring in Green Dairy Farm Systems.** — *Usama Mokashi, Wesam Rohouma, Shady Khalil, Rachid Benlamri, and Ajna Shukkoor*
ID063 [OL]	■ **Droop-Based Grid-Forming Inverter Control for a PV/BESS Microgrid Integrated with the IEEE 13-Bus Distribution System.** — *Muhammad Salim Khan, Amira Mohammed, Shady Khalil, Mostafa Farrag, and Mohamed Massaoudi*
ID096 [OL]	■ **Green Hydrogen Production from Harvested Rainwater: Evidence, Challenges, and Conceptual System Integration.** — *Mrinmoy Das, Ramesh Kumar, Gianluca Gatto, Suman Jana, and Amit Kumar*
ID114 [OL]	■ **Economic Evaluation of Utility-Scale Solar Photovoltaics, Onshore Wind, and Supercritical Coal Power Plants Using Annual Worth Analysis.** — *Aisyah Nadya Faradina, Shabriya Ridzqi Rahman, Filscha Nurprihatin, Hendy Tannady, Ali Vaezi, and Michael Nayat Young*
ID122 [OL]	■ **A Supervisory Zonal Multi-Agent Volt/VAR Controller for Active Distribution Networks.** — *Muhammad Usman*

13:15 - 14:30 TBD	**Power Converters and Energy Storage Systems** *Session Chair: TBA*
ID071 [IP]	■ **A Finite-Set Model Predictive Control Approach for Enhanced Energy Sharing in Fuel Cell / Supercapacitor Fed OEWIM Based Electric Drive.** — *Khaled Safsouf, Jean Sawma, and Hadi Kanaan*
ID024 [OL]	■ **Air-Lead Control of PEM Fuel Cells for EV Powertrains using Converter-Based OER Regulation and DC-Bus Stabilization.** — *Iyad Ameur, Mohamed Bougrine, Atallah Benalia, Abdelhamid Rabhi, Mohamed Trabelsi, and Mohammed Benmiloud*
ID093 [OL]	■ **FPGA-in-the-Loop Validation of P&O and Fuzzy Logic MPPT Controllers for Photovoltaic Systems.** — *Brahim Attia, Olfa Bel Hadj Brahim Kechiche, Hazem Dridi, Mahmoud Hamouda, and Xibo Yuan*
ID102 [OL]	■ **Thermal State Estimation-Assisted Self-Healing Interleaved SEPIC Converter for Reliable PV Battery Charging.** — *Manoharan S, Mahalakshmi B, Gianluca Gatto, and Amit Kumar*
ID048 [OL]	■ **NLADRC Control for Five Phase PMSM-Based Naval Propulsion Under Disturbances and Fault Conditions.** — *Khouloud Ben Hammouda, Mohamed Trabelsi, Trabelsi Trabelsi, and Riadh Abdelati*

14:30 - 14:45	**Coffee Break**

KUWAIT COLLEGE OF SCIENCE & TECHNOLOGY
كلية الكويت للعلوم والتكنولوجيا

14:45 - 16:00
Hall C

Artificial Intelligence for Engineering Applications (Part II)

Session Chair: TBA

- **ID116** [IP] ■ **Quantum Machine Learning for Efficient and Sustainable Robotics: A Review.** — *Ali Alfayly, Mahmood Seyyedzadeh, Amir Ghavifekr, Hooman Esmaielzadeh, and Amir Ghiasi*
- **ID117** [IP] ■ **An Attention-Based Neural Network Method for Acoustic Anomaly Detection in Pipelines.** — *Muhammad Zia Ur Rehman, Omar Khattab, Zurasha Zaryab, Ismail Kashif, Ahsan Ali Raza, and Murad Khan*
- **ID056** [OL] ■ **Vision-Based Cattle Weight Estimation: A Systematic Review of Computer Vision Models and Image Acquisition Methods.** — *Marcus Joshua Fortes, Ajna Shukkoor, Usama Mokashi, Shady Khalil, and Wesam Rohouma*
- **ID066** [OL] ■ **A Unified Vision-Based System for Cattle Weight Estimation and Behavior Detection.** — *Ajna Shukkoor, Wesam Rohouma, Shady Khalil, Rachid Benlamri, Amit Khandakar, and Usama Mokashi*
- **ID085** [OL] ■ **Comparative Analysis of Traditional Machine Learning and Graph Neural Network Models for Anti-Money Laundering Detection.** — *Aman Reddy, Areesha Hamad, Glen Anton Paul, Manikandan Pradeep, Mohammad Hasheer, Umar Syed, Syama Kurungodathil, and Asmath Famnaz*

14:45 - 16:00
TBD

Renewable Energy, Distributed Generation, and Smart Grids (Part II)

Session Chair: TBA

- **ID010** [OL] ■ **Impact of Bulk and Interface Defects on 2T Perovskite/ACIGS and 2J/3J All-Perovskite TSCs.** — *Nour El I Boukortt, Antonio Garcia Loureiro, Salvatore Patanè, Mohamed Trabelsi, Trupti Ranjan Lenka, and Ahmad Abushattal*
- **ID034** [OL] ■ **Validation of Acoustic Propagation Models for Wind Farm Noise Assessment: An End-to-End Reproducible Workflow Using Python and Power BI.** — *Emna Annabi, Mohamed Wassim Jmal, and Ali Wali*
- **ID043** [OL] ■ **Experimental Performance Comparison of Different Monofacial and Bifacial PV Modules in Outdoor Conditions.** — *Daniela Olivares Baqué, Alejandro González Moreno, Alberto Dolara, and Sonia Leva*
- **ID083** [OL] ■ **Innovative Economic Valorization Strategies for Industrial Photovoltaic Systems: Revamping, Repowering and Storage.** — *Giacomo Manzan, Sonia Leva, and Marco Mussetta*
- **ID111** [OL] ■ **Blockchain-Enabled Stackelberg Pricing for RealTime Decentralized Microgrid Energy Management:A Hardware-Validated Framework.** — *Abdullah Umar, Vijayakumar Gali, Prashant Jamwal, Shady Khalil, Mostafa Farrag, Ajay Kumar, and Nitin Gupta*

14:45 - 16:00
TBD

Cloud-Edge Computing, IoT, and Renewable Energy

Session Chair: TBA

- **ID107** [IP] ■ **A Behaviour-Level Zero-Trust Digital Twin for Secure ATM and Cash-Kiosk Maintenance.** — *Jumanah Alfailakawi, Anfal Almutairi, Aalyah Alotaibi, Khaled Allanqawi, and Basil Alothman*
- **ID121** [OL] ■ **Planning IoT-Enabled Road Management in Kuwait: Assessment, Technology Mapping, and a Pilot-Deployment Priority Index.** — *Haya Albader and Naser Alselahi*
- **ID127** [OL] ■ **Analysis of Expenditure and User Dissatisfaction Optimization in Residential Energy Management Systems Using Social Network Optimization.** — *Andrea Montuoro, Alessandro Niccolai, and Francesco Grimaccia*
- **ID064** [OL] ■ **Generalizable solar irradiance prediction from satellite data using spatial–temporal graph attention network.** — *Konin Pierre-Claver Kakou, Georg Frey, Daud Mustafa Minhas, and Jörg Trentmann*
- **ID069** [OL] ■ **Threshold-Based Sensitivity Assessment of Industrial PV-BESS Systems from Energy Balances to Economic Viability.** — *Nicoletta Matera, Riccardo Ramaschi, Emanuele Ogliari, Sonia Leva, Francesco Torchiaro, and Enrico Lamanna*

19:00 - 21:00 **Gala Dinner (Venue: TBD)**

Wednesday, October 14, 2026

09:00 - 09:45	**Keynote Speaker II** **“Micro-Grid Implementation: Challenges and Solutions”** *Hatem Zeineldin, Professor and Chair of the Department of Electrical Engineering at Khalifa University, UAE* *Hall C*

09:45 - 11:00 Hall C	**IoT-Enabled Green Energy Systems and Smart Infrastructure (Part I)** *Session Chair: TBA*
ID026 [IP]	■ **Adaptive and Context-Aware Zero Trust Security for Scalable IoT Ecosystems.** — *Omar Dib and Belal Ali*
ID062 [IP]	■ **Resilience Analysis of V2X Communication Protocols for IoT-Enabled Smart City Infrastructure: A Kuwait Environmental Case Study.** — *Othman Aldaher and Yousef Alsaqabi*
ID074 [IP]	■ **MCP-Driven Conversational AI Agent for Real-Time IIoT Energy Budget Control in Industrial Environments.** — *Omar Dib, Tarek Adi, Abdullah Shamout, Ahmed Amin, and Sulaiman Alrefai*
ID090 [IP]	■ **AgroSense: An Automated Smart Hydroponic System with Edge AI Diagnostics for Arid Regions.** — *Othman Aldaher, Muhammad Sohaib, and Omar Hammuda*
ID097 [IP]	■ **A Cloud–Edge PM2.5 Forecasting.** — *Huda Alrashidi and Ghita Berrada*

09:45 - 11:00 TBD	**Special Session 1: AI-Driven Smart Materials and Sensing Systems for Advanced Energy Technologies (Part I)** *Organizers: Inés Abdellaziz, Férid Chaffar Akkari, Naoufal Khemiri, Nabila Elbitri (Tunisia)*
ID045 [IP]	■ **Smart NanomaterialEnhanced ZincRich Coatings for Sustainable Corrosion Resistance in Energy Applications.** — *Abeer Alfarhan and Abdulmuhsen Akbar*
ID038 [OL]	■ **A Hybrid CNN–XGBoost for EEG-Based Alzheimer's Monitoring and Energy Prediction in next generation intelligent biomedical implants.** — *Yosra Ben Fadhel, Kamal Al-Haddad, Olfa Bel Hadj Brahim Kechiche, and Mahmoud Hamouda*
ID065 [OL]	■ **Physical properties of CuMgSnS Powders and Thin Films.** — *Oumaima Hammami, Mounir Kanzari, and Naoufel Khemiri*
ID081 [OL]	■ **Effect of a Multi-stage Cooling System on the Performance of a Mini Photovoltaic Panel.** — *Marwa Hamza, Olfa Bel Hadj Brahim Kechiche, and Habib Sammouda*

09:45 - 11:00 TBD	**Special Session 2: Real-Time Embedded Artificial Intelligence Systems (Part I)** *Organizers: Imen Werda, Amina Kessentini, Amna Maraoui (Tunisia)*
ID052 [IP]	■ **LLM-Powered In-Vehicle Assistants: IoT Architecture and Case Study.** — *Mai Alotaibi, Sameerah Alajmi, Aliah Alsulail, Retaj Alsaeedi, and Basem Elhalawany*
ID070 [IP]	■ **Real-Time Selective Harmonic Mitigation in Multilevel Inverters Based on Gradient-Based Search Algorithm and Artificial Neural Network.** — *Noman Liaqat, Amir Ghasemian, Sobhan Mohamadian, Maria Gabriella Cimoroni, Concettina Buccella, and Carlo Cecati*
ID061 [OL]	■ **Hardware acceleration in FPGA-SoC of a lightweight neural network for partition optimization in the ECM encoder.** — *Sourour Soltani, Amna Maraoui, and Imen Werda*
ID113 [IP]	■ **An IoT-Enabled and Machine Learning-Supported Framework for Irrigation Decision Support.** — *Ghaliah Alhajraf and Ismaeel Al Ridhawi*

11:00 - 11:15	**Coffee Break**

11:15 - 12:30 Hall C	**Artificial Intelligence for Engineering Applications (Part III)** *Session Chair: TBA*
ID053 [IP]	■ **Large Language Models for Low-Resource and Agglutinative Languages: A Systematic Review.** — *Adil Akhmetov, Mohammed Alaa Ala'Anzy, Malek Ali, and Adnan Saeed*
ID095 [IP]	■ **AI-Driven Intrusion Detection Framework for Industrial IoT in Smart Grid Environments.** — *Mohamed H. Saad, Basem Elhalawany, and Ramy Samy*
ID029 [OL]	■ **Multi-Sensing Deep Learning Framework for PEMFC Degradation Forecasting Under Heterogeneous Operating Conditions.** — *Ahmad Braydi, Elena Breaz, Issam Salhi, Fei Gao, and Qian Xun*
ID047 [OL]	■ **Mallat Scattering Transform -Based Multiscale 3 branch CNN–BiLSTM Framework for Accurate Solar Irradiance Forecasting.** — *Mariem Mallat, Olfa Bel Hadj Brahim Kechiche, Mahmoud Hamouda, Houcine Oudira, and Aissa Chouder*
ID055 [OL]	■ **Deep Learning-Based Cattle Behavior Classification from Visual Data: A Comparative Review.** — *Josh Allen San Diego, Ajna Shukkoor, Usama Mokashi, Shady Khalil, and Wesam Rohouma*

11:15 - 12:30 TBD	**Special Session 1: AI-Driven Smart Materials and Sensing Systems for Advanced Energy Technologies (Part II)** *Organizers: Inés Abdellaziz, Férid Chaffar Akkari, Naoufal Khemiri, Nabila Elbitri (Tunisia)*
ID105 [OL]	■ **Compositional Engineering of CuO-SnO2 Thin Films: Influence on Structural, Optical, Morphological and Electrical Properties.** — *Nourhéne Mansouri and Bitri Nabila*
ID106 [OL]	■ **Oblique angle of deposition enhanced the photo-response performance of CuSbS2 thin films.** — *Mouna Idoudi, Ferid Chaffar Akkari, and Mounir Kanzari*
ID119 [OL]	■ **Spray-Coated CuO/rGO Hybrid Thin Films for Room-Temperature H2S Gas Sensing.** — *Imen Hemmedi and Bitri Nabila*
ID120 [OL]	■ **Deposition Time-Controlled Structural, Optical, Electrical, and Carrier Transport of RF Magnetron Sputtered ZnO-SnO Thin Films.** — *Emna Ben Sassi and Fatma Chaabouni*

11:15 - 12:30 TBD	**Special Session 2: Real-Time Embedded Artificial Intelligence Systems (Part II)** *Organizers: Imen Werda, Amina Kessentini, Amna Maraoui (Tunisia)*
ID035 [OL]	■ **A Smart Solar Panel Cleaning System Using Machine Learning and Embedded Systems.** — *Milia Habib, Maria Tedros, Saraa El Hazzouri, Zaher Merhi, Abdelrazzak Merheb, and Mohamad Kanaan*
ID076 [OL]	■ **A Lightweight Embedded AI-Based UAV System for Real-Time Diseased Tree Detection.** — *Bouthaina Abdallah, Fatma Belghith, and Nouri Masmoudi*
ID078 [OL]	■ **Toward Edge-AI Precision Livestock Farming: A Low-Power Wearable LoRa Sensor Node for Heat Stress and Activity Monitoring.** — *Usama Mokashi, Wesam Rohouma, Rachid Benlamri, Shady Khalil, Amith Khandakar, Arselene Ayari, Mujtaba Fida Ul Haq, Reza Alam, and Ajna Shukkoor*
ID098 [OL]	■ **Real-Time Fall and Crash Detection for Electric Micromobility Using IMU Sensor Fusion and Edge-Deployed Machine Learning.** — *Asel Rohouma, Mohamed Rohouma, Ajna Shukkoor, and Shady Khalil*

