



SCHOOL OF ENGINEERING & TECHNOLOGY

UNIVERSITY of WASHINGTON | TACOMA

School of Engineering and Technology (SET) Research Showcase. May 1, 2026.

Agenda

11-11:50 Networking event (CP 106)

12:00 pm: Introductions by Dean Raj Katti

12:05 - 12:45: Lightning talks (WPH)

12:45 - 1:30 Posters round 1 (poster #1-#13)

1:30- 2:15 Posters round 2 (poster #14-#26)

2:15 - 2:45 Lab tours (concurrent) + Networking continued

Radio Frequency Wireless IC Laboratory in CP 206M (Professor Debasis Dawn,
Electrical and Computer Engineering)

Information Technology Networking Lab in CP206C (Professor Yan Bai and Associate
Teaching Professor Charles Costarella)

Civil Engineering lab tour in MLG 128 (Assistant Professor Lorne Arnold, Assistant
Teaching Professor Nara Almeida)

Mechanical Engineering lab tour (ME Senior: Sam Babin) basement of GWP: GWP G30,
GWP G40, BB G40, WCG G60.

2:45-3:00 Awards

Acknowledgements

Judges for the Best Poster Award: Dean Raj Katti, Professor Joel Baker, Andrew Cole (last year's winner)

SET staff: Marife Tabao, Gabi Crosby, Norma Robles

Steering Committee Members: Eyhab Al-Masri, Cassandra Donatelli, Wenjun (Eric) Fan, Seyyed Babak Hamidi, Ka Yee Yeung

Review Committee Members:

Civil Engr: Nara Almeida

CSS: Eyhab Al-Masri, Xutong Liu, Ka Yee Yeung

ECE: Seyyed Babak Hamidi, Thillainathan Logenthiran

IT: Wenjun (Eric) Fan, Deveeshree Nayak

ME: Cassandra Donatelli, Matthew Ford

LIGHTNING PRESENTATIONS AT-A-GLANCE

(ordered by poster numbers)

Poster #	Primary presenter	Title
1	Syed Huzaifa Bin Afzal	Are You Aware of Shadow AI? GovernAI for Addressing Emerging Risks
2	Rupali Singh	Just-In-Time Access for Machine Identities: Reducing Exposure Windows in Microservices
3	Ling-Hong Hung	STAR Suite: Integrating transcriptomics through AI-assisted software engineering
4	Rohan Avireddy	SPAgent: using an AI-agent to facilitate interpretation of spatial proteomics analysis
5	Patrick McKeever	Interactive data integration dashboard for the MorPhiC program
6	Xiaoling Wei	Mind the Attack: How security threats are shaping software deployed on Unmanned Aerial Vehicles
7	Kevin Michalson	Rules of Engagement: How Video Games Justify State Violence

POSTER PRESENTATIONS AT-A-GLANCE

(grouped to allow authors to present multiple posters)

1. **Are You Aware of Shadow AI? GovernAI for Addressing Emerging Risks** - Syed Huzaifa Bin Afzal, Justin L. Bowman, Renee Chiang, Ayush Singh, Wenjun Fan (MCL, Information Technology)
2. **Just-In-Time Access for Machine Identities: Reducing Exposure Windows in Microservices** - Rupali Singh, Corrazon Ogot, Wenjun Fan (MCL, Information Technology)
3. **STAR Suite: Integrating transcriptomics through AI-assisted software engineering** - Ling-Hong Hung, Ka Yee Yeung (Computer Science & Systems)
4. **SPAgent: using an AI-agent to facilitate interpretation of spatial proteomics analysis** - Naga Venkata Sai Rama Krishna Rohan Avireddy, Shrey Srivastava, Bryce Fukuda, Jocelyn Wright, Kimberly Smythe, Cecilia Yeung, Ling-Hong Hung, Ka Yee Yeung (Computer Science & Systems)
5. **Interactive data integration dashboard for the MorPhiC program** - Patrick McKeever, Jesse Flores, Bryce Fukuda, Ling-Hong Hung, Ka Yee Yeung (Computer Science & Systems)
6. **Mind the Attack: How security threats are shaping software deployed on Unmanned Aerial Vehicles** - Xiaoling Wei, Jacob Klymenko, Damiano Torre (Computer Science & Systems), Davide Fucci (Blekinge Institute of Technology, Sweden)
7. **Rules of Engagement: How Video Games Justify State Violence** - Kevin Michalson, Joanne Clarke Dillman (Computer Science & Systems, Global Honors)
8. **Real-Time Linear Prediction Coding of Speech on STM32F7 Discovery Board** - Kadyrzhan Tortayev, Jie Sheng (Electrical and Computer Engineering)
9. **Real-Time Power Quality Monitoring and Disturbance Detection using DSP Techniques on the STM32F746G Discovery Board** - Thillainathan Logenthiran, Jie Sheng, Tanmaya Sreeshekhar (Electrical and Computer Engineering)
10. **Software-Engineering Benchmarks for Explainable Intrusion Detection Systems in Autonomous Platforms** - Yanhong Miao, Damiano Torre (Computer Science & Systems)
11. **A Systematic Evaluation Of Accessibility Policies In U.S. Shared Micromobility Programs** - , Brandon Meyers (Mechanical Engineering), Jeff Walters (Civil Engineering), Alexander Brown (Mechanical Engineering), Mikell Smith (Civil Engineering), Robin Watkins (Civil Engineering), Angela Kitali (Civil Engineering), Heather Dillon (Mechanical Engineering), Anne Brown (College of Design - University of Oregon)
12. **Optimizing Hydrothermal Alkaline Treatment at Milder Conditions for PFAS Remediation** - Emese Hadnagy, Ansa V. Karim, Brian Pinkard, Annie Heble, Conrad Austin, Joel Baker (Civil Engineering)
13. **Tales of Scales: The Relationship Between Hydrodynamics and Antifouling in Pleuronectidae** - Ellen Misun Pak (Biomedical Sciences), Isabella Claeson (Mechanical Engineering), Megan L. Vandenberg (Biology - UW Seattle), Karly E. Cohen (Friday Harbor Laboratories), Cassandra M. Donatelli (Mechanical Engineering)

