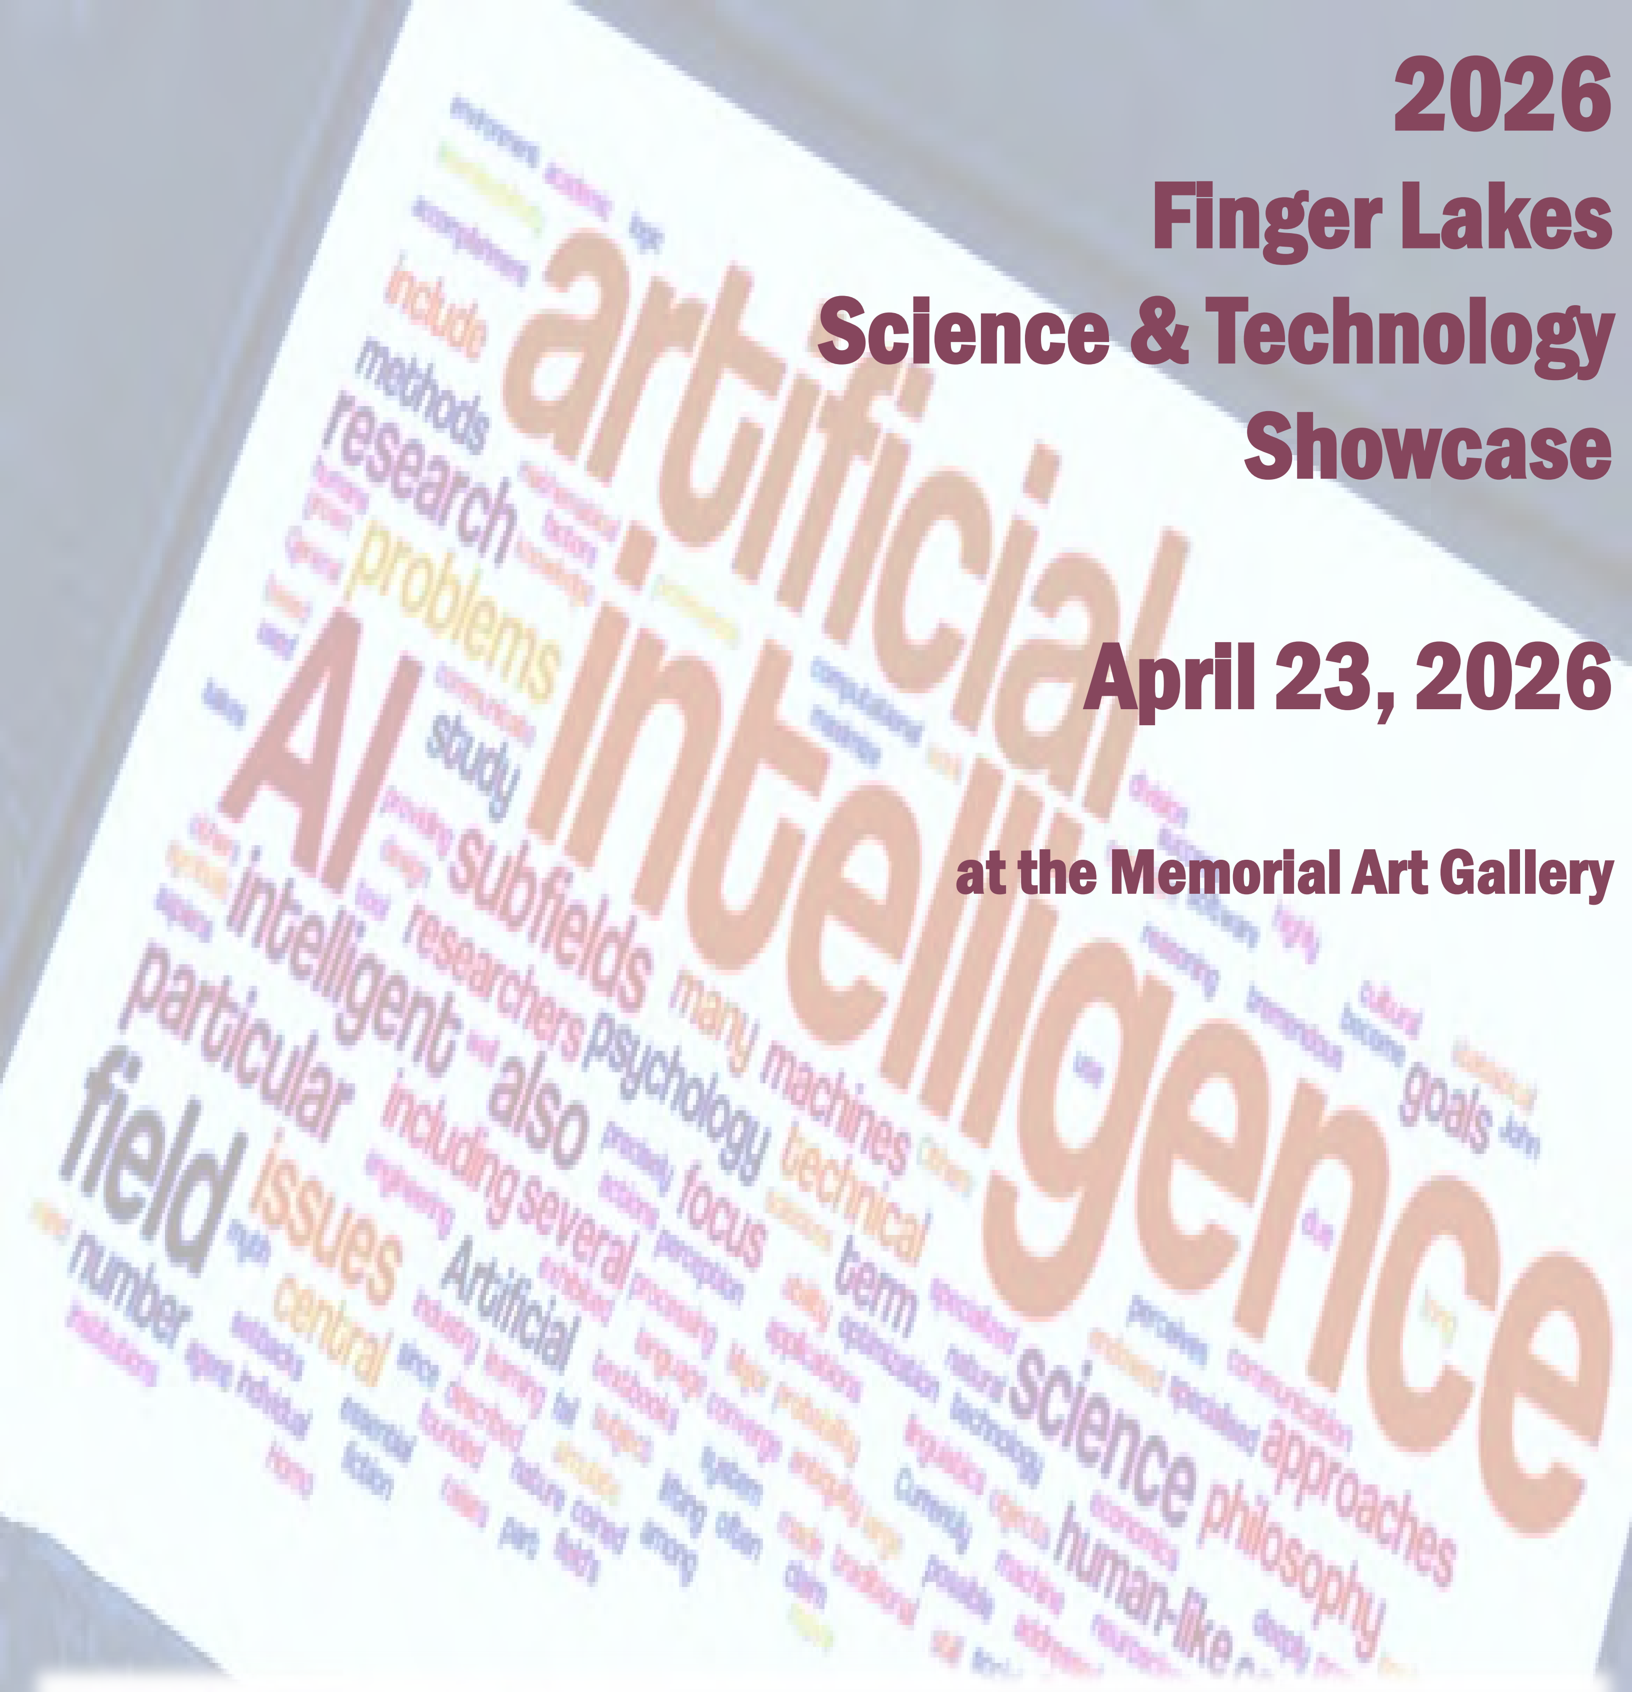

2026 Finger Lakes Science & Technology Showcase

April 23, 2026

at the Memorial Art Gallery

University of Rochester

Center for Emerging & Innovative Sciences

Center of Excellence in Data Science & Artificial Intelligence

Directors Welcome

Welcome to the 2026 Finger Lakes Science & Technology Showcase. This annual event is co-sponsored by the Center for Emerging and Innovative Sciences (CEIS) and the Center of Excellence (CoE) in Data Science and Artificial Intelligence, at the University of Rochester. With support from New York State, our centers encourage and support industry - university collaboration and technology transfer to create economic growth in New York State. Industry and university representatives working on a wide range of technology areas will be represented this year, including data science and AI, optics, imaging, photonics, sensors, semiconductors and electronics, biomedical technologies, acoustics, materials, and others.

This year our day begins with the 2026 Technology Woman of the Year Awards organized by TechRochester with CEIS and CoE as major sponsors. The Showcase events to follow include a Speed Networking Pitch Session, lunch, Keynote Speaker Barry Silverstein speaking on the topic: “Beyond Smart Display Glasses: Creating a path to Comfort, Context, and Depth in Spatial Computing,” and a poster session from researchers from the URochester and RIT, with people’s choice cash awards for the best posters - please be sure to vote!

We hope that you enjoy spending the day in the welcoming surroundings of the Memorial Art Gallery and use this opportunity to make new connections that we hope ultimately will contribute to regional economic expansion and job growth.

Both CEIS and the CoE in Data Science and AI provide NYS matching funds for company-sponsored research at NYS universities and the CoE also funds student internships at small companies and startups across NY. Please feel free to contact us to learn more about these programs and to discuss ways that CEIS and the CoE in Data Science and AI can help.

Warm Regards,

Mark I. Bocko

Mark Bocko, PhD
Director
Center for Emerging & Innovative Sciences

Mujdat Cetin, PhD
Director
COE in Data Science and AI

How are we doing?

Your feedback is important to us.

Please take a moment to scan the QR code to complete our event survey.