12:30 - 13:15	**Lunch Break**

13:15 - 14:30 Hall C	**Renewable Energy, Distributed Generation, and Smart Grids (Part III)** *Session Chair: TBA*
ID123 [OL]	■ **Assessing EV Penetration Impacts on Urban Air Quality: A Kigali Case Study.** — *Samuel Chukwuma, Kevi Jonathan Iradukunda, Temidayo Akanmu, and Jesse Thornburg*
ID124 [OL]	■ **Multi-Parametric Performance Evaluation of Six- and Twelve-Sector Direct Power Control for Grid-Connected Converters under Partial Load Conditions.** — *Mohamed Amine Fnaiech, Salwa Baserrah, Ebrahim Abdulrahman Hasan, Khaled Zehar, Mohamed Trabelsi, and Hani Vahedi*
ID134 [OL]	■ **Single-Switch MPPT Control of a PMSG Wind Turbine for Standalone and DC-Grid-Connected Residential Buses.** — *Maged Bauomy, Haytham Gamal, Adel Shaltout, Shady Khalil, and Morcos Metry*
ID075 [IP]	■ **Governor Valve Actuator Systems for Smart Grids: A Review of Control, Fault Diagnosis, and Predictive Maintenance.** — *Alamera Nouran Alquennah, Mohamed Trabelsi, and Naji Alqallaf*
ID023 [OL]	■ **From LLM-First to Governance-Bound Hybrid Execution: An Experimental Study of Reliability and Safety in Energy API Automation.** — *Swapneswar Sundar Ray*

13:15 - 14:30 TBD	**Artificial Intelligence for Automation of Energy Systems** *Session Chair: TBA*
ID036 [OL]	■ **Predictive Climate-Aware Intelligent Energy Management System for Fuel Cell Hybrid Electric Vehicles Under Kuwait Extreme Summer Conditions.** — *Ghaneemah Alasfour, Wael Farag, and Zakariya Dalala*
ID054 [OL]	■ **AI-Driven Identification of Compound Climate Hazard Pathways: Enhancing Energy Efficiency and Operational Resilience in Water Infrastructure.** — *Renalda El-Samra*
ID072 [OL]	■ **A Lightweigh Reproducible Multimodal Surrogate Learning Framework for EL-Free Photovoltaic Degradation Estimation from HDR Images.** — *Marcello Polenghi, Roberto Caldelli, Emanuele Ogliari, and Valsamis Ntouskos*
ID132 [OL]	■ **Comparing Large Language Models and Machine-Learning Methods for Short-Term Electricity Price Forecasting.** — *Seyed Ali Hosseini, Paria Valizadeh, Sven Tomforde, and Francesco Grimaccia*
ID125 [OL]	■ **Effect of Post-Training Quantization on Residential Electrical Load Demand Anomaly Detection.** — *Giuseppe La Tona, Massimiliano Luna, and Maria Di Piazza*

14:30 - 15:00	**Short Break before Tutorial 2**

15:00 - 16:00	**Tutorial 2 (Part I): Machine Learning Techniques for Energy Output Prediction** **Online Interactive Session** *Organized by: Lotfi Snoussi (Tunisia), Olfa Fakhfakh (Tunisia)* *Online*

16:00 - 16:15	**Short Break**

16:15 - 17:15	**Tutorial 2 (Part II): Machine Learning Techniques for Energy Output Prediction** **Online Interactive Session** *Organized by: Lotfi Snoussi (Tunisia), Olfa Fakhfakh (Tunisia)* *Online*

S4IoT'26 | KUWAIT COLLEGE OF SCIENCE & TECHNOLOGY كلية الكويت للعلوم والتكنولوجيا | ICAIGE

Thursday, October 15, 2026

09:00 - 09:45 **Keynote Speaker III**
"Electric Vehicle to Vehicle (V2V) Charging"
Vinod Khadkikar, Professor at Khalifa University, Abu Dhabi, UAE
Hall C

09:45 - 10:30 **Panel Session**
Hall C

10:30 - 10:45 **Coffee Break**

10:45 - 12:00 — IoT-Enabled Green Energy Systems and Smart Infrastructure (Part II)

Hall C
Session Chair: TBA

- **ID110** [IP] **Ergodic Capacity Analysis of Coordinated NOMA Transmission for LEO Satellite IoT Networks.** — *Anas Elsayed, Mohamed Elsayed, Ahmed Samir, Esraa M. Eid, Hala M. Abd Elkader, and Basem M. Elhalawany*
- **ID112** [IP] **Compute-Aware Adaptive Intrusion Detection for ATM Edge Networks.** — *Anfal Almutairi, Aalyah Alotaibi, Jumanah Alfailakawi, Khaled Allanqawi, and Basil Al Othman*
- **ID130** [IP] **Explainable Short-Term Photovoltaic Power Forecasting for PAIoT Energy Management.** — *Sara Salman and Zied Bouida*
- **ID030** [OL] **Energy-Aware Routing in Software-Defined Optical Networks via Reinforcement Learning.** — *Skander Abid, Mourad Menif, and Latifa Guesmi*
- **ID131** [IP] **IoT-Enabled AI Fusion for Early-Onset Alzheimer's Disease: A Survey and Preliminary Analysis.** — *Joicy John and Zied Bouida*

10:45 - 12:00 — Artificial Intelligence for Engineering Applications (Part IV)

TBD
Session Chair: TBA

- **ID089** [OL] **Robust Teacher Detection and Single-Target Tracking in Computer Science Laboratory Videos Using Spatial Context and Hybrid Vision Models.** — *Raoufi Elyousfi and Taoufik Ladhari*
- **ID094** [OL] **Accuracy Is Not Edge: An Honest, Leakage-Safe, Cost-Aware Benchmark of Machine-Learning Trade-Signal Selection on Indian (NSE) Equities.** — *Gauransh Kalla and Kartik Gupta*
- **ID109** [OL] **ML-based multi-layered hazard classification for runaway reactions in hybrid systems.** — *Amani Grati, Alamera Nouran Alquennah, Marcelo Castier, and Luc Vechot*
- **ID126** [OL] **Intelligent Learning and Assessment Ecosystems: A Multi-Agent AI Framework for Exam Analysis and Adaptive Learning.** — *Belhassen Guettat and Aymen Ben Brik*
- **ID128** [OL] **SEE: Prompt-Guided Structured Semantic Scene Understanding with LLaVA-OneVision for Surveillance Video.** — *Maria Khan and Hafiza Hina Kanwal*

12:00 - 12:30 **Closing & Award Ceremony**
Hall C

12:30 - 13:15 **Lunch**

Artificial Intelligence for Engineering Applications

ID022: A Danger-Aware Heuristic Pathfinding Algorithm for Safe Autonomous Navigation in 2D Engineering Environments

Hajar Alsulaili, Muneerah Alshammeri, and Manar Almutairi

Safety-aware navigation is an important artificial intelligence problem in engineering applications, where autonomous agents must balance efficient goal-reaching with reduced exposure to environmental hazards. Traditional grid-based algorithms such as A* can find shortest paths efficiently, but classical A* often ignores the risk level of cells along the path. As a result, it may produce routes that are optimal in length but less safe in environments containing hazardous or undesirable regions. Existing pathfinding methods mainly focus on reducing path length, runtime, or node expansions, while specialized risk-aware approaches may require complex environmental models and additional computational cost. This paper proposes A*-DZ, a danger-aware extension of A* for 2D grid-based engineering environments. The method adds a danger-zone penalty directly to the cost function, where each traversable cell contributes an additional cost according to its hazard level. This allows the algorithm to avoid risky cells when safer alternatives are available, while preserving the main heuristic structure of A*. Experiments were conducted on three 2D grid layouts: Room, Maze, and Random. The results show that A*-DZ maintains the same path length as classical A* in all tested maps, while reducing node expansions by up to 54.35% and improving runtime by up to 50.12% in open and semi-structured environments. These findings demonstrate that A*-DZ provides a simple, computationally efficient, and tunable AI-based approach for safer pathfinding in 2D grid-based engineering navigation tasks.

ID032: Artificial Intelligence Enabled Curriculum Transformation Through an Intelligent Ecosystem for Future Workforce Readiness

Ziad Salem, Hessa Alanjari, Frishtay Fatima, and Abdullatif Baba

Higher education institutions face increasing pressure to align academic curricula with rapidly evolving workforce requirements resulting from advances in artificial intelligence (AI), automation, and digital transformation. Conventional curriculum review cycles frequently lack the agility to address emerging industry demands, leading to persistent skill gaps between graduates and employers. This study introduces an Artificial Intelligence-Enabled Intelligent Curriculum Ecosystem (ICE), which integrates stakeholder feedback, educational analytics, machine learning classification, and explainable AI techniques to facilitate evidence-based curriculum transformation. The proposed framework integrates structured questionnaire data from students, faculty members, and industry professionals with predictive analytics employing Random Forest, Decision Tree, Support Vector Machine, and Logistic Regression models. SHAP-based explainability methods are incorporated to identify the most influential curriculum and workforce readiness indicators. The methodology introduces measurable indicators, such as the Curriculum Relevance Score (CRS), Workplace Readiness Score (WRS), and AI Integration Score (AIS). Experimental evaluation using stakeholder data demonstrates that ensemble learning approaches outperform baseline classifiers, achieving a maximum AUC score of 0.91. Statistical validation through Analysis of Variance (ANOVA) and post-hoc analysis further confirms significant perception gaps between academia and industry regarding graduate readiness. The proposed ecosystem contributes toward intelligent, adaptive, and workforce-aligned curriculum modernization in higher education.

ID099: AI-Driven Crime Hotspot Prediction and Energy-Aware Patrol Allocation for Sustainable Smart Cities

Mohamad Hoda

Urban patrol planning requires balancing publicsafety demand with operational efficiency. This paper presents a two-stage AI decision-support framework for crime hotspot prediction and energy-aware patrol allocation in a smart-city setting. Using 8.56 million Chicago crime records from 2001 to 2026, incidents are aggregated to a police beat-week panel and grouped into four categories: violent, property, public-order, and other. The first stage predicts whether a beat will be a nextweek hotspot, defined as falling in the top 20% of beats by next-week priority-crime count. The second stage uses predicted hotspot probabilities to allocate a fixed number of patrol units while penalizing unnecessary geographic spread; travel distance and emissions are estimated from beat centroids. On a chronological 2025 test set, logistic regression achieved an F1-score of 0.641 and ROC-AUC of 0.898, compared with 0.618 F1 for a static historical beat-rank baseline. In the allocation prototype, the proposed probability-and-distance method improved hotspot coverage to 0.268 while reducing mean weekly travel distance to 65.8 km and estimated emissions to 16.5 kg CO_2e. A sweep over the distance-penalty weight traces a coverage-emissions Pareto frontier, showing that this trade-off is tunable rather than fixed to a single chosen setting. A district-level spatial equity check further shows that risk-ranked allocation methods, including the proposed method, concentrate patrols more strongly than the underlying next-week crime-burden reference. The results show that hotspot prediction offers incremental value beyond persistent spatial risk, that energy savings and coverage can be jointly tuned along an explicit frontier, and that such gains currently come with a measurable spatial-equity trade-off.

ID086: Integrated Load-Flow and CCT-Based Transient-Stability Assessment of Photovoltaic Penetration in the IEEE 14-Bus System Under Variable Loading and Fault Conditions

Mustafa Alhariri, Rashad M. Kamel, and Sultan S. Alanzi

The integration of photovoltaic distributed generation (PVDG) can alter both the steady-state power-flow characteristics and transient stability performance of existing power systems. This paper evaluates PVDG integration into the IEEE 14-bus system using ETAP under different loading levels, PV penetration levels, fault types, and fault locations. PV penetration levels of 5%, 10%, 15%, and 20%, referenced to the 259 MW design load, are considered with the PVDG concentrated at bus 9. Steady-state performance is examined under design, summer, and winter loading conditions using bus voltages, branch currents, voltage drops, and total active power losses. A comparative critical clearing time (CCT)-based transient-stability assessment is performed under three-phase and single-line-to-ground (SLG) faults using the post-fault speed response of the largest synchronous generator, G1. Because the relative rotor-angle separation of all synchronous machines is not evaluated, the reported CCTs are interpreted as G1-response-based comparative stability margins rather than complete multimachine synchronism limits. The synchronous generators and condensers remain connected in all cases; therefore, synchronous inertia is maintained as PV penetration increases. The results show that increasing PV penetration improves the voltage profile and reduces active power losses under design and summer loading. Under winter loading, the minimum loss occurs at 15% PV penetration, after which network losses increase. Higher PV penetration also increases the calculated CCT under both fault types, while electrically remote faults produce larger CCTs. These findings demonstrate that the effects of PV integration depend on loading, fault location, fault type, generator dispatch, and the assumed treatment of synchronous generation.

ID017: Uncertainty-Aware Physics-Informed Neural Networks for Degree of Polymerization Estimation in Oil-Immersed Power Transformers

Alamera Nouran Alquennah, Mohammad Alshaikh Saleh, Ali Ghrayeb, Haitham Abu-Rub, Shady Khalil, Mohammed Al-Hajri, and Sunil Khatri

This paper proposes an uncertainty-aware physics-informed neural network (PINN) model for predicting the mean degree of polymerization (DP) and its standard deviation in the insulation paper of oil-immersed power transformers under uncertain future

thermal conditions characterized by variations in the hotspot temperature (HST). The proposed model produces DP predictions in the form of uncertainty-aware degradation trajectories rather than single-point estimates, thereby providing deeper insight into the impact of future load plans, represented as thermal stress, on insulation degradation. The PINN model integrates deterministic historical measurements with uncertain future operating conditions by learning the data-driven mapping, incorporating Monte Carlo sampling within the stochastic region, and enforcing the governing physical degradation law through Emsley's model. The model is trained and evaluated using synthetically generated DP profiles based on randomly distributed sensor measurements and a wide range of thermal stress scenarios. This approach, which incorporates a diverse set of randomly generated DP profiles, is intended to enhance robustness, reduce the risk of overfitting, and improve the model's generalization capability.

ID116: Quantum Machine Learning for Efficient and Sustainable Robotics: A Review

Ali Alfayly, Mahmood Seyyedzadeh, Amir Ghavifekr, Hooman Esmaielzadeh, and Amir Ghiasi

Quantum machine learning (QML) exploits quantum-mechanical phenomena superposition, entanglement, and interference to enrich the expressivity and, for selected tasks, the computational efficiency of learning models. Modern robotics depends increasingly on data-driven perception, planning, and control, yet remains constrained by computational cost, sample efficiency, robustness, safety, and the limited energy budgets of embedded and mobile platforms. This paper reviews QML techniques from a robotics perspective and highlights their implications for efficient and sustainable autonomous systems. We describe the principal QML algorithm families quantum neural networks, quantum kernel methods and support vector machines, quantum clustering, quantum generative models, and quantum reinforcement learning with their problem settings, advantages, and practical limitations on noisy intermediate-scale quantum (NISQ) hardware. We then survey robotics applications in perception, motion planning, control, and decision making, comparing evaluation criteria, experimental settings (hardware versus simulation), and reported performance and efficiency gains. Finally, we outline open challenges and directions toward reproducible, safe, energy-aware, and hardware-ready QML-enabled robotics.

