



Abstract Paper ID: 16

Digital Health Ecosystem: Architecture, Interoperability, and Predictive Analytics

Digital health is evolving from isolated clinical information systems toward intelligent, interconnected ecosystems integrating Internet of Medical Things (IoMT), artificial intelligence (AI), and cloud-edge computing infrastructures. This paper presents an Intelligent Digital Health Ecosystem (IDHE) framework enhanced with a digital twin paradigm for predictive healthcare monitoring. The proposed architecture integrates sensing devices, interoperable data layers, machine learning analytics, and clinical decision support into a unified system-of-systems. Unlike conventional reactive monitoring platforms, the frame-work enables continuous patient modelling through a virtual twin representation that predicts health risks using real-time multimodal data streams. A layered architecture is introduced covering sensing, interoperability, data integration, intelligent analytics, and healthcare service delivery. The framework supports remote monitoring, personalized treatment planning, public health analytics, and secure data exchange. Challenges related to interoperability, scalability, privacy, and AI governance are also discussed. The reviews provides a foundation for future digital twin-enabled healthcare platforms and intelligent population health systems.



Abstract Paper ID: 17

AI-Powered Cybersecurity Mesh for Real-Time Threat Detection in Financial Networks

Financial institutions process millions of transactions every day across digital banking platforms, payment gateways, mobile wallets, and open banking ecosystems. This scale has increased both operational efficiency and exposure to increasingly coordinated cyberattacks such as credential abuse, account takeover, transaction laundering, and cross-institution fraud. Traditional centralized Security Operations Centers (SOCs) and perimeter-based monitoring systems often struggle to detect these threats quickly because of data silos, delayed correlation, and limited visibility across distributed environments. This paper proposes an AI-powered cybersecurity mesh for financial networks that improves real-time threat detection using distributed intelligence instead of centralized monitoring alone. The architecture combines edge AI agents, federated learning, graph-based fraud correlation, and privacy-preserving threat sharing to help institutions detect and respond to threats faster without exposing sensitive transactional data. Transformer models are used for behavioral anomaly detection, while Graph Neural Networks (GNNs) identify suspicious transaction relationships across entities such as accounts, devices, and merchants. Experimental evaluation using synthetic financial transaction datasets and emulated fraud scenarios shows that the federated model achieves an F1-score of 0.92 compared with 0.87 for the centralized baseline, corresponding to a 5.75% relative improvement. Recall increases from 0.84 to 0.91, while median edge-inference latency remains 126 ms. The architecture also aligns with Zero Trust principles and regulatory expectations for privacy, explainability, and digital sovereignty. This work presents a practical implementation model for financial organizations seeking stronger cyber resilience while maintaining compliance and operational scalability.



Abstract Paper ID: 18

RMiT-Aligned Cyber Security Governance and Risk Management Disclosures: Evidence from Malaysia's Financial Sector

This paper studies how the Malaysian financial sector, which consists of banks and insurance companies describes cyber security governance and risk management in their recent corporate reports. The study uses Bank Negara Malaysia's Risk Management in Technology (RMiT) policy as an interpretive benchmark to assess the visibility of RMiT-related disclosure domains. RMiT is a comprehensive policy mandating strict technology risk and cyber security standards for all financial institutions operating within the country. A total of 290 coded excerpts from annual, integrated and sustainability reports for the 2023–2025 reporting period were analysed using Braun and Clarke's six-phase thematic analysis of qualitative methodology was adopted. The analysis yielded five themes: governance and oversight, operational capabilities, cyber resilience and emerging risks, culture, skills and resilience and impact and outcomes. Board-level or committee oversight is the most visibly reported topic in terms of the various RMiT components analysed, with technology and cyber risk viewed primarily as a risk factor. The primary focus is on cyber resilience frameworks, incident management practices, cyber security awareness initiatives and ongoing cyber risk assessments and testing routines. Less visible were disclosures relating to the risk relating to third parties, outsourcing partnerships and cloud computing environments. Based on the findings, the policy objectives of RMiT partially echo the language used in company reports, predominantly focusing on governance and technology risk management domains, yet, the evidence is less robust concerning reporting transparency and details on areas such as the scope of assurances provided to stakeholders, achieved results and in-depth information on third-party risks.



Abstract Paper ID: 20

Adaptive Pooling for Time-Efficiency Optimization in Cross-Modal Pedestrian Image-Text Matching

Cross-modal matching between pedestrian images and textual descriptions has emerged as a crucial technology for intelligent surveillance systems, facilitating natural language-based person retrieval from large-scale image databases. Nevertheless, existing approaches primarily emphasize matching accuracy while disregarding computational efficiency, which poses a significant bottleneck for real-world implementation in resource-constrained surveillance settings. This research systematically examines the optimization impact of adaptive pooling strategies on the temporal efficiency of cross-modal pedestrian image-text matching. Surveillance image-text paired data was gathered and curated, and comprehensive comparative experiments were carried out among six representative pooling methods: Global Average Pooling (GAP), Global Max Pooling (GMP), Spatial Pyramid Pooling (SPP), Generalized Mean Pooling (GeM), Attention-based Pooling, and Adaptive Pooling strategy proposed in this research. Through rigorous experimental analysis using the CUHK-PEDES benchmark dataset, it shows that the adaptive pooling method attains a Rank-1 accuracy of 74.6% while maintaining an inference latency of only 14.8ms per query, which represents a 2.4-fold speedup compared to attention-based pooling with merely a 0.5% degradation in accuracy. The proposed approach balances matching accuracy and computational efficiency by dynamically adjusting pooling parameters according to input feature characteristics. Comprehensive ablation studies verify the contribution of each component, and cross-dataset generalization experiments confirm robust performance across various scenarios. The findings conclusively indicate that adaptive pooling can effectively enhance the efficiency of cross-modal matching for pedestrian images and text, meeting the strict real-time requirements of practical surveillance applications.