14. **Numerical Simulation of Dielectric Rewarming to Prevent Thermal Runaway in Cryopreserved Organs** - Hwisung (Andrew) Kim, Rebekah Cureg, Ruidong Ma, Zhiquan (Andy) Shu (Mechanical Engineering)
15. **Towards Sustainable Energy-Storing Concrete: Zeolite as a Functional SCM for Supercapacitor Applications** - David Hernandez (Mechanical Engineering), Aspen Dyea (Civil Engineering), Nara Almeida (Civil Engineering), Cassandra Donatelli (Mechanical Engineering), Orlando Baiocchi (Electrical Engineering)
16. **Subsurface Mapping from Historical Construction Data** - Emmaly Larsson, Ryan Carson, Daniel Berton, David Hopkins, Raphaela Plameras, Lorne Arnold (Civil Engineering)
17. **Standardization of Power Distribution Centers (PDC) for Combined Cycle Power Plants (CCPP)** - Students: Qaturi Vaughn, Refaa Aziz, Alejandra J. Malca Perez, and Muna Mohamed. Faculty advisor: Thillainathan Logenthiran (Logan) Industry Sponsor: Kindle Energy (Joey Redmond) (Electrical and Computer Engineering)
18. **Case Study of Digital Predistortion for Power Amplifier Linearization** - Jie Sheng, Sheethal Christy Anthony, Mohammed Jasim (Electrical and Computer Engineering)
19. **Decentralized Communication and Scheduling of Microgrids Based PoolCoElectric Market** - Beserat Tewolde, Thillainathan Logenthiran, Vahid Dargahi (Electrical and Computer Engineering)
20. **MIRA: Multi-Vector Retrieval with Hybrid Indexing and Adaptive Query Processing for Web-Scale Search** - Solmaz Seyed Monir, Dongfang Zhao (Computer Science & Systems)
21. **LDP-MIC: Correlation-Aware Local Differential Privacy for Federated Learning under Untrusted Aggregation** - Wenjun Yang, Eyhab Al-masri, Olivera Kotevska (Computer Science & Systems)
22. **QoS-Aware AI Agents for Intelligent API Selection and Composition** - Ishwarya Narayana Subramanian, Eyhab Al-Masri (Computer Science & Systems)
23. **Enhanced Road Crack Detection** - Reetika Chauhan, Sixian Zheng, Wei Cheng, Juhua Hu (Computer Science & Systems)
24. **Agentic AI Serverless Code Generation: Towards Autonomous Improvement of Performance, Cost, and Code Quality** - Xinghan Chen, Robert Cordingly, Ling-Hong Hung, Wes Lloyd (Computer Science & Systems)
25. **Inference Profiler: Enabling Resource Utilization Profiling of LLM inferencing** - Austin Bomhold, Morteza Nabavinejad, Wes Lloyd (Computer Science & Systems)
26. **Impacts to Public Key Infrastructure from Post-Quantum Cryptography** - Sofian Abbasi, Jaren Thorsen (Computer Science & Systems, Information Technology, Master of Cybersecurity & Leadership (MCL))

ABSTRACTS OF PRESENTATIONS

Poster 1 Are You Aware of Shadow AI? GovernAI for Addressing Emerging Risks

Authors: Syed Huzaiifa, Bin Afzal, Justin L. Bowman, Renee Chiang, Ayush Singh

Advisor: Wenjun Fan

Program: MCL, Information Technology

Abstract: Shadow AI refers to employees using public or unsanctioned generative-AI tools outside enterprise IT and security control, increasing the risk of data leakage, compliance violations, and unobservable governance gaps. To reduce these risks without blocking productivity, we designed and evaluated a lightweight, on-premises generative-AI safe harbor that organizations can operate within their own security boundary. The platform combines Ollama for local large language model inference and Open WebUI for a browser-based interface, deployed on a Windows Server host using Docker. It enforces data locality by keeping prompts, outputs, and uploaded documents on-premises, and it supports role-based access control and administrative visibility for basic governance. We validated data residency and storage behavior using host-level monitoring and controlled firewall rules, and we assessed access boundaries across standard users, AI platform administrators, and system administrators. Initial results show the platform remains functional without outbound connectivity and stores user uploads and chat history in a known persistent volume, enabling explicit auditing and retention decisions. Usability testing across common workflows suggests response times can be comparable to leading public tools on suitable local hardware, improving the likelihood that users will adopt the sanctioned platform instead of external services. Ongoing work focuses on fine-grained auditing, stronger isolation, and integration with enterprise governance frameworks.

Poster 2 Just-In-Time Access for Machine Identities: Reducing Exposure Windows in Microservices

Authors: Rupali Singh, Corrazaon Ogot

Advisor: Wenjun Fan

Program: MCL, Information Technology

Abstract: Modern microservices commonly rely on stateless JSON Web Tokens (JWTs) for machine-to-machine authentication, but their non-revocable nature creates a critical security risk: once issued, a JWT remains valid until expiration, even if the associated client secret is compromised. Default configurations such as Keycloak's 300 second access token lifespan create multi-minute exposure windows that attackers can exploit. This work presents a practical Just-in-Time (JIT) access mechanism that significantly reduces this risk by enforcing extremely short-lived machine credentials without requiring token introspection, revocation lists, or architectural changes to existing services.

