CCSCNE '26

The Thirtieth Annual
Consortium for Computing Sciences in Colleges
Northeastern Conference

April 10 – 11, 2026

Hosted by Smith College
Northampton, Massachusetts

Undergraduate Student Poster Abstracts

CCSC
CONSORTIUM FOR
COMPUTING SCIENCES
IN COLLEGES

ccsc.org

Poster 5001

Personalized Interactive Calendar Program

Alan Catic, Colin Zouck
Computer Science
Utica University, Utica, NY

Faculty Advisor
Unnati Shah
Utica University, Utica, NY

Imagine relying on a digital calendar every day, only to discover that it cannot warn you about conflicting events, lacks meaningful personalization, forces you to navigate multiple screens to understand your own schedule, and provides limited support for recurring events. Modern calendar applications, such as Google Calendar and iPhone Calendar, are widely used, yet these unresolved limitations impact millions of users. These systems often restrict visual customization, fail to detect overlapping events, provide minimal real-time feedback, and make it cumbersome to manage multiple entries efficiently. Such shortcomings can lead to missed appointments, confusion, and reduced productivity, highlighting the need for more user-friendly and reliable solutions. To address these challenges, we developed a fully customized and interactive web-based calendar designed to enhance usability and reliability. Using HTML, CSS, and JavaScript, we created a personalized interface that incorporates a custom banner, hover effects, and a clean visual layout to improve user experience. A pop-up modal window allows users to add event details, including title, description, and time, with events stored dynamically during runtime. Scheduled events are displayed both inside the modal and directly on the calendar through day-specific indicator dots. One of the key contributions of our project is the event-overlap detection system. When users attempt to schedule multiple events at the same time, the program immediately alerts them and blocks the conflicting entry. This prevents scheduling errors and ensures accurate, conflict-free planning. Real-time feedback during event creation further enhances interaction, making the calendar smoother and more intuitive to use. For future development, we plan to implement persistent data storage, recurring event management, mobile responsiveness, automated notifications, and customizable visual themes. These improvements will expand usability and create a more adaptable, engaging, and reliable scheduling tool. By addressing the limitations of traditional digital calendars and improving personalization, accuracy, and interaction, this project provides a more dependable and user-friendly alternative for everyday scheduling needs.

Poster 5002

The birb poster

William Henry Mutschler

STEM

Landmark College, Putney, VT

Faculty Advisor

Karina Assiter

Landmark College, Putney, VT

My advisor and I are affiliated with a course on data visualization, a subject focused on transforming data into programs and graphics that present information in a clear and meaningful way. The final assignment for this course required students to design an original visualization using datasets that we either located from external sources or generated ourselves.

For my project, I selected a topic that I initially perceived as primarily "entertaining" and that some might even consider "gimmicky." The visualization depicts each U.S. state's official state bird and includes additional information, such as how each bird was selected as the state bird and current estimates of its population. To construct the dataset, I used online tools and resources to efficiently gather the substantial amount of information required. I then implemented the visualization using D3 and the JavaScript programming language to create an interactive display of the data.

At the outset, I did not regard the project as particularly useful or important; I even titled it "Birb Map" to reflect my view of it as something lighthearted or simply "neat," rather than consequential. My primary objective was to create a visual representation of a topic that I personally found engaging, rather than working with subject matter in which I had little interest. I chose birds because they are familiar to nearly everyone and are widely appreciated, making the topic more likely to appeal to a broader audience. While I could have chosen a topic more narrowly tailored to my own interests, I was concerned that it might not resonate with others.

Over time, however, I came to recognize that this project may hold greater value than I originally assumed, particularly for individuals who are curious about the world around them. Children, casual learners, or people exploring information for its own sake may find a state bird map both engaging and informative. It satisfies a kind of curiosity that can prompt someone to think, "I am glad this exists." I would be pleased to present and discuss the project in more detail in person.

Poster 5003

Smart Traffic Routing Using AI and Software-Defined Networking

Nathaniel Illingsworth, Alexander Gomez, Makayla Coleman
Computer Science
Utica University, Utica, NY

Faculty Advisor
Unnati Shah
Utica University, Utica, NY

Network congestion at key routers often leads to packet delays, uneven traffic flow, and reduced quality of service. Traditional static routing approaches struggle to respond dynamically, creating bottlenecks that impact performance across cloud infrastructures, 5G networks, and smart-city systems. This study explores how combining Software-Defined Networking (SDN) with Artificial Intelligence can improve routing performance during congestion. Through simulated scenarios, we observed packet buildup at specific routers and tested whether AI-enhanced SDN could respond more effectively by selecting smarter alternative paths. Using models such as LSTMs for traffic prediction, Random Forests for anomaly detection, PPO for real-time policy updates, and Graph Neural Networks for routing optimization, we compared traditional routing against AI-assisted strategies. Our analysis showed that AI-driven SDN consistently reduced packet buildup, distributed traffic more efficiently, and improved failover speed, leading to better overall network performance. These findings suggest that AI-enhanced SDN can enable more adaptive, stable networks and may provide valuable insights for emerging systems in 5G, cloud computing, and smart-city environments. Future studies could explore heterogeneous network architectures and address security concerns to further evaluate AI-assisted routing.

Poster 5004

Sustainable Networking: Energy-Efficient Routing for a Greener Internet

Patrick Penna, Connor Pirkl
Computer Science
Utica University, Utica, NY

Faculty Advisor
Unnati Shah
Utica University, Utica, NY

Modern networking consumes massive amounts of energy through routers, switches, and data centers. Inefficient transmission paths and idle routing components often result in unnecessary energy consumption. Notably, the active routing components, not the hardware of end devices themselves, are the primary targets for optimization. This means that not all end and networking devices need to operate simultaneously to maintain network functionality. To address this issue, this study investigates the use of Green Ethernet within standardized Star and Mesh topologies. Using Cisco Packet Tracer simulations in a Cisco routing environment, inactive links are no longer kept constantly active, with power savings influenced by cable size. Limiting the number of active nodes significantly reduces energy consumption from active routing components. Additional modifications to Star and Mesh topologies including structural adjustments, redundancy management, and node availability based on activity were explored to push the limits of energy optimization. Several challenges remain, such as reduced redundancy in Mesh topologies for the sake of efficiency, careful consideration of topology deployment and node placement, and limited standardization of some proposed techniques. Despite these trade-offs, this study aims to find a "middle ground," achieving a balance between Quality of Service (QoS) and energy efficiency.

Poster 5005

AURA: File Conversion Made Simple

Owain Lucas
STEM Department
Landmark College, Putney, VT

Faculty Advisor
Karina Assiter
Landmark College, Putney, VT

AURA is an audio conversion software built on a Java code base by Bachelors student Owain Lucas in the STEM department at Landmark College in Putney, VT. The goal of AURA is to be an easy tool for radio stations, production units, and newsrooms to manage audio files and their data effectively. The idea started with my work in the radio station and my desire to bridge that experience with my computer science background through a software enginnering project. I researched the current solutions for audio conversion available on the software market online and with our college station having a limited budget, I could not find a viable solution that kept the tag data and did not have a price tag. Originally known as “EchoMaster”, AURA improved on a lot of the problems that the original software had, such as performance, user experience, and versatility. This improved software uses numerous libraries including Swing, JAVE, and JAudioTagger to process files. One of the major hurdles with working with multimedia files is the amount of formats available now (ie. WAV, MP3, FLAC, WMA, AAC, OGG, AIFF). All these formats use different codecs for storing not only the audio contents, but also the text data associated with it, known as metadata. The different codecs mean data can sometimes be lost when converting between formats. AURA solves this problem by converting the file and transfering the metadata in separate steps for every file in its conversion queue. This approach always attributes that the only change between the two files is the format of the file, while the audio content and metadata remain the same. This makes it a reliable tool for many applications that regualrly work with audio. The poster will go over the key features of the program and the improvements it has over the old solution. Full descriptions and key features are presented in detail.

Poster 5006

Interactive Floor Plan Visualization for Energy Consumption Analysis

Eduardo Caceres
Computer Science
Eastern CT State University Willimantic, CT

Faculty Advisor
Garrett M. Dancik
Eastern CT State University Willimantic, CT

Energy consumption and efficiency are critical to consider in modern architecture; however, traditional energy analysis tools often lack user friendliness and customization. This research presents an interactive visualization tool made to assist in analyzing energy consumption in a floor plan. Implemented in Python using the Pygame library, the tool allows users to create a floor plan, assign energy use to each room in kWh, and visualize consumption through colors and a summary graph. Unlike complex traditional tools, this program presents simplicity and real-time interaction to make it accessible for early-stage design decision-making. A demonstration shows that the visualization tool highlights high energy consumption areas, allowing users to explore various alternative layouts and energy-saving strategies. Limitations include the lack of real-world data integration and predictive modeling that would take into account real-world energy
use and fluctuations due to weather and other environmental factors, which could be addressed in future updates. This tool represents a first step toward building needed energy visualization systems that support sustainable building design.

Poster 5007

ScrapDish: A Food Waste Minimizer App

Adelina Chocho, Megan Mohr, Reeya Patel
Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

Household food waste remains a significant environmental, financial, and health issue often driven by forgotten ingredients and unclear expiration labels. ScrapDish is an AI-powered mobile application, available on both Apple and Google stores, designed to help reduce food waste through ingredient tracking and personalized recipe generation. The application allows users to log food items using OCR (Optical Character Recognition) to read expiration dates and product details directly from food packaging, reducing manual data entry. For fresh produce and items without a printed label, the application uses the YOLOv8 (You Only Look Once) object detection model to recognize fruits and vegetables from images and automatically label them. Users can also log items using speech-to-text or manual input as an alternative. Items are stored in a digital pantry and organized by storage location, such as the fridge, freezer, or pantry.

Based on the contents in the pantry, ScrapDish generates recipe recommendations that prioritize ingredients nearing expiration while respecting personalized dietary preferences and restrictions. Each generated recipe also includes a detailed nutrition profile to encourage more health-conscious food habits while also reducing food waste. Additionally, ScrapDish provides timely notifications to alert users of soon-to-expire items and a nutrition-focused chatbot that offers guidance on topics such as safe food storage and ingredient substitutions. To increase long-term engagement, the application includes a community-driven recipe forum where users can upload their own recipes, rate these recipes, and leave written reviews. Through collaboration and user interest, this forum will bring inspiration to users. Contemporary literature suggests that users are often disengaged with existing food waste reduction applications due to cluttered interfaces and the effort required to log items. We kept this in mind when designing ScrapDish with an emphasis on a simple visual design and automated logging to reduce user effort and encourage long-term engagement.

Poster 5008

Tempo.AI - Emulating a DJ via Artificial Intelligence

Sebastian Salazar, Alex Power, Chris Anzilotti, Ryan Fennelly
School of Computing and Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal, Kruti Shah, Daniel Minogue
Quinnipiac University, Hamden, CT

We present our experience developing Tempo.AI in the School of Computing & Engineering at Quinnipiac University. The goal of our project is to make a AI driven DJ that mixes and is more accessible by using its artificial intelligence to support smooth transitions, clear song timing, and simple user control. During our fourth year at Quinnipiac, our team designed an interactive AI DJ application. We had numerous ideas to start with for this project, but moving forward we slowly narrowed our choices down until only one remained. To start with, our goal was to explore how different interface and audio decisions affect usability, including how quickly a new user can understand controls and produce a coherent mix. Before deciding upon this direction, we investigated a hardware-style controller interface and a playlist-only recommendation tool, but decided against both in favor of mixing AI assistance that keeps creativity with the user. After turning down those alternatives, we examined approaches suitable for building a small but complete DJ workflow. Having in mind a general idea of the experience we would like to design, we experimented with interaction flows and audio-processing strategies, ultimately deciding upon a three-model design. The first model supports AI-driven mixing by matching tempo, applying frequency filtering to reduce clashing elements, and adjusting pitch when needed so transitions remain musically coherent for listeners. The second model automatically segments songs into musically useful sections to support natural transition points, and the third model enables natural-language commands through a prompt box, including auto-queueing tracks based on metadata and similarity. Our poster will discuss our reasoning for choosing the AI DJ topic and the three-model design, and the challenges we encountered integrating real- time audio control with model outputs. The full description and implementation features of Tempo.AI, along with a planned evaluation comparing transitions to Apple’s and Spotify’s AutoMix, will be included in detail.