Abstract Paper ID: 21

Detecting Harsh Words in YouTube Podcast Content Using Bidirectional Encoder Representations from Transformers Algorithm

The development of digital media such as YouTube is increasing, presenting a wide variety of topics, including podcasts. However, the lack of content filtering in interactive digital media, such as YouTube podcasts, can raise concerns, including conversations that often contain harsh language. This harsh language can often cause discomfort for listeners. Therefore, this study aims to develop a model to detect harsh language in YouTube podcast content as a preventive measure against the spread of harmful content in society. The model was built using the Bidirectional Encoder Representations from Transformers (BERT) algorithm with a dataset consisting of YouTube podcast video transcripts. In addition, this study used the Synthetic Minority Over-sampling Technique (SMOTE) to balance the text transcription results from the collected videos. Experimental results show that the BERT model is able to achieve optimal performance on a data split of 80% for training and 20% for testing in 3 epochs, with an F1-score in the BERTScore evaluation of 0.829. In addition to contributing to the harsh language detection model, this study also produced a prototype web-based application that automatically processes YouTube video transcripts, detects harsh language, and displays the analysis results. The results of this research contribute to the development of machine learning-based text analysis technology to create more conducive digital communication for the community.



Abstract Paper ID: 22

Data-Driven Runtime Performance Profiling of NPC Crowd Density in a Standalone Virtual Reality Tawaf Simulation

This paper presents a data-driven runtime performance profiling study of non-player character (NPC) crowd density in a standalone virtual reality (VR) Tawaf simulation. The study responds to the need for reproducible performance evidence in immersive simulation systems, where frame stability is affected by animated crowd agents and hardware limitations. A Unity-based VR Tawaf scene was profiled on a PICO 4 Ultra standalone headset using six crowd-density levels and seven NPC processing conditions. The final dataset consisted of 126 CSV runs and 44,430 clean frame samples after startup logging samples were removed. The analysis focused on frames per second (FPS), mean frame time, frame-time percentiles, late-frame percentage, and practical near-90 Hz pass or fail status. The revised paper includes graphical FPS and frame-time trends, device specifications, and a discussion of the practical meaning of the 150 NPC upper test density. Results show that the baseline condition supports only 25 NPCs under the practical near-90 Hz rule, while animation-related optimisation conditions support the maximum tested density of 150 NPCs.



Abstract Paper ID: 29

Adversarial Machine Learning in Cybersecurity: A Bibliometric Review of Emerging Threats and Research Hotspots

This study presents a comprehensive bibliometric analysis of 3,705 publications on Adversarial Machine Learning (AML) within the cybersecurity domain, retrieved from the Scopus database. Spanning 2020 to 2025, the research maps the intellectual structure and evolutionary trajectory of this rapidly expanding field. The results reveal an exponential surge in scholarly output, with 2024 and 2025 alone accounting for over 50% of total publications, underscoring the critical urgency of securing AI-driven infrastructure. Geographic analysis highlights the dominant contributions of the United States and China, with Nanyang Technological University identified as the leading institution. Citation network analysis pinpoints Ian Goodfellow and Nicholas Papernot as seminal figures whose work on adversarial examples and defense mechanisms anchors the domain. Thematic clustering reveals a shift from foundational attack methodologies toward advanced topics such as federated learning, explainable AI, and robust intrusion detection systems. While the study is limited to English-language sources indexed in Scopus, it provides a crucial roadmap for researchers and practitioners, identifying emerging hotspots and strategic gaps in the defense against sophisticated cyber threats.



Abstract Paper ID: 31

Smart Debug AI for Strategic Manufacturing Decisions: TRAM Framework Approach in Batam

This study investigates the integration of Smart Debug technologies an emerging class of intelligent diagnostic systems into the strategic decision-making processes of electronic manufacturing industries in Batam City, Indonesia. As part of the broader shift toward Industry 5.0, these technologies offer real time fault detection, predictive maintenance, and automated decision support, enabling manufacturers to enhance operational efficiency and resilience. The study adopts the Technology Readiness and Acceptance Model (TRAM) to examine the psychological and organizational factors influencing the adoption of Smart Debug systems. A mixed-method approach was employed, combining quantitative surveys as a pilot study (N=40) and qualitative interviews with industry professionals at PT. Flextronics Technology Indonesia. Key constructs include perceived usefulness, perceived ease of use, and individual technology readiness dimensions such as optimism, innovativeness, discomfort, and insecurity. Organizational readiness factors such as leadership support, infrastructure, and employee training were also considered. Pilot findings reveal exceptionally high acceptance levels, with all positive constructs scoring above 3.9/4.0, and virtually no psychological barriers (discomfort=1.11, insecurity=1.10). Regression analysis confirmed all TRAM hypotheses, with Perceived Usefulness (PU) and Perceived Ease of Use (PEOU) strongly predicting Behavioral Intention ($\beta=0.51$ and $\beta=0.42$, respectively). Organizational Readiness demonstrated the strongest influence on Actual Use ($\beta=0.68$). The pilot study successfully validated the research instruments and provides preliminary insights, with a full-scale study planned for 200 respondents to confirm and extend these findings. The study contributes to the theoretical development of TRAM in industrial settings and offers practical recommendations for manufacturers seeking to implement intelligent diagnostic systems.



Abstract Paper ID: 32

Early Detection of Arterial Heart Disease via Quantum-Enhanced Explainable Deep Learning in Cardiovascular Imaging

Arterial related cardio vascular diseases are becoming a major cause of health issue all over the world. Early and accurate detection of medical issues are vital for effective handling and better results. Although machine learning has improved in processing of medical images and its analysis, but recent models face challenges in managing large and complex datasets efficiently and lack transparency which is essential for trust in clinical diagnosis process. This paper introduces a new system that combines quantum and classical computing methods to help detect arterial heart disease early. This system integrates explainable machine learning directly into the process. Unlike traditional step-by-step methods, this system combines classical techniques for extracting features, aligning different types of data and using quantum-enhanced processing to better understand complex and non-linear relationships between imaging and clinical data. The proposed system will use explainable AI with the integration of classical methods for explanation of model decisions along with quantum-based transparency procedures, to deliver strong, dependable explanations that match with the thinking of medical physicians. The system is tested for its ability to predict, process data efficiently, be clear and be ready for real-world use in clinical settings.