ID117: An Attention-Based Neural Network Method for Acoustic Anomaly Detection in Pipelines

Muhammad Zia Ur Rehman, Omar Khattab, Zurasha Zaryab, Ismail Kashif, Ahsan Ali Raza, and Murad Khan

Globally pipeline networks have significant sustainability and economic challenges stemming from leaks in oil, gas, and water transporting systems. Therefore, timely detection, location and reporting of abnormal pipeline conditions are required for safe and sustainable operation. However, the present leak detection methods are generally unreliable and robust under noisy industrial environments. This motivates the development of efficient data-driven methods. Most of the available datasets and Machine Learning studies are focused on image and speech recognition, and only a few studies are focused on acoustic anomaly detection for pipeline leakage based on industrial sensor signals. In this paper, we propose an attention-based neural network for acoustic anomaly detection in pipelines with flow, pressure and vibration features from sensors. In the proposed approach, we explore missing values, feature correlation and statistical behaviour using exploratory data analysis before the model is trained. The attention mechanism enables the model to assign higher importance to informative signal segments while reducing the effect of irrelevant or noisy observations. The suggested model is evaluated against artificial neural networks, convolutional neural networks, and long short-term memory benchmarks. Experimental results show that the attention-based neural network achieves superior performance, with 97.75% accuracy and 97.05% F1 score. These results demonstrate that attention-guided acoustic anomaly detection can provide a reliable and scalable direction for pipeline monitoring, leakage detection, and industrial infrastructure safety.

ID056: Vision-Based Cattle Weight Estimation: A Systematic Review of Computer Vision Models and Image Acquisition Methods

Marcus Joshua Fortes, Ajna Shukkoor, Usama Mokashi, Shady Khalil, and Wesam Rohouma

Accurate and non-invasive cattle weight estimation is critical for optimizing herd management, feed efficiency, health monitoring, and market readiness. Conventional livestock weighing practices are often constrained by high labor requirements, animal stress, and limited scalability. Computer vision (CV) technologies have emerged as a promising non-contact alternative, enabling automated weight estimation from camera-captured imagery. This paper presents a structured review of CV-based approaches for cattle weight estimation. Five primary experimental studies are reviewed in detail, including four cattle studies and one pig study retained for methodological comparison alongside several review and survey papers for broader contextual synthesis. Image acquisition methods are reviewed and compared, ranging from standard two-dimensional (2D) and closed-circuit television (CCTV) cameras to red, green, blue, and depth (RGB-D) systems. In addition, relevant machine learning models including linear regression, Convolutional Neural Networks (CNNs), and segmentation-based approaches are examined. The reviewed studies tend to report that depth-based systems can provide richer geometric information and have achieved strong reported performance, while standard 2D camera setups remain a potentially cost-effective and comparatively underexplored alternative. Major challenges including dataset scarcity, environmental sensitivity, and model generalization are discussed. Future research directions are identified, with particular emphasis on CCTV-based deployment in real farm environments.

ID066: A Unified Vision-Based System for Cattle Weight Estimation and Behavior Detection

Ajna Shukkoor, Wesam Rohouma, Shady Khalil, Rachid Benlamri, Amit Khandakar, and Usama Mokashi

Accurate and non-invasive monitoring of cattle weight and behavior is essential for modern precision livestock management. Traditional approaches, which rely on manual observation or contact-based weighing, are labor-intensive, stressful to animals, and impractical for large-scale implementation. This paper presents a unified, vision-based pipeline that jointly estimates cattle live weight and classifies behavioral states from standard CCTV footage. The proposed architecture employs a shared preprocessing stage-encompassing bounding-box detection, individual cow cropping, and normalization, followed by two independent deep CNN branches: a regression branch for weight estimation and a classification branch for behavior recognition. A viewpoint filter restricts the weight branch to side- and rear-view frames to improve geometric consistency, while temporal aggregation (majority voting for behavior and temporal averaging for weight) stabilizes predictions over time. Experiments on the CBVD-5 benchmark and an independent weight dataset demonstrate that the decoupled, modular design achieves competitive performance on both tasks while sharing a common feature extraction pipeline. The proposed system offers a practical, scalable solution for automated cattle monitoring in real-world farm environments.

ID085: Comparative Analysis of Traditional Machine Learning and Graph Neural Network Models for Anti-Money Laundering Detection

Aman Reddy, Areesha Hamad, Glen Anton Paul, Manikandan Pradeep, Mohammad Hasheer, Umar Syed, Syama Kurungodathil, and Asmath Famnaz

Anti-money laundering detection is a severely imbalanced classification problem in which random splitting, premature resampling, and future graph edges can produce misleading results. This study compares Random Forest, XGBoost, linear Support Vector Machine, Graph Convolutional Network, GraphSAGE, and Graph Attention Network on one common transaction-level task using the complete IBM AML HI-Small file. After nine exact duplicates were removed, 5,078,336 transactions were ordered by time and divided into approximately 60% training, 20% validation, and 20% testing partitions without splitting identical timestamps. Preprocessing, historical aggregates, negative sampling, hyperparameter selection, and threshold selection were restricted to training or validation data as appropriate. All models used the same seed-specific training transactions, while validation and test prevalence remained unchanged.

Across five seeds, XGBoost achieved the highest mean average precision of 0.4402 and F1-score of 0.4192, with precision 0.7798 and recall 0.2874. GAT had the highest GNN average precision at 0.0387, while GraphSAGE had the highest GNN F1-score at 0.0689. A timestamp-tail sensitivity analysis, measured runtimes, SHAP explanations, GNNExplainer cases, and graph ablations show both the value of leakage controls and the limits of static accountgraph message passing for this benchmark.

ID053: Large Language Models for Low-Resource and Agglutinative Languages: A Systematic Review

Adil Akhmetov, Mohammed Alaa Ala'Anzy, Malek Ali, and Adnan Saeed

Natural language processing has advanced remarkably with large language models (LLMs), yet this progress concentrates overwhelmingly on high-resource languages, leaving speakers of thousands of other languages at a systematic disadvantage. Agglutinative languages such as Kazakh present compounding challenges: morphological complexity drives extreme vocabulary sparsity, Cyrillic script handling introduces encoding barriers not addressed by standard preprocessing pipelines, and BPE tokenizers applied to Turkic languages produce token fertility rates of 3.2 to 4.1, substantially higher than rates observed for morphologically simpler languages. A systematic literature review of 33 studies published between 2020 and 2026 addresses these challenges for low-resource and agglutinative languages broadly, using Kazakh as a representative case study within the Turkic family. The review synthesizes evidence across four dimensions: tokenization failures in morphologically rich languages, comparative performance of massively multilingual models across language families, parameter-efficient fine-tuning strategies for resource-constrained settings, and cross-lingual transfer limitations. Analysis of five major multilingual models reveals that Turkic languages achieve F1 scores of 46.1 to 71.4, substantially below the 82.1 to 94.2 range recorded for Indo-European languages. On the basis of these findings, the paper proposes a development roadmap for a sovereign Kazakh LLM (KazLLM), specifying tokenization strategy, pretraining data composition, fine-tuning methodology, and evaluation benchmark design.

ID095: AI-Driven Intrusion Detection Framework for Industrial IoT in Smart Grid Environments

Mohamed H. Saad, Basem Elhalawany, and Ramy Samy

The increasing integration of Industrial Internet of Things (IIoT) devices into smart-grid infrastructures has introduced new cybersecurity challenges for operational technology networks. Ensuring secure and reliable grid operation requires intrusion detection mechanisms that combine high detection accuracy with low inference latency, particularly in renewable-energy-enabled power systems. Existing approaches often prioritize predictive performance without adequately addressing computational efficiency, limiting practical deployment in real-time smart-grid edge environments. This paper proposes an AI-driven intrusion detection framework specifically designed for IIoT traffic in smart-grid systems. Using the publicly available TON_IoT benchmark dataset, the framework integrates correlation-based feature optimization (reducing dimensionality from 49 to 29 features, a 41\% reduction) with comparative supervised learning using Random Forest, XGBoost, and Long Short-Term Memory (LSTM) models. Experimental evaluation employs $5 imes2$-fold cross-validation to ensure statistical robustness. XGBoost achieves the best overall performance with 98.1\% accuracy, 97.9\% F1-score, and an average inference latency of 13.4~ms on standard hardware. Feature optimization reduces inference latency by up to 21\% without degrading classification accuracy. The framework explicitly quantifies the trade-off between detection performance and real-time responsiveness, providing practical guidance for edge deployment in smart-grid environments that support green-energy sources. These findings establish a reproducible, latency-aware benchmark for AI-enabled intrusion detection in critical energy infrastructure and support scalable cybersecurity for next-generation smart grids.

ID029: Multi-Sensing Deep Learning Framework for PEMFC Degradation Forecasting Under Heterogeneous Operating Conditions

Ahmad Braydi, Elena Breaz, Issam Salhi, Fei Gao, and Qian Xun

Proton exchange membrane fuel cells (PEMFCs) are promising energy conversion systems for transportation and distributed power applications, yet their progressive degradation remains a major challenge for reliability and lifetime management. This paper proposes a multi-sensing deep learning framework for PEMFC degradation forecasting under heterogeneous operating conditions. The proposed approach combines long short-term memory (LSTM) networks with a self-attention mechanism to capture both temporal operational dynamics and long-term degradation behaviour from multi-variable operational data. The framework integrates multi-variable operational measurements to improve degradation prediction under steady-state, quasi-dynamic, cyclic-load, and start-stop operating regimes. Experimental evaluation on multiple PEMFC degradation datasets demonstrates strong predictive capability under diverse operating conditions, achieving an R^2 of 0.9768 under quasi-dynamic operation and long-term forecasting performance of 0.9919 and 0.9612 under cyclic-load and start-stop conditions, respectively. The results highlight the effectiveness of the proposed framework for robust PEMFC degradation forecasting and long-term prognostics.

ID047: Mallat Scattering Transform -Based Multiscale 3 branch CNN–BiLSTM Framework for Accurate Solar Irradiance Forecasting

Mariem Mallat, Olfa Bel Hadj Brahim Kechiche, Mahmoud Hamouda, Houcine Oudira, and Aissa Chouder

This paper proposes a hybrid forecasting framework that separates feature extraction from temporal modeling. The Mallat Scattering Transform is a feature extraction method that provides translation-invariant and deformation-stable representations that capture local, global, and cross-scale variations of irradiance. These features are then combined with a Multiscale 3 Branch CNN-BiLSTM model to leverage spatial patterns and bidirectional temporal dependencies. The model is trained and evaluated on one year of hourly irradiance data using standard performance metrics. Results show that the proposed model reduces the prediction error from approximately 21.27 W/m^2 to 3.2 W/m^2 and improves the coefficient of determination from 0.989 to 0.9996 compared to the baseline. These results demonstrate that the framework offers a highly reliable solution for photovoltaic energy management and forecasting.

ID055: Deep Learning-Based Cattle Behavior Classification from Visual Data: A Comparative Review

Josh Allen San Diego, Ajna Shukkoor, Usama Mokashi, Shady Khalil, and Wesam Rohouma

Automated cattle behavior classification using computer vision-based systems has emerged as a promising approach to precision livestock farming, offering a non-invasive alternative to sensor-based monitoring. This paper presents a comparative review of ten studies that employ deep learning to classify cattle behaviors from CCTV and camera-based video data. The reviewed studies are analyzed across three architectural categories: CNN-based approaches, CNN with temporal models, and spatiotemporal models. Their methodologies, datasets, camera setups, behavior coverage, and performance metrics including precision, recall, F1-score, mAP, and frames per second are examined. The results indicate that most reviewed systems achieved strong performance, exceeding 90% across reported metrics, with a general trend of decreasing performance as the number of classified behaviors increases. A clear architectural progression is evident, with spatiotemporal models demonstrating improved capability for recognizing behaviors defined by continuous movement. However, the review also identifies several recurring challenges including environmental variability, class imbalance, inconsistent evaluation protocols, and deployment constraints, which limit real-world applicability. Future directions toward standardized benchmarks, diverse datasets, and lightweight deployable architectures are outlined to guide further advancement in this field.

ID089: Robust Teacher Detection and Single-Target Tracking in Computer Science Laboratory Videos Using Spatial Context and Hybrid Vision Models

Raoufi Elyousfi and Taoufik Ladhari

Reliable teacher localization is a key prerequisite for classroom video analytics, but it remains challenging in realistic computer laboratory scenes due to multiple students, partial occlusions, perspective distortion and non-overhead camera viewpoints. This paper proposes a spatially guided hybrid framework for teacher detection and single-target tracking. The method exploits classroom-specific spatial cues, including board position and laboratory geometry, to initialize and preserve teacher identity over time. It combines periodic person detection, contextual teacher selection, local temporal tracking and motion-based identity correction. Experiments on ten real computer laboratory videos show that the proposed method achieves an average IoU of 0.72, precision of 0.87 and recall of 0.85, outperforming YOLO alone and a generic YOLO-based tracking baseline on average. These results indicate that integrating spatial classroom context improves teacher-centered tracking under realistic classroom conditions.

ID094: Accuracy Is Not Edge: An Honest, Leakage-Safe, Cost-Aware Benchmark of Machine-Learning Trade-Signal Selection on Indian (NSE) Equities

Gauransh Kalla and Kartik Gupta

Published machine-learning (ML) stock-prediction systems often claim very high accuracy, yet such figures rarely survive an evaluation that controls data leakage, tests strictly out of sample, and charges transaction costs. This paper reports exactly that kind of evaluation on the Indian market. For every (stock, trading-day) pair, a reproducible pipeline defines a concrete ATR-based trade (entry, target, stop) over a five-day horizon, grades it with the triple-barrier method, and trains classifiers to rank stocks by their calibrated probability of a winning trade. The data cover approximately 491{,}000 labelled examples from 474 National Stock Exchange of India (NSE) equities over roughly five years. Seven model variants (logistic regression, random forest, a multilayer perceptron, XGBoost, LightGBM, and histogram-based gradient boosting, the last with and without isotonic calibration) are compared under one walk-forward protocol with an embargo, giving 285{,}197 out-of-sample predictions over 625 test days. Accuracy and profit turn out to be nearly unrelated. Every model shows near-random ranking skill (ROC-AUC 0.517--0.537) and accuracy indistinguishable from the majority-class baseline, yet only the gradient-boosting family earns a positive after-cost expectancy on decisive (win/loss) trades, with profit factors up to 1.18; logistic regression, random forest, and the neural network are negative on the same measure. A by-day block bootstrap supports the boosting-versus-rest split, one-sided p = 0.005), and the boosting superiority still holds up under realistic transaction costs. No single model can carry significance alone, and the edge is concentrated in time. Headline accuracy is therefore the wrong target for trade-signal research; probability calibration and cost-aware economic metrics are the more informative yardsticks, and honest results in this setting are modest by construction. All code and the evaluation harness are released for reproducibility.

ID109: ML-based multi-layered hazard classification for runaway reactions in hybrid systems

Amani Grati, Alamera Nouran Alquennah, Marcelo Castier, and Luc Vechot

Exothermic reactions under runaway conditions pose catastrophic risks in the chemical processing industry. To prevent or mitigate these hazards, layers of protection are implemented, which include emergency relief systems (ERS) for the prevention of vessel explosions. Designing effective ERS requires a comprehensive understanding of reactive systems' behavior under venting conditions. This is especially true for hybrid systems, which generate both vapor and non-condensable gases during the runaway. For such systems, the extent of tempering phenomena, i.e., the consumption of latent heat of the liquid contents, directly impacts the runaway behavior of the reactive mixture during venting and must be understood for ERS sizing. This work assesses the performance of six supervised machine learning algorithms to identify the optimum model for predicting the behavior of a hybrid system under runaway conditions during venting. A total of three key operational parameters were utilized as input to classify hybrid systems into one of three distinct hazard levels. To ensure reproducibility, all models were trained and tested on the same randomized dataset. While all

algorithms achieved high performance in terms of accuracy, exceeding 90%, the model selection prioritized minimizing specific errors to ensure hazard reduction over overall accuracy.