Poster 5009

English To American Sign Language Robot

Morgan Montz, Evan Vastakis, Eric May, Shawn Acheampong
Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

In a world where verbal communication is the default, deaf and non-speaking individuals are often overlooked. Utilizing a combination of hardware, software, and AI, we have created a new way to effectively bridge the communication gap between deaf and speaking individuals. This project aims to build a robot that can accurately translate spoken English to ASL (American Sign Language), making spaces more accessible for those that are hearing-impaired.

With many studies suggesting that current forms of deaf to hearing communication are impersonal and non-inclusive, this robot aims to provide a more natural and efficient method of translation. Mimicking a human translator, the robot consists of a torso and subsequent limbs resting atop a mount that can be wheeled around or stay in one specific location. With the uncanny valley in mind, the robot features a facial screen that pairs ASL signs with animated expressions, which is installed on the head. The software is run on a Raspberry Pi 5, connected to two ESP32's each with their independent power supply, allowing for powerful computation to support the local AI models running on the hardware. To translate, the robot captures audio, converts it to text, and then the text translates to ASL gloss tokens via a local lightweight AI model. These tokens are then used to access the MongoDB database to access the motion scripts which allow the robot to execute the signs. Through several waves of testing with Hartford Healthcare, signs were verified and tweaked to ensure accuracy. For signs that are not contained in our database, we have a fast and reliable finger-spelling system.

Poster 5010

Optimizing Adaptive Blending Units

Joshua Feldman
Department of Computer Science
Skidmore College, Saratoga Springs, NY

Faculty Advisor
Michael Eckmann
Skidmore College, Saratoga Springs, NY

Adaptive Activation Functions (AAFs) offer a promising alternative to fixed nonlinearities in deep neural networks by allowing networks to learn activation shapes during training. One such method, Adaptive Blending Units (ABUs), constructs activation functions as trainable linear combinations of multiple base activations, enabling layer-specific optimization of nonlinear transformations. Although ABUs have demonstrated improved validation performance across several architectures, their adoption is limited by increased computational cost due to additional activation evaluations and trainable parameters.

This paper and poster introduce **Optimized Adaptive Blending Units (OABU)**, a computationally efficient variant of ABU designed to reduce training and inference overhead while preserving adaptive behavior. Three optimizations are proposed: (1) removal of the identity activation from the blending set, (2) gradient-based freezing of blending weights once stabilization is detected, and (3) dynamic hard-freezing to a single dominant activation when its learned weight exceeds a predefined threshold. These modifications reduce both trainable parameters and forward-pass computations over time.

A convolutional neural network was trained on the CIFAR-10 dataset to evaluate functionality and performance. Experimental results demonstrate that OABU successfully learns stable activation blends, frequently freezing within early epochs, while maintaining competitive classification accuracy. Benchmark comparisons show measurable reductions in training time, inference latency, and total trainable parameters relative to standard ABU implementations.
These findings suggest that adaptive activation functions can be made computationally practical without sacrificing flexibility, moving trainable nonlinearities closer to viable deployment in large-scale deep learning systems.

Poster 5011

InterQu: AI Enhanced Practice Interview Software

Aidan Armellino, Lucas Jenkins, Hayden Lacy, Connor Ryan
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

Interviewing skills are integral to furthering someone's career and a major hurdle for many. InterQu's goal is to help its users improve their skills in interviewing and to get our users to be comfortable for their next real interview. InterQu leverages AI and visual/audio analysis to simulate, grade, and coach users through their interview skills. InterQu uses facial recognition and emotion detection technology through LibraFace, allowing interviews to be dynamic. Mediapipe's Eye tracking and posture recognition can allow InterQu to detect whether a user is confident or uneasy. If a user performs well in their mock interview, harder and more thought-provoking questions will be given for the user to answer. In turn, if the user is struggling, easier questions that will aim to make the user more comfortable will be given instead. Through mock interview recordings, users will receive feedback from AI that will point out the specific skills that the user needs to improve so that they can practice their skills overtime. We are using OpenAI, which allows us to scale questions based on their responses to these questions, thus allowing users to have a comfortable experience while also improving their skills. A user's interview transcript can be seen to the AI model to evaluate their response. This means, after interviews are completed, users will receive actionable feedback that they can use to help improve their skills. With our backend Supabase database, each user will have their own profile where they can save their past interviews. By showing users feedback on their responses to questions, posture, eye movements, and emotions, users can then see these specific areas of where they can improve and incorporate strategies to fix their skills. Through repetitive use and practicing the feedback provided to them, users will hone their skills and see themselves grow.

Poster 5012

PathFinder: Autonomous Indoor Security Drone

Grant Foody, Zakaria Merzougi, Fahim Kalange, Daniel Vitashkevich
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

PathFinder is an autonomous indoor drone designed to address limitations in traditional indoor security systems, such as fixed camera blind spots and the high cost of those permanent systems. The project's focus was developing an affordable, self-contained drone capable of navigating commercial or residential indoor environments during times where humans aren't present, such as overnight or during after-hours security sweeps. Using onboard sensing, autonomous navigation, and pathfinding algorithms, PathFinder performs patrol routes and is able to monitor areas in a way that traditional systems cannot. Users receive alerts whenever the drone's camera detects a human. As part of PathFinder's initial setup, users scan their environment using their phone's camera. Users then choose a drone-accessible location for its dock, which the drone will return to when it is not active. Finally, small transceivers are placed throughout the space to allow accurate indoor positioning. Once set up, users only interact with the PathFinder website. User interactions include changing hours of operation, temporarily toggling off the drone, and forcing a sweep. The portable nature of the system allows it to be packed up and moved to another space. The path it takes throughout the space changes and the order in which rooms are visited is shuffled to avoid a predictable pattern that could be exploited. PathFinder shows how affordable autonomous indoor drones can be a viable alternative to conventional monitoring systems for homes and small commercial environments while providing a foundation for future work in indoor mapping and intelligent security systems.

Poster 5013

Compile Quest

Ellie Macdonald, Diego Matayoshi. Lehan Dharmatilake, Mark Pfister
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

Right now, the computer science industry is a vital and abundant job market, supplying lucrative opportunities for those interested. However, an issue has become apparent to us: education for this field is lacking, and there are no good platforms for students to educate themselves. We have found that platforms that tout themselves as educational, such as *Scratch* or *Code.org*, do not improve a student's cognitive ability or problem-solving skills. We have developed a solution to the problem called Compile Quest, and this study is to demonstrate that our platform educates students in a fun and effective manner. Compile Quest seeks to gamify and ease newer learners into coding through structured, level-based learning. Within Compile Quest, the player is presented with multiple planets that cover a vast array of coding concepts that would typically be taught to students from ages 10-13. Through the earlier planets, the player is informed about simpler coding concepts such as variable declaration and loop usage. As they proceed to the later planets and levels, they're introduced to more advanced concepts for their age group. Once we roll out our finished product to our two demo schools, *Saint Rita's* and *Hamden Middle School*, we aim to test how well students learn from our platform. A control group of students is asked to decipher a block of code before starting. Then, a group of students will clear the first two worlds of the game. Additionally, we will have a third group play Scratch to compare our results. Finally, the study group is asked to decipher a similar block of code. Our goal isn't to teach Python or Java but to teach concepts and reasoning. By the end of our game, the goal is for our users to understand core concepts and improve their reasoning.

Poster 5014

ReMind - A Convenient at Home Solution to TBI Rehabilitation

Michael Alvarado, Billy Godoy, Joshua Molin
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal
Quinnipiac University, Hamden, CT

Traumatic Brain Injury (TBI) is referred to as an injury that is caused by trauma to the head and can cause long term and short-term effects to regular behavioral functions. Based on the degree of TBI, some patients must schedule appointments with a neuropsychologist to assess cognitive and behavioral function. Although neuropsychologists can get data from these appointments, there is difficulty in assessing the patients in their day-to-day lives due to the lack of consistent monitoring of TBI symptoms and their impact. To address these concerns, our mobile application titled “ReMind” serves as an at-home therapy for patients to aid them in treatment. We designed and implemented ReMind as an iOS and Apple Watch application that delivers daily cognitive exercises while collecting physiological and behavioral data outside of clinical environments. The system was developed using a serverless cloud architecture to support secure authentication, scalable data storage, and long-term monitoring.

With our product, we aim to improve the rehabilitation process for people who have gone through a TBI by making consistent monitoring more accessible and convenient to fit in their daily lives. This would be made possible by prompting the user to complete three exercises daily to keep the users engaged in their recovery process. The data collected from these exercises will be compiled into one main score, while the provider will be able to access the detailed data and monitor them to provide their own insight. The exercises are also built with user adaptivity in mind, allowing the app to select levels based on performance and physiological signals such as heart rate variability. Additionally, the app has journal check-ins and clinical grade surveys incorporated to improve the quality of data collected and better inform providers of their patients' status. ReMind demonstrates how mobile and wearable computing can support continuous rehabilitation and long-term data collection outside clinical environments.

Poster 5015

I Know You AR Smart Glasses

Brady Galligan, Aislin Hayes, Jacob Levin, Andrew McCleary
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal, Soumyashree Sahoo, Aditya Gupta
Quinnipiac University, Hamden, CT

The ability to quickly and accurately identify an individual poses a significant challenge for many, including individuals managing memory or facial recognition impairments. To assist in 'putting name to face', we developed the I Know You Glasses: a pair of glasses that combines augmented reality and computer vision technology to provide real-time facial identification onto a Heads-Up Display (HUD).

Our system was designed to split the load of two heavy convolutional neural network models across devices, ultimately solving the known 'pose variation issue' without compromising the speed or accuracy of identification. We prioritized responsiveness and high-performance, allowing for the information on the HUD to remain up to date without the lag that plagues other real-time wearables. All data obtained for and used by the glasses was stored securely and will be deleted following the completion of the project.

To demonstrate the usefulness of the glasses, we collected publicly available web-scraped information on Connecticut residents, in accordance with the CTDPA. From just a LinkedIn profile picture and voter records, we were successfully able to create a device which can identify passersby ‘at a glance’. This project was developed entirely using open-source models and information, demonstrating the growing surveillance capabilities of modern technology, and will not be released.

Poster 5016

SmartFocus: A Novel Study on Stress Facilitated by a Smartphone Application

Wilson Chen, Luis Gilbuena, Hunter Pageau, MD Fayed Bin Salim
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Soumyashree Sahoo, Tobias Wasser
Quinnipiac University, Hamden, CT

Chronic stress is widespread among college students and contributes to academic difficulties and mental health challenges. Understanding how everyday behaviors relate to stress can support earlier awareness and prevention. Smartphones offer a scalable, low-burden means of observing such behaviors in real-world contexts.

SmartFocus is a cross-platform smartphone application designed to examine relationships between smartphone usage patterns and perceived stress. The application collects daily anonymized app-usage data and weekly self-reported stress measures. App usage data are aggregated into categories (e.g., social, entertainment, gaming), while stress is assessed using a weekly adaptation of the validated Perceived Stress Scale (PSS-10), referred to as PSS10-W.

The project is an interdisciplinary collaboration between the School of Computing and Engineering and the Frank H. Netter MD School of Medicine, guided by clinical psychiatrist Dr. Tobias Wasser to ensure appropriate interpretation, ethical oversight, and responsible handling of stress-related measures. Data collection will involve only students at Quinnipiac University, with no inclusion of external populations.

