




DDMLab Year in Review
2025 - 2026

CONTENTS


- 4-5 [ANNUAL SUMMARY FROM COTY GONZALEZ](#)
- 6-7 [CHANGING MEMBERS](#)
- 8-9 [CELEBRATING 25 YEARS](#)
- 10-17 [RESEARCH UPDATES FROM DDMLAB MEMBERS](#)
- 18-19 [RESEARCH HIGHLIGHT: C.O.R.A.](#)
- 20-21 [RESEARCH HIGHLIGHT: PHISHING TRAINING](#)
- 22 [PUBLICATIONS AND CONFERENCE PAPERS](#)
- 23 [RESEARCH HIGHLIGHT: AI-SDM](#)
- 24 [KEEPING IN TOUCH ALL YEAR!](#)

18-19

[Making a Difference: C.O.R.A](#)

Learn about C.O.R.A., the Collaborative Operations & Resource Management with Agentic AI game. This work, done in collaboration with the AI-SDM and American Red Cross, is designed to address the critical challenges of disaster management training, focusing on the intricate trade-offs inherent in distributing limited resources.





LETTER FROM THE DIRECTOR

Dear Friends and Collaborators:

Another year has gone by, quickly!! This has been a year of resilience, growth, and new adventures! In this newsletter, we are excited to share some of the activities, accomplishments, challenges, and adventures of the Dynamic Decision-Making Laboratory (DDMLab) during 2025–2026. I hope you enjoy reading about it!

This year is also a special milestone: 2026 marks 25 years of the DDMLab! Since the lab began in 2001, it has been shaped by an extraordinary community of students, researchers, collaborators, and friends. Throughout this newsletter, you will find a few pictures and reflections celebrating that journey. To all the people who have made these 25 years possible: Thank you!!

Over the last year, the research environment in the United States has changed dramatically, particularly for the behavioral and social sciences. One of the most difficult developments for our lab was the early termination of the IARPA ReSCIND program in September 2025, after we had been working on the project for only about ten months. This was especially disappointing because we had assembled a terrific team and were pursuing a fascinating and unusual idea: using our understanding of human cognition (cognitive biases, errors, and limitations) to develop new approaches for defending computer systems against human cyber attackers. Although the program ended much earlier than we expected, the ideas and science did not. Our collaboration with researchers at the University of Texas at El Paso, the University of Washington, and Peraton Laboratories generated important results, publications, data, and new questions that we continue to explore today. If there is one lesson from this experience, it is that good ideas can continue well beyond the lifetime of a particular program. And there is much brighter news!

Our lab continues to advance several exciting projects. With support from the National Science Foundation, we are collaborating with colleagues in Civil Engineering to investigate how cognitive models can help us understand and enhance social capital in physical spaces. Sensors have now been deployed in a coffee shop in CMU's Hunt Library, giving us an exciting opportunity to observe patterns of social interaction in a real community environment. Even better, one of our talented students has developed a dashboard that allows us to monitor these patterns in real time and eventually explore interventions. I am particularly excited about this project because it connects cognitive science, technology, and the possibility of making a tangible difference in our own community.

Through our MURI CATCH project, we are also continuing our work on phishing, cybersecurity, and Human-AI teaming. This year, we have focused particularly on the cognitive biases that are deliberately exploited in phishing emails and on a fascinating question: Can humans and AI agents recognize when those biases are being used against them? This work is allowing us to explore new ways of connecting humans, Cognitive AI agents, machine-learning systems, and large language models.

Much of our energy continues to go toward the NSF AI Institute for Societal Decision Making (AI-SDM), which is now entering its fourth year. Hard to believe!! The institute continues to thrive, bringing together an extraordinary interdisciplinary community working to understand how AI can help people make better decisions in consequential settings. You will find many AI-SDM highlights throughout this newsletter. One highlight I am especially looking forward to is our upcoming third workshop on Human-AI Complementarity for Decision Making. This year's theme is Dynamic Alignment: how humans and AI systems can continuously adapt to one another as environments, goals, knowledge, and capabilities change. I think this question will become increasingly important as AI becomes more deeply embedded in our lives and our decisions.