ID126: Intelligent Learning and Assessment Ecosystems: A Multi-Agent AI Framework for Exam Analysis and Adaptive Learning

Belhassen Guettat and Aymen Ben Brik

Higher education institutions face growing pressure to align examinations with course learning outcomes (CLOs) and accreditation frameworks such as AACSB and BGA, while exam design, review, and grading remain largely manual. This paper presents an agentic AI framework for exam analysis, course assessment review, and adaptive learning, built on a Model Context Protocol (MCP)-based multi-agent architecture. Twenty-three specialized agents, organized into five coordinated layers, ingest exam documents, classify questions by Bloom level and difficulty, map them to learning outcomes, flag coverage gaps, and generate actionable recommendations, while a retrieval-augmented generation (RAG) chatbot supports students with course-grounded explanations and personalized quizzes. We report a case study on a real mathematics examination and discuss design trade-offs observed across three architectural iterations, positioning the contribution with respect to SDGs 4, 9, and 10.

ID128: SEE: Prompt-Guided Structured Semantic Scene Understanding with LLaVA-OneVision for Surveillance Video

Maria Khan and Hafiza Hina Kanwal

Video anomaly detection in surveillance environ ments is often limited by visual representations that describe appearance and motion without explicitly organizing scene se mantics. This paper isolates and presents the SEE (Semantic Extraction Engine) module of a larger neuro-symbolic frame work. SEE transforms three consecutive surveillance frames into a structured semantic state using LLaVA-OneVision and a scene-specific prompt. The temporal window retains sufficient information to distinguish stationary and moving objects while avoiding full-video processing. The prompt restricts analysis to visible evidence, predefined surveillance objects, controlled location labels, and explicit motion states. It also requires valid JSON containing object identifiers, categories, locations, motion attributes, relationships, a concise caption, and a detailed scene description. Experiments use the NWPU Campus surveillance dataset. Of 242 processed videos, 202 produced valid, nondupli cate semantic records for controlled evaluation. To isolate the contribution of SEE prompt engineering, the proposed prompt is compared with a generic baseline while Knowledge Graph integration is disabled. The proposed configuration improves accuracy from 89.42% to 90.05%, precision from 50.00% to 63.64%, recall from 5.00% to 30.43%, and F1-score from 9.09% to 41.18%. These findings show that structured, scene specific prompting substantially improves semantic evidence for downstream anomaly interpretation.

Renewable Energy, Distributed Generation, and Smart Grids

ID087: Toward Smart Livestock Infrastructure: Sensor Benchmarking for Methane Monitoring in Green Dairy Farm Systems.

Usama Mokashi, Wesam Rohouma, Shady Khalil, Rachid Benlamri, and Ajna Shukkoor

The growing demand for smart agricultural infrastructure has accelerated the integration of Internet of Things (IoT) technologies into livestock farm management systems. Methane (CH?) emitted from enteric fermentation in dairy cattle represents a significant contributor to global greenhouse gas emissions, yet current measurement approaches remain expensive, labor-intensive, and unsuitable for continuous barn-scale deployment. This paper presents a sensor benchmarking study and IoT framework design for a low-cost wearable methane monitoring system applicable to both indoor and outdoor green dairy farm environments. Four sensing technologies Non-Dispersive Infrared (NDIR), Metal Oxide Semiconductor (MOS), Electrochemical, and Optical are systematically evaluated against key criteria including cost, accuracy, power consumption, environmental tolerance, and wearable IoT suitability. Benchmarking results indicate that NDIR sensors, specifically the MH-Z19B, offer the most favorable performance profile for continuous individual animal-level monitoring. Building on this, a lightweight wearable IoT sensing architecture is proposed, designed for scalable deployment within smart dairy farm infrastructure. This work contributes evidence-based sensor selection guidance and a practical IoT framework toward the development of sustainable, data-driven green livestock management systems.

ID063: Droop-Based Grid-Forming Inverter Control for a PV/BESS Microgrid Integrated with the IEEE 13-Bus Distribution System

Muhammad Salim Khan, Amira Mohammed, Shady Khalil, Mostafa Farrag, and Mohamed Massaoudi

Solar photovoltaic (PV) generation and battery energy storage systems (BESS) are increasingly interfaced to distribution grids through power electronic converters, driving these grids toward low-inertia, converter-dominated operation. Conventional grid-following (GFL) inverters rely on a phase-locked loop (PLL) to synchronize with an existing voltage reference; consequently, they cannot establish voltage or frequency independently once the upstream grid disconnects. This paper presents a droop-based grid-forming (GFM) control scheme for a renewable microgrid supplying 25% of the load on the IEEE 13-bus distribution test feeder. Of this, 20% is provided by a 693.2 kW PV array and 5% by a 173.2 kW lithium-ion BESS. A perturb-and-observe maximum power point tracking (MPPT) routine extracts peak power from the array, and a bidirectional DC-DC converter regulates battery charging and discharging. In place of a PLL, the proposed GFM controller obtains its operating point directly from P-f and Q-V droop laws. Four cases were modeled in MATLAB/Simulink R2024a: islanded PV-only operation, islanded PV with BESS, grid-connected operation, and a seamless transition between modes. In every case, the controller maintained the point of common coupling at 480 V RMS and 60 Hz while tracking the unbalanced three-phase loading characteristic of the IEEE 13-bus feeder. These results indicate that droop-based GFM control provides a practical, PLL-free path to stable microgrid operation under high penetration of inverter-based resources.

ID096: Green Hydrogen Production from Harvested Rainwater: Evidence, Challenges, and Conceptual System Integration

Mrinmoy Das, Ramesh Kumar, Gianluca Gatto, Suman Jana, and Amit Kumar

This paper evaluates harvested rainwater as an alternative source water for green hydrogen production through electrochemical, photocatalytic, photoelectrochemical, and renewable-powered pathways. Rainwater-specific and supporting studies are evaluated with a focus on water quality, pretreatment, electrode and catalyst design, electrolyte management, operating conditions and liquid-vapour-phase hydrogen production. Direct rainwater-specific evidence remains limited; therefore, indirect evidence from wastewater treatment, atmospheric moisture utilisation and vapour-phase photocatalysis is considered as supporting context. The variability of rainwater in ionic content, suspended matter, dissolved contaminants and microbial load presents risks of fouling, corrosion, catalyst deactivation and unstable electrolyser performance. Reported laboratory results are not readily comparable because of variations in reactor design, catalyst loading, illumination, electrolyte composition and definitions of efficiency. Major research gaps are the lack of a standard approach to characterise rainwater, limited durability data, limited pilot-scale demonstrations, and limited rainwater-specific techno-economic and seasonal assessment. Therefore, a conceptual architecture is proposed, including rainwater harvesting, treatment, renewable power, liquid and vapour phase utilization, and supervisory control, with experimental, safety, durability, seasonal-resource, and techno-economic validation identified as requirements for future development.

ID114: Economic Evaluation of Utility-Scale Solar Photovoltaics, Onshore Wind, and Supercritical Coal Power Plants Using Annual Worth Analysis

Aisyah Nadya Faradina, Shabriya Ridzqi Rahman, Filscha Nurprihatin, Hendy Tannady, Ali Vaezi, and Michael Nayat Young

This paper presents an evaluation of three electricity generator projects, consisting of utility-scale Solar Photovoltaics, Onshore Wind Turbines, and Supercritical Coal-based Steam-Cycle, using the engineering economy approach. An Annual Worth (AW) method is conducted to analyze the viability of the options, focusing on calculating the Equivalent Annual Cost from the extracted first cost and operating cost in a standardized per megawatt electricity (MWe) over each life cycle. A uniform calculation from the financial data that were taken specifically the 2023 values, and using the Minimum Attractive Rate of Return (MARR) of 5.5%, which follows the Bank Indonesia rate. Solar Photovoltaics technology emerges as the most economically viable with the lowest annual cost among others. The rest of the options are eliminated due to a higher annual cost result for Supercritical Coal-based Steam-cycle, followed by Onshore Wind Turbine. Based on the result, the comparison is aimed to stand as financial scheme-based evidence for a future decision on electricity generator investment, reinforcing Indonesia's commitment to the sustainable energy transition through the power-generating system.

ID122: A Supervisory Zonal Multi-Agent Volt/VAR Controller for Active Distribution Networks

Muhammad Usman

High penetration of renewable generation can cause fast voltage variations in active distribution networks, especially during periods of high generation and low demand. Conventional voltage control devices provide useful support, but they have slow response time. This paper proposes a supervisory zonal multi-agent Volt/VAR controller that coordinates fast voltage control devices with slower distribution system operator (DSO) owned device. The feeder is divided into zones, where each fast control device is modelled as a distributed-generator agent. Each agent uses semantic token-history information from its local and zonal operating conditions to determine reactive power action using a multi-agent soft actor-critic policy. The DSO supervises zone-level voltage conditions and updates DSO-owned control devices at a slower interval. The controller is evaluated on the IEEE 141-bus distribution feeder using one-minute resolution of load and PV data. The proposed controller demonstrates effective Volt/VAR coordination in the distribution feeder with high penetration of renewable generation.

ID010: Impact of Bulk and Interface Defects on 2T Perovskite/ACIGS and 2J/3J All-Perovskite TSCs

Nour El I Boukortt, Antonio Garcia Loureiro, Salvatore Patanè, Mohamed Trabelsi, Trupti Ranjan Lenka, and Ahmad Abushattal

In this work, the 2T perovskite/ACIGS and 2J/3J all-perovskite tandem solar cells are numerically analyzed with TCAD tools, taking the electrical and optical properties of each layer explicitly into account in the analysis. We employ two interconnection techniques, band-to-band tunnelling and series connection with ITO. The stacking process results in a Voc loss, and, therefore, a current matching at the maximum power point is necessary. One of the major challenges in such devices is to enhance Voc and FF without compromising Jsc, which is often attributed to interfacial defects in the sub-cells based on perovskite. These defects are due to the deposition methods, lack of additives or mismatches of the film solution during fabrication. The impact of defect states within the perovskite bulk and at its interfaces on tandem solar cell performance is systematically investigated. The study considers recombination processes at the ETL/PVK and PVK/HTL interfaces through their respective surface recombination velocities, aiming to quantify efficiency losses and support the optimization of advanced photovoltaic architectures.

ID034: Validation of Acoustic Propagation Models for Wind Farm Noise Assessment: An End-to-End Reproducible Workflow Using Python and Power BI

Emna Annabi, Mohamed Wassim Jmal, and Ali Wali

This paper presents a Python-driven end-to-end reproducible workflow for validating acoustic propagation models applied to nine Regulated Emergence Zones (ZERs) in onshore wind farms. Heterogeneous data streams (SCADA logs, sound level meters, meteorological masts, curtailment schedules) are harmonised at ten-minute resolution. Preprocessing includes circular wind direction averaging, hub-height wind speed normalisation, curtailment decoding, and per-ZER propagation condition classification following the French measurement protocol. Four visual diagnostics (trend curves, box plots, scatter plots, heatmaps) are generated by Python scripts and integrated into Power BI dashboards. Three error metrics (MAE, RMSE, Bias) are reported per stratified condition, with wind speed class identified as the dominant prediction-error driver, constituting the data-driven analytics dimension relevant to AI-assisted green-energy monitoring. Applied to nine ZERs, the framework yields global MAE of 1.04--2.20 dB and RMSE of 1.54--2.99 dB. The ZER1 case study (N=1 755) confirms close agreement at low wind speeds (classes 3--5 m/s, MAE ? 1.35 dB) and identifies the sole safety-critical underestimation stratum at night-time favourable conditions, class 6 (MAE = 1.96 dB, Bias = ?0.96 dB).

ID043: Experimental Performance Comparison of Different Monofacial and Bifacial PV Modules in Outdoor Conditions

Daniela Olivares Baqué, Alejandro González Moreno, Alberto Dolara, and Sonia Leva

Industrial and commercial rooftops offer significant potential for PhotoVoltaic (PV) installation; however, constraints related to module tilt, row spacing, and thermal management can limit system performance, particularly in bifacial PV applications. This study presents an experimental performance comparison of four commercial PV technologies installed on an industrial rooftop test bench and monitored under realistic outdoor conditions across three measurement campaigns conducted in 2025. The experimental setup enables sequential acquisition of I–V curves, paired with simultaneous measurements of plane-of-array irradiance on both the front and rear sides and back-surface module temperature. The collected data are used to evaluate daily efficiency profiles, efficiency–temperature correlations, performance under varying thermal conditions and bifacial gain, the latter assessed through temporary rear-side shading.

ID083: Innovative Economic Valorization Strategies for Industrial Photovoltaic Systems: Revamping, Repowering and Storage

Giacomo Manzan, Sonia Leva, and Marco Mussetta

As the global energy landscape transitions toward renewable sources to mitigate climate change, the economic viability of photovoltaic (PV) systems remains a critical driver for their widespread deployment. This paper explores advanced strategies to enhance the economic benefits of industrial-scale PV plants in Italy, based on real-world data and a collaborative study with a leading company in PV installations. The effectiveness of multiple valorization pathways is here analyzed: plant revamping, repowering, and the integration of utility-scale storage systems. The methodology encompasses a granular breakdown of Capital Expenditure (CAPEX), Operational Expenditure (OPEX), and lifecycle cash flow simulations over a 25-year period. Results demonstrate that Renewable Energy Communities incentives and the strategic revamping of aging "Conto Energia" PV plants significantly improve investment metrics, reducing payback times (PBT) to as little as 3 years and driving Internal Rates of Return (IRR) above 30%. Conversely, while standard utility-scale storage currently faces high CAPEX hurdles, trading-enabled systems participating in Terna's upward modulation services show distinct promise for future grid resilience and profitability.

ID111: Blockchain-Enabled Stackelberg Pricing for RealTime Decentralized Microgrid Energy Management:A Hardware-Validated Framework

Abdullah Umar, Vijayakumar Gali, Prashant Jamwal, Shady Khalil, Mostafa Farrag, Ajay Kumar, and Nitin Gupta

Abstract— Real-time coordination of distributed energy resources (DERs) in community microgrids requires decentralized, transparent, and incentive-compatible mechanisms capable of managing variable generation and heterogeneous loads. Centralized architectures suffer from communication bottlenecks and single-point failures, limiting their suitability for high-prosumer environments. This paper proposes a real-time decentralized management framework integrating Stackelberg game-theoretic pricing, blockchain execution, and hardware validation on a six-bus microgrid prototype. The Energy Management System (EMS) acts as the Stackelberg leader, computing equilibrium prices from aggregate local measurements, while prosumers and consumers optimize costs as followers within device and operational constraints. A smart contract on the Polygon proof-of-stake network records aggregate market decisions, validates equilibrium outcomes, and ensures tamper-proof settlements without disclosing individual data. Hardware experiments with PV emulators, inverters, smart meters, and embedded controllers demonstrate stable operation, robustness to measurement noise, and accurate equilibrium equilibrium under varying generation and load conditions. In the hardware implementation, the EMS price converges in seven iterations, while in the simulation it does so in six. The final equilibrium price is 0.427 USD/kWh on hardware and 0.412 USD/kWh in simulation (3.6% deviation). Most importantly, the prosumer allocations track the corresponding theoretical optima with a 3.2% to 3.7% error, and the on-chain settlement is finalized in less than 2 s.