Using Keycloak as the identity provider, we configure a 60-second Access Token Lifespan for OAuth 2.0 client-credentials-based machine identities and implement automatic token rotation. A fully containerized testbed including a protected resource server, a legitimate microservice client, and an attacker simulation demonstrates deterministic token invalidation at the 60 second boundary and instantaneous propagation of authorization policy changes. Experimental results

show an 80% reduction in the exposure window compared to default settings, while preserving the scalability benefits of stateless authentication.

This work validates sub-minute token lifetimes as a lightweight, standards-compliant mitigation strategy for securing non-human identities in distributed systems and provides an immediately deployable approach for production microservices environments.

Poster 3 STAR Suite: Integrating transcriptomics through AI-assisted software engineering

Authors: Ling-Hong Hung, Ka Yee Yeung

Program: Computer Science & Systems

Abstract: To accommodate rapid methodological turnover, bioinformatics pipelines typically consist of discrete binaries linked via scripts. While flexible, this architecture relies on intermediate files, sacrificing performance, and treats complex codebases as static silos. For example, the STAR aligner—the standard engine for transcriptomics—uses an external script for adapter trimming, necessitating the decompression and re-compression of massive files. These limitations presented scalability problems for uniform processing of data in the NIH Morphic consortium. We present our solution, STAR Suite, a human-engineered and AI-implemented overhaul that integrates functionality directly into the C++ source. In just three months, a single developer added 80,000 lines to the original 27,000-line codebase to produce four unified modules: STAR-core, STAR-Flex, STAR-perturb, and STAR-SLAM without introducing any new dependencies and can be installed as a pre-compiled binary. These four modules represent an integrated transcriptomics software package that can process diverse omics data types, including RNA sequencing (RNA-seq), single cell RNA-seq, perturb-seq, and SLAM-seq data. Using STAR as a case study, this work demonstrates an AI-guided development framework that facilitates community contributions to (and maintenance of) well-established bioinformatics software. Most importantly, our methodology can be used to streamline extension of high-performance bioinformatics software for new applications.

Poster 4 SPAgent: using an AI-agent to facilitate interpretation of spatial proteomics analysis

Authors: Rohan Avireddy, Shrey Srivastava, Bryce Fukuda, Jocelyn Wright, Kimberly Smythe, Cecilia Yeung, Ling-Hong Hung, Ka Yee Yeung

Advisor: Ling-Hong Hung, Ka Yee Yeung

Program: Computer Science & Systems

Abstract: Precision oncology relies on spatial proteomics to link high-resolution tissue images with molecular measurements and guide immunotherapy decisions. Despite automated analysis pipelines, researchers still manually reconcile large image sets with high-dimensional tables to produce interpretable reports. We present SPAgent, a web-based platform that automatically drafts structured scientific reports while enforcing claim-evidence traceability. SPAgent treats a report as a structured object backed by a registry of artifacts (images, tables etc.) and ingests outputs from our spatial proteomics pipeline (DAPI images, segmentation overlays, cluster maps, UMAPs, marker heatmaps, and quantitative tables). The pipeline performs artifact parsing and metadata extraction, table summarization for quantitative context, vision-language analysis for per-figure captions and insights, and provenance-aware assembly that preserves evidence links. The UI synchronizes report text with figures and tables, supports click-through verification, shows or hides AI reasoning, logs prompts for reproducibility, and can fetch relevant

literature with in-text citations. A VLM-assisted 'Query a report' mode answers user questions using report content. SPAgent reduces reporting bottlenecks and improves transparency, and is being refined with feedback from pathologists and biomedical scientists at the Fred Hutchinson Cancer Center (FHCC).

Poster 5 Interactive data integration dashboard for the MorPhiC program

Authors: Patrick McKeever, Jesse Flores, Bryce Fukuda, Ling-Hong Hung, Ka Yee Yeung

Advisor: Ling-Hong Hung, Ka Yee Yeung

Program: Computer Science & Systems

Abstract: Uniform processing of bulk RNA sequencing (RNA-seq) data from different studies is essential for integrated analysis, but it is time-consuming and difficult. Though the NCBI Sequence Read Archive (SRA) contains publicly available RNA-seq data, the process of identifying relevant samples, specifying experimental conditions, and harmonized processing across multiple datasets requires substantial expertise. We present a user-friendly dashboard which enables users to uniformly process and perform integrated analysis of multiple bulk RNA-seq datasets with a single click. The dashboard consists of an interactive graphical interface leveraging a standard semantic search and SRA study metadata. Users can select individual samples, assign samples to groups reflecting experimental design, and execute the approved RNA-seq pipeline of the MORPHIC project. This pipeline includes steps for adapter trimming (via cutadapt), alignment to hg38 (via STAR), and read quantification (via salmon), as well as an analysis step which outputs heatmaps and volcano plots identifying differentially expressed genes between user-specified experimental conditions. Our approach addresses several of the "pain points" of bulk RNA harmonization, including the standardized analysis of single and paired end samples. Additionally, our dashboard is integrated with the published datasets of the MORPHIC project.