SmartFocus' system architecture comprises a React Native + Expo mobile frontend, a Swift-based extension enabling access to Screen Time data on iOS, and a secure cloud backend hosted on a Linode virtual machine with SSH-restricted access. All data is transmitted exclusively over encrypted HTTPS connections. A lightweight Node.js/Express service manages the ingestion and storage of aggregated app-usage metrics and survey responses. Participant anonymity is preserved using pseudorandom identifiers, and no personal or device-specific information is collected or stored.

Poster 5017

Beyond Accuracy: A LIME-Based Comparative Analysis of Custom CNN and Pretrained ResNet50 for Lung Cancer CT Classification

Samarpan Poudel
Computer Science
Caldwell University, Caldwell, NJ

Faculty Advisors
Vladislav D. Veksler
Caldwell University, Caldwell, NJ

Lung cancer remains one of the leading causes of cancer-related mortality, which claims 1.8 million lives annually. While deep learning models have demonstrated high performance in CT-based lung cancer classification, model interpretability remains a critical challenge for clinical adoption. To overcome this issue, we examine not only classification accuracy but also the representational behavior of different neural network architectures using explainable artificial intelligence techniques (XAI). We developed a custom convolutional neural network (CNN) and compared it with a fine-tuned ImageNet-pretrained ResNet50 model on a lung CT dataset. Both models were trained and evaluated on a publicly available dataset using a stratified train–validation split to ensure balanced class representation. The custom CNN achieved 96.5% accuracy on unseen data with 98% precision, while the pretrained ResNet50 achieved 97.5% accuracy with 97% precision.

Furthermore, we applied Local Interpretable Model-Agnostic Explanations (LIME) visualizations to generate explanation maps for both correct and incorrect predictions. Our findings reveal distinct attention patterns between models. The custom CNN tends to rely on more localized lung regions, whereas the pretrained ResNet50 exhibits broader, context-driven attention across larger structural areas. Moreover, failure-case analysis shows that the two models misclassify images for different underlying reasons, reflecting architectural biases that are not apparent from accuracy metrics alone. These findings suggest that combining visual explanation with clinical metadata can notably enhance lung cancer diagnosis and support real-time decision-making in medical applications. Our ongoing work focuses on investigating the quantitative correlation of LIME-generated superpixels within decision-making regions to better characterize model attention patterns and enhance the reliability and precision of interpretability analysis.
.

Poster 5018

QChat: An AI Driven Conversational Assistant for Improving Information Access in Higher Education Portals

Thomas Rua, Tuana Turhan, Samantha Woodburn
School of Computing & Engineering
Quinnipiac University, Hamden, CT

Faculty Advisor
Chetan Jaiswal, Lynn Byers
Quinnipiac University, Hamden, CT

While universities increasingly depend on digital platforms to manage academic resources, administrative services, and student support, the expansion of these systems has led to challenges navigating information, causing user confusion and frequent reliance on support staff. This work introduces an AI-driven, conversation based, personal assistant designed to improve accessibility and efficiency when searching for information regarding university resources. This system will streamline information retrieval and enhance overall user experience within higher-education digital environments. This chatbot is being tested and implemented within a university environment but could be easily adapted to other organizations by modifying the underlying resource base.

This work examines common usability challenges in educational platforms and shows how an AI-based solution can address these issues. The proposed system leverages NLP, RAG, and secure integration with university infrastructure to deliver accurate, context-aware responses. Unlike traditional search bars or static help pages, the conversational interface interprets user intent and provides reliable, citation-supported answers. The system is also capable of referencing prior interactions with the current user to develop personalized responses based on the user's interests, while ensuring that all personal information is encrypted to protect students' privacy and comply with FERPA regulations. Encrypted student information and FAQ data is stored within MongoDB, where it can be monitored by the system administrators.

By simplifying routine inquiries and guiding users through complex digital portals, the chatbot reduces workload for administrative personnel while offering students and staff consistent, round-the-clock assistance. Ultimately, the goal is to demonstrate how an intelligent conversational assistant can improve satisfaction, efficiency, and digital accessibility across modern higher-education institutions.

Poster 5019

Roots in Resilience: Women of Color in STEM Garden

Sarah Jaimes Santos
Computer Science
Kalamazoo College, Kalamazoo, MI

Faculty Advisors
Sandino Vargaz-Perez
Kalamazoo College, Kalamazoo, MI

This project is a response to the systemic exclusion of women of color in STEM. We present the development of an interactive digital platform designed as a "counterspace"—an environment where women of color can reconstruct a sense of self-worth and belonging. The website functions as a living archive, combining a historical timeline, a user-generated story map, and a real-time community chat to foster visibility, connection, and collective memory.

The project is grounded in research on racial battle fatigue, troubled success, and the "equity ethic" among students of color in STEM, which frames their work as a drive for justice rather than personal advancement. The platform was designed to operationalize these concepts, transforming isolation into dialogue. The design phase in Figma focused on creating a visual identity that embraces femininity, warmth, and cultural resonance—deliberately challenging the sterile, neutral aesthetic of traditional STEM spaces. The technical implementation used HTML, CSS, and JavaScript, with Firebase for data storage and Leaflet for the interactive map, to build a functional space where these theoretical commitments are made tangible.

The outcome is a digital intervention that asserts the presence and power of women of color in STEM. By centering identity and lived experience, the platform reframes technology as a medium for cultural survival and community care—a modern extension of ancestral knowledge and resistance. This project demonstrates that computing can be a tool for social justice, and that creating spaces of recognition is essential for challenging the exclusionary narratives that have long defined the field. It is a tangible example of how technical work can be grounded in identity and aimed at collective liberation.

Poster 5020

Evaluating Privacy–Utility Trade-offs in Federated Learning for Heart Disease Prediction

Jessica Theodore

Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisors

Brian O'Neill
Quinnipiac University, Hamden, CT

Federated learning (FL) enables multiple institutions to collaboratively train machine learning models without sharing raw data, making it well-suited for privacy-sensitive environments such as healthcare. However, federated learning alone does not fully prevent privacy leakage. Additional protections, such as secure aggregation and differential privacy, improve confidentiality but may reduce model accuracy and increase computational cost. Understanding these trade-offs is essential for real-world deployment of privacy-preserving machine learning systems.

This project evaluates privacy–utility trade-offs in federated learning for heart disease prediction using tabular electronic health record-style data. A cross-silo federated learning environment will be simulated by partitioning the dataset into multiple institutional datasets with different distributions. Three federated optimization algorithms will be compared: FedAvg, FedProx, and SCAFFOLD.

Privacy protections will include secure aggregation and central differential privacy. Differential privacy will be implemented using gradient noise and clipping, and privacy guarantees will be measured using epsilon values. Secure aggregation overhead will be evaluated through runtime and communication costs.

Each configuration will be evaluated using prediction accuracy, training time, communication overhead, and privacy guarantees. The results will quantify how optimization algorithms and privacy mechanisms interact and will identify configurations that provide practical balances between privacy protection and model performance.

This work aims to provide implementation-focused guidance for healthcare institutions seeking to deploy federated learning systems that protect patient data while maintaining useful predictive performance.

Poster 5021

Evaluating ChatGPT's Ability to Answer First-Order Logic-Based Questions

Stephen Canavan
Computer Science Department
Eastern CT State University Willimantic, CT

Faculty Advisor
Garrett M. Dancik
Eastern CT State University Willimantic, CT

AI chatbots are increasingly being used in all fields of study and work. With the continuing growth and use of AI chatbots, it becomes important to understand the accuracy of these chatbots. Specifically, this study evaluates the performance of ChatGPT's ability to understand and answer first-order logic-based questions. Using the Logic-Bench dataset and custom questions about Eastern Connecticut State University (professor office hours, contact information, and class times), ChatGPT was prompted with different types of logical questions, and its accuracy was evaluated. Using the logic-based programming language Prolog as a control for perfect performance, ChatGPT was evaluated by its accuracy. We found that ChatGPT made mistakes across all models given all types of logic and questions. Compared to Prolog, ChatGPT fell short of the perfect standard and had accuracy results varying between 20% and 100% across different logic and question types. While Prolog did perform with perfect accuracy, its viability as an Eastern chatbot is lacking, since it cannot be interfaced with using natural language, giving an advantage in that space to ChatGPT. This study proved that ChatGPT is a versatile tool with the ability to answer some logic-based questions; however, it is still not perfect and should be treated as such. Utilizing ChatGPT to its full effectiveness requires a user to understand its limitations and adapt their usage accordingly.

Poster 5022

coLearn-AI: A Process-Constrained Intelligent Tutoring System for Preserving Collaborative Reasoning in AI-Mediated Classrooms

Godwin Idowu
Program in Computing
Kenyon College, Gambler, OH

Faculty Advisor
James Skon
Kenyon College, Gambler, OH

We present coLearn-AI, a process-constrained intelligent tutoring system designed to preserve collaborative reasoning in classrooms increasingly mediated by large language models (LLMs). We developed core components of the platform, including the collaboration engine, role-based interaction constraints, deterministic group generation, and the persistent epistemic trace architecture that captures structured records of student reasoning. Recent advances in LLMs decouple output quality from cognitive effort, allowing students to generate correct answers without engaging in the reasoning processes that produce durable learning. Building on the pedagogical framework of Process-Oriented Guided Inquiry Learning (POGIL) and constructivist learning theory, this project operationalizes learning as durable cognitive restructuring rather than information transfer. The system enforces structured collaboration and rotates authority among group members to ensure equitable participation. The platform was architected using a React frontend with role-based rendering, a Node/Express backend implementing deterministic state transitions, and a database that persists activity instances, group membership, responses, and AI probes. The AI component is wrapped in a constrained interface: it may issue prompts and reflective probes but cannot provide authoritative solutions or grading decisions. A persistent epistemic trace logs dialogue, revisions, and AI interactions, enabling process-level analysis rather than answer-level evaluation. Deployment in multiple undergraduate computer science courses (COMP118, COMP218, COMP318, COMP368) demonstrated improved structural collaboration, measurable participation equity, and instructor visibility into reasoning processes. Unlike conventional collaborative tools that allow divide-and-conquer behavior or AI shortcutting, coLearn-AI enforces cognitive participation at the system level. This work contributes a computational model of AI-mediated collaborative cognition, offering a framework for preserving epistemic integrity while integrating generative AI into higher education.

Poster 5023

How Algorithms Respond to Shares and Comments on Social Media

Bitania Hailemichael
School of Management
Wentworth Institute of Technology, Boston, MA

Faculty Advisor
Hollis Greenberg
Wentworth Institute of Technology, Boston, MA

Social media platforms use engagement driven algorithms to decide which content appears in users' feeds. While misinformation and online behavior have been widely studied, less attention has been given to how algorithms respond specifically to commenting and clicking the share button. This research examines how these two forms of user engagement influence content visibility and contribute to the wider spread of posts. Research shows that algorithms treat the amount and speed of engagement as signals that a post is important, without evaluating whether the interaction is positive or negative. When users comment on a post, whether they agree, criticize, or try to correct it, the system counts that activity as engagement and may increase the post's visibility. Clicking the share button has an even stronger impact because it places the content directly into new users' feeds and exposes it to new audiences beyond the original network. Importantly, the algorithm does not reliably distinguish between supportive and critical responses. If users comment negatively on a false or misleading post, that interaction can still cause the system to show the post to more people, spreading it just as much as supportive comments would. This creates a built-in incentive within the platform for content that generates strong reactions, disagreement, or controversy. Posts that spark more interaction are more likely to be promoted, regardless of accuracy. Although platforms attempt to limit misinformation through warning labels or reduced visibility, continued commenting and sharing can weaken those efforts. Understanding how engagement based ranking systems operate is essential for improving how platforms balance user interaction with credibility and accurate information.