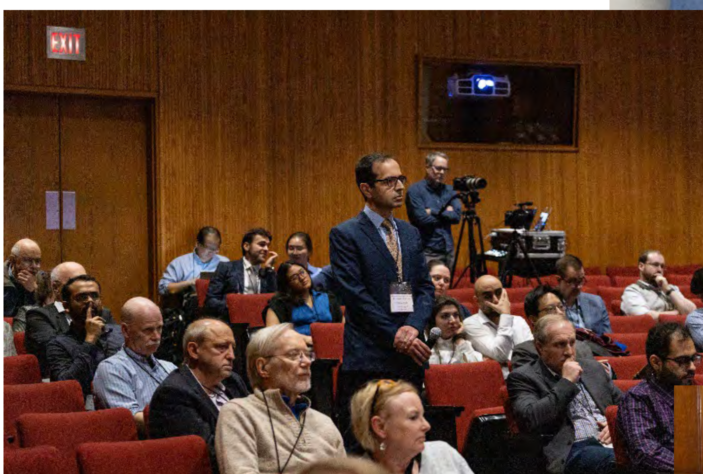

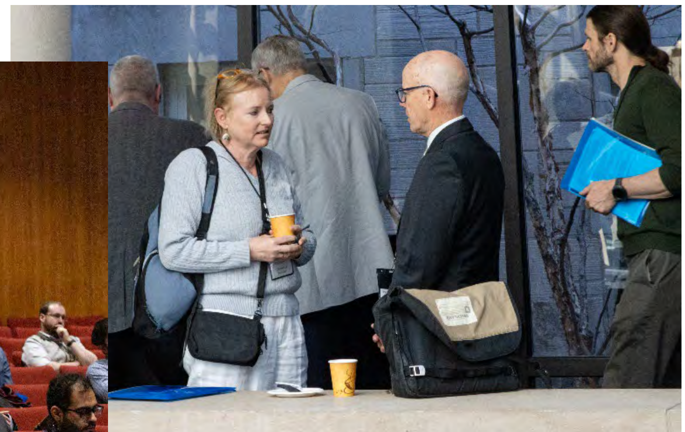

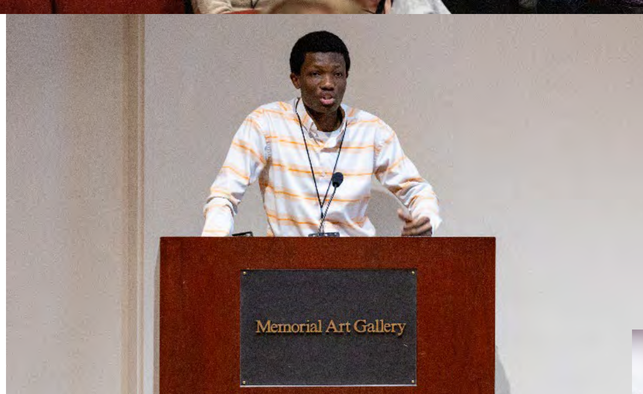

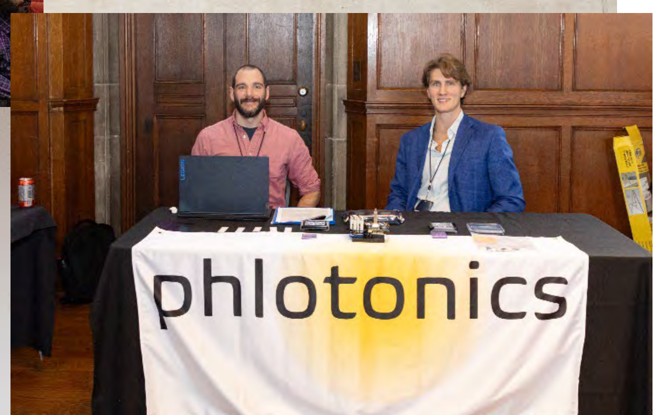

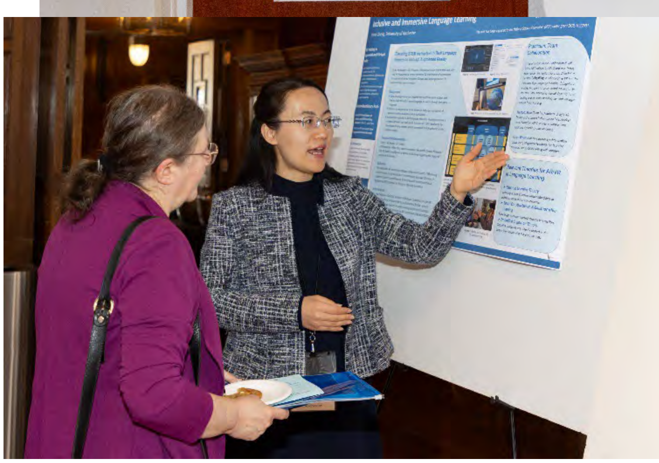

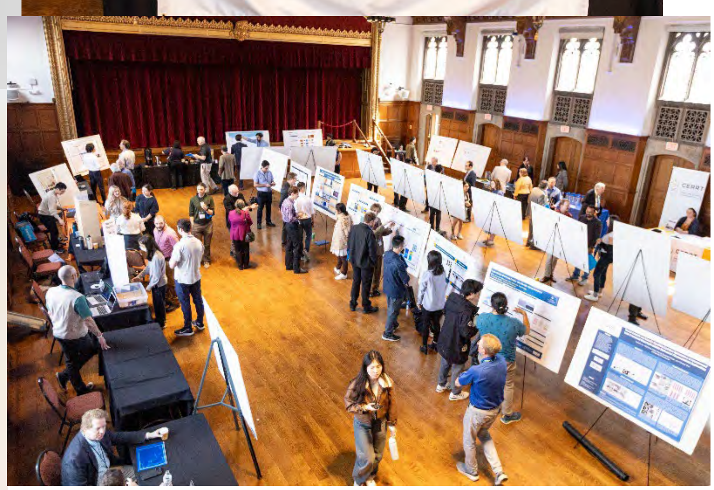

Agenda

9:00 AM - 1:00 PM	**Showcase Registration**
10:30 AM - 10:45 AM	**Showcase Welcome Remarks** *(Mujdat Cetin-CoE in Data Science & AI; Harry Bronson-NYS Assembly)*
10:45 AM - 12:00 PM	**Business/Faculty Networking Pitches & Partner Appreciation Awards**
12:00 PM - 1:00 PM	**Lunch**
1:00 PM - 2:00 PM	**Keynote Speaker** **Bary Silverstein – Director of the Center for eXtended Reality (CXR)** **Beyond Smart Display Glasses: Creating a Path to Comfort, Context, and Depth in Spatial Computing**
2:00 PM - 4:30 PM	**Open Poster Session/ Exhibitor Tables in Ballroom** ***(please check out the posters and vote before 4:00 – ballot may be found in your registration packet)***

Finger Lakes Science & Technology Showcase
Networking/Pitch Event

	Notes
Atma Anand | University of Rochester aanand6@ur.rochester.edu	
Lisa Belodoff | Bell Collaborative lisa.beledorf@bellcollaborative.com https://bellcollaborative.com/	
Michael Birnkrant | Freehold AESO, LLC support@freehold-aeso.com https://freehold-aeso.com	
Damon Diehl | NextCorps Luminate damon.diehl@nextcorps.org https://luminate.org	
Chen Ding | University of Rochester cding@cs.rochester.edu | 585.260.8182	
Paul Gebel | TechRochester/ITX pgebel@itx.com | 585.208.6799 https://www.techrochester.org	
Karla George | FLX AI, Inc. karla.george@flxai.com | 585.629.1286 https://flxai.com	

Finger Lakes Science & Technology Showcase
Networking/Pitch Event

	Notes
Michael Heilemann | University of Rochester mheilem@ur.rochester.edu | 716.870.3976 https://www.mchaudio.com/	
Kate Jordon | University of Rochester Kate_Jordon@urmc.rochester.edu | 315.250.5369	
David Levy | Mosaic Microsystems david.levy@mosaicmicro.com https://www.mosaicmicro.com/	
Haolin Liao | The Institute of Optics, University of Rochester hliao10@ur.rochester.edu	
Billy Manza | CDR GLOBAL, INC. bmanza@cdrglobal.com https://cdrglobal.com	
Susana Marcos | University of Rochester smarcos2@ur.rochester.edu https://www.urmc.rochester.edu/eye-institute/research/labs/marcos-lab	
Christopher McCarthy | Keller Williams Realty Greater Rochester crmccart@kw.com | 315.945.2688 https://greater-rochester.kw.com	

	Notes
Derel Meixell | Synergy IT Solutions dmeixell@synergyits.com https://www.synergyits.com	
Mark Wilson | Neworks, LLC mark@neworks.com https://www.neworks.biz/	

Finger Lakes Science & Technology Showcase Partner Appreciation Awards

CoE Recipient – Buffalo Solar

Buffalo Solar has been a valued partner of the Center of Excellence in Data Science and Artificial Intelligence for the past several years, participating three times in the Summer Internship Program and engaging University of Rochester students in high-impact, real-world data science projects. Through this collaboration, student interns supported the integration of business systems, development of data-driven tools, and generation of actionable insights that contributed to significant company growth. These efforts directly support the CoE-DSAI's mission to drive economic impact across New York State. Buffalo Solar's internship experiences have generated more than $1 million in new revenue and led to the creation of five new positions, including a full-time data analyst role. In addition, through its work in expanding access to clean, renewable energy solutions, Buffalo Solar contributes to a more sustainable energy future for New York State. This partnership exemplifies the power of applied data science and reflects the CoE-DSAI's commitment to connecting academic talent with industry to advance innovation, workforce development, and measurable impact.