Poster 6 Mind the Attack: How security threats are shaping software deployed on Unmanned Aerial Vehicles

Authors: Xiaoling Wei, Jacob Klymenko, Damiano Torre, Davide Fucci

Advisor: Damiano Torre

Program: Computer Science & Systems

Abstract: Context Autonomous unmanned aerial vehicles (UAVs) are increasingly used in safety-critical applications, making security a major concern. However, research on UAV attacks and defenses remains fragmented. This paper provides a systematic, software-oriented overview of security attacks against autonomous UAVs, identifying attack types, affected software components, and mitigation strategies. We also assess empirical rigor and reproducibility practices. We conduct a systematic mapping study of 140 peer-reviewed primary studies selected from an initial set of 1,557 papers. Guided by eleven research questions, we analyze attacks, targeted software components, exploited vulnerabilities, and detection and mitigation techniques. Research predominantly focuses on communication and navigation layer attacks, particularly GPS spoofing and denial of service, while autonomy centric components such as perception and decision logic remain underexplored. A hierarchical taxonomy mapping attacks to software functionalities reveals a strong reliance on machine learning for detection, with limited use of control-theoretic and formal methods. Evaluations depend heavily on

simulation, and the lack of open UAV datasets and artifacts significantly hinders reproducibility. We provide six actionable recommendations for the UAV software engineering community. Our findings emphasize that autonomy shifts the security unit of analysis to the end-to-end autonomy pipeline, requiring multi-layer, closed-loop, and reproducible evaluation approaches

Poster 7 Rules of Engagement: How Video Games Justify State Violence

Authors: Kevin Michalson, Joanne Clarke Dillman

Program: Computer Science & Systems, Global Honors

Abstract: We have shifted from primarily experiencing war through passive media such as news and history books to interactive simulations through video games. The Military-Entertainment Complex (MIME-NET) serves as a collaboration between entertainment companies and state defense agencies, and video games have begun serving as tools of soft power, designed to export national ideology and sanitize the reality of state violence for a global audience. I analyze how the gameplay of titles forces players to accept foreign policy worldviews. I conduct a comparative analysis of mechanics between Western AAA titles such as Call of Duty and Battlefield and non-Western titles, such as Russia's Syrian Warfare and China's Glorious Mission. I conduct this analysis through three specific rhetorical lenses: Distancing, the dehumanization of enemies via technology; Sanitization, the removal of consequences to create clean war; and Glorification, the gamification of violence through rewards. We are entering an era of memory wars in which historical truth is being defined by the nationality of the game studio whose games someone has played. It is manufacturing a generation of youth that consents to specific geopolitical narratives.

Poster 8 Real-Time Linear Prediction Coding of Speech on STM32F7 Discovery Board

Authors: Kadyrzhan Tortayev, Jie Sheng

Advisor: Jie Sheng

Program: Electrical and Computer Engineering

Abstract: This project investigates the implementation of Linear Prediction Coding (LPC) for speech compression on an embedded system. Using the STM32F7 Discovery board, a real-time system was developed to model speech production using the source-filter model. The system processes audio at 16 kHz, estimating the vocal tract transfer function using the Levinson-Durbin recursion algorithm with model orders up to $N = 30$. Experimental results demonstrate that LPC successfully approximates the spectral envelope of speech, with higher orders providing significantly better formant resolution. The implementation achieves a theoretical compression ratio of approximately 20:1 compared to raw 16-bit PCM audio, confirming the efficacy of parametric coding for bandwidth-constrained applications.

Poster 9 Real-Time Power Quality Monitoring and Disturbance Detection using DSP Techniques on the STM32F746G Discovery Board

Authors: Tanmaya Sreeshekhar, Jie Sheng, Thillainathan Logenthiran

Advisor: Jie Sheng, Thillainathan Logenthiran

Program: Electrical and Computer Engineering

Abstract: Power Quality disturbances such as voltage sags, swells, harmonics, and transient events significantly impact the reliability and efficiency of modern electrical systems. The project

presents a real-time power quality monitoring and disturbance detection system implemented on the STM32F746G Discovery Board using Digital Signal Processing (DSP) techniques.

The system acquires voltage signals and processes them using embedded DSP algorithms to extract key power quality metrics, including RMS voltage, Total Harmonic Distortion (THD), and frequency variations. Advanced signal processing methods such as Fast Fourier Transform (FFT) and threshold-based event detection are employed to identify and classify disturbances in real time. The implementation emphasizes computational efficiency and low-latency performance, enabling continuous monitoring on a resource-constrained embedded platform. Experimental results demonstrate accurate detection of common disturbances and reliable real-time operation. This work contributes to the development of cost-effective, portable, and scalable solutions for smart grid monitoring and industrial power systems. The proposed system can be extended for predictive maintenance and integrated into IoT-based energy management frameworks.

Poster 10 Software-Engineering Benchmarks for Explainable Intrusion Detection Systems in Autonomous Platforms

Authors: Yanhong Miao, Damiano Torre

Advisor: Damiano Torre

Program: Computer Science & Systems

Abstract: Intrusion Detection Systems (IDS) for autonomous platforms — such as drones, mobile robots, and IoT/edge devices — face increasing demands for explainability, reproducibility, and software-engineering reliability. Existing IDS and XAI research suffers from three key limitations: dataset inconsistency that prevents fair comparison, evaluation bias toward machine-learning-only metrics, and low-quality explanations that provide limited practical value to engineers deploying autonomous systems. This work-in-progress addresses these gaps by developing a research-grade benchmark for Explainable IDS (X-IDS) tailored to autonomous and IoT environments. The benchmark will include curated datasets, baseline IDS models (tree/ensemble methods, 1D CNN, and autoencoder), and XAI adapters using SHAP, LIME, and permutation-based interpretation. Evaluation criteria will combine machine-learning performance metrics with software-engineering metrics capturing maintainability, integration effort, explanation clarity, and reproducibility. Current progress includes replication studies of foundational X-IDS papers and initial dataset preprocessing using CICIDS2017, with full benchmark implementation ongoing.