Abstract Paper ID: 33

Forecasting Geomagnetic Induced Currents Using Neural Network During Severe Solar Storms

The Geomagnetic Induced Current (GIC) formation underground is dangerous due to its harmful effect to the surrounding electrical and pipeline appliances. The mitigation of the GIC level in Low Latitude and Equatorial region such as Malaysia is important since it might trigger unintended phenomenon such as sudden blackout or electrical equipment failure. This research study proposed the GIC prediction by using Neural Network based on the available GIC data measured during the Medium and Low-Level Solar Storm at the Tenaga Nasional Berhad (TNB), KUL South, Taman Pelangi, Petaling Jaya, Kuala Lumpur with geographic latitude and longitude ($3^{\circ} 4' 37''$ and $101^{\circ} 39' 55''$) by using HIOKI measurement equipment from 7th April 2022 to 5th May 2022. Based on the research it is found that while applying Neural Network, the GIC generated spikes between -9.72A to 12.905A during the Severe-Level Solar Storm and between -5.986A to 6.37A during Strong-Level Solar Storm from 1st November to 29th November 2025. The research investigation utilizing the data of the Magnetic Field collected from the SuperMAG website and the Solar Storm that was taken from OMNIWeb website.



Abstract Paper ID: 34

Digital Twin Security in Cyber-Physical Systems: A Systematic Literature Review Using PRISMA

Digital twin (DT) technology is emerging in cyber-physical systems (CPS) to create new attack surfaces that are not covered by traditional security models. This paper presents the PRISMA-compliant systematic literature review on digital twin security within CPS environments. From an initial Scopus extract of 20,000 records, a three-criterion keyword screening and six-item quality assessment yielded 750 confirmed studies spanning the years 2020–2025, representing a 523% publication growth over this period. A novel eight-category threat taxonomy is empirically derived; integrity/false data injection (FDI) and denial-of-service (DoS) attacks are the most explicitly addressed threats. The countermeasure landscape is fragmented, with security frameworks, classical machine learning, and anomaly detection each accounting for approximately 15% of studies, while zero trust architecture (ZTA) appears in only one paper. Six critical research gaps are identified, including the absence of dedicated DT synchronisation attack taxonomies, standardised benchmarks, and adversarial machine learning based defences. The findings offer evidence-based guidance for securing DT-CPS deployments across industrial, energy, and critical infrastructure domains.



Abstract Paper ID: 35

Security Vulnerabilities in Generative AI Systems: A Comprehensive Review and Future Directions

Automated content generation has taken a significant leap forward in various fields with the advent of Generative Artificial Intelligence, especially Large Language Models and deep generative architectures. The paper offers a comprehensive overview of the security threats facing generative AI systems, including prompt injection, data leakage, model poisoning, adversarial attacks, and prompt leakage. We systematically examine the threat landscape and current mitigations, including model alignment, privacy-preserving learning mechanisms, reinforcement learning from human feedback, and input validation. Furthermore, the key research gaps are identified, including the absence of standardized approaches to evaluate the quality of defense robustness. Lastly, we discuss future research opportunities to improve the robustness, reliability, and trustworthiness of generative AI systems. This review offers a structured foundation for understanding the current security challenges emerging for AI deployment and how to mitigate them.



Abstract Paper ID: 36

Glycaemic-Blind Machine Learning for Prediabetes Screening and Risk Stratification: A Focused Scoping Review of Validation and Reporting Gaps

Prediabetes is highly prevalent, frequently undiagnosed, and partly reversible, making it an attractive target for machine-learning (ML) screening. This focused evidence map compares 15 studies across model types, predictor sets, discrimination, calibration, decision-curve analysis (DCA), explainability, rapid methodological-concern ratings, and longitudinal design. The comparator studies were purposively selected from a broader structured literature synthesis to represent the principal prediabetes outcome classes, model families, cohort scales, and reporting practices relevant to glycaemic-blind screening. Gradient-boosting ensembles were the most frequent and usually performed best in head-to-head comparisons on tabular data. Discrimination varied by outcome: AUROC values of approximately 0.81–0.86 for cross-sectional detection, around 0.73 for incident prediabetes, and 0.67–0.87 for progression to type 2 diabetes. Only four studies were fully glycaemic-blind; calibration was reported in five and DCA in two. No model combined the full evaluation suite with prospective external validation. Six recurring gaps were synthesized, and a conceptual glycaemic-blind, calibrated, explainable, longitudinal two-stage architecture is proposed as a research agenda requiring future empirical validation.



Abstract Paper ID: 37

Intelligent Snapshot Scheduling in Storage Area Networks: A Machine Learning Approach

While Storage Area Networks (SANs) play a vital role in enterprise data infrastructures, backup snapshot scheduling in SANs is not a subject that is widely covered. Traditional scheduling solutions are rigid and do not consider the dynamic and unpredictable nature of enterprise workloads, which creates poor recovery windows, high storage costs, increased risk of data loss and network performance degradation during peak times. This paper introduces a novel adaptive predictive snapshot scheduler (APSS) which is based on machine learning to dynamically optimize backup snapshot scheduling in SAN environments. The hybrid ML system, APSS, is one that incorporates Random Forest classification and Gradient Boosting regression, enabling in real time prediction of workload intensity, storage utilization patterns, and the probability of hardware failure. The system dynamically adjusts snapshot frequency, timing and resource allocation based on these predictions. APSS delivers peak bandwidth reduction of 22.6%, storage overhead reduction of 28.3%, RPO compliance of 41.2%, and an average backup completion time reduction of 34.5% over a 90-day simulated SAN environment, compared to the best performing baseline.



Abstract Paper ID: 38

Evaluating Prompt Engineering Strategies for LLM-Based Synthetic Indonesian Text Data Generation

Synthetic data generation has become an important approach for overcoming data scarcity, privacy restrictions, and class imbalance in machine learning applications. Large Language Models (LLMs) provide new opportunities to generate synthetic textual datasets, but the quality of generated data depends strongly on the prompt engineering strategy used. This study aims to evaluate the effectiveness of several prompt engineering techniques for LLM-based synthetic text data generation. The research uses 13 prompting techniques for 5 text classification datasets in the Indonesian language. This research evaluates the relevance, diversity, consistency, fluency, format validity, and semantic similarity of the synthetic data generation results. Then, the synthetic text data results are used for Logistic Regression, Multinomial Naïve Bayes, and Support Vector Machine in the text classification task. The result is an shows that the prompt strategy produces synthetic text data with 0.52 of score quality for the labeled dataset. Overall, the synthetic data generation can enhance the accuracy, recall, precision, and F1-score. This paper contributes a reusable evaluation for selecting prompt strategies in synthetic data generation using LLMs.