CEIS Recipient – James McGrath – University of Rochester

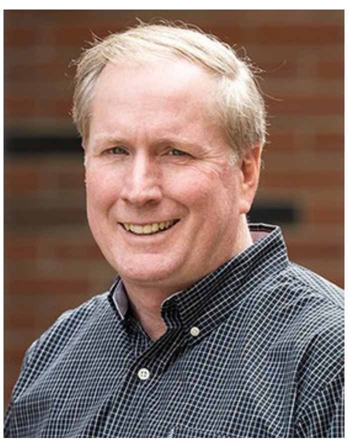

James McGrath is the William R. Kenan, Jr. Professor of Biomedical Engineering and the Director of the Translational Center for Barrier Microphysiological Systems ((TraCe-bMPS). Since 2001, Dr. McGrath has been on the Biomedical Engineering faculty at the University of Rochester and served the BME for over 10 years as the first director of its graduate program. Jim has led the Nanomembrane Research Group (NRG) since 2007. NRG is comprised of students, engineers, scientists, faculty and entrepreneurs at UR, SiMPore Inc., RIT, and an expanding network of academic and industry collaborators from around the world.

Jim co-founded three companies to advance the commercial development and application of nanomembranes: SiMPore (2007), Parverio (2020), and SiObex (2024). His familiarity with industry has motivated him to prepare his students by emphasizing the fundamentals of problem solving until the key steps become habit.

For over 20 years Jim has been one of the leading principal investigators funded by CEIS. He has been awarded numerous CEIS grants for his work with SiMPore on ultrathin silicon nanomembranes. His successive research contracts with SiMPore (one of CEIS's long standing partners) and his developing relationship with another local company, Bausch + Lomb, are testimony to the quality of his work and the economic impact it has generated.

Finger Lakes Science & Technology Showcase
Keynote Presentation

Beyond Smart Display Glasses: Creating a path to Comfort, Context, and Depth in Spatial Computing

Barry Silverstein
Director of the Center for eXtended Reality (CXR)
University of Rochester

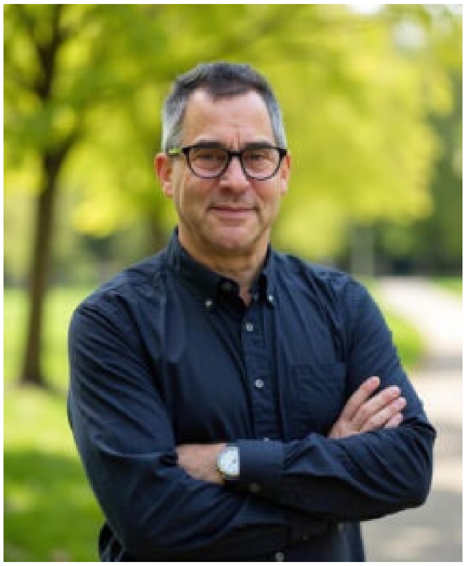

The first generation of smart glasses is validating consumer value by delivering individualized, context-aware information in a wearable form factor. However, a significant engineering challenge separates these early data displays from the milestone of comfortable, 'true AR' with a wide field of view and high-fidelity digital overlays. This transition point is defined by simultaneous technical hurdles: scaling the new display technologies, integrating prescriptions without compromising weight and yield, and ensuring visual comfort, particularly during interactions involving both near-field and distant objects (as comparative visual studies have demonstrated). This talk analyzes the critical decision-making framework and key technical and commercial benchmarks that define the road from present monocular single focal optics to more complex, variable-focus display architectures. We will present a phased strategic roadmap that balances the necessary codesign of optics and computational engines, detailing the current knowns and unknowns required to reach a context-aware, all-day wearable.