Poster 11 A Systematic Evaluation Of Accessibility Policies In U.S. Shared Micromobility Programs

Authors: Brandon Meyers, Jeff Walters, Alexander Brown, Mikell Smith, Robin Watkins, Angela Kitali, Heather Dillon, Anne Brown

Adviser: Jeff Walters

Program: Mechanical Engineering; Civil Engineering

Abstract: Shared micromobility programs (e-bikes, e-scooters, and similar low-speed devices), have expanded rapidly across the United States, yet concerns remain about uneven access, rising costs, and limited regulatory oversight. This study provides the first systematic, nationwide assessment of equity policies across 239 shared micromobility programs, drawing on hundreds

of policy documents, including MOUs, RFPs, regulatory guidelines, and operator agreements. Using a mixed deductive–inductive qualitative coding approach, we examined how agencies and operators define, implement, and monitor equity-oriented mechanisms such as smartphone-free access, cash payment compatibility, reduced-fare programs, multilingual services, adaptive vehicles for users with disabilities, and mandated geographic service areas. Our findings show wide variation in how equity is conceptualized and operationalized: while many programs require operators (i.e., Lime and Razor) to provide alternative access pathways and inclusive service features, the depth and enforceability of these requirements differ substantially. This gap between policy intent and measurable accountability underscores persistent challenges in ensuring that shared micromobility systems equitably serve underserved communities. By comparing programs across diverse geographic and institutional contexts, this study identifies emerging best practices, highlights structural shortcomings, and provides actionable guidance for strengthening equity policy design, monitoring, and implementation in future shared micromobility programs.

Poster 12 Optimizing Hydrothermal Alkaline Treatment at Milder Conditions for PFAS Remediation

Authors: Ansaf V. Karim, Brian Pinkard, Annie Heble, Conrad Austin, Joel Baker, and Emese Hadnagy

Program: Civil Engineering

Abstract: The increasing occurrence of per- and poly-fluoroalkyl substances (PFAS) in environmental matrices is a growing concern. Existing treatment technologies, such as biological processes, advanced oxidation, membrane filtration, adsorption, etc., are insufficient for complete degradation and defluorination of PFAS compounds. Hydrothermal alkaline treatment (HALT) is a proven destructive treatment approach for the complete mineralization of PFAS at high temperature (350 °C), pressure (25 MPa), and under highly alkaline conditions in concentrated waste matrices. The current study aims to optimize the treatment process under milder treatment conditions, specifically, lower temperatures (<250 °C) and reduced chemical dosage (<0.5 M NaOH) with the goal of creating an economical method for more dilute aqueous waste matrices. Batch experiments were conducted to determine the optimal reaction temperature, reaction time, and chemical dosage for maximizing the degradation of fluorotelomer sulfonate (6:2 FTS) and 5:3 fluorotelomer carboxylic acid (5:3 FTCA), two commonly occurring compounds in municipal wastewater. Experiments conducted at 175 °C and 2 M NaOH achieved 97.6% degradation of 6:2 FTS and 99.8% degradation of 5:3 FTCA within 60 minutes, demonstrating the high efficiency of the process.

Poster 13 Tales of Scales: The Relationship Between Hydrodynamics and Antifouling in Pleuronectidae

Authors: Ellen Misun Pak, Isabella Claeson, Megan L. Vandenberg, Karly E. Cohen, Cassandra M. Donatelli

Advisor: Cassandra M. Donatelli

Program: Biomedical Sciences; Mechanical Engineering

Abstract: The diverse surface textures of fishes serve many functions. Particularly, fish scales are known to enhance bioaccumulation resistance (antifouling) and hydrodynamic performance

in different ways via micro- and nano-scale surface structures, making them a good source of inspiration for bioinspired models. However, most studies observe these functions independently, so the relationship between hydrodynamics and antifouling features remains unclear. Insight into these interactions could guide more efficient and sustainable maritime technologies. Flounders swim in the water column and bury in the substrate, benefiting from having scales with both drag reducing and anti-fouling features. Our preliminary data showed that flounder models produced a wake pattern conducive to both. We used Keyence microscopy and μ CT to describe microscale topographies of flounders to obtain a standard for bioinspired 3D models based on shape and arrangement. A hydrodynamic flume was used to quantify drag with particle image velocimetry. Models were submerged across two sites in the Salish Sea, Tacoma and Friday Harbor, with differing wave action to evaluate fouling. We believe that models with a straight spike arrangement may excel in drag reduction, but with less antifouling capability, while models with a staggered spike pattern may have a better balance of drag reduction and antifouling.