Abstract Paper ID: 42

Machine Learning Approach for Type 2 Diabetes Risk prediction: A Comprehensive Review Approaches, Challenges and Future Directions

The global burden of type 2 diabetes continues to rise, and this places growing pressure on clinical practice. Many clinics still rely on diagnostic testing rather than automated risk-prediction tools that could flag the disease at an earlier stage. This review examines recent machine learning models for type 2 diabetes risk prediction, ranging from non-invasive and lifestyle-driven models to biostatistical frameworks built on classical biomarkers. It focuses on two problems that researchers face repeatedly: imbalanced datasets and high-dimensional clinical feature sets. The methods reviewed span ensemble techniques and deep learning architectures developed for this task. These advances carry three persistent challenges: models that generalize poorly, limited external validation, and low accessibility of results for healthcare professionals. The review also considers where automated machine learning can reduce heterogeneity between studies and improve translational readiness. Model robustness is discussed in relation to dimensionality reduction and hyperparameter tuning, and performance is assessed using accuracy, recall, and AUC-ROC. The main contribution is a structured synthesis of methods, datasets, and reported performance, supported by a transparent PRISMA-based search protocol and a comparative analysis that shows where each family of methods is most useful and where the evidence remains weak.



Abstract Paper ID: 44

Real-Time ASL Translation Using CNN-LSTM and MediaPipe Landmark Features

This paper presents the design and development of a real-time American Sign Language (ASL) translation system that integrates computer vision, deep learning, and natural language generation to support more inclusive communication for individuals with hearing and speech impairments. The proposed system employs MediaPipe to extract three-dimensional (3D) hand, body-pose, and facial landmarks, producing a dense representation of gesture dynamics comprising 126 features per frame. To support model training and evaluation, a custom dataset of 3,360 samples representing 16 ASL word classes was constructed using data captured from multiple signers under varied camera viewpoints. A hybrid Convolutional Neural Network–Long Short-Term Memory (CNN-LSTM) model is developed to learn both spatial landmark patterns and temporal dependencies across gesture sequences. Experimental results indicate that the proposed CNN-LSTM model achieves 98% accuracy, outperforming comparative architectures such as LSTM, Gated Recurrent Unit (GRU), Bidirectional LSTM, and ResNet-LSTM. The system further enhances translation quality by passing recognised gesture tokens to an OpenAI language model, which generates grammatically coherent English sentences from the predicted word sequence. The complete pipeline is implemented as a web-based application that integrates real-time video acquisition, landmark extraction, gesture classification, and sentence generation with low inference latency. Evaluation suggests that the system remains robust across variations in lighting conditions, user positioning, and inter-signer differences. Overall, this study contributes an end-to-end, vision-based ASL translation pipeline that offers a low-cost, non-invasive, and scalable alternative to wearable sensor-based sign recognition systems.



Abstract Paper ID: 46

Advancing Fungal Disease Detection in Crops Through Artificial Intelligence: A Technical Perspective

Fungal pathogens remain one of the costliest biotic stresses on global crop production, depressing yields, degrading produce quality, and inflicting billions of dollars in losses every season. Visual scouting by pathologists, though accurate in expert hands, scales poorly and is expensive to sustain across commercial operations. The arrival of deep learning and more recently transformer-based vision has altered how researchers approach this problem. We examine the state of AI-driven detection for three of the most consequential foliar disease groups: rust, powdery mildew, and downy mildew. The review walks through dataset construction, classical machine learning baselines, convolutional networks, transfer learning, object detection frameworks, and emerging transformer architectures, then benchmarks representative studies side by side. We pay particular attention to the gap between curated-dataset accuracy and field deployment performance, an issue that the literature frequently under-reports. Persistent obstacles around dataset diversity, model generalization, interpretability, and on-farm integration are discussed, together with the emerging toolset like explainable AI, multimodal learning, federated training, generative data synthesis, edge inference, and agricultural foundation models. The result is a strategic map intended for researchers and practitioners working to make AI-based crop disease management genuinely useful at farm scale.



Abstract Paper ID: 48

MoodTune: A Real-Time Emotion-Based Music Recommendation System Using Convolutional Neural Networks and Facial Expression Recognition

Most of the current music streaming platforms do not consider your current emotional state, as the majority of them still rely on your stated preferences and historical patterns of listening habits. In response to this lack of a personalized music experience, this paper presents MoodTune, an intelligent online web-based application that utilizes a combination of dynamically selected songs and facial recognition utilizing convolutional neural networks (CNN) for facial expression analysis. The application will use the user's webcam as a real time means of capturing his/her facial expressions, and will use a CNN model developed through TensorFlow.js's Face API to classify the user's expressions into categories such as happiness, sadness, anger, surprise, or neutral. Based on the identified emotion, the application queries the Spotify Web API to create a customized playlist based on the user's emotional state that will be displayed through an interactive frontend (developed using Angular). Additionally, the user is able to confirm or correct the detected emotion and the resulting recommendation through a human-in-the-loop confirmation system. The user's emotional data will be stored in Firebase for analytics and personalized experiences. The evaluation of the prototype shows strong performance under ordinary lighting conditions, including the ability to accurately classify a user's emotional expression and retrieve music without interruption. MoodTune bridges the gap from affective computing research to actual usage in everyday multimedia applications, and adds a modular and scalable architecture suitable for use in digital health and wellness applications.