ID123: Assessing EV Penetration Impacts on Urban Air Quality: A Kigali Case Study

Samuel Chukwuma, Kevi Jonathan Iradukunda, Temidayo Akanmu, and Jesse Thornburg

Estimating the air quality benefits of electric vehicle (EV) adoption is difficult in low-resource cities that lack emissions inventories, traffic activity data, and dispersion models. This study develops a scenario-based framework to estimate PM2.5 reductions from EV penetration in Kigali, Rwanda, anchored in a publicly available Rwanda Environment Management Authority analysis of the 2025 UCI Road World Championships, during which major roads were closed and traffic was substantially disrupted. The authors did not collect the underlying measure- ments; the contribution is the modeling framework built from the openly reported data. The observed reduction is treated as an empirical traffic-sensitive response factor rather than a source- apportioned traffic share, and is combined with electrification scenarios for motorcycles, buses, and private cars. Present-day penetration corresponds to an estimated reduction of 1.6%. Under Rwanda's 2030 targets the model projects 9–12%, the upper end assuming only tailpipe emissions are removed and the lower end accounting for non-tailpipe brake, tire, and road- dust emissions. A Monte Carlo propagation places the 90% interval at 6.0–12.8% across model variants. Holding the number of electrified vehicles constant, the category-targeted allocation delivers 2.6 percentage points more reduction than a uniform fleet-wide share, a 27.7% relative gain at identical fleet cost. Sensitivity analysis shows the dominant uncertainties come not from fleet composition but from the level of ambient PM2.5 that is reducible by electrification.

ID124: Multi-Parametric Performance Evaluation of Six- and Twelve-Sector Direct Power Control for Grid-Connected Converters under Partial Load Conditions

Mohamed Amine Fnaiech, Salwa Baserrah, Ebrahim Abdulrahman Hasan, Khaled Zehar, Mohamed Trabelsi, and Hani Vahedi

This paper presents a multi-parametric performance assessment of direct power control (DPC) strategies for grid-connected voltage-source converters under partial-load conditions, in which three switching-table configurations are investigated: the conventional twelve-sector DPC, an asymmetric six-sector DPC, and a symmetric six-sector DPC. Unlike conventional comparisons based on fixed controller settings, this study evaluates the combined influence of sampling frequency and hysteresis bandwidth on transient response, active/reactive power regulation accuracy, and grid-current THD. The analysis highlights the strong dependence of DPC performance on digital implementation parameters, particularly at low power levels where current ripple has a higher relative impact

on power quality. The presented simulation results show that the asymmetric six-sector DPC provides the best overall trade-off, achieving faster dynamics, improved active/reactive power decoupling, lower energy fluctuation, and reduced THD compared with the twelve-sector and symmetric six-sector strategies. The proposed evaluation framework offers practical guidelines for calibrating DPC sectorization and tuning parameters for grid-connected converters operating under challenging partial-load conditions.

ID134: Single-Switch MPPT Control of a PMSG Wind Turbine for Standalone and DC-Grid-Connected Residential Buses

Maged Bauomy, Haytham Gamal, Adel Shaltout, Shady Khalil, and Morcos Metry

Small-scale variable-speed wind turbines (VSWT) based on permanent-magnet synchronous generators (PMSG) offer a practical renewable-energy option for residential DC systems. This paper presents a single-switch maximum power point tracking (MPPT) strategy for a wind-energy residential local branch system (WER-LBS). The PMSG is connected to the DC bus through an uncontrolled rectifier and a DC–DC buck converter. Using the measured generator speed, the controller determines the optimum electromagnetic-torque and converter-current references, eliminating direct wind-speed measurement. The system is modeled in MATLAB/Simulink and evaluated in standalone and DC-grid-connected operating modes under step changes in wind speed and load demand. Results show convergence of generator speed and torque to their optimum operating points at 12, 10, and 9 m/s. In standalone operation, the DC-bus voltage varies according to available wind power and load demand, whereas the grid-connected configuration maintains 120 V while balancing generation and load. The results demonstrate the feasibility of the low-switch-count interface for maximum-power extraction in residential DC wind-energy applications.

ID075: Governor Valve Actuator Systems for Smart Grids: A Review of Control, Fault Diagnosis, and Predictive Maintenance

Alamera Nouran Alquennah, Mohamed Trabelsi, and Naji Alqallaf

Governor Valve Actuators (GVAs) are key elements in turbine governing systems, directly affecting speed regulation, power-output control, frequency stability, and plant reliability. This paper reviews GVA systems with emphasis on control strategies, fault diagnosis, and predictive maintenance. Conventional, model-based, fault-tolerant, robust, hybrid, and AI-based approaches are compared in terms of objectives, advantages, and limitations. Recent advances in data-driven diagnostics, explainable AI, and predictive maintenance are also discussed. The review highlights the growing role of hybrid and AI-based methods in improving adaptability, fault awareness, and maintenance decision-making, while identifying key research needs, including real industrial datasets, physics-informed AI, digital-twin-enabled monitoring, and integrated control–diagnosis–maintenance frameworks.

ID023: From LLM-First to Governance-Bound Hybrid Execution: An Experimental Study of Reliability and Safety in Energy API Automation

Swapneswar Sundar Ray

Large language models (LLMs) are increasingly being connected to energy automation workflows that coordinate forecasting services, tariff engines, battery controllers, smart meters, and market-facing operational APIs. Although LLM-first orchestration improves flexibility and natural-language programmability, it can also introduce unstable contracts, schema drift, hallucinated parameters, policy violations, and unsafe execution paths in critical infrastructure contexts. This paper proposes Governance-Bound Hybrid Execution (GBHE), an architecture that separates semantic planning from deterministic control. In GBHE, an LLM generates candidate intent and action plans, while policy validation, schema normalization, runtime verification, audit logging, and rollback are handled by bounded deterministic components. We evaluate the approach using an implemented controlled failure-injection benchmark over representative energy API automation scenarios, including demand-response dispatch, battery charge scheduling, tariff-aware load shifting, telemetry repair, and device-command execution. The controlled simulation results show that, under the evaluated fixed failure schedule, the hybrid approach improves validated action completion, reduces unsafe execution, and increases recovery rates

compared with LLM-first automation, while introducing moderate latency overhead. These benchmark results are not guarantees of live-grid performance. The study argues that LLMs can be useful in energy automation when they are treated as planners inside a governed execution envelope rather than as unconstrained operators.

IoT-Enabled Green Energy Systems and Smart Infrastructure

ID026: Adaptive and Context-Aware Zero Trust Security for Scalable IoT Ecosystems

Omar Dib and Belal Ali

As IoT deployments rapidly expand in enterprise environments, traditional static verification methods for Zero Trust Security (ZTS) are no longer sufficient. This paper introduces an Adaptive Zero Trust Automation (AZTA) Framework that integrates machine learning (ML) classifiers with automated policy enforcement. The model enables security postures to adapt in real time. It includes technical tools such as dynamic decision trees and risk-scoring pseudocode to demonstrate how contextual data (e.g., device behaviour, location, and traffic patterns) can trigger automated micro-segmentation and quarantine actions. This ensures that Zero Trust implementations remain scalable and resilient while avoiding excessive administrative burdens.

ID062: Resilience Analysis of V2X Communication Protocols for IoT-Enabled Smart City Infrastructure: A Kuwait Environmental Case Study

Othman Aldaher and Yousef Alsaqabi

The transition toward automated driving requires robust Vehicle-to-Everything (V2X) communication to support both low-latency Basic Safety Messages (BSMs) and high-capacity LiDAR payloads. However, V2X performance is highly sensitive to extreme environmental stressors. This study presents a cross layer analytical framework to evaluate the Physical (PHY) and Medium Access Control (MAC) performance of IEEE 802.11p DSRC, 5G C-V2X, and a prospective multi-band 6G architecture under severe desert conditions, specifically modeling 55°C ambient heat and dense sandstorms. Grounded in empirical atmospheric attenuation data, the MATLAB-based model calculates theoretical bounds for Packet Delivery Ratio (PDR), Air-Interface latency, and data rates. The results demonstrate a fundamental trade-off between propagation resilience and channel capacity. While lower-frequency links (5.9 GHz) penetrate dust effectively via Rayleigh scattering, their narrow bandwidths cannot support massive sensor payloads. Conversely, high-frequency mmWave and sub-THz bands provide necessary gigabit capacities but experience severe Mie scattering, hitting a physical “dust wall” that reduces their effective range to under 75 meters during severe weather. The analysis indicates that future autonomous vehicular networks cannot rely on a universal technological solution; instead, they require adaptive, multi-band architectures and advanced multiple-access schemes to maintain reliability in harsh environments.

ID074: MCP-Driven Conversational AI Agent for Real-Time IIoT Energy Budget Control in Industrial Environments

Omar Dib, Tarek Adi, Abdullah Shamout, Ahmed Amin, and Sulaiman Alrefai

Abstract—Industrial facilities running Industrial Internet of Things (IIoT) networks often blow past their energy budgets because their monitoring systems are stuck in the past. They rely on fragmented data, fixed thresholds, and device- level alerts that don’t understand what’s happening across the wider system. That’s the problem. This paper presents a conversational AI agent that combines GPT-4o with the Model Context Protocol (MCP), a standardised tool-invocation layer introduced by Anthropic in 2024. The agent watches real-time IIoT telemetry, detects weekly energy budget violations, and generates conditional ptimisation plans. But it doesn’t get to touch the hardware on its own. Every proposed action needs mandatory human approval first, which is exactly how industrial AI control should work. The prototype uses a Raspberry Pi 4 as the IIoT gateway, connected to a PZEM-017 DC power sensor and Athom 16A smart plugs. A Flask backend coordinates the system, TimescaleDB stores the time-series data, and a React.js dashboard gives operators a clear view of what's happening. It gets messy. Industrial energy data always does. But this setup shows that an LLM-based agent can make sense of it without pretending the human operator is optional. Validation across ten structured test cases showed more than 90% prompt-response accuracy, 100% MCP command execution success, and steady 15-second telemetry sampling with zero data loss. That matters. These results suggest MCP-standardised LLM-to-IIoT control isn't just an interesting lab idea. It's a practical architecture for industrial energy management, especially in facilities that need smarter decisions without handing full control to automation.

ID090: AgroSense: An Automated Smart Hydroponic System with Edge AI Diagnostics for Arid Regions

Othman Aldaher, Muhammad Sohaib, and Omar Hammuda

This paper presents AgroSense, an automated, closed-loop Deep-Water Culture (DWC) hydroponic system that integrates Internet of Things (IoT) sensing, edge artificial intelligence, cloud-based monitoring, and targeted crop treatment through automated medicinal delivery for arid environments. The system employs a distributed multi-microcontroller architecture comprising ESP32 controllers and an Arduino UNO Q vision coprocessor connected to a Firebase Realtime Database. It continuously monitors pH, TDS estimate (ppm), water level, temperature, and humidity, while closed-loop dosing algorithms regulate reservoir conditions with minimal manual intervention. To mitigate electrical interference between the pH and TDS sensors, a time-division power-isolation method was implemented, reducing the measured pH variance from 0.2041 to 0.00063. For crop-health monitoring, INT8-quantized MobileNetV2 models were deployed at the edge to classify crop growth stages and visible health conditions. The models achieved held-out test accuracies of 97.4% for strawberry classification and 91.33% for mint classification. Detections above the confidence threshold trigger the spray actuator with a cooldown safeguard; prototype testing verified actuation using a milk–water test fluid, while treatment chemicals, dosage, efficacy, and phytotoxicity remain future work. Experimental evaluation further demonstrated automated TDS recovery from 400 ppm to 685 ppm within two dosing cycles and 20 minutes, without exceeding the target operating range. The results demonstrate the feasibility of integrating environmental regulation, edge-based crop diagnostics, and localized intervention within a modular hydroponic platform designed for water scarce and resource-constrained regions.

ID097: A Cloud–Edge PM2.5 Forecasting

Huda Alrashidi and Ghita Berrada

Air pollution is a critical public-health challenge in Kuwait, where PM?.? concentrations routinely exceed World Health Organization thresholds due to dust storms, dense traffic, and oil-based industry. Existing monitoring infrastructure records historical measurements but lacks predictive capability and an operational pathway from raw sensor data to actionable public-health decisions. This paper makes two distinct contributions, evaluated separately. First, we propose a five-layer cloud–edge system architecture for smart environmental monitoring in Kuwait; this architecture is conceptual and is presented as a deployment blueprint, not an implemented system. Second, we develop and experimentally validate the forecasting engine intended for the architecture's cloud AI layer: a Long Short-Term Memory (LSTM) network that predicts the next-hour PM?.? Air Quality Index (AQI) from a 24-hour historical window of multivariate pollutant and meteorological observations. Only this forecasting engine is trained and evaluated in this paper. It is evaluated on 409,267 hourly records collected from 14 Kuwait Environmental Public Authority (EPA) monitoring stations across six governorates (2018–2022), with a multi-stage gap-aware imputation pipeline handling severe real-world missingness (overall missing rate 32.95%). Under a strict chronological split, the LSTM model achieves a scaled MAE of 0.452, RMSE of 0.660, and R^2 of 0.592, demonstrating that deep sequential modeling can capture the daily and seasonal dynamics of dust-driven Gulf pollution. These results validate the forecasting core of the proposed framework; the surrounding cloud–edge architecture remains a design proposed for future operational deployment and has not been implemented or benchmarked in this work.

ID110: Ergodic Capacity Analysis of Coordinated NOMA Transmission for LEO Satellite IoT Networks

Anas Elsayed, Mohamed Elsayed, Ahmed Samir, Esraa M. Eid, Hala M. Abd Elkader, and Basem M. Elhalawany

Low Earth Orbit (LEO) constellations increasingly serve Internet-of-Things (IoT) terminals while reusing the same frequency bands across neighbouring beams, which raises the question of how to serve those terminals efficiently where beams overlap. We study the downlink ergodic capacity (EC) of two coordinated LEO satellites that jointly serve three terrestrial IoT terminals using power-domain non-orthogonal multiple access (NOMA) under Shadowed-Rician fading. Closed-form series expressions are derived for the EC of the two near users. The shared far user lies in the beam-overlap region and is served by both satellites through selection combining (SC); For the shared far user, a semi-analytical ergodic capacity expression is derived together with a maximum-ratio-combining (MRC) upper bound. We identify an irreducible SINR ceiling that limits the shared user's capacity, extend the analysis to imperfect SIC, and formulate a power-allocation algorithm that maximises the sum capacity subject to per-user rate constraints. Monte Carlo simulation confirms the analysis. Finally, comparisons with a benchmark OMA-based system demonstrate that coordinated NOMA consistently achieves superior ergodic-capacity performance across the entire SNR range. Furthermore, SC achieves nearly the same diversity performance as MRC while avoiding the inter-satellite phase synchronization that low-cost IoT terminals cannot support.

ID112: Compute-Aware Adaptive Intrusion Detection for ATM Edge Networks

Anfal Almutairi, Aalyah Alotaibi, Jumanah Alfailakawi, Khaled Allanqawi, and Basil Al Othman

Automated teller machines (ATMs) are unattended IoT edge devices that require continuous intrusion monitoring under constrained compute budgets. We present ATM-EAIDS, a confidence-gated cascade in which logistic regression screens every event and uncertain samples are escalated to a Stage-2 detector (the heavier detector invoked after escalation). CICIoT2023 is used as a behavioral proxy, not ATM telemetry. The corrected evaluation supersedes the originally submitted numerical results, which could not be reproduced because the original subset-construction pipeline was not fully preserved. The revised evaluation uses validation-only recall-constrained margin selection, measured gate calibration, 100 timing repetitions after warm-up, leakage-oriented split stress tests, and a 30-seed descriptive stability analysis. At the legacy margin m = 0.15, retained only as an a priori comparison point, the corrected Random-Forest cascade reaches F1 0.686 with FNR 0.440 and normalized modeled cost 0.184, versus Random Forest F1 0.772 and FNR 0.314. An unpruned Decision Tree reaches F1 0.719 and FNR 0.284 at normalized modeled cost 0.007. At the legacy margin, the cascade has higher precision and lower false-positive rate than the Decision Tree, but higher FNR and modeled compute cost. Electrical energy was not measured and the physical kiosk was not used experimentally; the results establish a resource-aware proof of concept for smart-infrastructure edge intrusion monitoring, not production ATM readiness.