Poster 5024

Robo-Advisors and the Investor Experience

Anuj Khadka
Computer Science
Caldwell University, Caldwell, NJ

Faculty Advisor
Vladislav D. Veksler
Caldwell University, Caldwell, NJ

Global assets managed by robo-advisors (automated investment platforms) are projected to grow at a compound annual growth rate of 7.14%, reaching $2.8 trillion by 2030 [1,2]. While recent research has cataloged the prevalence of certain algorithms within these systems [3], there seems to be a lack of comparative analysis to determine meaningful differences and benefits.

This project addresses this gap by simulating popular portfolio construction algorithms, including MPT, Sample Portfolio, Risk Parity, Full Scale Optimization, and the Black-Litterman Model, along with a baseline measure (Equal Weight). A common asset universe consisting of U.S. equities, international equities, metals, aggregate bonds, and inflation-protected bonds was used, and all strategies were applied to identical monthly market data, starting wealth, and rebalancing structure to isolate the effect of algorithmic choice. Analysis focuses on internal portfolio behavior through time-varying weights and turnover, and comparison of investor outcomes across full-sample and crisis periods, focusing on drawdown paths, volatility exposure, and recovery dynamics, rather than terminal wealth alone.

Preliminary results show that algorithm choice noticeably changes what investors experience. During the 2022 rate-hike crash, drawdowns ranged from -16.4% (Naive Risk Parity) to -20% (Equal Weight) to -26.1% (MPT), a difference in 9.7% in peak-to-trough decline. These findings suggest that more complex algorithms offer no better protection in a crisis while generating more instability. More importantly, these results indicate that the proposed analysis framework provides a useful comparative robo-advisor analysis tool, filling a major gap in this domain.

[1] Jonathan Walter Lam. Robo-Advisors: A Portfolio Management Perspective. (Visited 01/2026).
[2] Statista. Robo-Advisors. https://www.statista.com/outlook/dmo/fintech/digital-investment/robo-advisors/worldwide. (Visited 01/2026).
[3] Mikhail Beketov, Kevin Lehmann, and Manuel Wittke. “Robo Advisors: quantitative methods inside the robots”. Journal of Asset Management 19.3 (2018).

Poster 5025

Simulation of Trade Between Capitalist and Labor-Time Economies

Zhejun Zhang
Department of Computer Science
Boston College, Boston, MA

Faculty Advisor
Amittai Aviram, Alexander Creiner
Boston College, Boston, MA

A labor-time economy (LTE) is a proposed economic system that uses labor time credits instead of money. This project develops a dynamic simulation of trade between a multi-commodity capitalist economy (CE) and an LTE. The system consists of (1) an event-based capitalist production model that generates prices, supplies, and reproduction requirements over time, and (2) a trade interface that mediates exchange between the two economies under strict value equivalence.

The capitalist backend is constructed using a Leontief input–output framework with labor coefficients grounded in classical political economy and Marxian value theory. At each time step, the model computes total reproduction demand as intermediate production requirements plus worker and capitalist consumption. Equilibrium is detected when prices and supplies stabilize; a 40-step post-equilibrium window then allows stochastic shocks and trade interaction. Trajectories of prices, supplies, reproduction demand, and reproduction gaps are recorded and exported for analysis.

The trade interface begins at the detected equilibrium and operates in discrete time. Capitalist exports are constrained to surplus above reproduction requirements, while LTE shortages are randomly generated. Trade is strictly balanced in monetary terms. Labor-time values are converted to money using a dynamically recomputed Monetary Expression of Labor Time (MELT), calculated from GDP per capita (PPP) and average annual hours worked. Relative prices are empirically calibrated using Producer Price Index data and anchored to a 2024 steel benchmark to obtain absolute monetary values.

Under current assumptions, the capitalist economy remains largely self-reproducing, and trade volumes are structurally limited by reproduction feasibility. Because exports occur only when surplus exists and trade value is symmetrically balanced, neither side accumulates persistent monetary gains. Currency trajectories remain stable over the 40-step window, demonstrating that the algorithm produces a reproduction-constrained, value-consistent, and dynamically sustainable trade regime.

Poster 5026

Evaluating Large Language Models for Phishing Email Identification

Regan Cunningham
Computer Science and Mathematics Department
Eastern Connecticut State University Windham, CT

Faculty Advisor
Garrett M. Dancik
Eastern Connecticut State University Windham, CT

Phishing attacks continue to be one of the most common and damaging forms of social engineering, due to the fact that they exploit a human's vulnerability to obtain sensitive information. As large language models (LLMs) become increasingly integrated into an average user's daily life, their untapped potential to assist a user in identifying malicious emails is significant. This study evaluates the effectiveness of three popular LLMs – OpenAI's ChatGPT-4, Google's Gemini 1.5 Flash, and Microsoft's Copilot with GPT-4 – in classifying phishing and non-phishing emails. Using a pre-written prompt, each model was tested with fifty emails, split between twenty-five phishing and twenty-five non phishing emails, taken from two different data sets. Accuracy bar charts and confusion matrices were used to measure the LLMs' performance. Results show that Gemini achieved 100% accuracy, ChatGPT achieved 94% accuracy, and Copilot achieved 92% accuracy. These findings show that popular LLMs can effectively detect phishing emails and can potentially be used by everyday users as an acceptable and reliable tool.

Poster 5028

Facial Expression Recognition and Response with Tension Mapping and LLM-Assisted Dialogue

Satyam Regmi
Computer Science
Caldwell University, Caldwell, NJ
Faculty Advisor
Vladislav D. Veksler
Caldwell University, Caldwell, NJ

This research showcases the design and development of an adaptive facial expression recognition system. The work comprises a browser-based system that uses real-time facial geometry analysis to calculate a "facial tension score" and incorporates an LLM to facilitate contextual conversations when there is a possibility of emotional ambiguity.

Conventional facial emotion recognition systems are highly dependent on static datasets and pre-defined emotion categories (such as happy, sad, and angry), which frequently overlook individual differences in facial morphology and intensity of emotional expressions. To overcome this drawback and design a system that is more personalized, interpretable, and adaptive for emotion analysis, this work focuses on deviation from a personalized neutral state instead of predicting a fixed emotion category.

Using MediaPipe Face Mesh (468-478 landmarks), geometric features such as eye aspect ratio, mouth aspect ratio, eyebrow movement, jaw movement, and head tilt were extracted. A per-landmark deviation model was developed to calculate real-time displacement based on a captured neutral expression. These displacements were then combined into a continuous "Tension Score" from 0 to 100 and represented as a facial tension heatmap to indicate localized muscular activity. When confidence in prediction was reduced or tension levels exceeded a threshold, an external LLM was activated to start dialogue-based clarification to support contextual emotional feedback.

Preliminary testing indicated that deviation from the baseline offered a more personalized assessment of emotional intensity than traditional classification methods based on the dataset. The Tension Score positively correlated with the user's emotion intensity. The combination of geometric modeling, heatmap, baselines, incremental learning, and dialogue-based AI appears to offer a promising approach to more human-centric emotion recognition systems.

Future research will focus on the development of emotional companionship systems for loneliness relief and infant stress detection systems that can start interventions to calm the infant or alert caregivers.

Poster 5029

Simulating Load Sharing in an Internet of Things Network

Evelyn Ray, Lucy Kreipke, Olohi John
Computer Science
Smith College, Northampton, MA

Faculty Advisor
Michael Robson
Smith College. Northampton, MA

We present the initial findings of our work on the project: Simulating Load Sharing in an Internet of Things Network. Over the past year we have experimented with several Parallel Discrete Event Simulators (PDES) to determine one that would be best suited for our project. We were evaluating PDES based on their maintenance status, functionality, and ability to model parallel events in order to run experiments modeling an Internet of Things network. Throughout this project we examined a variety of PDES, including ROSS, NS-3, Devastator, RootSim, and Simian. This involved reading documentation, tutorials, and research publications, building and running their example models on our local and remote machines, and working through any bugs we discovered. This involved directly debugging pre-existing models, reading literature that critiques the PDES, and, in some cases, communicating directly with the creators of the software. We compared and debated the benefits and detriments of each PDES and we have ultimately determined that Simian fits all of our requirements. We have begun writing models to run on Simian and our poster will discuss our analysis of each PDES and the comparisons that ultimately led us to decide on Simian for our project. It will also detail the next steps that we will take in our project to experiment with the efficiency of various load sharing algorithms in IoT systems.

Poster 5030

Emerging Technologies and the Design of Long-Term Adoption

Jahnavi Sharma
Information and Computer Sciences
University of Massachusetts – Amherst, Amherst, MA

Faculty Advisor
Ella Tuson
University of Massachusetts – Amherst, Amherst, MA

Emerging technologies often arrive surrounded by bold claims about transformation, efficiency, and the future of everyday life. Still, many never move beyond early curiosity. The gap between what a product promises and how people actually live with it points to a deeper design challenge. How does a technology move from demonstration to habit? This study investigates how early design choices shape that trajectory.

Focusing on wearable augmented reality systems such as Google Glass, spatial computing platforms like Apple Vision Pro, social smart glasses including Ray-Ban Meta Smart Glasses, and immersive metaverse environments such as Meta Horizon Worlds, this research explores how interface design, embodied interaction, and public presence influence long-term engagement. Each case reflects a different attempt to bring digital systems directly into shared social space.

The methodology combines close interface analysis with a synthesis of early adopter reviews, developer materials, and media coverage. Cognitive complexity, onboarding flow, and task structure were assessed in relation to patterns of sustained use and disengagement. Sentiment analysis across user forums and social media traced emotional responses, including excitement, hesitation, frustration, and curiosity. These insights were examined alongside observable behaviors such as public reception, frequency of use, and shifts in product framing.

Across cases, similar pressures surface: unclear everyday purpose, social unease during visible use, and initial excitement that fades without lasting relevance. Adoption emerges as a cultural and emotional process shaped by design. This project develops a user-centered framework for recognizing early signals of whether emerging technologies will find a durable place in daily life.

Poster 5031

GenMark: A New Testing Ground for Procedurally Generated Game Content

Maximus Riat

Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisor

Brian O'Neill
Quinnipiac University, Hamden, CT

I present my experience in developing GenMark, an evaluation tool for procedural content generator (PCG) algorithms for games. The goal of GenMark is to create a benchmark that evaluates PCG algorithms on a number of different qualities including the computational strains the algorithm has and the overall quality of the content generated. PCG developers can use GenMark to optimize the algorithm and generate more novel, high-quality content. Procedural content in games is primarily used to create dynamic maps and is usually limited to such. Through GenMark, I hope to be able to improve these systems by creating more diverse, high-quality, and controllable game content. The first part of the project was to determine the need for an overall benchmark for PCG systems. While other approaches do exist for evaluating algorithms, they either are limited to how they evaluate the content generated or are targeted towards different kinds of algorithms altogether. By combining aspects from all of these approaches, GenMark serves as an excellent analysis tool for the evaluation of procedural game content generation and will help bring these generators to a more creative state. To start with, I created a simple two dimensional game map generator that creates a simple map with several parameters fed in by the user, these being the diversity, quality, and controllability of the content generated. During generation, GenMark evaluates the map periodically using an A* algorithm to determine if the level is solvable and analyzes how performance intensive the generation is. The goal of this generator is to find the optimal combination of parameters that produces the most effective and most engaging game map based on "rewards" given to the game agent completing the level.