Abstract Paper ID: 49

AI-Based Interactive Fitness Coaching System for Children Using a Humanoid Robot

RoboMove is an interactive fitness coaching application developed for integration with the Daman humanoid robot to promote physical activity among children aged 6 to 12. As children increasingly fail to meet the recommended 60 minutes of daily physical activity, RoboMove addresses this gap by delivering an engaging, robot-assisted exercise experience with real-time interaction and feedback. The system guides users through 14 exercises across three levels; Warm Up, Active, and Challenge through four modules: Exercise Guidance, Pose Detection, Feedback, and Interaction. Pose detection uses Google ML Kit to track body landmarks via the robot's built-in RGB camera, evaluating exercise form and counting repetitions. Voice commands are handled by the offline Vosk speech recognition engine restricted to a fixed grammar for fast, accurate recognition without internet connectivity, spoken feedback is delivered through Android's Text-to-Speech engine, and robot head movement is controlled via the Reeman AIDL service interface. All processing runs locally on the robot's built-in Android 6.0 tablet, requiring no external sensors, cloud connectivity, or additional hardware. This paper presents the requirement analysis and system design of RoboMove, demonstrating that a fully integrated, child-friendly fitness coaching system can be built within the constraints of existing humanoid robot hardware.



Abstract Paper ID: 54

Precision Weed Detection in Paddy Fields Using UAV Remote Sensing and Deep Learning

Accurate weed detection in paddy fields remains a major challenge in precision agriculture because weeds and paddy often exhibit similar visual and spectral characteristics, particularly in UAV RGB imagery captured under complex field conditions involving wet soil, standing water and uneven illumination. This study aims to develop a deep learning-based technique for pixel-level weed detection in paddy fields using UAV remote sensing imagery. The proposed methodology employs manually selected region-of-interest training samples representing weed, paddy, soil and water classes. Red, Green and Blue pixel intensity values are extracted as spectral feature vectors and classified using a Multi-Layer Perceptron (MLP) neural network. To improve the spatial quality of the detected weed regions, morphological post-processing operations, including area opening, hole filling and closing, are applied to reduce noise and refine the final weed mask. Experimental results show that the proposed technique achieved an overall classification accuracy of 97.64%. For the weed class, the model obtained a Precision of 99.01%, Recall of 98.08% and F1-score of 98.54%, demonstrating strong separation between weed and paddy pixels. The resulting weed detection map provides a practical basis for site-specific weed management and enabling more sustainable paddy field management through UAV-assisted precision agriculture.



Abstract Paper ID: 55

Rule-Based Insider Threat Detection System via User Behavior Monitoring for Small-Scale Desktop Environments

Insider threats remain difficult to detect because authorized users can misuse legitimate access while their actions may appear similar to normal behavior. Existing security mechanisms such as firewalls, antivirus tools and traditional intrusion detection systems mainly focus on external or network-level threats, while advanced user behavior analytics and machine learning solutions often require large datasets, model training and enterprise-level resources. This paper presents Insider Watch, a lightweight rule-based insider threat detection prototype that monitors selected Windows desktop user activities. The system observes session behavior, file access, USB activity, application usage, URL access and working-hour violations. A customizable rule engine evaluates captured events against predefined thresholds and generates risk-based alerts with contextual information. The prototype also provides a live dashboard, configurable settings and structured incident report export in PDF and Excel formats. The proposed approach emphasizes explainability, low implementation complexity and suitability for small-scale environments. Future work will focus on broader testing with realistic datasets and deployment in controlled organizational scenarios.



Abstract Paper ID: 56

Serene Chat: An Edge-AI Privacy Architecture for Localized Mental Disorder Screening in Peer-Support Systems

Mental health disorders among Malaysian youths have become an increasing public health issue, while privacy concerns, social stigma, and fear of identity disclosure continue to discourage timely help-seeking. Existing AI-powered mental health applications typically rely on cloud-based inference, requiring sensitive conversations to be transmitted to external servers and raising privacy and security concerns. As a result, personal information is exposed beyond the user's device, increasing the potential risk of data leakage and unauthorized access. This paper presents Serene Chat, a privacy-preserving, serverless peer-support architecture that successfully mitigates network-layer data exposure through localized client-side artificial intelligence execution. Shifting the computational paradigm entirely to the browser sandbox environment, the framework integrates TensorFlow.js with the Universal Sentence Encoder (USE) to run high-dimensional semantic vector calculations directly at the edge. To account for localized linguistic variations, a custom client-side normalization dictionary scales out to interpret code-switched phrases, regional shorthand, and Manglish distress signals before semantic inference. To protect user privacy, the platform uses a zero-knowledge approach that processes raw text entirely on the device, while relying on Google Firebase for fast, real-time message delivery. Empirical validation using a representative 120-synthetic evaluation dataset demonstrated a 95.0% overall classification accuracy, capturing 100% of explicit crisis indicators with zero high-risk false negatives. These outcomes prove the feasibility of decentralized, browser-based Edge-AI for executing secure, privacy-preserving mental health screening frameworks.



Abstract Paper ID: 57

Design of a Multi-Layered Cybersecurity Architecture for Small and Medium Enterprises (SMEs)

Small and Medium Enterprises (SMEs) face increasing cybersecurity threats while operating with limited resources, flat networks, and inadequate monitoring. This paper presents a multi-layered cybersecurity architecture based on Defense-in-Depth that integrates pfSense firewall enforcement, VLAN segmentation, DMZ isolation, Suricata intrusion detection, and Telegram alerting. A five-phase Design Science Research process is used, including requirements-based analytical validation of representative attack paths. The validation traces Internet attacks, lateral movement, DMZ compromise attempts, and IDS-to-administrator alert delivery against explicit control responsibilities and acceptance criteria. A quantitative design-coverage comparison shows that the proposed architecture supports 14 of 14 assessed features, while reviewed studies cover 3 to 7 features; this score represents architectural breadth rather than measured detection effectiveness. A common experimental protocol is therefore specified for future measurement of detection rate, false-positive rate, alert latency, throughput, computational overhead, and deployment cost. The resulting design provides SMEs with a structured, modular, and cost-conscious reference architecture while clearly separating current design-level evidence from future empirical evaluation.