This year also brought a very meaningful personal academic milestone. After exactly 20 years on the research faculty track at CMU, during which my research program was supported through external funding, I became a tenured Full Professor in January 2026. The department, college, and university voted unanimously in support of the appointment. This is meaningful for me, following many years of building the DDMLab, collaborating with wonderful students and colleagues, and pursuing sometimes unconventional research questions about human decision making. It also came after being named an AAAS Fellow, becoming the 15th AAAS Fellow in the history of CMU's Dietrich College of Humanities and Social Sciences. I am enormously grateful to the many students, collaborators, friends, and family members who have been part of this journey.



As we begin the 2026–2027 academic year, I look forward with hope and curiosity. I hope for renewed strength in U.S. science and research, for new discoveries and collaborations, and for thoughtful reflection about the extraordinary advances in AI and how they are changing our world. I also hope for something equally important: time for family, friends, cycling, painting, travel, and the people who make all of this worthwhile.

And to those friends and colleagues out there whom I have not seen or talked with in far too long: reach out! I would be delighted to hear from you. Let us continue bringing our ideas, energy, and optimism together to make a meaningful impact, and, in whatever ways we can, to make this a better world.



A special professional experience this year was returning to Texas Tech University as a Rawls Distinguished Scholar Speaker, where I gave a presentation about my career journey and current research. It was especially heartwarming to walk around campus again and to reconnect with my former advisor, George Kasper, who is now retired. George made the trip all the way to Lubbock just to attend this special event, which meant a great deal to me. Thank you, George!



January 2026 took me to Singapore, where I co-taught a tutorial at the AAAI Conference with Aarti Singh. In June, I traveled to Montreal, where I had the honor of giving a keynote address at the International Conference on Cognitive Modeling. Preparing an important talk always pushes me to think more deeply about what I am doing, why I am doing it, and where science should go next. For that reason alone, these experiences were tremendously rewarding.



Science also gave me a good excuse to return to Rio de Janeiro for the Cognitive Science Conference. The conference was excellent, and Rio is still Rio: beautiful, energetic, crowded, surprising, and a lot of fun! It had changed considerably since my previous visit, which reminded me that places, like people and research programs, never stop evolving.



Outside of work, cycling again took me somewhere new. This year I traveled to Croatia, where I cycled roughly 400 miles. It was my first visit, and Croatia exceeded all my expectations: the cycling, landscapes, people, food, and history were wonderful. I also visited Sardinia for the first time, where wine, water, and family made for another memorable experience.



SAYING HELLO To!



Yuchen Li graduated CMU with a Master's degree in Electrical and Computer Engineering. Her work focuses on the American Red Cross (ARC) project, where she develops AI-powered, game-based disaster response simulations to support training and decision-making in dynamic environments.



Kibo Schaffer is an independent cybersecurity researcher.

His interests fall in the the intersection of feedback structures, organizational design patterns, and cognitive limits that lead to security vulnerabilities, from which he derives heuristics for what parts of a system are most likely to harbor bugs.



Jet Jacobsen is a recent graduate from Duquesne University.

As a new research assistant he is working on exploring optimal stopping task models and how these models can help us understand how humans learn from experience. His work aims to shed light on optimal stopping to advance the field of behavioral science.


SAYING GOODBYE TO!



Don Morrison retired from CMU in 2026 after over 15 years supporting the DDMLab as well as the Psychology Department's FMS Group. In his retirement, Don will continue his hobby of change ringing, an obscure art form that arose in Renaissance England and combines music, sport, and group theory.

Maria Ferreira and **Anu Aggarwal** continue their work on decision making and online security.



Marko Morrison is continuing his PhD in the societal computing program.




Research assistants **Arya Bhawe**, and **Raghav Sharma** all made valuable contributions during their time with our lab.