ID130: Explainable Short-Term Photovoltaic Power Forecasting for PAIoT Energy Management

Sara Salman and Zied Bouida

Photovoltaic agricultural Internet of Things (PAIoT) systems face critical energy-management challenges due to solar intermittency and remote deployment constraints. This paper presents an explainable short-term photovoltaic (PV) forecasting framework engineered as a core edge-intelligence and decision-support component for PAIoT infrastructure. Utilizing a one-year, one-minute SolarTechLAB dataset containing 525,601 records, XGBoost is benchmarked against a persistence baseline at 15- and 60-minute horizons. At the 60-minute horizon, XGBoost reduces mean absolute error (MAE) and root mean square error (RMSE) by 41.6% and 34.2%, respectively, providing the anticipatory foresight required for predictive energy management. SHAP (SHapley Additive exPlanations) analysis identifies cyclical time features, current power, and short-term lags as dominant drivers, offering transparent interpretability for downstream IoT resource scheduling, such as battery dispatch and precision irrigation. A subsequent 40-trial Optuna search improved the 60-minute model to MAE 17.72 W, RMSE 27.39 W, and R^2 0.811, corresponding to 44.4% and 35.8% reductions in MAE and RMSE relative to persistence.

ID030: Energy-Aware Routing in Software-Defined Optical Networks via Reinforcement Learning

Skander Abid, Mourad Menif, and Latifa Guesmi

Energy-efficient optical transport is essential for green AI-enabled infrastructure, yet classical shortest-path routing ignores energy and static consolidation heuristics can fail under non-uniform demand. This paper presents a reproducible packet-level framework for energy-aware routing in software-defined DWDM optical networks. The framework combines a lit/dark per-channel energy model, an SDN controller, a proactive energy-aware baseline, and a tabular Q-learning policy on the 14-node National Science Foundation (NSF) reference topology. Energy is computed from actual packet trajectories, so savings are directly tied to the number of wavelength channels kept dark. Across five traffic scenarios, four routing policies, and five independent random seeds, Q-learning matches the heuristic on uniform workloads, reducing energy by 14.3% relative to shortest-path at unit Packet Delivery Ratio (PDR). Under hotspot traffic, the learned policy adapts to the many-to-one structure, recovers a compact 14-channel routing tree, and improves energy by 12.5% relative to the heuristic while maintaining 0.999 PDR.

ID131: IoT-Enabled AI Fusion for Early-Onset Alzheimer's Disease: A Survey and Preliminary Analysis

Joicy John and Zied Bouida

Early-Onset Alzheimer's Disease (EOAD), defined by symptom onset before age 65, is a clinically and biologically distinct form of Alzheimer's disease characterized by heterogeneous and rapid progression. Unlike late-onset forms, EOAD commonly presents with non-amnestic symptoms such as executive dysfunction, language impairment, and behavioral changes which frequently precede memory decline and lead to misdiagnosis. While advances in neuroimaging and cerebrospinal fluid biomarkers have improved pathological detection, traditional episodic clinical assessments often fail to capture the subtle, daily micro-fluctuations of cognitive decline. To address this, this paper presents a comprehensive survey of EOAD mecha-nisms and proposes a multidimensional edge-cloud Internet of Things (IoT) architecture. The proposed framework features an Artificial Intelligence (AI) Data Fusion Layer designed to synthesize continuous biometric data from wearable sensors with static clinical, genetic, and neuroimaging profiles. To validate the clinical baseline of this architecture, preliminary machine learning experiments were conducted using an open-source multidimensional dataset. A Random Forest classifier achieved an overall predictive accuracy of 87.78% across discrete diagnostic categories, demonstrating an exceptional recall rate of 0.90 for identifying transitional Mild Cognitive Impairment (MCI). Feature importance analysis confirms that integrating cognitive assessments with structural MRI volumes provides a robust predictive foundation. Ultimately, this research highlights how IoT-enabled AI fusion can shift EOAD management from reactive clinical evaluations to predictive and personalized healthcare.

Special Session 1: AI-Driven Smart Materials and Sensing Systems for Advanced Energy Technologies

ID045: Smart NanomaterialEnhanced ZincRich Coatings for Sustainable Corrosion Resistance in Energy Applications

Abeer Alfarhan and Abdulmuhsen Akbar

Corrosion and cracking of carbon steel components remain a critical challenge in the oil sector, particularly in electrical submersible pumps used for water injection and pressure enhancement. Conventional epoxy coatings have provided limited protection, with failures occurring within a year under high temperature, high-pressure sweet corrosion environments containing CO? and saline water. This study introduces a smart nanomaterial approach by integrating zinc-rich epoxy coatings with eco-friendly, water-soluble inhibitors to enhance durability and sustainability. The synergistic effect of zinc particles and nanomaterial-based inhibitors was systematically evaluated using computerized corrosion monitoring methods, electrochemical impedance spectroscopy (EIS), autoclave testing, and salt spray exposure. Results demonstrate that combining a zinc-rich epoxy matrix with 1000 ppm of selected inhibitor significantly

reduces corrosion rates, extending service life under aggressive operating conditions. The findings highlight the potential of smart nanomaterial coatings as cost-effective, environmentally responsible solutions for protecting energy infrastructure, ensuring operational reliability, and advancing next generation materials for oilfield applications.

ID038: A Hybrid CNN–XGBoost for EEG-Based Alzheimer's Monitoring and Energy Prediction in next generation intelligent biomedical implants.

Yosra Ben Fadhel, Kamal Al-Haddad, Olfa Bel Hadj Brahim Kechiche, and Mahmoud Hamouda

Alzheimer's disease is a progressive type of neurodegenerative disease that demands constant monitoring and effective assessment of brain activities for an early diagnosis and treatment of patients. On the other hand, intelligent biomedical implant devices should consume less energy during operation but ensure accurate signal processing and wireless communication at the same time. In this paper, we develop a hybrid artificial intelligence framework involving a one-dimensional convolutional neural network (CNN) together with extreme gradient boosting (XGBoost) for Alzheimer's monitoring based on EEG data and prediction of the energy consumption during wireless communication. Our framework includes advanced electroencephalography (EEG) signal pre-processing, feature engineering, normalization, and multimodal feature extraction to utilize both neurological biomarkers and communication parameters. CNN automatically extracts discriminative features from EEG and additional system-level features, whereas XGBoost conducts accurate nonlinear regression of the energy consumption. As shown by our experimental results, the proposed algorithm is highly effective, providing MAE=0.021, RMSE=0.037, and R^2 = 0.972. Comparative analysis demonstrates the superiority of the proposed hybrid approach compared to pure CNN and XGBoost algorithms. Finally, residual and feature importance studies further validate the effectiveness and robustness of the proposed solution.

ID065: Physical properties of CuMgSnS Powders and Thin Films

Oumaima Hammami, Mounir Kanzari, and Naoufel Khemiri

Copper, magnesium, and tin sulfide Cu_2MgSnS_4; (CMTS) is being studied as a non-toxic, naturally abundant quaternary semiconductor absorber intended to replace conventional zinc kesterite-based solar cells. As part of this work, polycrystalline CMTS compounds were synthesized via solid-state reaction under vacuum, and the corresponding thin films were processed and then subjected to post-sulfurization annealing at 250 ° C. X-ray diffraction (XRD) and micro-Raman spectroscopy confirm successful crystallization in the tetragonal kesterite phase (space group I4 ?) with strong (112) preferential orientation and a highly resolved local Mg-S lattice vibration mode at 285 cm -1 (E symmetry). Thermogravimetric analysis and differential thermal analysis (TGA/DTA) demonstrate solid thermal stability below 300 °C, with an exothermic peak associated with matrix crystallization identified at 721 °C. Diffuse reflectance spectroscopy yields an ideal direct optical bandgap of 1.54 eV. Finally, mapping of the surface topography using AFM and SEM reveals that, while substitution with magnesium effectively eliminates the high density of pits and porosity inherent in basic kesterites, it induces localized mechanical stress in the crystal lattice, which is released through the formation of microcracks on the surface during cooling. These results provide fundamental insights for the design of defect-free, stress-controlled CMTS solar energy conversion layers intended to optimize operational performance.

ID081: Effect of a Multi-stage Cooling System on the Performance of a Mini Photovoltaic Panel

Marwa Hamza, Olfa Bel Hadj Brahim Kechiche, and Habib Sammouda

Abstract—This paper presents an experimental investigation of a multi-stage thermoelectric cooling system designed to improve the thermal regulation and electrical performance of a mini photovoltaic (PV) panel. The proposed experimental setup integrates three successive cooling stages consisting of a thermoelectric (TEC) module, a CPU heat sink, and an electric fan mounted beneath the mini-PV panel. The cooling system was experimentally evaluated by progressively integrating each cooling stage while maintaining a constant distance of 35 cm between the light source and the PV module. During the experiments, the operating temperatures of the

mini-PV panel and the TEC module, together with the electrical characteristics of the PV module, were monitored to assess the effect of each cooling stage. The obtained results show that the operating temperature of the mini-PV panel has a significant influence on its electrical performance. The first and second cooling stages progressively reduced the panel temperature from 31.5 °C to 27.8 °C, increasing the maximum output power from 52.21 mW to 53.53 mW and the conversion efficiency from 6.02% to 6.18%. The integration of the third cooling stage further stabilized the panel temperature at approximately 25 °C, resulting in the highest electrical performance, with a maximum output power of 59.34 mW and a conversion efficiency of 6.85%. These results demonstrate the effectiveness of the proposed multi-stage cooling system in enhancing the thermal stability and electrical performance of mini photovoltaic modules.

ID105: Compositional Engineering of CuO-SnO2 Thin Films: Influence on Structural, Optical, Morphological and Electrical Properties

Nourhéne Mansouri and Bitri Nabila

This study investigates the synthesis and characterization of pure CuO, pure SnO2, and CuO–SnO2 mixed oxide thin films deposited on glass substrates by a low-cost spray pyrolysis technique. By varying the CuO–SnO2 molar ratios (75%:25%, 50%:50%, and 25%:75%) and applying post-deposition annealing at 500 °C, thin films with tunable structural, morphological, optical, and electrical properties are obtained. X-ray diffraction analysis confirms the coexistence of monoclinic CuO and tetragonal rutile SnO2 phases in the mixed films without the formation of secondary phases, indicating the successful incorporation of both oxides within the thin-film matrix. Optical investigations reveal a strong dependence of the optical response on composition, with the optical band gap increasing progressively from 1.76 eV for pure CuO to 3.71 eV for pure SnO2. The mixed films exhibit intermediate band gap values, demonstrating the possibility of tailoring the electronic structure through compositional control. Electrical characterization performed using the four-point probe technique shows a gradual decrease in resistivity from 9.19×10?³ ?.cm for pure CuO to 0.79×10?³ ?.cm for pure SnO2, highlighting the significant influence of SnO2 content on charge transport behavior. Morphological analysis further reveals notable changes in surface organization, grain distribution, and film compactness with varying composition, while thickness measurements indicate enhanced film growth in the mixed oxide systems compared to the corresponding pure oxides. Overall, the results demonstrate a strong correlation between composition and the physical properties of CuO–SnO2 thin films, confirming that compositional engineering provides an effective approach for tuning their structural, optical, and electrical characteristics.

ID106: Oblique angle of deposition enhanced the photo-response performance of CuSbS2 thin films

Mouna Idoudi, Ferid Chaffar Akkari, and Mounir Kanzari

This study investigated the amorphous structure of CuSbS? thin films and found that their morphology is strongly dependent on the tilt angle of the substrate (0°,85°). SEM analysis reveals the formation of inclined columnar structures accompanied by increased porosity and reduced film thickness as the deposition angle increases. The optical bandgap varies from 1.68 to 2.21 eV and the films predominantly exhibit p-type conductivity. The I–V characteristics exhibit quasi-linear behaviour, confirming ohmic contact and a stable photoelectric response within the voltage range of ?2 V to +2 V. Time-resolved measurements under periodic light illumination demonstrate reproducible photoconductive behaviour characterised by cyclic resistance variations. Furthermore, photodetection performance significantly improves with increasing deposition angle, sensitivity reaches 76.09%, and spectral responsivity increases to 18.31 ??¹·W?¹·m?² at 85°. Additionally, response and recovery times are substantially reduced at high tilt angles, reaching 7.87 s and 20.40 s respectively, confirming the suitability of these thin films for efficient photodetector applications.

ID119: Spray-Coated CuO/rGO Hybrid Thin Films for Room-Temperature H2S Gas Sensing

Imen Hemmedi and Bitri Nabila

This study investigates the fabrication, structural characterization, electrical properties, and room-temperature hydrogen sulfide (H2S) sensing performance of spray-coated CuO/reduced graphene oxide (rGO) hybrid thin films. Chemically reduced graphene oxide, prepared from graphene oxide using hydrazine hydrate, is combined with CuO at different volume ratios to examine the influence of hybrid composition on the properties of the deposited films. Grazing-incidence X-ray diffraction confirms the coexistence of monoclinic CuO and rGO without detectable secondary phases, while electrical measurements show that increasing the rGO content progressively enhances the electrical conductivity of the films. The gas-sensing capability is evaluated toward H2S concentrations of 5, 10, and 20 ppm at room temperature. All sensing films exhibit reversible resistance changes, with the response magnitude increasing systematically with both H2S concentration and rGO content. Among the investigated compositions, the CuO/rGO (25/75) thin film exhibits the highest response of 56.57 % toward 20 ppm H2S, demonstrating the beneficial synergistic effect between CuO and rGO. These results highlight spray-coated CuO/rGO hybrid thin films as promising candidates for simple, scalable, and low-cost room-temperature H2S sensing.

ID120: Deposition Time-Controlled Structural, Optical, Electrical, and Carrier Transport of RF Magnetron Sputtered ZnO-SnO Thin Films

Emna Ben Sassi and Fatma Chaabouni

The properties of RF magnetron-sputtered ZnO-SnO2 thin films are strongly influenced by the deposition conditions. This work investigates the effect of deposition time (30, 60, and 120 min) on the structural, optical, and electrical properties of ZnO-SnO2 thin films and the corresponding Ag/ZnO-SnO2/p-Si/Al heterojunctions. Increasing the deposition time progressively enhances the crystalline quality of the films, as evidenced by sharper diffraction peaks, larger crystallite size, and lower microstrain and dislocation density, with diffraction features attributed to the ZnSnO3 phase observed under the investigated deposition conditions. The films maintain an average visible transmittance of approximately 75%, while the optical band gap decreases from 3.59 to 3.46 eV. Hall-effect measurements reveal an increase in carrier concentration and a decrease in electrical resistivity with increasing deposition time. The fabricated heterojunctions exhibit a decrease in series resistance from 9.60 to 4.37 ? and in ideality factor from 1.60 to 1.03. The effective minority carrier lifetime increases from 2.19 ? 0.09 ?s at 30 min 2.84 ? 0.50 ?s at 60 min and remains nearly unchanged at 2.85 ? 0.49 ?s for 120 min. These correlated structural and electrical changes show that deposition time is an effective processing parameter for controlling the properties of the films and the electrical characteristics of the resulting heterojunctions. Compared with previously reported ZnO/p-Si heterojunctions, the 120 min device shows favorable diode characteristics. The film deposited for 120 min exhibits the best overall combination of crystalline quality, electrical transport, and heterojunction performance, making these structures of interest for oxide/Si-based optoelectronic applications.