Abstract Paper ID: 60

Artificial Intelligence-assisted Communication and Emotional Health Monitoring System for Autistic Individuals and Their Caregivers

Autism Spectrum Disorder (ASD) can affect communication, emotional expression, and social interaction, which may make it difficult for caregivers to recognize changes in emotional and behavioral conditions. Existing assistive applications commonly focus on Augmentative and Alternative Communication (AAC), with less emphasis on combining communication support with continuous emotion and behavioral monitoring. This paper presents Lumi, an AI-assisted communication and emotional health monitoring system that integrates text input, speech-to-text communication, and AAC icons with a fine-tuned DistilBERT model for 18-class emotion recognition and a weighted fuzzy assessment for behavioral risk score. The classifier confidence score is retained as classification-certainty information and is excluded from the behavioral risk calculation. To differentiate emotional contributions, the 18 emotion classes remain distinct during classification and are mapped to five relative concern bands for fuzzy assessment. Lumi also provides caregiver recommendations, automated alerts, emotional-trend visualization, and PDF reports. Its main contribution is the integration of multimodal communication, fine-grained emotion recognition, differentiated emotion contribution, and behavioral monitoring that extends communication-focused support with structured caregiver decision-support information.



Abstract Paper ID: 61

Effect of Transmitter and Receiver Coil Turns on Signal Response in Magnetic Induction Spectroscopy

Magnetic Induction Spectroscopy (MIS) is a non-contact measurement technique with wide applications in biomedical diagnostics and industrial non-destructive testing, in which the system sensitivity is critically determined by transmitter and receiver coil designs. The effect of the number of circular turns of transmitter and receiver coils (N) on the performance MIS system with respect to inductance, phase shift and which determines the magnitude of magnetic field developed has been modeled in this work. To this end, three transmitter coils designs (Tx- 5, Tx-8 and Tx-12) were analyzed in a systematic manner with COMSOL Multiphysics based numerical simulations, while the number of receiver coils was fixed at five turns ($N= 5$) to isolate the effect of transmitter turn variation. The results show that an increasing number of turns enhances inductance ($L \propto N^2$) and the amplitude of the phase shift, the greatest sensitivity of which belongs to Tx-12 (-0.04548° at 10 MHz) compared to Tx-8 (-0.04484°) and Tx-5 (-0.04382°). Quadratic sensitivity dependence on the number of turns and magnetic flux linkage ($\Phi = B \cdot A$) explains the superior sensitivity growth, particularly at higher frequencies (2–10 MHz). These trade-offs occur because excessive turns create parasitic capacitance and thermal noise, which reduce the operational bandwidth. Tx-12 configuration was found most suitable for high-sensitivity applications such as biomedical imaging, whereas less turns (Tx-5/Tx-8) may be sufficient for broadband industrial applications. The findings are aligned with theoretical models (e.g., Ampère's law, Biot-Savart principles) and prior literature, offering actionable suggestions for tailoring MIS systems to specific diagnostic or industrial contexts.



Abstract Paper ID: 62

AI Based Traffic Congestion and Accident Risk Prediction

Traffic congestion and road accidents represent critical challenges for modern urban transportation systems, leading to increased travel delays, economic losses, and safety risks. With the rapid growth of vehicular volume and complex road networks, traditional traffic management approaches are often inadequate in providing accurate short-term congestion forecasts and timely accident risk assessment. This study focuses on the design and implementation of an AI-based framework for short-term traffic congestion prediction and accident risk estimation using real-world spatio-temporal data. The proposed system integrates a Graph Neural Network combined with Long Short-Term Memory (GNN–LSTM) model for congestion forecasting and an XGBoost-based classifier for accident risk prediction. Traffic congestion prediction is performed using the PeMS dataset, where the road network is modeled as a graph to capture spatial dependencies among sensors, while temporal traffic patterns are learned through sequential modeling. Accident risk prediction is implemented using the US Accidents dataset, incorporating traffic, environmental, and temporal features to estimate the probability of crash occurrence. The experimental results demonstrate that the proposed GNN–LSTM model achieves effective short-term congestion forecasting with low prediction error, while the XGBoost model provides reliable accident risk classification with strong discriminatory performance. Evaluation metrics such as Mean Absolute Error (MAE), Root Mean Square Error (RMSE), accuracy, F1-score, and ROC-AUC confirm the effectiveness of the proposed approach. Overall, the results indicate that congestion forecasting with data-driven accident risk prediction can enhance situational awareness and support proactive traffic management and road safety decision-making.



Abstract Paper ID: 63

Classification of Skin Disease using Convolutional Neural Network

Skin disease classification is a challenging medical image analysis task because several lesion categories have similar colour, texture, and boundary characteristics. This article presents a convolutional neural network approach for classifying skin disease images into seven pigmented lesion categories: Actinic Keratosis, Basal Cell Carcinoma, Seborrheic Keratosis, Dermatofibroma, Melanoma, Melanocytic Nevi, and Vascular Lesion. The study is based on the HAM10000 dermatoscopic image dataset and compares MobileNetV1 with Inception-ResNet-v2 for use in a web-based skin disease classification system. Images were resized to 224 by 224 RGB input, divided into training and validation/testing subsets, and augmented to improve representation of minority classes. Model performance was evaluated using accuracy, loss, AUC-ROC behaviour, model size, and training time. MobileNetV1 achieved 0.9905 accuracy, 0.3306 loss, a 26.9 MB model size, and 32 minutes 33.9 seconds training time, while Inception-ResNet-v2 achieved 0.9851 accuracy, 0.3320 loss, a 427 MB model size, and 80 minutes 34.4 seconds training time. The findings indicate that MobileNetV1 provides the better balance between performance and deployment efficiency for the MediSkin web application.