Poster 5032

WikiProject Africa: Analyzing the Relationship between User Engagement and Article Content

Hadassah Mwazemba
Computer Science
Wellesley College, Wellesley, MA

Faculty Advisor
Eni Mustafaraj
Wellesley College, Wellesley, MA

Wikipedia is a very popular website due to its extensive topic coverage and free access. However, compared to its millions of content consumers, only thousands actively contribute to it, since editing is voluntary. Given this imbalance, how should contributors allocate their time to improve topic coverage? I explored this in WikiProject:Africa, which hosts 100,000+ Africa-related articles. Articles on Wikipedia vary in quality labels, assigned by editors based on a scale ranging from Featured-Articles, which are comprehensive and detailed such as the one on Cleopatra, to Stubs, which provide basic descriptions of topics like the Cape-Verdian Diaspora. I set out to investigate the interaction between editors, who assign importance based on the aforementioned scale, and users, whose measure of an article's importance can be analysed through the pageviews it attracts. As a proxy for user engagement, I employed DPDP, a dataset that captures the number of daily pageviews for all the articles on Wikipedia, organized by country and language. My initial hypothesis was that high quality articles have high user engagement, and the analysis of 2023-2025 pageviews data supports that. However, I found that more than 80% of the articles in the wikiproject never got any pageviews in the three-year period. Furthermore, there is a large number of non-high quality articles which received pageviews above the median of their quality class, which I believe editors can focus on when thinking of improving topic coverage. With the insight from this analysis, I employ quantitative methods to find the most popular categories, excluding high quality articles, to discover content areas in which editors can concentrate on. The poster will discuss the details of my analysis and the challenges encountered during the project. I hope to inspire other CS students to build similar monitoring tools to improve free access to knowledge.

Poster 5034

Educational Tool for Psychology: Development of Interactive AI Chatbots Representing Mental Health Conditions

Sana Seth, Nicholas Barragato, Sean Smith, Sunny Miah
Computer Science, Psychology
Siena University, Loudonville, NY

Faculty Advisor
Jami Cotler, Dmitry Burshtyen
Siena University, Loudonville, NY

We developed an Artificial Intelligence (AI)–based educational chatbot designed to simulate conversations with research-grounded personas representing individuals with mental health conditions and psychological professionals. The goal of this project is to create an interactive learning tool that supports psychology and computer science students in developing both conceptual understanding and empathy through engagement with lived experiences, symptoms, and treatment approaches.

Traditional psychology education often relies on lectures and written case studies, which may not fully capture the emotional and contextual complexity of lived experience. Advances in artificial intelligence and natural language processing now allow for more immersive and responsive educational tools. Our project builds on prior conversational AI systems but focuses specifically on psychology education and evidence-based persona modeling. We were motivated by the need for scalable, interactive tools that allow students to explore sensitive psychological topics in a structured, academically grounded environment while maintaining consistency with peer-reviewed research.

The chatbot integrates speech recognition, large language model processing, and speech synthesis into a unified framework. User voice input is transcribed using Microsoft Azure and processed by ChatGPT, which generates persona-consistent responses guided by structured prompts derived from psychological literature and DSM-informed criteria. Responses are then converted into natural-sounding speech using ElevenLabs, creating a fully interactive voice-based experience. Multiple personas were developed to represent diverse mental health conditions and professional perspectives. The system architecture is modular, allowing additional personas to be incorporated without redesigning the framework as the project expands.

The chatbot promotes active learning and empathy development through interactive engagement rather than passive instruction. Potential applications include classroom instruction, training environments, and empathy research. Future work will focus on formal evaluation of learning outcomes and broader deployment through a web-based interface. This project highlights the potential of interdisciplinary AI systems to enhance psychology education through structured, evidence-based persona simulation.

Poster 5035

Print Media Database and Analysis

Pierre Field
Computer Science
Ithaca College, Ithaca, NY

Faculty Advisor
Ali Erkan
Ithaca College, Ithaca, NY

I present my research work and corresponding database for country-wide print media statistics that I synthesized for the Computer Science Department at Ithaca College. The goal of my project is to mold information about newspaper circulation into an accessible database and to compare the circulation information with other country-wide statistics to understand why print circulation increases or decreases in areas, while illustrating my findings with QGIS mapping systems. I determined this to be my topic as the differences between digital news and print news continually is widening, thus looking at what demographics can influence the choice of medium can provide insight into why different areas behave the way they do. The initial spreadsheet was provided to me, which I uploaded onto an Ubuntu server that can be remotely accessed by anyone so that this information is as accessible to the public as possible. I then created a GQIS representation of the density of print media purchases organized by zip code, which I then contrasted with different demographic statistics. One of our results illustrated that there is a general trend that states located on the western side of the country will have a decrease in utilizing print media, whereas states near the Great Lakes may see an increase in pockets. This quantifies not only why certain attitudes may be held in each of those regions, but also may be helpful for those who are unsure of which news medium to use when launching their own news outlet. My poster will show all of the different demographics tested in graph form and detail what conclusions can be made from each comparison.

Poster 5036

Simulating Sample Size Re-estimation for Crossover Clinical Trials

Mollie McDonald
Department of Computer Science
Kalamazoo College, Kalamazoo, MI

Faculty Advisors
Pam Cutter
Kalamazoo College, Kalamazoo, MI

This project involves the simulation of a new statistical method that was developed as a research project at the Colorado Summer Institute of Biostatistics in 2025. The goal of the project was to simulate a method of sample size re-estimation for crossover clinical trials across several scenarios. Clinical trials are crucial for advancing medical knowledge, and sample size estimation is critical to the success of a trial. However, determining the appropriate sample size depends on uncertain parameters. Therefore, sample size re-estimation, a type of clinical trial adaptation, can be highly beneficial. It is necessary to conduct many simulated trials with virtual observations in order to test the efficacy of the statistical method. The simulation study was written in the R programming language. Overall, the finished code proved adaptable for several different scenarios with varying crossover effects, with the possibility of easily modifying the code based on changes in the statistical method and goals. The simulation study also proved effective in testing the statistical method and generating clear results; the method ultimately performed as expected across all scenarios.

Poster 5037

EthicsBot: Scaling Ethical Support for Computer Science Research

Zijing Zhao, Jessica Chen
Department of Computer Science
Wellesley College, Wellesley, MA

Faculty Advisors
Eni Mustafaraj, Julie Walsh
Wellesley College, Wellesley, MA

We present a user study on EthicsBot, an LLM-powered ethical analysis tool for computer science researchers. Our interviews with 16 computer science researchers explored their perceptions of EthicsBot and assessments of the hypothetical risks and benefits of its widespread deployment. Currently, computer science researchers face complex ethical challenges in a rapidly evolving digital landscape, yet often receive inadequate institutional support for navigating these issues. EthicsBot aims to fill this gap by providing researchers with immediate feedback on the ethical implications of their proposed research.

To select a base Large Language Model (LLM) for our prototype, we experimented with various open-source models locally to guarantee transparency and replicability, ultimately selecting Qwen3-14B. We then engineered a comprehensive baseline prompt through an iterative process. After giving EthicsBot a research protocol, it followed the prompt and returned a structured ethical analysis surfacing potential ethical concerns in research design, data collection and analysis, and potential impact of study.

To prepare for the user study, we built a prototype of EthicsBot on Figma to help users gauge basic functionalities. We conducted 1-hour interviews with 16 computer science researchers to assess their critiques of the structured LLM output, the prototype itself, and the conception of an LLM-powered tool itself. A reflexive thematic analysis of the results reveals researchers' cautious optimism: while researchers valued EthicsBot as a timely and accessible resource for ethical guidance, they expressed concerns about the specificity and reliability of its recommendations. These findings highlight both the promise and limitations of AI-assisted ethics support in research contexts.

Poster 5038

C.A.R.E.: A Governance Framework for Robust and Equitable AI in Oncology Imaging

Brooks Jackson
Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisors
Brian O'Neill
Quinnipiac University, Hamden, CT

The growing integration of deep learning tools in oncology practices has highlighted a significant "validation gap" between these tools' performance in research environments and clinical implementations. The research performed in this thesis address the weaknesses of 3D Convoluted Neural Networks (CNNs) in cancer imaging, specificallywith the detection and classification of lung nodules when subjected to distribution shifts. To address these vulnerabilities, the C.A.R.E. Framework is proposed to operationalize four distinct pillars (Clinical Accountability, Reliability, Equity, and Explainability) and evaluate the clinical and ethical readiness of AI models.

The implementation of this framework was executed using a 3D Squeeze-and-Excitation Residual Network (SE-ResNet50) architecture from MONAI, the leading open-source framework for medical imaging AI. After an initial training phase on the LUNA16 and LIDC-IDRI lung cancer datasets, the model was subjected to a stress test on various, diverse external cohorts such as CPTAC-LUAD (Lung Adenocarcinoma) and VAREPOP-APOLLO (Lung Adenocarcinoma and Squamous Cell Carcinoma). This evaluation revealed significant results, including processability failure rates from clinical data fragility and "shortcut" learning behaviors. Internal attention was also observed to identify spatial artifacts in the bottom-right quadrants in image crops, as opposed to locating biological morphology. Furthermore, an equity audit showed a variance in marked sensitivity across different CT scanner manufacturers, particularly with regards to hardware neutrality deficiencies.

The C.A.R.E. Framework creates a measurable rubric to quantify these failures, shifting from simple aggregate accuracy to a multi-faceted safety audit. By performing supervision on non-standardized inputs and utilizing isotropic resampling strategies during preprocessing, this thesis demonstrates how all parties involved in technical governance can locate and mitigate algorithmic biases. Ultimately, the framework serves as a mechanism to ensure that AI-driven oncology tools are not only high-performing in the laboratory but will be equitable and deployment-ready for complex clinical environments.

Poster 5039

Is Computer Science inclusive? Rethinking Accessibility in CS Education for neurodivergent students

Emily Cruz
Computer Science Department
Smith College, Northampton, MA

Faculty Advisor
Hia Ghosh
Smith College. Northampton, MA

Current institutional structures – such as accommodations, accessibility technology, teaching methods, and curricula – do not adequately address all needs of neurodivergent students. The present study is an examination of current interventions and applications of accessibility services for neurodivergent students in higher education within the Computer Science (CS) department. My research started with a general examination of accessibility in higher education where results pointed out that higher education institutions are not fully equipped to support all kinds of students and accommodations focus on diagnoses rather than addressing contextual and functional needs. Building on this foundation, the research narrowed its focus to neurodivergent students pursuing CS in higher education. Neurodivergent students in CS have challenges with collaboration, executive function, and visual or cognitive barriers to processing code, all of which makes their academic journey more difficult than their neurotypical peers. This informs my research methodology that is a Semi-Systematic Literature Review with a total of 9 articles and journals reviewed from ACM SIGCSE proceedings. Three research areas were identified for classification: Barriers to facilitators of accessibility in higher education, pedagogical strategies to improve accessibility, and Assistive Technologies in Computer Science education. These areas emerged frequently throughout literature from repeated themes, challenges, and solutions. The poster will discuss the findings in each of these areas and address gaps in current research. In addition, this review integrates findings and proposes recommendations for faculty and institutions aiming at ensuring equitable access and retention of neurodivergent students in computing education.

Poster 5040

Semantic Edge Detection for Structurally Meaningful Boundary Extraction

Tristan Wolenski, Ngan Huynh
Computer Science
New Jersey City University, Jersey City, NJ

Faculty Advisor
Moitrayee Chatterjee
New Jersey City University, Jersey City, NJ

Russell Carpenter
NASA Goddard Space Flight Center, Greenbelt, MD

We developed an edge detection framework capable of supporting high-precision boundary extraction for analyzing complex terrain from rover cameras, satellites, orbiters etc.. The goal is to contribute toward more reliable computer vision pipelines for aerospace applications and autonomous systems operating in visually complex environments.

Classical gradient-based edge detection methods, such as Canny and Sobel operators, are widely used in computer vision but frequently fail due to their inherent inability to distinguish structural object boundaries from common artifacts like texture patterns or illumination-induced gradients. Spacecraft imagery also struggles with a plethora of artifacts, leading to producing fragmented and noisy edge maps. This limits the effectiveness of essential downstream tasks such as damage detection and autonomous docking.

We investigated the use of deep learning-based edge detection models like HED and DexiNed, to generate cleaner, semantically meaningful structural edge maps. Unlike traditional gradient-based methods, these architectures leverage hierarchical feature representations in order to suppress texture edges while preserving coherent object contours. We benchmarked classical and deep learning-based approaches on labeled datasets and synthetic CAD rendered spacecraft imagery. Furthermore, we performed structural alignment comparisons against ground-truth CAD wireframes to assess how improved edge representations impact downstream applications like spacecraft pose estimation and robotic navigation in simulated orbits.

