CSEE RESEARCH DAY

Friday May 2, 2025

9 a.m. – 4 p.m.

BWTECH@UMBC SOUTH

MAIN SEMINAR ROOM

Virtual attendees join via Webex (Linked)

A celebration of selected research accomplishments from 2024 – 2025 by UMBC faculty and students in the Department of Computer Science and Electrical Engineering. Open to faculty, researchers, students, and visitors.

Schedule

9 a.m. – 9:20 a.m. Continental Breakfast

9:20 a.m. – 9:30 a.m. Opening Remarks

9:30 a.m. - 10:40 a.m. Session I
(talks at 9:30 a.m., 9:55 a.m., 10:20 a.m.)

10:40 a.m. – 11 a.m. Break

11 a.m. – Noon Poster Session

Noon – 1 p.m. Lunch

1 p.m. - 2:10 p.m. Session II
(talks at 1 p.m., 1:25 p.m., 1:50 p.m.)

2:10 p.m. - 2:30 p.m. Break

2:30 p.m. - 4 p.m. Session III
(talks at 2:30 p.m., 2:55 p.m., 3:20 p.m.)

4 p.m. Adjourn

TALKS

Session I (9:30 a.m. – 10:40 a.m.)

- **9:30 a.m.** Dr. Edward Zieglar
- **9:55 a.m. Best B.S. Research** - Bharg Barot, LOADS: LiDAR-based Privacy-Preserving Queue Monitoring and Analysis
- **10:20 a.m.** Dr. Cynthia Matuszek - Artificial Intelligence and Human-Robot Interaction

Session II (1 p.m. – 2:10 p.m.)

- **1 p.m. Best M.S. Research** - Pratik Shukla, DUNE: A Machine Learning Deep UNet++ based Ensemble Approach to Monthly, Seasonal and Annual Climate Forecasting
- **1:25 p.m.** Dr. Maksim Eren, "Tensor Decomposition for AI: Applications in Cyber-security, Data Privacy, Model Compression, and Hallucination Reduction"
- **1:50 p.m.** Dr. Riadul Islam, Event-Based Vision Meets Compute-In-Memory: A Path to Scalable, Energy-Efficient AI

Session III (2:30 p.m. – 3:40 p.m.)

- **2:30 p.m.** Dr. Sam Lomonaco, My Struggle with Quantum Entanglement
- **2:55 p.m.** Dr. Seung-Jun Kim, Flexible Functional Magnetic Resonance Imaging Data Analysis Using Machine Learning Methods
- **3:20 p.m. Best Ph.D. Research -** Yuechun Gu, Calibrating Practical Privacy Risks for Differentially Private Machine Learning

MENU

Continental Breakfast

9 a.m. - 9:20 a.m.

**Variety of Pastries*
**Muffins*
**Assorted Bagels with Spreads*
**Fresh Fruit Salad*
**Assorted Juices*
**Coffee and Tea Service*

Lunch

Noon - 1 p.m.

**assorted sandwiches & wraps*
**tossed salad*
**2 cold salads*
**desserts*
**drinks*

POSTER PRESENTERS

Improving Shift Invariance in Convolutional Neural Networks with Translation Invariant Polyphase Sampling
Presenter: Sourajit Saha

Shaping Perception of Emotional Storytelling with Synthesized Speech
Presenter: Arya Honraopatil

Plasmonically Enhanced 2D Material based Phototransistor with Efficient Heat Management
Presenter: Raonaqul Islam

Formation of Multiple Stable Regions for Single Solitons in the Presence of an Avoided Crossing
Presenter: Logan Courtright

A New Benchmark Solution for T-cell Receptor Epitope Binding Prediction
Presenter: Chen Wu

Leveraging LLMs for the Construction and Validation of an IoT Knowledge Graph
Presenter: Gia Oriana Santos

The Dual Role of Student and Creator: Exploring the TikTok Experience
Presenter: Saquib Ahmed

Studying Multi-Color Solitons using 3-Wave Equations
Presenter: Pradyoth Shandilya

An Efficient Algorithm for Modeling the Slow Evolution of Solitons Due to Interaction and Noise in Microresonators
Presenter: Sanzida Akter

Calibrating Practical Privacy Risks for Differentially Private Machine Learning
Presenter: Yuechun Gu

Adaptive Domain Inference Attack with Concept Hierarchy
Presenter: Jiajie He

Continuous Blood Pressure Monitoring Using Smartphones and Wearables: A Survey of Recent Advances and Challenges
Presenter: Riishav Guptaa

UltraControl: Capturing Ultrasound Leaked from Home Appliances for Smartphone-Based Gesture Control Leakages
Presenter: Kodilinye Mkpasi

From Anomaly to Novelty: Active Detection and Adaptive Response in Smart Grids
Presenter: Leann Alhashishi