Poster 14 Numerical Simulation of Dielectric Rewarming to Prevent Thermal Runaway in Cryopreserved Organs

Authors: Hwisung (Andrew) Kim, Rebekah Cureg, Ruidong Ma, Zhiquan (Andy) Shu

Advisor: Zhiquan (Andy) Shu

Program: Mechanical Engineering

Abstract: The rewarming process of tissues and organs in cryopreservation requires a rapid rewarming rate and uniform temperature distribution to avoid ice recrystallization and temperature gradient-induced thermal stress. Dielectric heating utilizes the electric field to rewarm cryopreserved samples fast; however, the thermal runaway effect limits the use of electric field heating, where parts of the sample that have higher temperature can absorb more energy resulting in an increasing temperature gradient between the hot and cold regions, i.e., a phenomenon called thermal runaway. This leads to escalated uneven warming and severe damage to the biological samples, highlighting the need for understanding of this thermal runaway phenomenon and precise control of the rewarming process. Thus, this study utilizes numerical simulation to study the thermal runaway effect and design optimized cryoprotective agents (CPAs) to mitigate this impact. We first study the temperature profiles during heating in a normal cryopreserved sample (i.e., dielectric loss increases with temperature) when either the electromagnetic field or sample dielectric properties are nonhomogeneous, which explains the origins of thermal runaway phenomenon. We then proposed novel CPAs, which present an inverted U-shaped correlation between dielectric loss and temperature. Later, we compare the temperature profiles during heating of samples for the optimized CPAs and other normal CPAs. The study reveals the mechanism of thermal runaway effect during dielectric heating and provides guidance on the design of CPAs to mitigate this effect.

Poster 15 Towards Sustainable Energy-Storing Concrete: Zeolite as a Functional SCM for Supercapacitor Applications

Authors: Aspen Dyea, David Hernandez, Nara Almeida, Cassandra Donatello, Orlando Baiocchi

Advisor: Nara Almeida

Program: Civil Engineering; Mechanical Engineering; Electrical Engineering

Abstract: This study evaluates zeolite as a sustainable alternative to carbon black in cement-based supercapacitors, as recent research has shown that conductive carbon additives can enable cementitious materials to function as electrodes in energy storage systems. Zeolite is investigated as a supplementary cementitious material (SCM) to increase internal surface area and pore tortuosity - key parameters governing ion transport and energy storage capacity. Potassium chloride (2M KCl) is incorporated into the mixing water to serve as the electrolyte. Three mix designs were prepared: (1) control (no zeolite, no KCl), (2) 15% zeolite as SCM without KCl, and (3) 15% zeolite as SCM with 2M KCl. Mechanical performance was assessed using nine 25.4 mm cube specimens per mix tested in compression at 7 and 28 days. Electrical performance was evaluated using cylindrical specimens (25 mm diameter × 15 mm height) fabricated from the same mixes and subjected to charge-discharge testing. Resistance and time constants were measured to calculate capacitance. Preliminary results are promising, showing measurable capacitance in specimens containing 15% zeolite with 2M KCl and 7-day compressive strengths exceeding 3,000 psi across all mixes. Ongoing CT scanning will quantify porosity and tortuosity to evaluate microstructural evolution and long-term performance.

Poster 16 Subsurface Mapping from Historical Construction Data

Authors: Emmaly Larsson, Ryan Carson, Daniel Berton, David Hopkins, Raphaela Plameras, Lorne Arnold

Advisor: Lorne Arnold

Program: Civil Engineering

Abstract: Geotechnical engineers analyze subsurface data and foundations to ensure soil resistance is adequate for the design loads of structures. For driven pile foundations, common in waterfront structures, construction logs record installation energy throughout the entire installation. Yet standard practice uses only the end-of-drive data to characterize the final bearing condition of each pile. The intermediate data, if properly interpreted, can reveal significant information about subsurface stratigraphy. This untapped resource is particularly valuable given that subsurface uncertainty remains one of the primary drivers of risk in civil engineering projects. This study presents a 3D subsurface characterization of the Port of Tacoma Pier 4 Reconfiguration project site, developed by digitizing hundreds of handwritten pile driving logs. Penetration resistance and applied energy data recorded throughout pile installation serve as a proxy for soil resistance at depth, enabling spatial mapping of capacity across the site. Our data processing evolved from manual data entry to AI-assisted digitization with human review. The resulting model is consistent with the geologic environment and resolves stratigraphic units at higher spatial resolution than any previous subsurface investigation at the Port, forming the basis for new methods to predict long-term pile capacity with reduced uncertainty.

Poster 17 Standardization of Power Distribution Centers (PDC) for Combined Cycle Power Plants (CCPP)

Authors: Students: Refaa Aziz, Alejandra J. Malca Perez, Muna Mohamed, and Qaturi Vaughn. Faculty advisor: Thillainathan Logenthiran (Logan) Industry Sponsor: Kindle Energy (Joey Redmond)

Advisor: Thillainathan Logenthiran (Logan)

Program: Electrical and Computer Engineering

Abstract: As electricity demand increases due to rapid growth of data centers, electrification of transportation, and industrial expansion, combined cycle power plants (CCPPs) must be designed and constructed more efficiently to meet energy needs. A critical component of these facilities are Power Distribution Centers (PDC), which house electrical equipment such as switchgear, motor control centers, and protective devices responsible for distributing power throughout the plant. Currently, PDC layouts are custom designed for each project despite commonalities among electrical load requirements, leading to repeated engineering effort, increased costs, and extended project timelines. This project aims to develop standardized PDC configurations for Kindle Energy's CCPP facilities to reduce design redundancy and enable earlier equipment procurement. The project focuses on three primary PDC systems: the Main, Heat Recovery Steam Generator (HRSG), and Air Cooled Condenser (ACC). Existing electrical drawings, load lists, and one line diagrams were analyzed to identify common load patterns and equipment requirements while ensuring compliance with industry codes and internal design standards. By standardizing PDC designs, the project seeks to support pre purchasing strategies and shorten overall plant construction timelines, potentially reducing project completion from approximately three to four years to around eighteen months while maintaining system reliability, safety, and scalability.