Quantitative evaluation using pixel-wise precision showed that classical operators emphasize local gradients and texture, whereas deep learning models suppress background noise and produces semantically cleaner object contours. The results contribute toward more reliable computer vision pipelines for aerospace applications and autonomous systems operating in visually complex environments.

Poster 5041

Deploying a Flexible Remote Sensor Platform for Biological Monitoring

Alanna Town, Buddy Luong, Skye Worster
Department of Computer Science
Smith College, Northampton, MA

Faculty Advisor
Michael Robson
Smith College. Northampton, MA

We present our research on deploying a remote sensor platform for biological monitoring at Smith College. Our goal is to design a configurable framework so that field scientists can easily obtain data from remote sites without technical expertise. The system involves a field-deployed sensor module comprised of a Raspberry Pi microcomputer and an attached sensor. Data is transferred back to the user's host machine via a pipeline which connects over either a wi-fi, cellular, or radio network. On the user's host machine, a Python backend allows for communication with the remote Raspberry Pi. We are in the process of establishing a Java-based graphical user interface (GUI) to make the backend more adaptable and accessible to field scientists. We have implemented this system in two projects. The first is the Nocturnal Flight Call project which utilizes the wi-fi pipeline and a microphone sensor to record nighttime bird migration calls. In this project, the host system and remote sensors communicate directly using Secure Shell (SSH), which easily supplies encrypted and authenticated remote access to the Raspberry Pi. However, network configurations can pose risks. This has been solved in our second project which integrates Tailscale, a virtual private network (VPN) based on the WireGuard Protocol, which creates a mesh network that securely connects all participating devices using encrypted peer-to-peer tunnels. This is the Mothitor project, which is adaptable not only in its implementation of the VPN, but also can utilize any of the proposed systems, connecting over either wi-fi, cellular, or radio networks. This project involves a deployed light sensor which records images and light sensor data to track moth populations. Moving forward, we hope to make our remote sensor platform adaptable and accessible with the hopes of assisting a wide range of biological field scientists in smoothly collecting and accessing data.

Poster 5042

Examining AI's Impact On Novice Web Design Developers

Matthew Jones
Computer Science
Quinnipiac University, Hamden, CT

Faculty Advisors
Brian O'Neill
Quinnipiac University, Hamden, CT

Artificial intelligence (AI) tools have become increasingly integrated into modern web design workflows, allowing anyone to access automated support for ideation, coding, debugging, and layout generation. While these tools may improve efficiency and reduce technical barriers, there is a concern on how these tools could affect the abilities and mentalities of students. In regards to student web designers, concerns remain regarding their potential impact on creativity, cognitive engagement, and sense of authorship. This study has investigated how undergraduate students enrolled in web design courses perceived and experienced AI-assisted technologies within their coursework.

We used a qualitative research approach with semi-structured interviews that were conducted with students who have used and not used AI tools in their design-related assignments. We had asked specific interview questions that explored how students incorporated AI into their workflows, the perceived benefits of AI assistance such as faster iteration and problem-solving, and the potential risks including overreliance, design fixation, and diminished creative ownership. The gathered participants were also asked to show their previous completed design work and describe the extent to which AI influenced their decision-making processes. With these works, we were able to analyze the designs to determine any connection within the pieces.

Interview transcripts have undergone inductive thematic analysis. In doing this, it has identified recurring patterns related to learning outcomes, workflow efficiency, originality, and trust in AI-generated outputs. The findings of this research aim to provide insight into how AI technologies shape the educational experiences of novice web designers and to inform responsible integration of AI tools within web design curricula without undermining foundational creative and technical skill development.

Poster 5043

BportGo: An Indoor and Outdoor Navigation Mobile Application for a University Campus

Joseph Abdo
Computing Sciences
SUNY Brockport, Brockport, NY

Faculty Advisor
Adita Kulkarni
SUNY Brockport, Brockport, NY

Navigating the SUNY Brockport campus is currently challenging for new students, employees, and visitors because the only resource available to them is a static, high-level PDF map provided on the university website that lacks room-level detail. As a result, people new to the campus often struggle to efficiently find their destinations. BportGo was developed to address this issue by providing a comprehensive mobile navigation solution that supports both indoor and outdoor routing for Android and iOS. Built with Flutter, the application uses a single, shared codebase to run seamlessly on both Android and iOS platforms. Outdoor navigation is powered by real-time location tracking combined with the Google Maps API, enabling precise directions across the campus. To improve usability, the location search experience was redesigned using the Google Places Autocomplete API, allowing users to quickly and accurately find buildings and points of interest. For indoor navigation, the application converts existing static floor plans into interactive maps through Mappedin, making it possible to guide users through building interiors. The application ensures a smooth transition between outdoor paths and indoor routes, creating a continuous navigation experience. Beyond core functionality, the project development emphasized performance optimization via search efficiency improvements, enhanced navigation logic, and scalable, maintainable design practices. Significant effort was also dedicated to resolving iOS-specific deployment challenges, including location permission configuration, API key management, and dependency management using CocoaPods. BportGo represents a significant advancement in campus digital infrastructure, providing enhanced mobility, accessibility, and overall navigation efficiency.

Poster 5044

Normal Forms for Regular Expression writing

Oleg Efimov

Math, Computer Science and Statistics Department
St. Lawrence University, Canton, NY

Faculty Advisors

Kevin Angstadt
St. Lawrence University, Canton, NY

We present our research project on Regular Expression (regex) Normalization Rules. To support a more robust set of patterns (i.e., beyond theoretical regular expressions), common CPU-based processing algorithms generally support Perl-compatible regular expressions (PCRE), which are more complex than theoretical regexes, and these engines are often implemented using backtracking algorithms with exponential worst-case runtimes. Something as innocuous as parenthetical grouping of parts of an expression can result in additional computation and data storage overhead. In the worst case, a carefully crafted input to a poorly written regex can stall the engine, creating a denial of service vulnerability called Regular Expression Denial of Service (REDoS). Many Domain-Specific Languages, code translation and synthesis tools, and regex repair tools have been developed to aid in the construction of correct and performant regex, but these tools can still produce regexes with low readability and maintainability.

Poster 5045

Exploring College Students' Needs and Acceptance of AI-Based Stress Management Support Tools: A Pilot Study

Ankita Banik, Abhi Chatterjee
Business Analytics and Data Science
New Jersey City University, Jersey City, NJ

Faculty Advisor
JD Jayaraman
New Jersey City University, Jersey City, NJ

We developed an AI-based stress management support tool and evaluated the usability of this student-informed AI conversational prototype. This tool is intended to supplement existing support services rather than replacing professional care.

College students frequently experience elevated levels of academic, social, and financial stress, which can negatively affect well-being, academic performance, and persistence. Although campus counseling services provide essential support, many students face barriers such as limited availability, time constraints, and stigma. As a result, AI-based stress management tools are increasingly explored as complementary, non-clinical resources that may offer accessible and scalable support for student well-being.

Our study is grounded in data from a survey conducted by our team, with college students. We assessed perceived usefulness, trust, ethical concerns, perceived risks, and willingness to use AI-based stress management tools. Our findings directly informed the design priorities and safety considerations of the AI system. Based on these insights, we fine tuned an existing AI language model to function as a non-clinical stress management support tool. The system is designed to provide emotionally supportive conversational interaction, general stress coping strategies, and guidance toward appropriate campus and external resources. Explicit safety guardrails are incorporated to prevent the system from offering clinical diagnoses, treatment recommendations, crisis intervention, or responses to harmful or inappropriate prompts. The prototype is evaluated through a small-scale pilot usability study with college students. Post-use questionnaires assess usability, perceived usefulness, user comfort, trust, and overall acceptance, with an emphasis on user experience rather than clinical outcomes. Expected results include empirical evidence on student acceptance of AI-based stress management tools, usability insights to inform responsible AI design in higher education, and recommendations for future research.

Poster 5046

Monitoring and Detecting Anomalies during Space Exploration

Natalia Criollo, Xyanna Fuentes
Department of Computer Science
New Jersey City University, Jersey City, NJ

Faculty Advisor
Moitrayee Chatterjee
New Jersey City University, Jersey City, NJ

Russell Carpenter
NASA Goddard Space Flight Center, Greenbelt, MD

This research project develops a model application designed to monitor and detect anomalies within space habitats during space missions. The goal is to enhance the autonomy and reliability of life-support and environmental monitoring systems in deep-space environments.
As space missions travel farther from Earth, communication delays and reduced ground support require spacecraft to operate with greater independence. Space habitats must be able to identify, respond to, and mitigate anomalies without relying on immediate human intervention. Prior research has emphasized the need for fully autonomous monitoring and maintenance systems capable of supporting long-duration space missions. This project contributes to those efforts by focusing on machine learning–driven anomaly detection for habitat safety and crew well-being.

Machine learning models were developed using datasets including the Astronaut and Mission dataset from Kaggle and the Air Quality Prediction dataset from HuggingFace. These datasets supported training and testing for various environmental and life-support parameters such as temperature, humidity, CO_2 and O_2 levels, radiation, pressure, sleep quality, stress levels, and power usage. A Streamlit-based application is being built to simulate and visualize real-time space habitat conditions. The interface displays model predictions and anomaly detection outputs, enabling users to explore system behavior under different environmental and operational scenarios.

The resulting system provides an interpretable and dynamic platform for monitoring spacecraft environments and detecting anomalies. By modeling and visualizing behavior across key habitat metrics, the application supports the development of self-sufficient systems for deep-space missions, complementing existing efforts to increase spacecraft autonomy and improve mission safety.

Poster 5047

Mouse Model of Disease - Language, Metadata, and Cross-Species Comparison

Yuzhang Fu, Ayesha Sheikh
Department of Computer Science
Smith College, Northampton, MA

Faculty Advisor
Halie Rando
Smith College, Northampton, MA

Mice are widely used to model human diseases due to their genetic similarities, yet few pharmaceutical developments are successfully translated across species. This project develops a cross-species semantic alignment framework to examine how experimental metadata describe disease models and to support systematic comparison between human and mouse studies. Our broader motivation is to understand whether semantic representations can provide a foundation for evaluating the biological similarity.

We first introduced MONDO, a unified disease ontology that integrates multiple disease resources into a global hierarchy of disease concepts. Human and mouse GSE (GEO Series) metadata were analyzed separately using txt2onto 2.0, which maps controlled vocabularies onto free text and outputs prior probability scores for disease models. Applying a threshold greater than 1 allowed us to retain predictions with high confidence for cancer-related MONDO terms and assign consistent disease labels to individual studies.

To investigate ontology behavior, we examined prediction patterns across six hierarchical levels within the cancer branch of MONDO. Although txt2onto does not encode hierarchical structure explicitly, most studies showed stronger predictions at child nodes and weaker scores at parent terms. This pattern indicates that disease specificity present in experimental metadata aligns with the internal structure of the ontology.

We then analyzed human–mouse GSE pairs sharing the same predicted MONDO term. After preprocessing and lemmatization, we applied TF-IDF and generated sentence embeddings to compare studies. Cosine similarity scores were consistently higher for matched human–mouse pairs than for random pairs, suggesting structured alignment at the textual level.

To further characterize biomedical entities, we used SciSpaCy (en_ner_bionlp13cg_md) with its AbbreviationDetector and UMLS EntityLinker to normalize differently expressed phrases through shared CUIs, grouping related terms such as gene expression under unified concepts. Ongoing work maps these normalized entities back to the corpus and applies TF-IDF to enable refined cross-species comparison.