BoardVision: Real-Time Motherboard Defect Detection using YOLOv7 and Faster R-CNN
Presenter: Brandon Hill

Self-Calibration of LLMs for Robust Classification
Presenter: Christian Angel

Understanding data chart images with LLMs
Presenter: Sadia Tisha

The Digital Nutrition Label
Presenter: Aijaz Shaik

EViT-CiM: A Hybrid SRAM/DRAM Compute-in-Memory Architecture for Embedded Vision Transformers
Presenter: Dhandeep Challagundla

EvoLVE: Event-optimized Lightweight Vision Engine
Presenter: Joseph Mule

Filter Bank Common Spatial Pattern on Motor Imagery EEG Data
Presenter: Nathan Dayie

Copula Based Statistical Technique for Intrusion Detection System in CAN
Presenter: Rachit Saini

“LOADS: LiDAR-based Privacy-Preserving Queue Monitoring and Analysis”

Presenter: Bharg Barot

Abstract: "Long queues in retail and public environments can frustrate customers and negatively impact user experiences. Traditional camera-based monitoring systems are effective in analyzing queues, however, the potential for identification raises privacy concerns. Other queue-counting methods (such as WiFi or RFID) depend on user-carried devices or tags. In contrast, LiDAR sensors strictly measure distances and angles, which drastically reduces privacy risks and does not require users to carry specialized hardware. We present LOADS, an end-to-end, single-sensor, LiDAR-based IoT solution for queue-occupancy and wait-time estimation. LOADS incorporates Hierarchical Density-Based Spatial Clustering of Applications with Noise (HDBSCAN) to provide a robust method of accurately separating people from noise in real time. We employ a SARIMAX model to predict future queue lengths from historical data stored in a time-series database. A web interface shows real-time and historical queue information, enabling users to make informed decisions. We demonstrate the feasibility of LOADS in practical retail and conference scenarios, highlighting its privacy-preserving nature, accurate crowd estimation, and simple deployment."

“Artificial Intelligence and Human-Robot Interaction”

Presenter: Dr. Cynthia Matuszek

Abstract: We live in a time of tremendous advancements in artificial intelligence and robotics. Robots are rapidly becoming smaller, more capable, and less expensive; at the same time, advances in natural language processing, particularly in transformer-based architectures, are enabling computers to answer questions and support people's needs in a variety of novel ways. However, there is still a gap in what robots can actually be deployed to do, largely because modern systems are not flexible in the face of dynamic environments and changing tasks—a robot that is designed to clean the kitchen cannot handle a new task, even a closely related one such as picking up the living room. This is partly a failure of Human-Robot Interaction, or HRI; a person should be able to express what they want done in an intuitive way and have the robot follow those directions. In this talk, I will discuss how the Interactive Robotics and Language lab works on bringing robotics and natural language together to address this gap, enabling us to build robots that understand and act on language-based instructions and descriptions of the world.

DUNE: A Machine Learning Deep UNet++ based Ensemble Approach to Monthly, Seasonal and Annual Climate Forecasting

Presenter: Pratik Shukla

Abstract: Capitalizing on the recent availability of ERA5 monthly averaged long-term data records of mean atmospheric and climate fields based on high-resolution reanalysis, deep-learning architectures offer an alternative to physics-based daily numerical weather predictions for subseasonal to seasonal (S2S) and annual means. A novel Deep UNet++-based Ensemble (DUNE) neural architecture is introduced, employing multi-encoder-decoder structures with residual blocks. When initialized from a prior month or year, this architecture produced the first AI-based global monthly, seasonal, or annual mean forecast of 2-meter temperatures (T2m) and sea surface temperatures (SST). ERA5 monthly mean data is used as input for T2m over land, SST over oceans, and solar radiation at the top of the atmosphere for each month of 40 years to train the model. Validation forecasts are performed for an additional two years, followed by five years of forecast evaluations to account for natural annual variability. AI-trained inference forecast weights generate forecasts in seconds, enabling ensemble seasonal forecasts. Root Mean Squared Error (RMSE), Anomaly Correlation Coefficient (ACC), and Heidke Skill Score (HSS) statistics are presented globally and over specific regions. These forecasts outperform persistence, climatology, and multiple linear regression for all domains. DUNE forecasts demonstrate comparable statistical accuracy to NOAA's operational monthly and seasonal probabilistic outlook forecasts over the US but at significantly higher resolutions. RMSE and ACC error statistics for other recent AI-based daily forecasts also show superior performance for DUNE-based forecasts. The DUNE model's application to an ensemble data assimilation cycle shows comparable forecast accuracy with a single high-resolution model, potentially eliminating the need for retraining on extrapolated datasets.