About the speaker

Barry Silverstein '84 is an optics engineer and technology leader whose career has spanned industry-defining work in imaging, projection, and augmented and virtual reality. A University of Rochester optics alumnus, he spent 28 years at Eastman Kodak, later led hardware R&D at IMAX where he helped develop the company's advanced laser projection system, and then served at Meta Reality Labs as senior director and chief technology officer for optics and display research. He now directs the University of Rochester's Center for Extended Reality, where he is helping shape interdisciplinary work at the intersection of AR/VR, artificial intelligence, and human-centered technology.

Augmented Reality/Virtual Reality (1)

1. **BrainScale VR: Transforming Medical Education through Immersive VR and AI-Enhanced Interactive Learning**

Adma Gama-Krummel[1] , Akhil Kasturi[1], Ziming Li[2], Roshan Peiris[2], Axel Wismueller[1], Raffaella Borasi[1] , Mujdat Cetin[1]
[1]University of Rochester, [2]Rochester Institute of Technology

BrainScale VR is an innovative educational platform designed to enhance medical training in brain tumor diagnosis through the integration of immersive Virtual Reality (VR) and Artificial Intelligence (AI). This project aims to develop and evaluate an interactive VR environment that combines advanced tumor segmentation with AI-driven conversational avatars to support clinical reasoning. Using the BraTS 2020 dataset (2,349 MRI scans), the system employs the UNETR deep learning model to generate accurate 3D tumor visualizations, achieving a Dice score of approximately 76%. These models are embedded in a VR interface, allowing users to explore complex anatomical structures in an intuitive, spatially rich environment. Additionally, a knowledge-augmented conversational AI avatar facilitates diagnostic inquiry and supports learners through realistic clinical scenarios. Preliminary findings indicate that the integration of immersive visualization and AI interaction improves anatomical understanding, engagement, and independent decision-making among medical trainees. Clinician feedback further suggests the platform's potential to bridge the gap between theoretical knowledge and clinical practice. While results are promising, further validation is needed to improve segmentation accuracy and assess long-term educational and clinical outcomes. Overall, BrainScale VR demonstrates strong potential as a scalable tool for advancing medical education and improving patient care.

Biomedical Technology (2-5)

2. **µSiM-Outer Blood Retinal Barrier: A Versatile Tool for Modeling Retinal Drug Uptake, Toxicity, and Choroidal Neovascularization In Vitro**

Kevin C. Ling, Rufaro Gamariel, Gram Hepner, Jordan Jones, Andres Muriel, Arthur Masiukiewicz, Ahmet Gurcan, Meng Hsu-Chun, Mehran Mansouri, Sami Farajollahi, Vinay Abhyankar, Ruchira Singh, Danielle S.W. Benoit, and James McGrath

Dysfunction of the outer blood retinal barrier (OBRB) is implicated in retinal degeneration, a leading cause for blindness. The OBRB is comprised of the retinal pigment epithelium, the choriocapillaris, and the Bruch's membrane. The OBRB maintains retinal homeostasis by regulating nutrient transport. Clinically approved drugs such as digoxin are linked with toxicity induced retinal degeneration. Animal models of the retina are severely limited; mice lack maculae, and primates are costly and low-throughput. Current in vitro OBRB models lack the

specificity and sensitivity to reproduce key pathological responses in retinal degeneration or toxicity. We developed an OBRB-on-a-chip, which leverages ultrathin (100 nm) nanoporous silicon nitride membranes and the µSiM device platform to model the interface between the retina and the blood (µSiM-OBRB). Arising retinal pigment epithelial cells (ARPE-19) and human umbilical vein endothelial cells (HUVECs) were seeded on opposite sides of a nanomembrane to mimic the RPE-choriocapillaris interface. ARPE-19/HUVEC co-formed tight barriers: average transepithelial electrical resistance (TEER) was 136.6 ± 20.9 $\Omega\cdot cm^2$, and permeability was 5.2×10^{-4} cm/min. Immunohistochemistry showed that µSiM-OBRB ARPE-19 expressed tight junctions with cobblestone morphology (zonula occludens-1) and were highly polarized (ezrin). µSiM-OBRB exhibited co-deposition of fibronectin and collagen IV at the RPE-HUVEC interface, emulating the Bruch's Membrane. To further validate the model, liquid chromatography with mass spectroscopy was used to show that ARPE-19-HUVEC co-barriers regulate retinal drug uptake and actively exclude toxic compounds. To stimulate drug-induced retinal toxicity, ARPE-19/HUVEC co-barriers were inoculated with 10 ng/ml digoxin, causing significant cell death and increased permeability, indicating barrier disruption. Invasive choroidal neovascularization (CNV) of the retina is another key feature of retinal degeneration. Injectable fibrin hydrogels were used to form 3D choriocapillaris-like microvascular networks beneath ARPE-19 cell barriers. VEGF stimulation caused HUVEC invasion into the retinal chamber, mimicking CNV. Collectively, these findings show that µSiM-OBRB is a predictive tool for screening retinal drug uptake and toxicity. Further, by incorporating 3D microvasculature with injectable fibrin hydrogels, we show that µSiM-OBRB can be used to capture CNV in vitro, emulating the pathology of macular degeneration.

3. The µSiM-MVB – A Microfluidic System for Sepsis Drug Development and Immune Cell Studies

Arthur Masiukiewicz, Emily Reitz, Isabelle Linares, Anthony Pietropaoli, Richard Waugh, Danial Ahmad, James McGrath

Sepsis remains a leading cause of organ injury and mortality, yet drug development is limited by the lack of applicable human-relevant models. We developed the µSiM-MVB, a microphysiological system incorporating a silicon membrane to model immune-cell behavior during sepsis. The platform is designed to replicate a vascular endothelial barrier while enabling controlled analysis of neutrophil migration, barrier disruption, and inflammatory responses. This system supports development and testing of therapies targeting dysregulated immune activity. We are beginning to explore this capability via aquaporin-inhibition strategies pursued with ApoGlyx. To improve throughput and reproducibility, we are integrating automated computer vision and machine-learning analysis for image-based quantification of neutrophil transmigration among other assay outputs. A cloud-enabled workflow supports data upload, high-performance processing, result archival, and standardized reporting. In parallel, assay scaling is being advanced through “MultiSiM” device formats, including 24-device designs and plate-based layouts that increase experimental capacity while maintaining optical clarity and fluid compatibility. Protocol development and technology transfer are underway to support multi-institutional validation. Thus, as the µSiM-MVB platform provides a human-relevant disease modeling and drug development tool for studying sepsis-related immune dysfunction, this work aims to evaluate antisepsis therapeutics and hopefully improve translation from preclinical testing to patient-relevant outcomes.

4. **Portico Health: From CAD to Cadaver**

   PI: Greg Gdowski. Other Contributors: Nick Drury, Lucas Lassinger, Alex J. McMullen, Martin Gira, Nat Ordway, William Lavelle

Portico Health LLC is a Rochester, New York–based medical device startup developing innovative solutions for spine surgery. The company's patented Portico Spine System addresses key limitations in cervical laminoplasty, a motion-preserving alternative to fusion that remains underutilized due to technical challenges with existing implants. Unlike traditional systems requiring multiple small screws, the Portico system enables a unique single-screw approach that simplifies the procedure while directing instrumentation away from the spinal cord and other important soft tissue structures. Through a CAT-funded partnership with the Center for Emerging and Innovative Sciences (CEIS), the team advanced the technology from early prototypes to clinically relevant designs. The proposed work focused on fabrication of implant components using surgical-grade materials (e.g., titanium alloys) to enable clinically relevant preclinical evaluation, with initial plans for testing in a porcine model. All milestones were completed, including multiple design iterations and fabrication of functional titanium prototypes. The project exceeded its original scope by enabling evaluation in a human cadaver model through participation in the Central New York Biotech Accelerator Medical Device Innovation Challenge. This provided earlier, more clinically relevant surgeon feedback, accelerated feasibility assessment, and highlights the value of coordinated academic–industry partnerships in advancing early-stage medical technologies toward clinical translation.