Poster 5048

Development of a Course Planning Tool and Syllabus Generator

Audrey Windrow, Cam Senna, Sencere Rabel, Jessica Day
Computer Science
Eastern Connecticut State University Willimantic, CT

Faculty Advisor
Garrett M. Dancik, Lora Lee, Courtney Broscious
Eastern Connecticut State University Willimantic, CT

We present the design and implementation of a Course Planning Tool developed for faculty at Eastern Connecticut State University. The goal of this project is to automate syllabus generation and to help assist professors when designing a course. Effective syllabus planning is important in acadamia, as it plays a critical role in ensuring student success within the classroom.

Our contributions to the project were as follows. We first developed a Figma prototype of the tool, and created graphics for the site. We then developed the front-end in React, JavaScript, and TypeScript, and a JavaScript Object Notation (JSON) rendering engine in order to dynamically generate pages from JSON files. This architecture promotes modular feature extension while mainting scalability and long-term maintainabilty for future development. This will ultimately make it easier for individuals without a CS background to develop their own pages so the tool can be adopted at other institutions. We developed the backend using Python/Flask and use Google Firestore to store course information. An API call to the backend uses the docx Python module along with a template syllabus in order to generate the syllabus for the current course. Our presentation will demonstrate the features that we have implemented, as well as highlight key challenges in developing this tool, including handling the interaction between the frontend and backend, switching between the local and the production version to test new changes to the codebase, managing the login page during development, handling merge conflicts with Github.

A preliminary version of the tool is available from: https://courseplanningtool2.pages.dev/

Poster 5049

Gamifying the Campus Experience: A 3D Interactive Campus Exploration

Joshua Swanson, Sophia Lin
Computer Science
St. Mary’s College of Maryland, St. Mary’s City, MD

Faculty Advisor
Simon Read
St. Mary’s College of Maryland, St. Mary’s City, MD

We present our experience completing the St. Mary's Project (SMP), a capstone requirement at St. Mary's College of Maryland that encourages independent research and creative work. As virtual tours become essential for prospective students unable to visit in person, existing solutions often lack interactivity. Our project uses Unity to address this gap by creating an immersive digital twin of the campus for use at recruitment events, aiming to gamify exploration and promote user accessibility. This project's multidisciplinary aspects include not only computer science but also art and media production, allowing us to integrate digital art, voice-acting, soundtrack, and 3D environments, contributing to a cohesive and immersive user experience. Due to the intensive 3D modeling requirements, we followed a Waterfall development framework to prioritize the production of complex assets in Blender. The campus was reconstructed using Google satellite imagery and photographic references to ensure authenticity. Key technical challenges included managing high-volume repository constraints; we implemented Git LFS via Git Bash to effectively handle large binary assets. By parallelizing team roles, we simultaneously developed intuitive UI systems and voice-acted dialogue to enhance immersion and accessibility. Our results include a playable demo that serves as an innovative contribution to the college’s outreach efforts. We will discuss the benefits of a multidisciplinary approach, the trade-offs of our methodology, and the technical hurdles of translating a physical campus into a digital environment.

Poster 5050

What Do You Mean? Detecting LLM-Generated Texts using Meaning Representations

Jiyuan Ji, Thu Hoang, Brendan Shen
Department of Computer Science
Amherst College, Amherst, MA

Faculty Advisor
Shira Wein
Amherst College, Amherst, MA

Large language models (LLMs) have exhibited remarkable ability in generating texts resembling human authorship, posing challenges in distinguishing LLM outputs from human-written content and raising concerns about misinformation and linguistic diversity. Existing detection methods often rely on surface-level features (Muñoz-Ortiz et al., 2024), which can be ambiguous or prone to overfitting. As LLMs advance in producing human-like text, syntactic cues may no longer reliably indicate LLM generation. Park et al. (2025) explore comparing Abstract Meaning Representation (AMR; Banarescu et al., 2013) graphs between human- and LLM-generated rewrites to measure semantic gaps, achieving 96.5% accuracy and surpassing many state-of-the-art models. Therefore, we ask: Does incorporating semantic information help identify LLM-generated texts? If so, to what extent?

Using two meaning representations, AMR and Discourse Representation Structure (DRS), we test three approaches on short LLM-generated news texts from Muñoz-Ortiz et al. (2024). (1) Baseline: We run DetectGPT using the same setup as Mitchell et al. (2023). (2) Rephrase & Detect: We rephrase texts twice using LLMs, parse them into AMR and DRS graphs, compute semantic gaps, and train a decision tree on these features. (3) AMR Motifs: We extract statistically significant 3-node subgraphs from canonicalized document-level AMR graphs, adapt TF-IDF to score motifs, and finetune a Longformer model on these representations.

The baseline achieves a best macro F1 of 0.4506 (average 0.4206). Rephrase & Detect improves performance, reaching a best test macro F1 of 0.6952 (average 0.6528). AMR motifs achieve a best F1 of 0.8285, though we attribute much of this gain to the Longformer architecture rather than motifs alone. Overall, semantic information aids detection, particularly with simple, interpretable models. Qualitative analysis shows that LLM-generated texts frequently include a pronoun referring to a person performing an action, while human-generated texts more often feature adjectives modifying nouns, corroborating findings in Muñoz-Ortiz et al. (2024).

Poster 5051

BlueSweep: A Mobile App for Coastal Cleanup

Elisabeth Jacob
Computer Science
Rivier University, Nashua, NH

Faculty Advisor
Mihaela Malita
Rivier University, Nashua, NH

As a regular volunteer for beach cleanups, I am pleased to share BlueSweep which is a mobile application used to visualize data for waste collected along New Hampshire state beaches. This application was developed in Android Studio using Kotlin and Jetpack Compose. With the simple click of a button, volunteers can initiate a new cleanup event to record and upload waste items collected in real-time. The overall goal of BlueSweep is to provide a digital alternative to the traditional pencil and paper collection methods used which are often tedious, time-consuming, and subject to calculation errors.

BlueSweep was designed to allow data to be collected offline without the need for cloud-based services or login systems, as many waterways in New Hampshire are remote. BlueSweep uses a single centralized CSV file, allowing data to be stored locally with the ability to append new entries and aggregate results for visualization. Following cleanups, saved data is utilized to generate visual dashboards with charts and graphs of waste collected. This data can then be exported and shared with conservation organizations to support their mission of protecting waterways while educating the general public. BlueSweep provides a simple, touch-friendly interface which allows users to easily switch between screens to initiate new events, access summaries for past events, and recognize trends across all cleanups. The data that was collected using the BlueSweep app focuses on New Hampshire state beaches both along the seacoast and inland, recognizing that all bodies of water are interconnected and that pollution in one area directly impacts the entire ecosystem.

Poster 5052

Systems-Level Profiling of Bayesian Optimization Frameworks on HPC Infrastructure

Meron Oumer, Nargiz Akhmetova
Department of Computer Science
Smith College, Northampton, MA

Faculty Advisor
Michael Robson
Smith College, Northampton, MA

Hyperparameter optimization is a major computational bottleneck in modern machine learning workflows, often requiring hours to days of distributed computation. While Bayesian Optimization (BO) is designed to reduce the number of evaluations needed to find good configurations, the systems-level performance characteristics of existing BO frameworks on high-performance computing (HPC) infrastructure remain poorly understood.

We present a comprehensive profiling study of four widely used optimization frameworks: BoTorch, DeepHyper, Optuna, and Ray Tune, alongside baseline methods including random search and grid search. Experiments were conducted across synthetic benchmark functions (Branin, Hartmann, and Rastrigin) as well as real-world machine learning tasks such as neural network and XGBoost hyperparameter tuning. We instrumented each framework to measure per-iteration wall-clock time, Gaussian process (GP) fitting latency, acquisition function optimization cost, and model evaluation time on an HPC cluster.

Our results show that model evaluation consistently dominates total runtime for real-world workloads, confirming that training cost is the primary practical bottleneck. While surrogate modeling and acquisition optimization introduce measurable overhead, their impact remains secondary for expensive tasks. In higher-dimensional settings, GP fitting latency becomes increasingly noticeable in GP-based implementations. Distributed systems such as DeepHyper reduce wall-clock time through parallel worker pools, though they introduce communication and orchestration overhead.

These findings establish quantitative performance baselines and highlight opportunities for improvement at both the algorithmic and systems levels. Our ongoing work explores the design of a high-performance optimization framework aimed at reducing end-to-end wall-clock time while preserving the sample efficiency of Bayesian optimization.

Poster 5053

A Comparative Performance Evaluation of Modern Full-Stack Architectures for Gamified Web Applications

Luis Vazquez
Computer Science
Eastern Connecticut State University, Willimantic, CT

Faculty Advisor
Sarah Tasneem
Eastern Connecticut State University, Willimantic, CT

The rapid expansion of digital productivity platforms has accelerated the integration of gamification mechanisms to enhance user engagement and retention. While gamification has been widely studied from a design perspective, limited empirical research examines how different full-stack architectures affect the backend performance of gamified web applications. This study addresses this gap by comparing two functionally equivalent applications built with distinct modern technology stacks and evaluating their backend performance. Two functionally equivalent web applications were developed. The first application was constructed using a *React.JS* frontend, with a *Node.JS* backend, and a *MySQL* database hosted on Amazon Web Services' *Relational Database Service*. The second application was developed using an Angular frontend with a Python FastAPI backend, and a SQLite database. Both implementations incorporated identical gamification features such as progress tracking, and reward mechanisms, and were designed to support structured storage and management of user data. Backend performance was evaluated using Apache JMeter under controlled simulated user loads. Metrics included request completion times, request throughput, and request size. Results indicate that the FastAPI-based implementation achieved faster average request completion times compared to the Node.js implementation within the tested environment. The comparative analysis identified architectural tradeoffs affecting scalability, resource efficiency, and system responsiveness. These findings provide empirical evidence on how backend framework and database selection influence the performance characteristics of gamified web applications. By delivering a controlled, side-by-side evaluation of two modern full-stack architectures, this study offers practical guidance for developers selecting technology stacks for performance-sensitive productivity platforms.

Poster 5054

Implementing an AI-Assisted Feedback Pilot for Structured Student Learning

Hoang Huy, Yubraj Khatri
Computing Sciences
SUNY Brockport, Brockport, NY
Faculty Advisor
Ning Yu
SUNY Brockport, Brockport, NY

Giving precise and timely feedback in higher education is becoming increasingly challenging due to the increasing size of the classes and the grading workloads. Although there are many available tools for generative AI feedback, they are mostly independent of the grading rubrics and lack instructor control and answer leakage concerns. The project aims to bridge this gap by integrating AI feedback with the coursework.

We designed and deployed a multi-course configurable AI-assisted feedback platform used across three semesters at SUNY Brockport. The system has been implemented in eight courses spanning Education, Biology, Computer Information Systems, and Computer Science, supporting over 500 students, 600+ test submissions, and 10,881 student answers.

The architecture consists of a FastAPI backend and Next.js frontend, replacing an earlier Streamlit prototype to improve scalability. Course materials (46 large documents across 38 modules) are processed using LlamaExtract and structured chunking techniques. Rubrics and content are embedded using text-embedding-3-large and stored in Supabase (pgvector). A dynamic Retrieval-Augmented Generation (RAG) pipeline retrieves module-specific content and rubric constraints to generate feedback. Each module maintains isolated embeddings and rubric logic, enabling multi-tenant, instructor-controlled customization.

Unlike generic AI tools, the system is embedded directly within assignments and prevents answer leakage. Instead of revealing solutions, it references relevant slide or document sections automatically, guiding students to revisit source material. Instructors control feedback tone, strictness, and revision limits.

Performance testing shows average feedback generation times of 3–7 seconds per question using parallel workers. The system includes a 45-second hard timeout and automatic job recovery to handle failures, ensuring reliability at scale. Across deployments, instructors reported a significant reduction in repetitive grading workload, while most students indicated improved study reinforcement and more frequent revision cycles. These results demonstrate that AI can function as a dependable, instructor-governed component of structured learning systems in higher education.