Special Session 2: Real-Time Embedded Artificial Intelligence Systems

ID052: LLM-Powered In-Vehicle Assistants: IoT Architecture and Case Study

Mai Alotaibi, Sameerah Alajmi, Aliah Alsulail, Retaj Alsaeedi, and Basem Elhalawany

Large Language Models (LLMs) transform the rigid command-response nature of in-vehicle Assistants (IVAs) into natural and context-aware conversational agents. However, most LLM-based IVAs operate with zero vehicle coupling. That means that IVAs respond to generic queries without access to the actual live sensor data of the vehicle. This paper addresses this gap by providing a three-axis IoT taxonomy (deployment model, vehicle coupling depth, and interaction modality). Additionally, a comparative survey of several systems spanning the period from 2020 to 2025 is provided. Moreover, we present AutoMate as a bilingual Flutter/GPT-4 IVA

developed at Kuwait College of Science and Technology that fuses real-time OBD-II telemetry with natural language interaction. The survey results confirm the coupling deficit as the most significant open problem, alongside other challenges related to latency, privacy, offline resilience, and multilingual support.

ID070: Real-Time Selective Harmonic Mitigation in Multilevel Inverters Based on Gradient-Based Search Algorithm and Artificial Neural Network

Noman Liaqat, Amir Ghasemian, Sobhan Mohamadian, Maria Gabriella Cimoroni, Concettina Buccella, and Carlo Cecati

This paper presents an Artificial Neural Network (ANN)-based approach for real-time implementation of Selective Harmonic Mitigation (SHM) using a Gradient-Based Optimization (GBO) framework. Unlike conventional optimization-based SHM methods, the proposed formulation treats the fundamental frequency equation as a constraint rather than an optimization objective, thereby reducing the search space dimension and significantly improving convergence speed. This enables efficient generation of training data within a feasible computational time. The trained ANN is capable of predicting switching angles for varying DC input voltages and modulation indices, overcoming the limitations of existing methods that are typically restricted to fixed operating conditions. The proposed approach is validated through both simulation and FPGA-based hardware implementation, demonstrating real-time performance. The results show strong agreement between simulation and experimental data, confirming effective harmonic mitigation.

ID061: Hardware acceleration in FPGA-SoC of a lightweight neural network for partition optimization in the ECM encoder

Sourour Soltani, Amna Maraoui, and Imen Werda

This article describes the complete Field Programmable Gate Array-System on a Chip (FPGA-SoC) implementation of a lightweight neural network (LNN) accelerator designed to reduce the computational complexity of the block partitioning module in a post-VVC(Versatile Video Coding)video encoder with an enhanced compression model (ECM). The LNN anticipates and deletes redundant ternary tree (TT) partitioning evaluations during intra block encoding, so reducing complexity.The average encoding time is reduced by 40.14 % with negligible impact on compression efficiency and video quality . The accelerator hardware was developed in C++ using Vivado HLS 2019.1 and implemented on a Xilinx Zynq-7020 (PYNQ-Z2) SoC using an integrated hardware and software design process. Key optimization techniques such as segmentation, loop unrolling, and full-array partitioning reduce inference latency from 1.871 to as low as 156 clock cycles at 105 MHz, using 10.18 % of the LUTs and 23.18 % of the DSPs, confirming the feasibility of the integrated real-time implementation.

ID113: An IoT-Enabled and Machine Learning-Supported Framework for Irrigation Decision Support

Ghaliah Alhajraf and Ismaeel Al Ridhawi

Abstract—Efficient irrigation is important for sustainable agri- culture, especially in regions with limited water resources and harsh weather conditions. . This paper presents a framework for predicting irrigation needs in the immediate future using three classes, namely, None, Low, and High, as well as a binary version. Twenty-seven features are generated from current and previous soil-moisture readings, moisture changes, recent irrigation ac- tivity, and time information. Seven ML models are compared using the same Synthetic Minority Over-sampling Technique (SMOTE)-based training pipeline. The models are evaluated using stratified five-fold cross-validation and a chronological hold-out test that represents prediction on later data. XGBoost achieved the best three-class cross-validation performance, with a macro-F1 score of 0.891 ± 0.008 and an accuracy of 90.8%. However, macro-F1 decreased to 0.39–0.47 under chronological testing, showing that model performance can decrease as con- ditions change over time. Feature analysis showed that recent irrigation activity and soil-moisture history were the most useful inputs. A working prototype was also developed using an ESP32, Raspberry Pi 5, environmental sensors, online weather data, a Random Forest model, a monitoring dashboard, and automatic pump control. The results highlight the value of tree-based models and the importance of time-based

evaluation for practical smart irrigation systems. Index Terms—Internet of Things, smart irrigation, machine learning, XGBoost, Random Forest, soil-moisture sensing.

ID035: A Smart Solar Panel Cleaning System Using Machine Learning and Embedded Systems

Milia Habib, Maria Tedros, Saraa El Hazzouri, Zaher Merhi, Abdelrazzak Merheb, and Mohamad Kanaan

The efficiency of the solar panels can be compromised due to dust particles accumulation resulting in higher costs for maintenance. The traditional methods for cleaning the solar panel involve manual processes that require lots of water and are inefficient. This paper presents a smart automated solar panel cleaning system with a use of machine learning algorithms and embedded IoT devices. This system utilizes a Logitech C270 HD camera linked to the Raspberry Pi 4 and captures live images of the surface of the solar panels. A Convolutional Neural Network (CNN) model is trained on a labeled dataset of images to classify the state of the panel surface as either dusty or clean. Once dust is detected, the Raspberry Pi 4 provides a signal to the ESP32 board, which then manages the relay that activates the water pumping system and mechanical cleaning system. Standards metrics such as accuracy, precision, recall, and F1-score are used to measure the performance, achieving values of 94%, 95%, 94%, and 94%, respectively. The presented solution demonstrates the feasibility of an embedded, condition-based approach for intelligent automated solar panel cleaning systems.

ID076: A Lightweight Embedded AI-Based UAV System for Real-Time Diseased Tree Detection

Bouthaina Abdallah, Fatma Belghith, and Nouri Masmoudi

Early detection of diseased trees is essential for precision agriculture and targeted crop management. This paper presents a lightweight drone-based system for automatic diseased-tree detection and geolocation using embedded computer vision and GPS technologies. The proposed platform integrates a Raspberry Pi 4 and a YOLO11s deep-learning model for onboard real-time processing. Experimental results show that the system achieves a mAP@50 of 81.5\% while processing 6–7 frames per second without external computing resources. The prototype also satisfies the main design requirements, including portability and a flight autonomy exceeding 8 minutes. The obtained results demonstrate the feasibility of low-cost embedded artificial intelligence for real-time agricultural monitoring and precise localization of diseased trees. This work highlights the potential of combining embedded AI, UAVs, and precision agriculture to support sustainable and efficient orchard management.

ID078: Toward Edge-AI Precision Livestock Farming: A Low-Power Wearable LoRa Sensor Node for Heat Stress and Activity Monitoring

Usama Mokashi, Wesam Rohouma, Rachid Benlamri, Shady Khalil, Amith Khandakar, Arselene Ayari, Mujtaba Fida Ul Haq, Reza Alam, and Ajna Shukkoor

Undetected heat stress and subtle activity changes in cattle are linked to reduced productivity and increased veterinary costs, motivating continuous, low-cost individual-animal monitoring. This paper presents a collar-mounted IoT wearable that inte-grates, on a single custom PCB, an ESP32-C6 microcontroller, a BME280 sensor for heat-stress index estimation, a BMA400 accelerometer for activity monitoring, and an E22-900T22D LoRa (Long Range) transceiver. Unlike existing LoRa-based collars, which typically address either environmental or activity sensing in isolation and report sleep currents in the milliampere range, the proposed node combines both sensing modalities with a duty-cycled firmware architecture achieving a measured deep-sleep current of ~25 μA (microamperes). The firmware acquires fixed-length accelerometer windows (125 samples at 25 Hz) and on-demand environmental readings, transmitting chunked LoRa packets before returning to deep sleep. Power profiling with a Nordic PPK2 confirms transmission peaks of ~155 mA and projects a battery life of 150–160 days on a 2500 mAh cell at a 30-minute duty cycle, meeting the 5–6 month target for unattended deployment in arid desert dairy condi-tions. These results demonstrate that integrated dual-modality sensing and sub-30μA sleep current are jointly achievable on low-cost, off-the-shelf hardware a combination not previously reported in custom-PCB LoRa collar designs.

ID098: Real-Time Fall and Crash Detection for Electric Micromobility Using IMU Sensor Fusion and Edge-Deployed Machine Learning

Asel Rohouma, Mohamed Rohouma, Ajna Shukkoor, and Shady Khalil

Electric micromobility vehicles such as electric scooters (e-scooters) and e-bikes have seen rapid global adoption, yet rider safety remains a critical challenge. This paper presents a real-time fall and crash detection system that fuses multi-axis inertial measurement unit (IMU) data with a lightweight edge-deployed Decision tree classifier to identify fall and collision events. The hardware platform comprises an MPU6050 six-axis IMU connected to an ESP32 microcontroller and sampled at 50 Hz. Tri-axial acceleration data (ax, ay, az) are collected and labelled across three classes: Normal, Fall, and Crash. A windowed statistical feature extraction stage feeds a decision tree classifier that achieves 86.4% cross-validated test accuracy under a leave-one-crash-event-out evaluation designed to prevent data leakage; a Random Forest ensemble evaluated for comparison reaches 93.2% under the same protocol. An onboard alert module signals detected events via LED and serial notification, while a local SD card logs timestamped incident records. Results from controlled simulation trials confirm the system's reliability across diverse riding maneuvers and validate its suitability for autonomous incident detection on personal electric vehicles.

Artificial Intelligence for Automation of Energy Systems

ID036: Predictive Climate-Aware Intelligent Energy Management System for Fuel Cell Hybrid Electric Vehicles Under Kuwait Extreme Summer Conditions

Ghaneemah Alasfour, Wael Farag, and Zakariya Dalala

Energy management for fuel cell hybrid electric vehicles (FCHEVs) is almost always developed at nominal ambient temperature, leaving the behaviour of supervisory controllers under sustained extreme heat largely uncharacterised. This paper quantifies that gap for hot-climate operation, taking Kuwait as a representative case, and proposes a lightweight climate-aware intelligent energy management system (IEMS). Over a real GPS-derived Kuwait Drive Cycle, an ambient temperature of 45°C raises the HVAC auxiliary load to 4.69 kW and increases battery state-of-charge (SOC) depletion 2.7× relative to a 20°C European baseline (3.98% to 10.82%), while an Arrhenius-form thermal stress index rises 2.30×. The proposed IEMS pairs an interpretable five-node Decision Tree strategy selector with a four-mode supervisory controller whose Kuwait-Hot mode transfers the HVAC burden from the battery to the fuel cell. Comparative benchmarking against a rule-based baseline shows that the proposed controller sustains proactive Kuwait-Hot operation for the full cycle where the baseline reacts only after SOC has fallen below 50%, at a cost of 35% higher hydrogen consumption (1270 g versus 940 g) and with no measurable SOC or stress-index advantage in the present lumped thermal model. That null result is reported directly, and its cause traced to the thermal model's structure, rather than minimised. The classifier splits on ambient temperature alone at 24.99°C and 39.93°C, reaching 99.50% ± 0.68% cross-validated accuracy in 26.5 ?s and 3.59 KB as measured in MATLAB.

ID054: AI-Driven Identification of Compound Climate Hazard Pathways: Enhancing Energy Efficiency and Operational Resilience in Water Infrastructure

Renalda El-Samra

Climate Risk and Vulnerability Assessments (CRVA) are widely used to evaluate the impacts of climate change on water infrastructure. However, conventional approaches typically assess hazards individually and may overlook compound interactions that amplify infrastructure vulnerability. This study proposes a hybrid artificial intelligence (AI) framework for identifying compound climate hazard pathways affecting critical water infrastructure. The framework is applied to the Basalt Wellfield (BWF), Jordan, a strategic groundwater supply system located in one of the world's most water-stressed regions. By mitigating climate-induced stressors on pumping operations, the framework further supports the optimization of energy efficiency in water distribution systems. Climate-risk

indicators derived from the BWF CRVA were analyzed using Principal Component Analysis (PCA) and unsupervised hierarchical clustering. PCA was used to identify the dominant drivers of climate vulnerability and reduce the dimensionality of the hazard space, while hierarchical clustering was employed to classify recurring patterns of interacting climate hazards. The analysis revealed two dominant climate-risk dimensions associated with groundwater stress and hydro-meteorological infrastructure stress, which were subsequently confirmed through the identification of distinct compound climate hazard archetypes. The results demonstrate the potential of AI techniques to uncover hidden climate-risk structures, support targeted adaptation planning, and enhance infrastructure resilience. By translating complex climate-risk data into actionable decision-support information, the proposed framework enables infrastructure managers to prioritize adaptation investments, optimize operational planning, and improve energy-efficient system management under increasing climate uncertainty. The methodology provides a transferable foundation for smart water infrastructure and climate-informed decision-making.

ID072: A Lightweigh Reproducible Multimodal Surrogate Learning Framework for EL-Free Photovoltaic Degradation Estimation from HDR Images

Marcello Polenghi, Roberto Caldelli, Emanuele Ogliari, and Valsamis Ntouskos

This work proposes a reproducible, computationally efficient, surrogate modeling framework for photovoltaic (PV) degradation estimation from visible-domain HDR imagery, reducing the need for electroluminescence (EL) acquisition, which is accurate but costly and difficult to scale. The framework learns synthetic EL representations from HDR images through supervised image-to-image translation and exploits them as intermediate degradation-related descriptors rather than final reconstruction targets. Compact statistical features extracted from the generated EL images are then used for supervised power loss regression. The pipeline is trained and evaluated on a publicly available dataset containing a limited number of paired HDR images, EL images, and synchronized degradation measurements. Performance is assessed using both image reconstruction metrics and downstream regression criteria, including prediction errors and rank-based correlations, to evaluate the feasibility of the proposed surrogate modeling framework under the available data regime. Experimental results show that the proposed method consistently outperforms a naive baseline and approaches the performance of an oracle model, using real EL images within the investigated dataset while preserving good degradation ranking consistency. These findings suggest that synthetic EL representations may serve as surrogate descriptors for physical inference under the investigated experimental conditions, enabling computationally efficient and reproducible EL-free PV degradation estimation from HDR imagery.