5. **Epigenetic Regulation of Brain Aging GWAS Variants within Transposable Element Sequences**

   Sreejato (Sree) Chatterjee, Shuangshuang Feng, Dean Zhang, Hongbo Liu

Genome-wide association studies (GWAS) have identified many loci associated with brain aging, yet most lie in noncoding regions, complicating functional interpretation because their effects must be inferred through regulatory mechanisms. In the brain, where long-lived cells rely on stable epigenetic regulation to maintain transcriptional homeostasis, transposable elements (TEs) represent a major component of the regulatory DNA landscape and a potential substrate through which noncoding variation may be interpreted. However, it remains unclear whether brain-aging GWAS variants are preferentially localized within specific TE sequences, whether these variants are linked to biological pathways relevant to brain aging, and through which regulatory mechanisms and cell-type-specific chromatin contexts they may act. Here, we applied an integrative genomic framework combining RepeatMasker TE annotations, single-nucleus ATAC-seq chromatin accessibility maps across major brain cell types, ENCODE candidate cis-regulatory elements (cCREs), and partitioned LD Score Regression. We found that GWAS variants frequently overlap transposable element sequences, but this overlap does not indicate global enrichment relative to genomic background. Instead, enrichment is non-random and specific to TE families, as SINE elements, especially Alu, show the strongest genomic enrichment, whereas enrichment within cell-type-specific accessible chromatin is driven by other repeat families. Simple repeat elements exhibit consistent enrichment across multiple brain cell types, while TcMar-Tigger DNA transposons show strong, oligodendrocyte-specific enrichment. Brain-aging GWAS variants overlapping TE sequences are preferentially localized within transcription factor–associated cCREs and depleted from promoter-like regions, indicating bias toward transcription factor–mediated regulatory contexts. Functional enrichment analyses revealed limited and context-specific pathway associations. In contrast, partitioned heritability analyses showed no independent enrichment of TE annotations after conditioning on baseline functional features, indicating that TE-overlapping GWAS signals are largely explained by shared regulatory architecture. Together, these results distinguish between genomic localization, regulatory activity, and heritability contribution, and support a model in which brain-aging GWAS variants localized within transposable element sequences reflect underlying epigenetic regulatory architecture rather than independent TE-specific genetic effects.

Data Science (6-15)

6. Measuring Global Structural Oppression in Free-Text Identity Data via Rule-Guided Large Language Models

Sreejato (Sree) Chatterjee, Linh Tran, Hanjia Lyu, John Nguyen, Timothy Dye

Traditional global health metrics often overlook identity-based structural disadvantage by relying on material indicators and rigid, census-style categories that flatten cultural and historical nuance. As a result, lived experiences of marginalization, shaped by colonization, systemic racism, and institutional exclusion, are difficult to measure across countries and languages. We introduce a scalable framework that uses Large Language Models (LLMs) to transform free-text identity data into a quantitative measure of structural oppression, enabling cross-national health equity research grounded in sociological theory. Using a multilingual global COVID-19 survey (n = 11,019, 170+ countries, 8 languages), we task LLMs with assigning ordinal oppression scores (1-5) and explanations from open-ended ethnicity and country descriptions. We develop a rule-guided prompting framework that encodes sociological principles, such as grounding judgments in historical and institutional exclusion, avoiding globalized stereotypes, and interpreting national labels as dominant groups by default. We benchmark three models (Gemini 1.5 Pro, GPT-4o-mini, GPT-3.5-turbo) against expert human annotations (n = 334 across 10 countries). Rule-guided prompting substantially improves performance across all models, achieving strong agreement with expert ratings (r = 0.85) and reducing large scoring errors by 35-50% relative to vanilla and chain-of-thought baselines. The approach stabilizes model behavior across countries and mitigates common failure modes, such as overgeneralizing global narratives of oppression or misinterpreting ambiguous identities. To evaluate construct validity, we apply LLM-derived scores to the full dataset and test associations with theoretically relevant epidemiological indicators. When grouped into low (scores 1-2) and high (scores 3-5) oppression categories, AI-derived classifications show significant associations with colonial legacy, household material asset ownership, external health locus of control, and healthcare privilege during COVID-19. Individuals classified as higher oppression consistently experience greater structural disadvantage across domains. These results demonstrate that LLMs, when constrained by theory-driven rules, can reliably convert open-ended identity data into a valid, interpretable epidemiological variable. This framework enables regression, stratification, and causal modeling of inequality in multilingual surveys, offering a new pathway for measuring structural oppression at global scale.

7. Accessible Video Conferencing for Blind and Low-Vision Users: A Multimodal, Role-Aware Assistive Framework

Tasin Khan, Luke Liu, Yukang Yan (PI)

Most video conferencing platforms still fall short in terms of meaningful accessibility for blind and low-vision (BLV) users, particularly in presentation settings. Sighted presenters routinely rely on visual cues from their audience, such as facial expressions or visible reactions, to adjust pacing, clarify points, or gauge engagement. For BLV speakers, this feedback is largely inaccessible, creating a gap in participation and awareness. In this ongoing project, we are designing an assistive system that supports BLV users in both presenter and audience roles during live online meetings. When a BLV user is presenting, the system uses computer vision and large language models (LLMs) to interpret audience facial expressions and summarize high-level emotions, such as confusion, interest, or agreement. This information is conveyed through real-time audio cues, allowing the BLV presenter to adapt as they speak. The system also offers camera-positioning guidance to help presenters remain centered and visible on screen. For BLV users attending as audience members, the system provides audio descriptions of slide content at select moments, timed around natural pauses in speech or slide transitions to reduce overlap with the presenter’s voice. Rather than replacing existing video conferencing platforms, our approach is designed as a lightweight add-on accessibility layer. By supporting both presenting and listening roles, this work moves BLV accessibility beyond basic accommodations and supports more active, two-way participation, helping BLV users follow presentations more easily and present with greater confidence.

8. Mixed Precision Quantization of Vision Transformers

Navin Ranjan and Andreas Savakis

Vision Transformers (ViTs) achieve strong performance across visual tasks but are computationally expensive, limiting deployment on resource-constrained platforms. Quantization reduces model size, but most methods use uniform bit-width across layers, which is suboptimal. Mixed-precision quantization (MPQ) offers better performance-efficiency tradeoffs, yet existing MPQ methods typically rely on heuristics or exhaustive search, without systematically considering each layer’s functional role and sensitivity. We propose Mix-QViT, an explainability-driven MPQ framework that assigns bit-widths based on two principled criteria: (i) Layer Importance, estimated via Layer-wise Relevance Propagation (LRP), which quantifies each layer’s contribution to the final prediction; and (ii) Quantization Sensitivity, which measures how susceptible each layer is to performance degradation when quantized to lower precision. This joint strategy enables informed bit allocation without exhaustive search, balancing accuracy and efficiency. We evaluate Mix-QViT on ViT, DeiT, and Swin across multiple benchmarks. Our results show that Mix-QViT significantly outperforms prior methods at 3-, 4-, and 6-bit precision.

9. VideoSeek: Long-Horizon Video Agent with Tool-Guided Seeking
   Jingyang Lin, Jiebo Luo

Video agentic models have advanced challenging video-language tasks. However, most agentic approaches still heavily rely on greedy parsing over densely sampled video frames, resulting in high computational cost. We present VideoSeek, a long-horizon video agent that leverages video logic flow to actively seek answer-critical evidence instead of exhaustively parsing the full video. This insight allows the model to use far fewer frames while maintaining, or even improving, its video understanding capability. VideoSeek operates in a think-act-observe loop with a well-designed toolkit for collecting multi-granular video observations. This design enables query-aware exploration over accumulated observations and supports practical video understanding and reasoning. Experiments on four challenging video understanding and reasoning benchmarks demonstrate that VideoSeek achieves strong accuracy while using far fewer frames than prior video agents and standalone LMMs. Notably, VideoSeek achieves a 10.2 absolute points improvement on LVBench over its base model, GPT-5, while using 93% fewer frames. Further analysis highlights the significance of leveraging video logic flow, strong reasoning capability, and the complementary roles of toolkit design.