Poster 5055

Smart Energy Load Balancer: A Scalable Campus Energy Management Prototype using Machine Learning and parallel-optimized Scheduling

Alex Gomez, Kyle Hartness, Connor Pirkl
Computer Science Department
Utica University, Utica, NY

Faculty Advisor
Rutal Mahajan
Utica University, Utica, NY

University campuses face significant financial and environmental costs due to midday energy spikes, driven by the convergence of peak outdoor temperatures and high-occupancy classroom clusters. This project presents a Smart Energy Load Balancer prototype designed to flatten these demand peaks by integrating predictive machine learning with a high-performance, parallelized scheduling engine. Our methodology combines real-time environmental data via Weather APIs with synthetic campus activity models based on Utica University facilities. We trained an ensemble machine learning (ML) pipeline to estimate hourly usage based on factors including humidity, cloud cover, and school activity. Using model explainability analysis (SHAP), we identified ambient temperature as the primary driver of HVAC load. This predictive model generates the energy-cost curve used to guide optimization decisions.

To solve the complex task of scheduling classes against this curve, we developed a scheduling engine in C++ that balances room capacity with projected thermal demand. To ensure scalability across multiple buildings and thousands of potential permutations, energy calculations are parallelized using OpenMP. This significantly reduces runtime, enabling rapid schedule evaluations that would be computationally prohibitive in standard serial processing. For administrative practicality, we developed a Java-based GUI that allows users to visualize recommended schedules and energy impacts. Results demonstrate a significant reduction in midday spikes and a smoother daily energy profile compared to traditional schedules. While currently a technical prototype, this decision-support tool provides a scalable foundation for energy-aware campus management, with future iterations planned to incorporate soft constraints like faculty availability and student transit times.

Poster 5056

The Generative Generation? A Qualitative Investigation into Undergraduate AI Tool Use and Perspectives

Lance Cheng
College of Information and Computer Sciences
University of Amherst, Amherst, MA

Faculty Advisor
Ella Tuson
University of Amherst, Amherst, MA

Over a short few years, large language models have become ubiquitous academic partners-in-crime, annoyances, saving graces, fears, and everything in between. However, what everyday people think of these programs is frequently obfuscated behind extreme opinions, which often come with social pressure. In real life, professors may introduce no AI blanket terms for their classes that students cannot challenge or express their feelings on due to the power dynamic; students also may feel awkward opining around their classmates due to general peer pressure. Online, tech startup culture is extremely prevalent on career-focused websites such as LinkedIn, and has further spread to “mainstream,” general-purpose social media sites, especially Instagram and Twitter; influencers in this sphere depict themselves using AI for all kinds of everyday tasks, and are often themselves developing AI tools. Due to these social dynamics, which often result in mutual mistrust, there is a severe need for clarity on young people's true thoughts and motivations to use AI.

This study seeks to characterize UMass CICS undergraduates' use of LLMS and other generative AI tools, in particular their motivations for using or not using them, with an eye towards the role of complex academic social dynamics. Participants are sourced from UMass Computer Science and Informatics majors via social media outreach, physical flyers, and snowball sampling, first completing a screening survey, then an interview with questions targeted towards their experiences using AI in the contexts of different tasks and social groups. The researchers hope this study will provide recommendations to CICS staff on how to move forward crafting an educational experience with AI awareness and a student-first lens.

Poster 5058

Discrepancies in the Geotargeting of Spanish vs English Google Search Results During the 2024 US Elections

Nina Howley, Dianna Gonzalez
Computer Science
Wellesley College, Wellesley, MA

Faculty Advisor
Eni Mustafaraj
Wellesley College, Wellesley, MA

We present our experience investigating the difference between Google search results for election related queries in English vs. Spanish during the 2024 US election. In the context of this research, geotargeting can be defined as the accuracy to which a domain appears for the location of a search. It is important that results for election related queries are correctly geotargeted in order to provide relevant information and keep voting accessible. This is especially the case for locational queries such as “where to vote near me”. Our research questions were proposed to investigate the difference in the appearance and geotargeting of government domains between English and Spanish. We follow the Voting Rights Act section 203, which states that any jurisdiction where 5% or more of the population speaks a minority language other than English must supply official voting materials in that language. Studies show that people rely on search engines to access election and voting related information, so a search audit is an appropriate tool for investigating the access that these minority language voters have to official sources. We conducted a Google search audit using locations covered by this section along with English election related queries presented in related literature, and their Spanish translation. Our dataset for scraping consisted of 223 jurisdictions and 266 queries (133 in each language). From this we gathered 1,106,096 unique search results. Next, we utilized two open-source government domain databases to identify results from official sources and compare across languages. Through this audit, we conclude that despite the information existing for the chosen locations, Google’s search algorithm often mistargets or doesn’t surface them at all. Our poster will discuss the specifics of data collection, key findings including how we interpret the data, and what’s next for future work.

Poster 5060

Microsatellites Made Accessible: An Open-Source Tool for Microsatellite Data Analysis

Nargiz Akhmetova, Ruize Qin
Department of Computer Science
Smith College, Northampton, MA

Faculty Advisor
Halie Rando
Smith College, Northampton, MA

We present a Python pipeline that decodes binary ABIF files from Applied Biosystems capillary sequencers and exports raw fluorescence traces, size standards, and run metadata into a stable CSV format for downstream quantitative analysis and machine-learning–based allele calling.

ABIF files encode light from fluorescent dyes attached to DNA fragments, allowing a single run to measure fragment sizes at many microsatellite loci and making them central to wildlife and conservation genetics. In current practice they are processed in proprietary software such as GeneMapper, followed by time-consuming manual inspection of peaks. These closed, platform-specific workflows are expensive, hard to automate, and limit reproducibility, while existing open-source tools only partially decode or visualize the data. Our aim is to provide an open-source that exposes the full electropherogram and metadata for integration with standard data-science and machine-learning workflows. Methodologically, we treat each ABIF file as a structured binary stream with a header, directory entries, and data blocks. The program parses the header, walks the directory to locate data, and decodes blocks into numeric and character values, retaining intensity traces, size-standard channels, and technical metadata such as batch identifiers. The extraction logic is modular so that the same code can run across many files in a project, and the decoded structures are assembled into data frames, written to CSV, and optionally plotted to visualize peak shapes and size-standard alignment.

The current implementation successfully decodes representative ABIF files and exports machine-readable tables suitable for automated analysis. We are now using these outputs to engineer features from individual peaks and to train and evaluate machine-learning classifiers that predict alleles directly from raw fluorescence patterns. By replacing proprietary, manual allele calling with an open, automated pipeline, this work has the potential to reduce cost and human effort and to improve reproducibility in microsatellite-based conservation genetics studies.

Poster 5061

VulneraX: Drift-Aware Time-Series Forecasting for Software Vulnerability Risks

Patrick Penna, Vonnie Magnussen
Computer Science
Utica University, Utica, NY

Faculty Advisor
Unnati Shah
Utica University, Utica, NY

Pavan Manukonda
DataEQ Consulting, PA

Cybersecurity encompasses many necessities for companies to invest in to ensure their information and services are compliant with confidentiality, integrity, and availability. Notably, all organizations must understand the relevant risks to their interests and assess them based on their impact. Cybersecurity efforts must be willing not only to perform these risk assessments on valuable infrastructure but also make decisions based on the volatility in a certain vulnerability's prominence over time. This study presents a framework for how risk forecasting can strengthen risk assessment efforts. The solution, VulneraX, is a time-series forecasting model that predicts trends in high-severity vulnerabilities (CVSS $\geq$ 7.0) and formulates realistic predictions of those vulnerabilities based on inbound scanner outputs and the National Vulnerability Database (NVD). The solution was developed with Prophet and, through extensive research into forecasting methodologies and the inclusion of relevant key metrics, it serves to preprocess data, mitigate outliers, generate forecasts, detect drift, visualize trends and metrics, and issue alerts when predictive reliability is compromised. Moreover, this study builds on existing implementations and ideologies by introducing the concept of a “preliminary evaluation” during its preprocessing phase. Given a dataset, CVSS scores in two-day windows are juxtaposed with one another to calculate the standard deviation between the two; should that deviation exceed a predetermined threshold (<=50), the abnormal outlier will be ignored for the model’s training and placed into a CSV file containing relevant fields (date, vulnerability, vulnerability count, percentage from deviation, and the rate of change in between) for anomalies to be assessed separately. The goal of this solution is to build on existing, standardized drift-detection safeguards by reinforcing data preprocessing to enhance cybersecurity efforts and enable the model to produce a generalized, accurate, and stable forecast. Through this, VulneraX enables organizations to turn from reactive vulnerability remediation to proactive, data-driven risk management.

Poster 5062

Experience Report: Software Development Teams with Human-AI Collaboration

Adrianna Frazier
Computer Science Department
Worcester State University, Worcester, MA

Faculty Advisor
Karl Wurst
Worcester State University, Worcester, MA

Stoney Jackson
Western New England University, Springfield, MA

Kits are educational tools built to provide an authentic environment to learn Computer Science concepts. Kits consist of an open-source project with student activities and instructor guides. The MicroservicesKit allows students to work with and understand modern complex program integration and containerization while the GitKit provides students with experience using Git, the dominant version control system. The Kits project is currently refactoring, improving, and generalizing the tool used to build and deploy the kits, developing a top-level design that will decompose GitKits' monolithic architecture. The goal is an architecture where each task is a single, atomic-function building block that together form a larger cohesive build. With this reconstruction, the aim is to have the system move from working only with GitHub to one that also has GitLab integration. We are developing some of the ~80 components needed. By using Agentic AI to help develop our code, it has helped us develop the individual components, create a plan, implement changes, and test to make sure everything is integrated correctly and works as desired. In conversation with the AI tool we can brainstorm, design, plan implementation, write and run tests, and implement the code. We can do this in a fraction of the time, with higher quality, than it would take a human team member. As we continue to use the AI tool, we are learning how to make it work better in this project and in future projects.

Poster 5063

Teaching Advanced Computing Through Experiential Labs

Gabrielle Addotey, Luther Barreiro Roxo, Juidane Thomas
Golisano College of Computing and IS/Kate Gleason College of Engineering
Rochester Institute of Technology, Rochester, NY

Faculty Advisor
Daniel Krutz
Rochester Institute of Technology, Rochester, NY

When the conversation of making accessible and inclusive software presents itself, many developers and students find themselves awfully quiet. This typically stems from insufficient exposure to related topics early in education such as accessibility, artificial intelligence, ethics in computing, machine learning, and even cybersecurity. The solution to this problem? The Accessible Learning Labs (ALL) within the Golisano College of Computing and Information Sciences at the Rochester Institute of Technology seeks to address this gap through open-access experiential labs. In providing an interactive learning experience, our labs are intentionally designed to educate users on the challenges and impacts lack of accessibility and exposure can have within computing.

ALL has created labs on these crucial topics, each with four sections: a conceptual reading section, an interactive simulation, an assessment quiz, and a reinforcement section containing supplementary instructional videos.

This poster will provide an overview of two key labs: Lab 7, AI Cybersecurity, and Lab 14, Quantum Cryptography. Lab 7 seeks to educate users on the fundamentals of machine learning under cybersecurity. This was achieved through introducing key terms like unsupervised and supervised learning, autonomous systems, phishing and related cyber attacks. The experimental lab carries students through an autonomous system's decision on restricting file access based on certain criteria. Through the activity, users learn the ways in which AI/automated systems can be flawed and how they can rectify this. In Lab 14, users learn the core concepts of quantum computing, progressively building their knowledge through an informative simulation. The interactive component provides users with the opportunity to solve ciphers and visualizes how a quantum and classic computer differ in decryption methods. Instead of relying solely on static, written instructional material, the lab allows students to manipulate encryption parameters, observe computational scaling behavior, and compare classical and quantum encryption and decryption performance.

CCSCNE 2026 Host

SMITH COLLEGE

Student Poster Abstract Booklets Courtesy of