ID132: Comparing Large Language Models and Machine-Learning Methods for Short-Term Electricity Price Forecasting

Seyed Ali Hosseini, Paria Valizadeh, Sven Tomforde, and Francesco Grimaccia

Very-short-term electricity-price forecasting increasingly uses both conventional models and large language models (LLMs), although their accuracy, resource demands, and numerical sensitivity are rarely compared under a common protocol. This study evaluates persistence, ARIMA, XGBoost, LSTM, and three LLM families for 30-minute-ahead German electricity-price regression and direction classification using rolling-origin evaluation. The comparison also considers monetary cost, energy use, carbon emissions, zero-shot and few-shot prompting, and a counterfactual price-level shift. Few-shot Claude Haiku achieved the lowest regression error but improved only slightly over persistence, while few-shot GPT-5-mini provided the most balanced classification performance. Few-shot prompting generally increased classification accuracy across the LLMs. Some conventional models achieved competitive accuracy by predicting mainly the majority class. Although LLMs required substantially greater computational and financial resources than local baselines, their forecasts responded consistently to shifted numerical inputs, suggesting reliance on the supplied context rather than memorised price levels. These findings remain preliminary because of the limited evaluation period.

ID125: Effect of Post-Training Quantization on Residential Electrical Load Demand Anomaly Detection

Giuseppe La Tona, Massimiliano Luna, and Maria Di Piazza

Anomaly detection (AD) of electrical load demand is an important component of smart grids that enables predictive maintenance and facilitates timely intervention in case of faults. Recurrent neural network autoencoder models have shown good performance for electrical load demand AD. However, such deep learning models might require powerful hardware or cloud computational resources, which might be inappropriate for residential household applications due to cost in one case or privacy concerns in the other. Model optimizations techniques, like post-training quantization (PTQ) have been applied to enable the execution of deep learning models on resource-constrained devices like edge devices. This paper evaluates the impact of PTQ on the performance of a gated recurrent unit (GRU) autoencoder for residential electrical power consumption AD using a public dataset. Three quantization schemes are compared against the non-quantized baseline in terms of accuracy degradation, inference latency, and memory footprint. The results show that more aggressive quantization can reduce latency up to 41% at the cost of up to tenfold degradation in accuracy and 50% increase in model size. On the other hand, results showed that intermediate quantization schemes can improve deployment efficiency while preserving AD performance for edge-based smart grid applications.

Power Converters and Energy Storage Systems

ID071: A Finite-Set Model Predictive Control Approach for Enhanced Energy Sharing in Fuel Cell / Supercapacitor Fed OEWIM Based Electric Drive

Khaled Safsouf, Jean Sawma, and Hadi Kanaan

This paper proposes a Finite Set Model Predictive Control (FS-MPC) strategy for managing power sharing between a Fuel Cell (FC) and a Supercapacitor (SC) in an Open-End Winding Induction Motor (OEWIM) drive system. The proposed controller optimizes overall system efficiency while maintaining the supercapacitor State of Charge (SoC) within permissible limits. By extending the prediction horizon beyond a single step, the FS-MPC evaluates a broader range of future system states, enabling more informed control decisions and yielding improved dynamic performance and energy efficiency. A comparative analysis is conducted between single- and multi-step prediction horizons to quantify the performance gains associated with increased horizon depth. Additionally, the study benchmarks the efficiency of a standalone fuel cell configuration against a hybrid FC–SC topology, demonstrating the advantages of hybridization in electric vehicle drivetrain applications in terms of transient response and energy management.

ID024: Air-Lead Control of PEM Fuel Cells for EV Powertrains using Converter-Based OER Regulation and DC-Bus Stabilization

Iyad Ameur, Mohamed Bougrine, Atallah Benalia, Abdelhamid Rabhi, Mohamed Trabelsi, and Mohammed Benmiloud

This paper proposes an air-lead control architecture for Proton Exchange Membrane Fuel Cell (PEMFC) systems in electric vehicle (EV) powertrain applications, where fast load variations can compromise oxygen management, stack durability, and DC-bus stability. Conventional current-lead strategies allow practically the stack current to vary freely with the traction load, then force the compressor to follow the resulting oxygen demand. However, the slow pneumatic dynamics of the air-feed subsystem make this approach vulnerable to transient oxygen starvation during sudden power demand changes. To overcome this limitation, the proposed approach actively forces the stack current to follow a time-varying reference determined by the available cathode air flow. The downstream DC/DC converter switching action is used as the primary actuator for Oxygen Excess Ratio (OER) regulation, while the compressor motor voltage is reassigned to DC-bus voltage stabilization. This reformulation transforms OER control into a current-tracking problem, enabling a simple first-order bang-bang Sliding Mode Control (SMC) strategy. A robust outer Model-Free Control (MFC) loop is then used to regulate the DC-link voltage under uncertain nonlinear air-feed dynamics. The presented simulation results demonstrate that the proposed controller maintains the OER around its optimal value, prevents oxygen starvation, coordinates the stack current with the available oxygen supply, and restores the DC-bus voltage smoothly under severe load variations. These results highlight the relevance of the proposed air-lead architecture for fuel-cell EV powertrains, where stack protection and powertrain reliability are critical under highly dynamic driving conditions.

ID093: FPGA-in-the-Loop Validation of P&O and Fuzzy Logic MPPT Controllers for Photovoltaic Systems

Brahim Attia, Olfa Bel Hadj Brahim Kechiche, Hazem Dridi, Mahmoud Hamouda, and Xibo Yuan

This article presents an experimental validation of Perturb and Observe (P&O) and fuzzy logic (FL) controllers for maximum power point tracking (MPPT), using an FPGA-in-the-Loop (FIL) methodology. The photovoltaic (PV) system and the DC-DC Boost converter are simulated in MATLAB/Simulink, whilst the MPPT controllers are executed in real time on a Xilinx Zynq UltraScale+ MPSoC ZCU104 FPGA, using automated fixed-point conversion and HDL code generation. The FIL approach enables real-time co-simulation that accurately captures hardware-related effects, such as computational latency and fixed-point quantization. Both controllers achieve comparable steady-state tracking efficiency (97.05 per cent for the FL versus 96.05 per cent for the P&O), but the FL controller exhibits significantly reduced steady-state oscillations and lower dynamic power consumption, with an estimated total on-chip power of 0.609 W compared with 0.691 W for the P&O controller under the operating conditions considered. FIL co-simulation reduces simulation wall-clock time by 68.6 per cent for the FL controller and by 82.6 per cent for the P&O controller compared with a purely software-based simulation. This work confirms that FIL co-simulation provides a reliable framework for rapid prototyping and validation of MPPT strategies, thereby bridging the gap between simulation and practical FPGA implementation.

ID102: Thermal State Estimation-Assisted Self-Healing Interleaved SEPIC Converter for Reliable PV Battery Charging

Manoharan S, Mahalakshmi B, Gianluca Gatto, and Amit Kumar

This paper proposes a Thermal State Estimation -assisted Self-Healing Interleaved Single-Ended Primary-Inductor Converter (SEPIC) for reliable photovoltaic (PV) battery charging systems operating under dynamic irradiation and switch-fault conditions. The proposed architecture integrates thermal state estimation, temperature-dependent Mean Time Between Failures (MTBF)-based reliability assessment, rule-based anomaly prediction, and self-healing Pulse Width Modulation (PWM) redistribution. The converter is simulated for irradiation levels of v 400-1000 W/m2, shading conditions, MOSFET open-circuit faults, and thermal-stress conditions using MATLAB/Simulink. The simulation results show a 74.3% reduction in output-voltage ripple, a 63.5% reduction in total harmonic distortion (THD), a 58.23% reduction in the temperature metric from 79.0 °C to 33.0 °C, an efficiency of 93%, and an increase in the calculated MTBF from 15.3×10^3 h at 79 °C to 490.7×10^3 h at 33 °C (approximately 3107%, increase). The proposed method also recovers the output following the injected switch fault within 0.038 s. These results demonstrate that combining thermal-state estimation, fault detection, self-healing PWM redistribution, and temperature-dependent reliability assessment can improve the fault tolerance and thermal reliability of PV battery charging converters.

ID048: NLADRC Control for Five Phase PMSM-Based Naval Propulsion Under Disturbances and Fault Conditions

Khouloud Ben Hammouda, Mohamed Trabelsi, Trabelsi Trabelsi, and Riadh Abdelati

This paper investigates the application of Non- Linear Active Disturbance Rejection Control (NLADRC) to a five-phase Permanent Magnet Synchronous Machine (PMSM) intended for naval propulsion systems. Marine environments are characterized by highly nonlinear dynamics, significant hydrodynamic and aerodynamic disturbances, load variations, and parameter uncertainties. Moreover, the propulsion system may operate in both healthy and degraded modes, where phase faults or imbalances further increase control complexity and challenge conventional control strategies. To address these issues, a NLADRC-based control approach is developed using an extended state observer to estimate and compensate for total disturbances in real time, without requiring an accurate system model. The controller performance is evaluated under healthy and degraded operating conditions, as well as under realistic ship

propulsion disturbance profiles. A comparative study between classical vector control and NLADRC is carried out in healthy and degraded modes under Open Circuit Fault (OCF) scenarios.

Cloud-Edge Computing, IoT, and Renewable Energy

ID107: A Behaviour-Level Zero-Trust Digital Twin for Secure ATM and Cash-Kiosk Maintenance

Jumanah Alfailakawi, Anfal Almutairi, Aalyah Alotaibi, Khaled Allanqawi, and Basil Alothman

Maintenance sessions of cyber-physical assets, for instance, ATMs and self-service cash-exchange kiosks, connected via Internet of Things (IoT) connectivity, present a high-value post-login attack surface. A valid or compromised technician might authenticate and then open the cabinet, connect service media, modify system configurations, modify firmware, or interact with cash-handling commands. This paper proposes a lightweight Zero-Trust behaviour-level digital-twin abstraction that mirrors the maintenance event stream, contextual risk factors, and trust trajectory of a kiosk session. A bounded Trust-Score model assigns explainable penalties to concrete maintenance events and classifies each session as Normal, Suspicious, or High-Risk using a transparent single-pass algorithm. The evaluation combines 1,500 synthetic sessions, a 130-session scripted trace-capture testbed, an independent CERT r4.2 insider-threat benchmark, threshold sensitivity, ablation, cross-validated ML separability baselines, and a five-strategy adversarial evaluation. The synthetic study achieves 85.9% accuracy and 0.78 macro-F1, the testbed escalates 97.5% of malicious sessions with no benign High-Risk alarms, and the CERT mapping escalates 93.5% of flagged malicious sessions. The results provide initial evidence that the proposed mechanism may be useful as an explainable proof-of-concept approach for continuous maintenance-session monitoring, not as production ATM field performance.

ID121: Planning IoT-Enabled Road Management in Kuwait: Assessment, Technology Mapping, and a Pilot-Deployment Priority Index

Haya Albader and Naser Alselahi

Internet of Things (IoT) Connected sensing and digital twins offer road decision-makers new means to monitor traffic, detect hazards, and support operational decisions. In Kuwait, however, public transportation data are reported at different spatial and temporal scales and have not been integrated into a smart management system. This study addresses that gap through a three-stage framework for Internet of Things (IoT) –enabled road management. Official transportation data for 2015–2024 were analyzed to examine vehicle and network growth, crash burden, fatality severity, accident-type distribution, traffic violations, and emerging monitoring needs. The identified challenges were then matched to relevant sensing, analytical, and digital-twin functions. A Kuwait Pilot Priority Index was specified, and its measurable safety component—the Area Pilot Priority Index—was applied to 39 areas using total crashes, deaths, pedestrian or trampling crashes, and overturning crashes. Registered vehicles increased by 35.5% between 2015 and 2024, compared with 12.6% growth in paved-road length. Another finding, Crash volume and deaths showed only a weak association across the listed areas, while accident-type distributions differed significantly among governorates. Abu Halifa and Al-Fintas ranked first under the base weighting. The findings support IoT-based digital twin (DT) site-specific deployment, especially for traffic and pedestrian monitoring in high-volume urban and coastal areas, and for speed, road-condition, visibility, and emergency-response systems in severe-crash corridors. The framework provides a reproducible basis for screening candidate pilot areas, while final deployment requires exposure, feasibility, and field-validation data.

ID127: Analysis of Expenditure and User Dissatisfaction Optimization in Residential Energy Management Systems Using Social Network Optimization

Andrea Montuoro, Alessandro Niccolai, and Francesco Grimaccia

This paper investigates the operation of a residential Energy Management System (EMS) for demand-side manage- ment under a time-of-use electricity tariff. The considered res- idential system includes controllable and non-controllable loads with different flexibility characteristics: thermal loads, namely a water heater and a refrigerator; shiftable appliances, such as washing machine, dryer, and dishwasher; electric vehicle charg- ing; and a user-driven baseload. Each controllable load is mod- eled according to its physical behavior, operational constraints, and user-related dissatisfaction criteria. The EMS is analyzed through two independent single-objective optimization problems: minimization of electricity expenditure and minimization of user dissatisfaction. The optimizations are performed using the Social Network Optimization (SNO) algorithm over a 24-hour simulation horizon, with multiple independent trials to assess convergence robustness. The results show that expenditure mini- mization shifts flexible loads toward low-tariff periods, reducing energy cost but increasing dissatisfaction due to reduced thermal comfort, violation of preferred appliance operating windows, and reduced power availability margin. Conversely, dissatisfaction minimization produces smoother load profiles and improves user comfort, but requires less economically favorable scheduling and higher energy use for thermal loads. The comparison demonstrates that the two objectives lead to structurally different optimal schedules, confirming their conflicting nature even when optimized separately under the same residential scenario.

ID064: Generalizable solar irradiance prediction from satellite data using spatial–temporal graph attention network

Konin Pierre-Claver Kakou, Georg Frey, Daud Mustafa Minhas, and Jörg Trentmann

Despite advances in solar resource estimation, significant biases relative to ground-based measurements remain in some regions. Accurate forecasts are essential for filling observational gaps and reducing reliance on dense monitoring networks. This study presents a generalizable framework for solar irradiance forecasting based solely on satellite observations. A Spatial–Temporal Graph Attention Network (STGAT) is proposed to model spatial interactions among neighboring locations, while a Bayesian layer provides uncertainty estimates. To enhance spatial generalization, the model is evaluated using Leave-One-Station-Out cross-validation. Results demonstrate strong performance at short forecasting horizons. In Côte d'Ivoire, STGAT achieves lower RMSE than the corresponding SARAH-3 satellite irradiance estimates for forecast horizons up to 13 h. In Germany, where SARAH-3 performs better overall, the model trained on only five stations achieves comparable accuracy, with an RMSE difference of approximately 10 W/m2. From 1 to 24 hours ahead, RMSE increases from 102 to 170 W/m2 in Côte d'Ivoire and from 113 to 159 W/m2 in Germany. These findings highlight the model's potential as a reliable and transferable tool for solar energy planning and management across diverse climatic regions.

ID069: Threshold-Based Sensitivity Assessment of Industrial PV-BESS Systems from Energy Balances to Economic Viability

Nicoletta Matera, Riccardo Ramaschi, Emanuele Ogliari, Sonia Leva, Francesco Torchiaro, and Enrico Lamanna

This paper presents a threshold-based sensitivity assessment of an industrial photovoltaic-battery energy storage system under evolving electricity market conditions. The analysis considers a 600 kW PhotoVoltaic (PV) system coupled with a 600 kWh Battery Energy Storage System (BESS) installed at an industrial facility in Italy. An hourly simulation framework is used to evaluate the energy exchanges among PV generation, BESS, load demand, and the grid, and to calculate the levelized cost of electricity, the levelized cost of use, and the levelized cost of storage. Results show that PV size and annual energy consumption are the dominant drivers of system performance, whereas BESS capacity mainly affects the redistribution of surplus PV energy between BESS and grid export. The economic attractiveness of storage strongly depends on market conditions, with the most favourable conditions observed in years with higher electricity prices. The proposed approach provides a replicable framework for the assessment and design of industrial PV-BESSs.