The DDMLab turns 25!





Over the past 25 years, the DDMLab has served as a vibrant home to dozens of post-doctoral fellows, graduate students, research assistants, and staff members. Our alumni have made major contributions to dynamic decision-making research and are leaders in academia, industry, public policy, and defense worldwide. We are immensely grateful for the dedication, innovation, and diverse perspectives each member has brought to our lab family throughout our history.



RESEARCH UPDATES BY DDMLAB MEMBERS

From Grace Roessling

I have recently completed the first year of my postdoctoral fellowship in the DDMLab. Throughout the year, my research has focused on human-AI complementarity—



how we can design AI systems that coordinate effectively with humans to achieve performance neither could accomplish alone. Across different domains, including cooperative decision making, cybersecurity, and phishing detection, my work has focused on identifying the mechanisms that facilitate robust human-AI teams.

In collaboration with Chase McDonald, I have continued his work on human-AI collaboration in the Overcooked platform. Our research has investigated two research questions: whether increasing temporal alignment by slowing self-play reinforcement learning agents to human-like speeds improves human-AI complementarity, and whether Chase's Behavior Shaping (BeSH) algorithm better supports collaboration than conventional self-play agents. Over the past year, I deployed experiments, developed analysis pipelines, and led manuscript preparation. Our work on temporal alignment has recently been submitted to HCOMP, and we are now preparing a journal manuscript examining how BeSH, an AI policy that adapts to human preferences, improves human-AI complementarity.

Working with Tailia Malloy and Aarya Bhawe, I have helped develop a novel experimental platform investigating whether humans and large language models exhibit similar cognitive biases during phishing attacks. We first designed an annotation platform to identify cognitive bias triggers across a corpus of legitimate and phishing emails before constructing an experiment that directly compares the susceptibility of humans and LLMs to these

biases. The project is currently in the data collection phase, after which we will examine both decision accuracy and the cognitive biases that drive responses in both humans and LLMs.

In collaboration with Yinuo Du, I designed a cybersecurity experiment using the CHART framework to investigate how role interdependencies influence human-AI team performance during intrusion detection. The experiment examines how manipulations of authority and information resolution affect complementarity within mixed triad teams consisting of a human analyst, a heuristic intrusion detection model, and a large language model. Over the past year, I developed the experimental protocol, task design, and theoretical framework, as well as an IDS Event generation program. We expect to begin deploying the experiment in the coming months.

Finally, I joined an NSF-supported collaboration with researchers in Civil Engineering to study what features of a given community space facilitates social capital (i.e. the distribution of resources within a social network). We are deploying LiDAR sensors within a coffee shop in Hunt Library to collect naturalistic human trajectories. This data will support the development of a hybrid cognitive model that combines Instance-Based Learning (IBL) to represent high-level decision making with a behavioral dynamics model governing continuous movement, bridging the tasks of decision-making and locomotion.

Looking ahead to my second year, my primary goals are to publish the ongoing Overcooked manuscripts, deploy and analyze both cybersecurity experiments, and begin modeling data from the LiDAR project. More broadly, I aim to continue developing general principles for designing AI systems that complement human cognition by adapting to human behavior, supporting effective coordination, and improving team performance across a wide range of complex decision-making environments.

From Xiaohong Cai

This was my first year as a postdoctoral researcher in the Dynamic Decision Making Lab. My research has focused on understanding how people respond when AI makes systematic mistakes and how trust in AI develops through repeated interaction.



In one project, we examined how people respond to AI recommendations that consistently overestimate or underestimate resource needs in a simulated disaster-management task. Our findings suggest that people can identify systematic AI biases and adjust the recommendations in the appropriate direction. However, their reliance on the AI depends on the direction of the bias. In particular, participants appeared to trust the AI less when it systematically underestimated resource needs than when it systematically overestimated them. I am now extending this work to dynamic environments in which the AI's behavior changes over time. This project examines how people update their trust as an AI shifts between unbiased and systematically biased recommendations.