“Tensor Decomposition for AI: Applications in Cyber-security, Data Privacy, Model Compression, and Hallucination Reduction”

Presenter: Dr. Maksim Eren

Abstract: Tensor decomposition is a powerful unsupervised machine learning technique for discovering hidden patterns in large-scale data. This presentation will explore key applications of tensor methods in the domains of cybersecurity and data privacy. These works highlight the diverse utility of tensor decomposition in detecting anomalies in network traffic and power grids, identifying SPAM emails, mitigating credit card fraud, detecting and classifying malware (including novel families), user behavior analysis, and implementing data privacy via federated learning frameworks. In addition to these applications, the presentation will also cover the use of tensor decomposition techniques for compressing large language models (LLMs), as well as strategies for reducing hallucinations, two critical aspects of deploying LLMs efficiently and reliably in real-world systems.

Event-Based Vision Meets Compute-In-Memory: A Path to Scalable, Energy-Efficient AI

Presenter: Dr. Riadul Islam

Abstract: Event-based vision revolutionizes traditional image sensing by capturing asynchronous intensity variations rather than static frames, enabling ultrafast temporal resolution, sparse data encoding, and enhanced motion perception. While this paradigm offers significant advantages, conventional event-based datasets impose a fixed thresholding constraint to determine pixel activations, severely limiting adaptability to real-world environmental fluctuations. Lower thresholds retain finer details but introduce pervasive noise, whereas higher thresholds suppress extraneous activations at the expense of crucial object information. To mitigate these constraints, in this talk, I will discuss our two recently released event datasets for edge artificial intelligence (AI), the Smart Event Face Dataset (SEFD) and the Event-Based Crossing Dataset (EBCD).

Besides event vision, the evolving compute-in-memory (CiM) paradigm tackles this issue by facilitating simultaneous processing and storage within static random-access memory (SRAM) elements. Numerous design decisions taken at different levels of hierarchy affect the figures of merit (FoMs) of SRAM, such as power, performance, area, and yield. The absence of a rapid assessment mechanism for the impact of changes at different hierarchy levels on global FoMs poses a challenge to accurately evaluating innovative SRAM designs. This talk will present an automation tool designed to optimize CiM's energy and latency.

My Struggle with Quantum Entanglement

Presenter: Dr. Sam Lomonaco

Abstract: What is quantum entanglement (QE), and how can it be quantified?

This question has led to a surprising research journey.

For n-qubit systems, a cryptic answer is that QE is a fragile symmetry of the local group L(n)= SU(2)^[(X)n] lying in the special unitary group SU(2^n), where the Lie algebra of L(n) induces vector fields that in turn define a system of partial diU erential equations, the solutions of which give a complete set of QE invariants. The actual state space of an n-qubit system is the Kahler manifold complex projective space CP^((2^n)-1), which is the base space of the circle bundle S^[(2^(n+1))-1] CP^((2^n)-1). The non-zero De Rham cohomology classes associated with this bundle provide a global insight into the structure of QE.

Flexible Functional Magnetic Resonance Imaging Data Analysis Using Machine Learning Methods

Presenter: Dr. Seung-Jun Kim

Abstract: Functional magnetic resonance imaging (fMRI) is a non-invasive neuroimaging technique used to capture neural activity. Data-driven approaches based on latent variable models and matrix/tensor factorizations are increasingly being employed in fMRI data analysis. At the same time, there is a growing need to conduct large-scale analyses involving hundreds to tens of thousands of individuals. In this talk, we will present a progression of research addressing this important challenge, leveraging powerful machine learning techniques such as dictionary learning and deep learning. This body of work is the culmination of efforts by two Ph.D. students, one Master's student, and both internal and external research collaborations.

Calibrating Practical Privacy Risks for Differentially Private Machine Learning

Presenter: Yuechun Gu

Abstract: Differential privacy quantifies privacy through the privacy budget ϵ, yet its practical interpretation is complicated by variations across models and datasets. Recent research on differentially private machine learning and membership inference has highlighted that with the same theoretical ϵ, setting, the likelihood ratio-based membership inference (LiRA) attack success rate (ASR) may vary according to specific datasets and models, which might be a better indicator for evaluating real-world privacy risks. Inspired by this practical privacy measure, we study the positive correlation between the ϵ, setting and ASR. We also find that for a specific dataset and a specific task we can lower the attack success rate by modifying the dataset. As a result, we may enable flexible privacy budget settings in model training. One dataset modification strategy is selectively suppressing privacy-sensitive features without significantly damaging application-specific data utility. We use the SHAP (or LIME) model explainer to evaluate features' privacy sensitivity and utility importance and develop an optimized feature-masking algorithm. We have conducted extensive experiments to show (1) the inherent link between ASR and the dataset's privacy risk in terms of a specific modeling task; (2) By carefully selecting features to mask, we can preserve more data utility with equivalent practical privacy protection and relaxed ϵ, settings. The implementation details are shared online at https://github.com/RhincodonE/ On-sensitive-features-and-empirical-epsilon-lower-bounds.