Abstract Paper ID: 64

TG-GCN: A Tongue Graph Convolutional Network for Multi-Disease Classification Using Heterogeneous Tongue Feature Graphs

Medical diagnosis using tongue image classification has advanced substantially with deep convolutional neural networks (CNNs), yet most existing methods treat tongue images as flat, grid-based data and ignore spatial relationships between diagnostic regions and interactions among key tongue features such as color, coating, morphology, and sublingual vasculature. To address these limitations, we propose TG-GCN (Tongue Graph Convolutional Network), a graph-based framework that represents tongue images as heterogeneous, multi-relational graphs and models inter-region dependencies through graph convolution. TG-GCN integrates four complementary graph views: a Regional Feature Graph, a Morphology Superpixel Graph, a Sublingual Vein Graph, and a Multi-scale Patch Graph through a learned heterogeneous adjacency within a three-layer heterogeneous Graph Convolutional Network. We also introduce Tongue Clinical Graph (TCG), a new clinically annotated dataset of 12,480 tongue images spanning Type 2 Diabetes, TCM multi-organ syndromes, Gastric/Gastrointestinal Cancer, and cardiovascular disease. TG-GCN achieves 94.7% accuracy, 0.971 AUC, and 94.2% F1-score on TCG, improving over the best CNN/Transformer baseline by 3.0 accuracy points, and generalizes robustly to external datasets with only a 2.3% accuracy drop compared to 4.1-8.2% for CNN baselines. These experimental results demonstrate that graph-based learning produces more accurate and transferable tongue-diagnostic representations.



Abstract Paper ID: 65

Vision on Wheels: A YOLOv8 and OCR-Powered Smart Trolley for Frictionless Retail Checkout

As retail continues to grow, intelligent automation is required to combat the inefficiencies with manual checkout systems and barcode scanning. Because alternative technologies like RFID have high implementation costs, existing visual detection systems can face challenges such as product occlusion and changes in lighting conditions. To overcome these obstacles, we introduce a smart trolley system comprising AI integrated for an automated, secure self-checkout process. This architecture adopts a dual-stage computer vision pipeline that allows for YOLOv8n for real-time object detection and EasyOCR for secondary verification of objects. On the trolley, a camera tracks items as they are both added and removed, dynamically updating a shopping cart. Since this model prevents shoplifting and billing errors, the OCR module extracts textual data from product packaging and cross-references it with a central store database. Created using the CRISP-DM framework, the system's object detection model has been trained on its own dataset of 10,000 retail images. Experimental testing on 1,967 unseen images showed an outstanding performance, where the precision was 99.9%, the recall was 100%, and the mAP@50–95 was 93.9%. In addition, the model demonstrated a fast inference time of 3.7 milliseconds per frame, which allowed seamless real-time edge processing. Functional testing also verified the robustness of the integrated detection, cart management, and checkout modules. In the end, this multidimensional solution represents a reliable and secure structure for future smart retailing.



Abstract Paper ID: 67

A Comprehensive Review of FACTS-Based Low-Voltage Ride-Through Enhancement in Wind Energy Systems

Ensuring stable wind energy integration during grid faults necessitates robust Low-Voltage Ride-Through capability. This review examines FACTS-based techniques for improving LVRT in wind power systems, specifically analyzing various turbine configurations, FACTS technologies, control methods, and protective measures, while taking into account newly developed intelligent control techniques. Evidence suggests that STATCOMs and SVCs are highly effective for rapid reactive power injection, voltage restoration, and power quality enhancement, whereas Fault Current Limiters notably bolster transient stability. Smart algorithms improve performance, but physical hardware is still the ultimate shield for equipment during voltage drops. Improvements in intelligent fault-mitigation techniques and real-world validation are essential. Future scholarly efforts should prioritize the integration of FACTS devices with energy storage systems and predictive forecasting models, as these are essential to ensuring grid stability given the rising prevalence of wind energy.



Abstract Paper ID: 68

A Security Control Assessment Guidelines for AI-Powered Mental Health Companion Applications

AI-powered mental health companions are increasingly used to provide emotional support, mood tracking, and mental health assistance through conversational interfaces. However, these applications process highly sensitive personal and psychological data, raising significant concerns regarding privacy, security, and ethical data handling. This study investigates the implementation of security controls in AI-powered mental health companion applications and evaluates their effectiveness in protecting user information. A structured methodology based on the Waterfall Model was employed, comprising application selection, requirement analysis, hands-on testing, security control analysis, and recommendation development. Ten widely available AI mental health applications, including Replika, Earkick, Abby, Yuna, and Ash, were analyzed using a security control checklist covering technical and administrative controls such as authentication mechanisms, privacy policies, consent management, data deletion options, session management, and user support features. Scenario-based testing was conducted to examine how the applications responded to privacy-related situations, including crisis detection and data collection consent. The findings reveal that while all applications implemented basic administrative controls, critical technical safeguards such as multi-factor authentication, session timeouts, and bot protection were largely absent. A binomial probability scoring model was used to quantify the level of security control implementation, with the application Ash achieving the highest score by implementing 10 out of 20 evaluated features. The study concludes that current AI-powered mental health companions exhibit significant security gaps that may expose users to privacy and cyber risks. Practical guidelines are proposed to help developers strengthen security measures and improve the protection of sensitive mental health data.



Abstract Paper ID: 69

Smart-Sort: An AI-Based Image Classification System for Automated Household Waste Management

Global recycling programs are severely hindered by means of waste contamination because of "wish cycling". This paper proposes "Smart-Sort," an intelligent system designed to categorize household waste in real-time. Utilizing Convolutional Neural Networks and the Mobile Net framework via Transfer Learning, the system identifies Plastic, Metal, and Paper with high accuracy. The proposed prototype is realized using the use of Google Teachable Machine, providing an evidence-of-concept for computerized sorting at the source.



Abstract Paper ID: 70

U-Net Versus Swin U-Net for Flood Mapping: A Comparative Study Using Multiple Flood Datasets

Flood mapping with remote sensing imagery is crucial for disaster response and assessment of hydrological risk. Even though convolutional neural networks (CNN) based on U-Net architecture and transformer-based segmentation models have both been adopted for flood segmentation, prior works commonly evaluate different architectures on different datasets, preprocessing pipelines, training protocols, and evaluation metrics. Such heterogeneity in evaluation makes it difficult to attribute performance improvements to architectural design. This study aims to provide a controlled comparison of U-Net and Swin U-Net models for flood segmentation, using two distinct flood datasets: the well-known Sen1Floods11 benchmark and the recent STURM-Flood dataset. Both models were trained with the same preprocessing pipeline, optimizer, loss function, learning rate, batch size, number of epochs and metrics. The results reveal high performance of both models: on STURM-Flood, Swin U-Net achieved a Dice score of 0.7243, Intersection over Union (IoU) of 0.6476, Recall of 0.8093 and Precision of 0.7642 after 50 training epochs.