In another project, we examined how people use information from multiple alarms when the alarms are dependent. When alarms are dependent on one another, their errors can become correlated, making agreement among the alarms less informative. We found that people use alarm information differently depending on the level of dependency, although their decisions are not fully optimal for that dependency level. We are now extending this research to examine how people respond when alarm dependency changes dynamically over time. By improving our understanding of how people use dependent alarms, this work can help inform the design of safer monitoring systems.

I am also involved in developing the CORA game, a research platform for studying decision making in disaster-response settings. This year, I visited the American Red Cross to observe a disaster-response management simulation designed to prepare Red Cross employees and volunteers to serve effectively as leaders during disaster operations. The visit allowed us to see how the training was actually conducted and gave us first-hand insight into how participants communicated, coordinated their activities, and allocated resources.

The DDMLab's newly minted PhD's!



In 2025, Grace and Xiaohong earned their PhD's! Grace in Cognitive Science from Rensselaer Polytechnic Institute and Xiaohong in Psychological and Brain Sciences from Indiana University Bloomington. The DDMLab has been very fortunate to have them as they embark on this new chapter of their careers.



From Anu Aggarwal

This year at DDMLab, I worked on an experiment under the IARPA ReSCIND program, conducted on CyberVAN, involving a cyber kill chain exercise with defined objective. Participants included industry penetration testers as well as knowledgeable cybersecurity students, and the kill chain scenarios were designed to reflect real-life situations, mirroring the steps a potential attacker would typically take to breach a network and access sensitive data. Participants were given access to CyberVAN and completed these scenarios over a 12-hour span. I was responsible for designing the kill chain scenarios for CMU and managing the smooth execution of this experiment, which required coordinating a complex experimental design across team members from multiple institutions — UTEP, UW, and CMU — as well as our industry partner, Peraton. This experience gave me a strong sense of how to manage the full lifecycle of running an experiment, while also building my confidence in handling complex situations and thinking quickly to resolve issues as they came up.



In addition, I worked on a journal paper in collaboration with the S3D lab, based on an experiment we conducted on the picoCTF platform, which provided a realistic environment and helped capture ecological validity. We designed CTF challenges to trigger cognitive biases in participants and studied how these biases influenced their decision-making while solving the challenges. Our focus was on the Availability Heuristic and the Sunk Cost Fallacy.

I've been genuinely excited all year to gain exposure to such impactful experiments. I also had the opportunity to present our paper, "Evidence of Cognitive Biases in Capture-the-Flag Cybersecurity Competitions" at HICSS 2026. It was a great experience to present our work and an even better opportunity to meet other researchers in the same domain and exchange ideas with people working on similar problems.

This was my last year at DDMLab, and I want to take a moment to thank Coty for mentoring me over the past two years. I'm truly grateful to have worked alongside everyone in the lab — especially Maria and Jeffrey — and for all that I learned, in ways I never imagined I would. Beyond the research, DDMLab has just been so much fun, and I'll always remember the parties, the biking, and the friends I've made here.

Where have our members been this year?

In January, Anu represented the DDMLab at The 59th Hawaii International Conference on System Sciences. That paper, "Evidence of Cognitive Biases in Capture-the-Flag Cybersecurity Competitions" examined how cognitive biases influence adversarial decision-making when developing cyber defenses. In this work, we analyze over 500,000 submission logs from picoCTF, a large educational CTF platform, to identify behavioral signatures of cognitive biases with defensive implications. Our findings show that participants often submitted flags with correct content but incorrect formatting (availability bias), and persisted in attempting challenges despite repeated failures and declining success probabilities (sunk cost fallacy). We outline a framework for bias-informed adaptive defenses that anticipate, rather than react to, adversarial actions.



From Maria Ferreira


I completed my postdoctoral appointment at the DDMLab in December 2025, after nearly three years. During this time, my research focused on human decision-making, cognitive modeling, and cybersecurity.