Poster 18 Case Study of Digital Predistortion for Power Amplifier Linearization

Authors: Sheethal Christy Anthony, Jie Sheng, Mohammed Jasim

Program: Electrical and Computer Engineering

Abstract: Power amplifiers used in modern wireless transmitters exhibit nonlinear behavior when operated near saturation, resulting in spectral regrowth and signal distortion. This research presents a case study to demonstrate how Digital Pre-distortion (DPD) algorithm can improve signal quality using both simulation and embedded hardware implementation. Specifically, DPD coefficients are trained through MATLAB, and further the DPD algorithm gets implemented on an STM32 board.

Our approach generated a baseband signal first using MATLAB. It is then sent to the STM32 board and passed through a nonlinear analog circuit emulating power amplifier behavior. The output signal exhibits nonlinear distortion, which is observed and analyzed using spectral plots and nonlinear characteristics such as AM-AM and AM-PM.

Using the measured input and output signals, DPD coefficients got trained through MATLAB, and then applied to the input signal sending to the STM32 board before transmission through the nonlinear circuit. As a result, the distortion introduced by the amplifier is reduced, and the output signal demonstrates improved linearity.

Poster 19 Decentralized communication and Scheduling Microgrid Based PoolCo Electric Market

Authors: Beserat Tewolde, Thillainathan Logenthiran

Advisor: Thillainathan Logenthiran

Program: Electrical and Computer Engineering

Abstract: The increasing integration of distributed energy resources (DERs) has led to the development of microgrids as flexible and localized energy systems. This paper presents a comprehensive economic dispatch and financial settlement analysis for two scenarios: (1) a standalone microgrid consisting of three generators and five customers, and (2) a collection of interconnected microgrids operating under coordinated power exchange. A 24-hour simulation is conducted using a PoolCo market mechanism with merit-order dispatch and Market Clearing Price (MCP) determination. The study evaluates generator dispatch, unmet load allocation, and financial settlements during peak and off-peak conditions. Results demonstrate price volatility during peak demand and improved reliability under interconnected operation. The proposed framework provides insights into market-based microgrid management and capacity adequacy planning.

reference

[1] H. Bevrani and T. Ise, "Microgrid dynamics and control," in Proc. IEEE Power Energy Soc. Gen. Meeting, Detroit, MI, USA, 2011, pp. 1–5.

[2] N. Hatziargyriou, Microgrids: Architectures and Control. Hoboken, NJ, USA: Wiley-IEEE Press, 2014.

Poster 20 MIRA: Multi-Vector Retrieval with Hybrid Indexing and Adaptive Query Processing for Web-Scale Search

Authors: Solmaz Seyed Monir, Dongfang Zhao

Advisor: Dongfang Zhao

Program: Computer Science & Systems

Abstract: Similarity search over high-dimensional embeddings enables web-scale applications such as recommendation, semantic search, and retrieval-augmented generation. Existing approximate nearest neighbor (ANN) methods are constrained by static indexing, single-vector queries, and fixed parameters, leading to inefficiencies for adaptive and multi-vector search. We introduce Multi-Vector Information Retrieval with Adaptivity (MIRA), a self-optimizing ANN framework that integrates multi-vector fusion, hybrid indexing, and query-aware parameter tuning within a memory-efficient pipeline. Evaluated on text- and vector-based benchmarks spanning millions to billions of embeddings, MIRA maintains stable latency across embedding dimensions and similarity thresholds, achieving over $3\times$ lower memory usage, up to 86% lower latency, and nearly $10\times$ higher query efficiency than state-of-the-art methods, enabling scalable, adaptive, and efficient retrieval for web-scale applications.

Poster 21 LDP-MIC: Correlation-Aware Local Differential Privacy for Federated Learning under Untrusted Aggregation

Authors: Wenjun Yang, Eyhab Al-masri, Olivera Kotevska

Advisor: Eyhab Al-masri

Program: Computer Science & Systems

Abstract: Federated learning (FL) enables collaborative model training across distributed clients while keeping data local, yet privacy protection typically relies on a trusted aggregator for noise injection — a single point of failure. Local Differential Privacy (LDP) eliminates this

dependency, but existing LDP schemes suffer severe utility degradation due to correlation-agnostic noise injection, particularly under non-IID data distributions.

We present LDP-MIC, a correlation-adaptive LDP framework that co-designs privacy and integrity in federated learning. LDP-MIC quantifies feature–target dependencies locally using the Maximum Information Coefficient (MIC), then asymmetrically allocates a fixed privacy budget across dimensions based on these correlations. Perturbation is applied entirely client-side; model updates are computed from privatized data and clipped before transmission, requiring no trusted curator.

This asymmetric allocation preserves model utility while disrupting the cross-dimensional coherence that stealthy integrity attacks exploit. Evaluated on tabular and vision benchmarks under non-IID partitions, LDP-MIC consistently outperforms standard and prior LDP baselines — improving accuracy by up to 10 percentage points on Adult Census Income and up to 15 points on MNIST variants under strict privacy budgets ($\epsilon \leq 1$). LDP-MIC also reduces gradient-level information leakage and scales linearly with client count and feature dimensionality without additional communication overhead, establishing correlation-adaptive local privacy as a practical foundation for trust-free federated learning.