Abstract Paper ID: 71

SmartFruit AI: A CNN-Driven Teachable Machine Solution for User-Friendly Fruit Recognition

This project aims to develop an image-based automated fruit recognition system by leveraging artificial intelligence to address the inefficiencies and error-prone nature of manual fruit sorting. In commercial environments such as supermarkets and fresh produce markets, rapid and accurate classification of different fruit types is essential for quality control and operational efficiency. To tackle this challenge, Teachable Machine, a user-friendly web-based platform, was utilized to train a machine learning model capable of classifying eight distinct fruit categories. The model is built upon a Convolutional Neural Network architecture and employs supervised learning and transfer learning techniques to improve classification accuracy and reduce training time. The system workflow involves capturing fruit images using a standard webcam or smartphone camera, processing them through the trained model, and generating real-time classification predictions. By automating the fruit identification process, the proposed system reduces human error, increases sorting efficiency, and ensures consistent quality assessment, highlighting the practical value of accessible AI tools in addressing everyday logistical challenges.



Abstract Paper ID: 72

Generative AI-Style Sentiment Analysis for Understanding Customer Reviews

While customer reviews are a valuable record of customer perceptions, experiences, and satisfaction, they tend to be emotionally charged, complex, and hard to scale and interpret. This study focuses on the application of sentiment analysis with the help of generative AI to gain insights from customer reviews and generate actionable insights. The study analyzed 1597 product reviews for Amazon products from the Kaggle Amazon Product Reviews Dataset using a three-step approach that involved assistant GenAI-style sentiment and theme coding, VADER baseline sentiment analysis, and a rating-based sentiment classification. The findings reveal a strong positive sentiment in the data, with the assistant GenAI-style coding identifying a genuine mixed sentiment (not captured by VADER or star ratings). Product quality and performance, recommendation and satisfaction, ease of use, display/experience to read and price/value were the most common themes. The results of the agreement comparisons revealed that VADER was generally found to be more likely to label reviews as positive, whereas the rating-based sentiment was only somewhat similar to the text-based sentiment. Based on these results, sentiment analysis using a generative AI approach has the potential to enhance the understanding of customer reviews by capturing nuances in customer sentiment, identifying recurring themes, and offering practical advice. The study brings a clear methodological approach to integrate the use of generative AI sentiment analysis with baseline and rating-based comparison.



Abstract Paper ID: 73

Explainable Machine Learning Framework for Technical Loss Prediction in Power Distribution Systems

Although machine learning has demonstrated strong predictive performance in energy applications, many existing models operate as black boxes, limiting transparency and trust in engineering decision-making. Although conventional analytical techniques have a physical foundation, they are frequently computationally demanding and challenging to scale for vast networks. By utilizing operational data, this study proposes an explainable machine learning framework for technical loss prediction in power distribution systems that integrates Least Absolute Shrinkage and Selection Operator (LASSO) regression with SHapley Additive exPlanations (SHAP). Support Vector Regression (SVR) is included as a comparative benchmark to evaluate predictive performance. The proposed framework is validated using a real-world distribution network dataset comprising operational, electrical, and feeder-related characteristics. Performance is evaluated against analytical calculations, simulation results, and measured system losses. Experimental results show that the proposed LASSO model predicts total technical losses of 231.50 MWh, corresponding to an estimation error of 1.5% relative to measured losses, outperforming the analytical approach (7.1%) and slightly exceeding the accuracy of SVR (3.1%). SHAP analysis further reveals that maximum demand, peak power losses, and infeed energy are the most influential features contributing to model predictions, providing transparent explanations that align with established power system principles. The findings demonstrate that explainable machine learning can provide accurate, computationally efficient, and interpretable technical loss prediction, supporting trustworthy AI deployment in critical energy infrastructure.



Abstract Paper ID: 74

A Deep Learning Approach For Automated PCB Defect Detection Using FPN-R-FCN

Accurate and reliable defect detection in Printed Circuit Boards (PCBs) is an indispensable requirement for ensuring long-term product longevity and minimizing manufacturing losses. While current deep learning research has advanced significantly, focuses primarily on tiny defect detection and using the mAP as a primary benchmark, a critical flaw persists standard mAP metrics treat all defect classes equally, weighting functionally critical defects and minor cosmetic flaws with the same importance. This research addresses the issue by establishing a systematic taxonomy of PCB defects, classifying them into critical, major, and minor categories based on their definition and functional impact. Following this classification, a specialized detection framework is proposed, utilizing a R-FCN integrated with a multi-scale feature fusion strategy inspired by FPN principles. Quantitative validation and qualitative visual analysis reveal that the proposed method achieves a high mAP of 94.0% across six defect classes and 93.1% specifically for critical defects. Comparative evaluations demonstrate that the proposed model is 14.3% higher than the YOLO v4-Tiny baseline. Although its accuracy is 2.0% lower than the Extended FPN model, the proposed architecture delivers an inference speed of 35.4 FPS, which is approximately three times faster than Extended FPN.



Abstract Paper ID: 75

Evaluating ROI Feature Extraction Methods in BF- CloU Faster R-CNN for Steel Surface Defect Detection

Detection of low-contrast steel surface defects, particularly crazing and rolled-in scale, remains a challenge on the Faster R-CNN architecture standard. Previous research introduced the BF-CIoU Faster R-CNN that integrates selective preprocessing, custom anchor, ROI Align, and CloU Loss, but the evaluation was limited to two classes of defects without a systematic ablation study of the contribution of each component. This study evaluated the architecture in all six classes of NEU-DET defects through seven controlled experiments, focusing on the comparison of ROI Align and ROI Pooling as feature extraction methods. The results showed that both methods produced almost equal performance (mAP@50 57.31% compared to 58.63%), with a difference of 1.32 points, much smaller than the contribution of the anchor configuration of 9.57 points and the preprocessing strategy of 6.18 points. These findings indicate that in low-resolution datasets such as NEU-DET, the choice of ROI extraction method is not the main determinant of detection performance. The allocation of optimization efforts is more effectively directed at anchor configuration and preprocessing strategies that have a greater impact on overall detection accuracy.