10. SpecEyes: Accelerating Agentic Multimodal LLMs via Speculative Perception and Planning
    Jiebo Luo

Agentic multimodal large language models (MLLMs) (e.g., OpenAI o3 [36] and Gemini Agentic Vision) achieve remarkable reasoning capabilities through iterative visual tool invocation. However, the cascaded perception, reasoning, and tool-calling loops introduce significant sequential overhead. This overhead, termed agentic depth, incurs prohibitive latency and seriously limits system-level concurrency. To this end, we propose SpecEyes, an agentic-level speculative acceleration framework that breaks this sequential bottleneck. Our key insight is that a lightweight, tool-free MLLM can serve as a speculative planner to predict the execution trajectory, enabling early termination of expensive tool chains without sacrificing accuracy. To regulate this speculative planning, we introduce a cognitive gating mechanism based on answer separability, which quantifies the model's confidence for self-verification without requiring oracle labels. Furthermore, we design a heterogeneous parallel funnel that exploits the stateless concurrency of the small model to mask the stateful serial execution of the large model, maximizing system throughput. Extensive experiments on V* Bench, HR-Bench, and POPE demonstrate that SpecEyes achieves $1.1 - 3.35\times$ speedup over the agentic baseline while preserving or even improving accuracy (up to +6.7%), boosting serving throughput under concurrent workloads.

11. BUILD with precision: Bottom-Up Inference of Linear DAGs

Gonzalo Mateos, Samuel Rey, Geert Leus, Antonio G. Marques

Learning the structure of directed acyclic graphs (DAGs) from observational data is a central problem in causal discovery, statistical signal processing, and machine learning. Under a linear Gaussian structural equation model (SEM) with equal noise variances, the problem is identifiable and we show that the ensemble precision matrix of the observations exhibits a distinctive structure that facilitates DAG recovery. Exploiting this property, we propose BUILD (Bottom-Up Inference of Linear DAGs), a deterministic stepwise algorithm that identifies leaf nodes and their parents, then prunes the leaves by re- moving incident edges to proceed to the next step, exactly reconstructing the DAG from the true precision matrix. In practice, precision matrices must be estimated from finite data, and ill-conditioning may lead to error accumulation across BUILD steps. As a mitigation strategy, we periodically re-estimate the precision matrix (with less variables as leaves are pruned), trading off runtime for enhanced robustness. Reproducible results on challenging synthetic benchmarks demonstrate that BUILD compares favorably to state-of-the-art DAG learning algorithms, while offering an explicit handle on complexity.

12. Towards a Unified Evaluation Protocol for Link Prediction Interpretation over Knowledge Graphs

Beshani Weralupitiya, Dr. Carlos R. Rivero

Link prediction is a critical task in knowledge graph applications, enabling the inference of missing information within graph-structured data. However, existing evaluation protocols for link prediction are computationally expensive, lack statistical soundness, and suffer from various challenges, such as limited interpretability, inconsistent graph splitting methods, leading to unreliable model assessments. This research aims to establish new link prediction interpretation methods for knowledge graphs that are robust and efficient with statistical guarantees. It also aims to create a unified evaluation protocol that integrates variations like subgraph selection, graph splitting, and redundancy reduction. As the initial step, we have developed a knowledge graph splitting mechanism, which we present in this paper. Our research focuses on designing a robust evaluation framework, scalable rule-mining techniques, and reliable statistical estimators to enhance the interpretability of link prediction models.

13. Assessing Difficulty of Link Prediction on Benchmarking Knowledge Graphs

Lydia Klecan, Carlos Rivero

Knowledge graphs represent real-world data in a directed graph format where two entities, connected by a predicate, represent one fact. Link prediction models predict new relationships using existing entities and predicates, and are trained using benchmarking knowledge graphs. Redundancies exist within benchmarking knowledge graphs that artificially improve the results of link prediction. If a model scores well on a highly redundant knowledge graph, it may not directly correlate to good performance on other more complex knowledge graphs. This research creates new analysis methods and evaluation metrics for measuring redundancies in knowledge graphs. The two main goals include: (1) reporting levels of redundancy in entire knowledge graphs and (2) offsetting link prediction results based on predicate-specific redundancy levels within knowledge graphs. The entire knowledge graph metric (1) confirms high levels of redundancy in FB15k and WN18, which are known to be highly redundant. Offsetting link prediction results using predicate-specific metrics (2) displays less successful link prediction results for knowledge graphs with high redundancy, and more successful link prediction results for knowledge graphs with low redundancy, as expected. These methods will allow link prediction models to be realistically evaluated which will lead to more accurate link prediction models.

14. Evaluation Methods for Knowledge Graphs

Diba Masihi, Carlos R. Rivero

Knowledge graphs store real-world information as facts, where each fact connects two entities through a labeled relationship. For example, (Paris, capitalOf, France) states that Paris is the capital of France. Evaluating the quality of these graphs matters, especially when link prediction models generate new facts that may appear structurally correct but are semantically nonsensical. Traditional graph metrics like PageRank treat all connections equally. They measure structural importance but ignore what the labels mean. If you swap "capitalOf" with "languageOf," PageRank will not notice the difference. We propose Domain/Range PageRank (DRPR), which incorporates semantics into the PageRank computation. Our approach looks at which predicates share similar entities as subjects or objects. We evaluate DRPR on seven benchmark datasets which we intentionally corrupted the graphs by swapping labels or deleting connections to change the semantics. For each dataset and mutation level, we generate 25 independently mutated graphs. Our results show that DRPR detects semantic corruption more effectively than standard PageRank.

15. The Choir Without a Conductor: Phase Transitions in Sparse Predictive Multi-Agent Synchronization

Atma Anand, Eric Blackman

How does a large choir stay in time without a conductor? Listening to everyone scales as N^2 computations and breaks down for large groups. However, each musician can hold an internal sense of the beat and sample only a few neighbors to stay aligned. We formalize this intuition with a population of Kuramoto oscillators augmented by Kalman-filter internal models and ask how sparsely distributed agents can coordinate when each maintains a predictive model about the group state. Two results emerge. First, predictive agents achieve synchronization quality comparable to fully-connected baselines while using only ~5% of the interaction bandwidth, reducing total coordination cost to roughly 5–15% of the $O(N^2)$ all-to-all baseline at N=2000. Whether this reflects an asymptotic O(N) vs $O(N^2)$ separation requires further finite-size scaling analysis. The internal model successfully substitutes for direct observation across a wide range of model precisions. Second, this efficiency terminates in a sharp phase transition: below a critical sampling rate (ε_c), coordination collapses abruptly, and neither richer internal models nor longer integration times recover it in the regimes tested. Model sophistication affects convergence timescale but not the location of the threshold. The transition is consistent with a simple information-balance argument: observations must replenish belief faster than it drifts. This suggests a lower bound on communication cost for this class of coordination problems. The findings have practical implications for sensor networks, swarm robotics, and federated learning operating under bandwidth constraints, where adding more compute stops rescuing coordination past a hard floor.

Optics, Photonics, Imaging (16-22)

16. Modelling of Multi-Photon Events of a Superconducting Single-Photon Nanostripe Detector for Photon Number Resolution Experiments

Quang Nguyen, Ivan Komissarov, Roman Sobolewski

Superconducting nanostripe single-photon detectors (SNSPDs) [1] have quantum efficiencies approaching unity [2], gigahertz count rates [3], picosecond jitter [4], sub-hertz dark count levels [5], and relatively high operating temperatures (2-4 K). Although SNSPD operation was initially considered primarily in the threshold regime, studies of the photoresponse pulse shape dependence on the number of absorbed photons allowed implementation of a readout scheme that enabled photon number resolution (PNR) [6]. Time tracking of the pulse rising-edge arrival as a method for PNR in intrinsic (single SNSPD) schemes has been implemented using semiconductor electronics, but it has scalability limitations due to high energy dissipation per switching event [7]. A readout scheme using superconductor electronics, such as, e.g., a single-flux-quantum (SFQ) circuitry should address the above issue. SFQ-based PNR readout schemes had already been implemented in multi-pixel SNSPD systems [8]. Here, we investigate the possibility of designing a novel SFQ-based PNR readout scheme for a single SNSPD. We study in the LTSpice environment, an SNSPD's photoresponse pulse rise time dependence on the number of absorbed photons, ranging from 1 to 5. For this purpose, we have modified the known dynamic model for SNSPD simulation [9], in the way that the number of photons is specified by the number of series-connected resistors corresponding to hot spots formed by absorbed photons. In addition, we calculated the time difference between the heralded pulse and the time-moment corresponding to generation of the SFQ pulse by the DC-to-SFQ converter.