Poster 22 QoS-Aware AI Agents for Intelligent API Selection and Composition

Authors: Ishwarya Narayana Subramanian, Eyhab Al-Masri

Advisor: Eyhab Al-Masri

Program: Computer Science & Systems

Abstract: This study evaluates Large Language Models (LLMs), including open-source Hugging Face models, for automating service discovery and composition using a structured multi-agent architecture. The framework decomposes user queries into subtasks, retrieves candidate APIs through a RAG-based retriever, ranks them using an LLM-based agent, and composes executable workflows via a planner. A selector layer supports both semantic ranking and QoS-aware optimization using TOPSIS, enabling controlled comparison between functional and performance-driven selection strategies. The system is assessed across diverse queries using metrics such as functional alignment, completeness, technical validity, and Quality of Service (QoS). Preliminary results reveal a clear trade-off: pure semantic selection achieves higher task correctness and completeness, while strict QoS-driven optimization improves performance metrics but can reduce semantic alignment. These findings motivate a balanced multi-objective approach where semantic correctness is enforced as a primary constraint and QoS is applied as a secondary optimization layer.

Poster 23 Enhanced Road Crack Detection

Authors: Reetika Chauhan, Sixian Zheng, Wei Cheng, Juhua Hu

Advisor: Juhua Hu

Program: Computer Science & Systems

Abstract: Road surface cracks contribute to infrastructure deterioration and pose significant safety risks if left undetected. While deep learning models achieve high accuracy on benchmark crack datasets, their performance often declines when applied to different real-world road conditions. This gap highlights the challenge of cross-domain generalization in practical crack detection systems. This project develops a domain-aware crack detection pipeline combining

public datasets with real-world dashcam imagery. Dashcam videos can be efficiently collected for a targeted region and are processed to extract frames for crack detection. To enhance crack visibility, we proposed to apply the Sobel operator to extract high-frequency components, followed by SAM-based road segmentation to isolate pavement and suppress background clutter. Thereafter, each frame is divided into 200×200 grids, retaining only patches that overlap with the edge-enhanced road mask to focus on crack-like regions. Finally, we train a binary classifier for crack identification from other patches (e.g., shaded area and watermarks) to improve robustness. The proposed method enhances robustness, reduces false positives from background noise, and scales efficiently for large volumes of dashcam data in smart street management.

Poster 24 Agentic AI Serverless Code Generation: Towards Autonomous Improvement of Performance, Cost, and Code Quality

Authors: Xinghan Chen, Robert Cordingly, Ling-Hong Hung, Wes Lloyd

Advisor: Wes Lloyd

Program: Computer Science & Systems

Abstract: Large Language Models (LLMs) are transforming cloud native development by automating serverless function generation. However, current workflows often rely on blind iteration focused solely on functional correctness, neglecting critical non-functional attributes. Our preliminary evaluations reveal a hidden challenge: while LLMs often generate functionally correct code, the code can be very slow, as much as 100 times slower than human baselines or use more memory, which directly inflates serverless cloud hosting costs.

This research aims to address these issues by proposing an Agentic AI framework to autonomously improve the performance, cost, and quality of LLM-generated serverless functions. We introduce a closed loop system where intelligent LLM agents iteratively generate, deploy, profile, and optimize code. By leveraging interoperability protocols like the Model Context Protocol (MCP) and observability tools like the Serverless Application Analytics Framework (SAAF), we will develop agents to actively capture live runtime metrics and static code evaluations to inform their optimizations.

Ultimately, this work bridges the generative AI gap by arming agents to deduce and enhance diverse non-functional attributes. This will yield novel, performance aware tools that optimize execution efficiency, facilitate cross cloud migration to reduce vendor lock-in, and ensure serverless applications are cost effective.

Poster 25 Inference Profiler: Enabling Resource Utilization Profiling of LLM inferencing

Authors: Austin Bomhold, Morteza Nabavinejad, Wes Lloyd

Advisor: Morteza Nabavinejad, Wes Lloyd

Program: Computer Science & Systems

Abstract: The widespread deployment of Large Language Models (LLMs) requires substantial computational resources, often resulting in high operational costs and energy consumption. Optimizing these complex workloads necessitates precise visibility into how hardware accelerators and host system resources are utilized. This poster presents the InferenceProfiler, a software tool that measures and records the resource usage of containerized Large Language Model (LLM) inference tasks. The InferenceProfiler is designed to profile the CPU, memory,

disk, network, GPU, and power utilization of various tasks, collecting system metrics at the virtual machine, container, process, and hardware levels. InferenceProfiler leverages the NVIDIA Management Library (NVML) to capture real-time metrics across varying NVIDIA GPU architectures. InferenceProfiler integrates with the vLLM, an open-source high-throughput inference engine for serving LLMs, to correlate application-level telemetry directly with underlying metrics. InferenceProfiler supports performing time-series profiling to enable continuous monitoring of the resources consumed by LLM inference. Using the InferenceProfiler, we demonstrate resource requirements for LLM inferencing for a variety of user queries processed by different LLMs on alternate GPUs.

Poster 26 Impacts to Public Key Infrastructure from Post-Quantum Cryptography

Authors: Sofian Abbasi, Jaren Thorsen

Program: Computer Science & Systems, Information Technology, Master of Cybersecurity & Leadership (MCL)

Abstract: With the advent of quantum computing, the ability to secure data in-transit and data at-rest has become paramount. Current safeguards for secure digital communications consist of implementing Public Key Infrastructure (PKI) environments, which is the comprehensive set of policies, processes, and technologies to ensure confidentiality, integrity, and authenticity for digital communications. Post-Quantum Cryptography (PQC) is a set of algorithms developed to be resistant against attacks from quantum computers. This paper presents an overview of current PKI architectural elements and potential future integration of PQC algorithms within current PKI environments. Index Terms—cryptography, PKI, PQC, encryption, digital signatures, digital certificates, public/private keys.