This work was funded in part by the Center for Emerging and Innovative Sciences at the University of Rochester, a New York State Center for Advanced Technology funded by the Empire State Development Division of Science, Technology and Innovation (NYSTAR). Additional funding came from the National Science Foundation under the Expedition project: DISCoVER (Design and Integration of Superconducting Computation for Ventures beyond Exascale Realization) grant #2124453.

1) C.M. Natarajan et al. Superconductor Sci. Technol. Vol. 25 (6), p063001, 2012
2) G.N. Gol'tsman et al. Appl. Phys. Lett. Vol. 79 (6), p705, 2001
3) D.V. Reddy et al. Optica Vol. 7 (12), p1649, 2020
4) W.-J. Zhang et al. IEEE Transactions Appl. Supercond. Vol. 29 (5), p1, 2019

5) B. Korzh et al. Nature Photon. Vol. 14 (4), p250, 2020
6) T. Rambo Quantum Computing, Communication, and Simulation IV. – SPIE. 2024
7) O. Mukhanov et al. Fundamentals and Frontiers of the Josephson Effect. p611-701, 2019
8) H. Myoren et al. IEEE Trans. Appl. Supercond. Vol. 28, p2500304, 2018
9) K. K. Berggren et al. Supercon. Science and Technology Vol. 31(5), p055010, 2018
10) K.L. Nicolich et al. Phys. Rev. App. Vol. 12.3, p034020, 2019

17. Polarimetric imaging of peripheral nerves: an intraoperative aid

Haolin Liao, Zongyuan Lu, David J. Mitten, Wayne H. Knox*

To minimize the risk of iatrogenic nerve injury, an automatic intraoperative identification method is essential. Leveraging the intrinsic birefringence of peripheral nerve tissue, Rotating Crossed Polarization Imaging (RXPI) enables tissue differentiation based on the time-domain analysis of backscattered intensity signals. In this work, we present real-time nerve imaging results in hand surgeries and further explore the potential of our method as a preprocessing strategy to facilitate machine learning in low-data regimes.

18. A dual-use apparatus for infrared thermal defocus and refractive index determination

James A. Monroe, Paolo Masini, James M. Zavislan

The infrared (IR) optics supply chain currently does not certify the refractive index of IR transparent materials or the thermal performance of optic assemblies due to the high cost and long lead time for such measurements and certification. At the lens material level, refractive index certification using the minimum deviation method requires prisms that take weeks to manufacture and measurements that take up to a day to set up and collect. At the optic assembly level, performance is usually determined with thermal vacuum chamber testing in custom built facilities that are expensive to operate and have long wait times due to high demand. This work demonstrates that a wedged shear plate interferometer system can determine the refractive index of input IR transparent materials and the thermal defocus of output optic assemblies. Advantages include using easily manufactured flat polished samples for refractive index determination and simple thermal defocus measurement techniques. The interferometer setup is presented and the measurement principles discussed. Refractive index measurement and thermal defocus results using an infrared demonstration system are presented and discussed. This work demonstrates that a single interferometer can have dual uses to certify IR lens materials and full optic assemblies.

19. New VIS-VUV radiometric calibration facility for instrumentation at the RIT

Wyatt Robinson, Zoran Ninkov and Ross Robinson

The Laboratory for Advanced Instrumentation Research (LAIR) at Rochester Institute of Technology (RIT) is developing a radiometric calibration facility focusing on expanding into Vacuum UV (VUV) capabilities. The current configuration for 120nm-1100nm coverage has dual monochromators illuminated with Deuterium lamps. The monochromator for shorter wavelengths operates under vacuum. Potential applications include assisting in developing sensors and optics for space astronomy, defense and semiconductor manufacturing.

20. Neural network Design for Effective Index Prediction for Waveguides

Jaymit Surve, Martin Sanchez, Christer Everly, Lukas Weituschat, Sophia Korono, Pablo Postigo

Due to its optical properties and compact footprint, a single mode strip waveguide is widely utilized for vast range of applications. Designing a waveguide requires several simulations as the no. of modes supported by a waveguide depends on several structural parameters, therefore, as an alternative, we proposed an efficient approach by employing a neural network. We developed a Forward neural network that takes the dimension and materials of waveguide as input and predicts the effective index of fundamental modes for TE and TM polarization with an R2 score of 0.9957. Next, a Physics Informed Neural Network will be implemented by incorporating the eigenvalue equation into loss function, reducing the reliance on data driven methods.

21. Development of high-Q microresonators on thin-film lithium niobate platform

Zhengdong Gao, Qiang Lin

Thin-film lithium niobate (LN) microresonators have emerged as an important platform in frontier photonics research, owing to their outstanding linear and nonlinear optical properties, ultrahigh optical confinement, and broad transparency window. In this poster, we present recent advances in the development of high-Q lithium niobate microresonators, including essential fabrication techniques and a detailed overview of their applications, ranging from frequency conversion and electro-optic modulation to frequency comb generation, quantum photonics, and lasing.

22. Photonics Instrument and Assay Development for Biomolecule Analysis in Complex Fluids

Mert Gurcan, Ethan Yospin, Dan Steiner, Michael Bryan, Benjamin Miller

Photonic sensors are making inroads across a wide range of applications, from medical diagnostics to environmental monitoring and beyond. These systems offer the ability to directly monitor complex environments in real time, providing insight into system state, dynamics, and potential inefficiencies that are difficult to capture with conventional approaches. To address these opportunities, we are adapting a photonic sensing system originally developed for medical diagnostics into a format suitable for broader analytical use. First, we describe optical probe design modifications to achieve the required form factor and reduce system cost. Second, we outline assay formats for detecting representative analytes relevant to complex biological environments.

This approach builds on an integrated platform capable of automated optical alignment, synchronized spectral data acquisition, and correlation with imaging data, enabling high-fidelity, time-resolved measurements within a unified workflow.

23. Substrate-independent Microdroplet Transport via Simple Superhydrophobic Tweezers

Tianshu Xu, Ran Wei, Yichen Deng, Kunyang Cao, Luheng Tang, Subhash C. Singh, and Chunlei Guo

We developed a substrate-independent droplet transfer technique by creating superhydrophobic tweezer tips using femtosecond laser processing. By optimizing laser parameters with machine learning and adjusting real-time tip separation, adhesion forces are precisely controlled. The SH tweezers transfer 1–35 μL of various liquids with no content loss or cross-contamination, capable of volume-selective and simultaneous multi-droplet transfers. Bypassing the need for complex external controls, this simple, cost-effective tool provides a highly effective solution for droplet transfer across various biomedical applications.

24. Development of a Low-Cost, Low-Power Integrated Machine Health Monitoring Sensor

Benjamin Thompson, Aaron Bundy, Michael Heilemann

This project, in conjunction with ADVIS Inc., is focused on developing a device to cost-effectively bring machine health monitoring (MHM) to a broad spectrum of Department of Defense (DoD) assets (vehicles, pumps, rotating machinery), where the implementation of conventional monitoring systems is cost prohibitive. During this Phase II effort, we have focused on minimizing power consumption. A new processing board model was developed that can achieve a battery life of ~1 year while maintaining a form factor of ~1 cubic inch. We have also developed a mechanical test bench that allows us to systematically introduce faults during real-time data collection and analysis.

25. Toward Data-Driven Thermoelectric Power Generation for Oceanographic Equipment

Zhiheng Xu, Bobby Kovach

Thermoelectric generators (TEGs) provide a promising approach for powering oceanographic sensors by harvesting energy from natural temperature gradients. However, their practical use in marine environments is limited by low thermal differences and varying ocean conditions. This work evaluates the feasibility of TEG-based power generation for ocean sensing applications through modeling and system-level analysis. The results highlight the influence of temperature gradients and device design on power output and demonstrate the potential for continuous, low-power operation in stable ocean environments. This study also outlines future directions toward data-driven optimization by integrating environmental data to improve system performance under dynamic conditions.

26. Solenoid Actuated Liquid Air Lock

Mike Kaminski, Dan Phinney, Cash Close, Abby Che, and Max Morris

A novel high pressure injection system for injecting an additive (aqueous or gaseous) into a pressurized pipe by means of MCU actuated solenoids to create a temporary low pressure injection window. Such a window allows a low power pump to inject additive into a pressurized line, because the solenoids remove the pressure and fluid from the pipe. Having neutralized the line pressure, the pump no longer needs to exceed the pressurized flow, thereby reducing energy consumption and system cost significantly.

27. On-chip Flux Bias Driver for Superconducting Quantum Circuits

Yerzhan Mustafa, Oleg Mukhanov, and Selçuk Köse

As quantum computing and quantum sensor technologies move toward scalable and energy-efficient implementations, the need for compact, high-fidelity, and cryogenically compatible readout architectures is becoming more urgent. Certain superconducting quantum circuits such as Josephson digital phase detector need to be excited with a specifically timed and shaped current pulse that serves as a flux bias. In this work, we developed an on-chip flux bias driver circuit, which offers more scalable and energy-efficient architecture as compared to supplying the current from room temperature electronics. Several flux bias driver designs have been fabricated using SEEQC 10 μA/μm^2 process. Experimental results at 4 K temperature show the correct switching of bias current ranging from hundreds of μA to few mA over a large inductive load.

28. Universal Optimal Transport-Based Voice Conversion

Anton Selitskiy and Mark Bocko

The problem of voice conversion (VC), which aims to transform speech from a source speaker to sound as if it were produced by a target speaker while preserving linguistic content, is studied in this work. We propose a flow-matching–based neural network that approximates the optimal transport map between speakers in a pretrained speech embedding space. In our approach, speech signals are first mapped into high-level representations using the self-supervised speech model WavLM Large. Training pairs are constructed using a kNN-VC strategy. To capture speaker identity, an arbitrary voice sample from the target speaker is incorporated via an attention mechanism. The transformed embeddings are then converted back into waveform audio using a neural vocoder. This framework enables effective voice conversion without requiring parallel training data and achieves high quality even for short audio segments. Experimental results show that the proposed method is competitive with kNN-VC, SinkVC, and MKL, and outperforms them in the low-resource setting of short input utterances. The model is trained on the LibriSpeech Clean-100 dataset.

29. Characterization of a Stirling Cycle Heat Pump for Waste Heat Recovery

MA Garlatti, FD Kelley, EH Chimowitz

The objective of this research project is to design, build and experimentally investigate the efficacy of a thermodynamic engine for converting waste heat into electrical energy. This will be accomplished by implementing a Carnot cycle within a Stirling engine-based mechanical device for the purpose of waste heat recovery. Engine regenerator materials will be investigated as a means to improve engine efficiency. Engine frequency is measured using a piezoelectric film sensor and a charge amplifier with a counter to measure the number displacements of the film. Frequency count ensures the engine is operating at resonance.

30. Enhancement of Sensitivity and Yield in Integrated Optical Biosensing by Machine Learning

Lukas Weituschat, Martín Sanchez, Christer Everly, Pablo A. Postigo

Recovering information from noise-dominated signals typically involves amplification to enhance signal visibility, especially in regimes approaching fundamental noise limits. Here, we demonstrate a deep neural network (DNN) that accurately extracts Lorentzian resonance positions from optical ring resonator spectra across noise levels down to the noise-dominated regime limit (Signal/Noise < 1). The DNN is trained on synthetically generated spectra with controllable noise statistics, removing the necessity of time-expensive collection of experimental data and its labeling. We show that the trained model maintains accurate resonance extraction even at low signal-powers and generalizes to experimentally measured data not included in the training, enabling accurate resonance inference and tracking in the weak-signal regime. This approach provides a scalable route to improving measurement fidelity at low signal powers in integrated photonic platforms for quantum technologies.

31. Structure representation learning for music with GNN

Huiran Yu, Yiyang Wang, Wei-Cheng Lin, Zhiyao Duan

Musical structures such as motifs, grooves, forms, and relational dependencies between events are fundamental to music analysis and generation. Learning representations that are sensitive to such structures remains a core challenge in symbolic music modeling. While recent sequence-based models have made good advances, they are designed to process linearized event streams and still struggle to model the rich structures underlying polyphonic music. In this work, we propose a Graph Neural Network (GNN)–based approach for symbolic music representation learning, where a sparse, interpretable graph among musical notes is learned through a masked node attribute reconstruction objective. Over this graph, note embeddings are also learned to carry musical structure information. We argue that this graph-based approach is natural and efficient in modeling simultaneous musical events and their relations in polyphonic music. We evaluate the learned note embeddings of our approach on multiple structure-required downstream tasks, including pattern matching, music segmentation, and harmony recognition. Results show that, with a significantly smaller model size, our approach achieves superior or competitive performance over state-of-the-art Transformer-based models. To facilitate future research on structure-aware music modeling, we additionally release a symbolic music dataset with fine-grained melodic pattern annotations.

We Remember

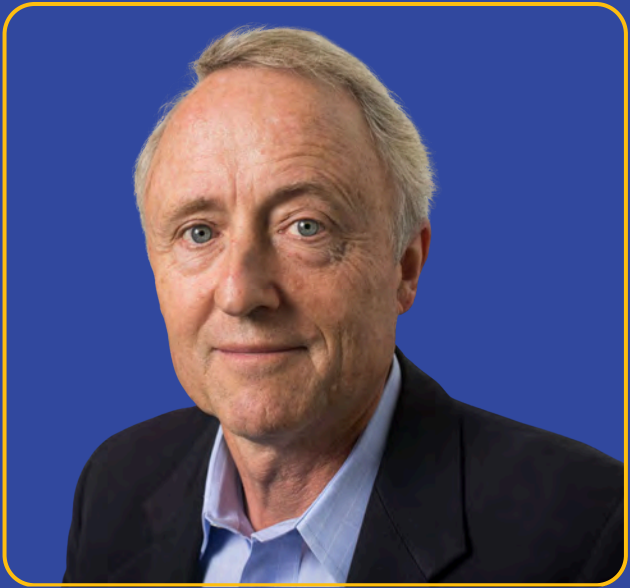

Kevin J. Parker, PhD

1954 -2025

A pioneering scholar, inventor, entrepreneur, mentor, and leader whose work transformed ultrasound imaging and touched countless lives. A world-renowned researcher, beloved teacher, and Dean Emeritus of Engineering and Applied Sciences at the University of Rochester, Professor Parker helped shape modern medical imaging through innovations in elastography and related diagnostic technologies. He was also a founder of VirtualScopics and a valued partner of CEIS and the Center of Excellence in Data Science and AI, reflecting his deep commitment to translating research into real-world impact. Through his leadership, generosity, and dedication to students, colleagues, and collaborators, Kevin leaves a legacy that endures in the science he advanced, the companies and partnerships he helped build, and the many people whose lives and careers he influenced.

TECHNOLOGY SUPPORTERS

BELL COLLABORATIVE

BUSINESS INSIGHT CENTER- CENTRAL LIBRARY OF ROCHESTER & MONROE COUNTY

CENTER FOR INTEGRATED RESEARCH COMPUTING (CIRC)

GOERGEN INSTITUTE FOR DATA SCIENCE AND AI (GIDS-AI)

GWEN M. GREENE CENTER FOR CAREER EDUCATION & CONNECTIONS

MONROE COUNTY ECONOMIC DEVELOPMENT

MONROE COUNTY UPSTATE NEW YORK APEX ACCELERATOR

NEWORKS, LLC

NEXTCORPS

PHLOTONICS, INC.

STEMPOWER STUDIO

SYNERGY IT SOLUTIONS

TSG SECURITY

URVENTURES