Much of my last year's work was dedicated to the ReSCIND project, which unfortunately ended early following changes in U.S. government funding. Together with collaborators from CMU and partner institutions (UW, UTEP, and Peraton Labs), we conducted three human-participant studies examining how cognitive biases influence cybersecurity decisions. The first study adapted well-known cognitive biases to cybersecurity scenarios, the second embedded these scenarios within the Cyber Kill Chain framework, and the third translated them into CyberVAN, a realistic cybersecurity environment. Although the project concluded earlier than anticipated, it generated valuable insights into human behavior in cyber contexts.

In parallel, I collaborated with Carolina Carreira, Anu Aggarwal, Alejandro Cuevas, Hussein Hibshi, and Coty Gonzalez on research investigating cognitive biases in Capture-the-Flag (CTF) competitions. Our first study, Evidence of Cognitive Biases in Capture-the-Flag Cybersecurity Competitions, presented at the 59th Hawaii International Conference on System Sciences (HICSS 2026), analyzed more than 500,000 picoCTF submission logs and identified behavioral signatures consistent with cognitive biases during attacker decision making. Building on these findings, we completed a second study that embedded psychological manipulations into four fully functional picoCTF challenges to experimentally evaluate the causal influence of cognitive biases on attacker behavior. This research is currently being prepared for submission to a journal.

I also continued my research on adaptive decision-making using the Binary Choice Task. By combining behavioral experiments with cognitive modeling based on Instance-Based Learning theory, I investigated when adaptive support can improve decisions in dynamic environments. Across two studies, we found that selective model-guided interventions substantially improved adaptation, with personalized guidance based on IBL model tracing proving especially effective in environments that systematically hinder learning. This work is also currently being prepared for submission to a journal.



Marko really is that tall!

Looking back on my time at DDMLab, I am incredibly grateful for the opportunity to work with such talented and supportive colleagues. I am especially thankful to Professor Coty Gonzalez for her mentorship, guidance, and encouragement throughout my postdoctoral appointment. The collaborative and welcoming environment of the lab not only strengthened my research but also helped me grow as an independent scientist. I leave DDMLab with lasting collaborations, valued friendships, and a deep appreciation for everything I learned during this journey.

From Marko Morrison

This was my first year at CMU and DDMLab. During this time, I've been focusing on autonomous cybersecurity systems and how we can find out where, when, why, and how they fail.



In the winter, I focused on developing MHBench (<https://github.com/cylabcyberautonomy/MHBench>), an open-source cybersecurity experimentation platform for quickly deploying vulnerable environments and hacking them with LLM-powered attackers. Because MHBench emulates full-scale computer networks, the initial system was incredibly slow, which prevented rapid iteration on my research. However, through collaboration with Lakshmi Adiga, we were able to speed up experimentation setup by over 10x, cutting the average time to fully execute an experiment down from ~11 hours to under 40 minutes.

During this summer, I've been thankful to work with four undergraduate interns through the REUSE program at CMU thanks to collaboration between Coty and Hanan Hibshi. The interns have been working on a variety of projects centered around improving the beginner experience of the PicoCTF platform; however, I've been most involved in the CyBlocks project, a visual coding IDE for creating and testing agentic cybersecurity systems. With this tool, someone with no code experience could construct their own hacker, introducing a greater diversity in expertise to the design of such systems. We plan to release CyBlocks to the public shortly, and encourage people to try building their own hacker - no coding experience required!

Overall, it's been a lot of work, but even more fun. Being in DDMLab has introduced me to such a variety of people and topics that I never would have experienced otherwise, and I'm thankful to have spent my first year with such a wonderful team. Excited to see what the future holds! :)

From Yuchen Li

This year, I worked on advancing CORA, an AI-powered disaster response training game, from a prototype into a more practical and usable tool. In particular, I focused on redesigning core gameplay systems, improving the user experience through iterative playtesting, and developing the game's infrastructure to support human-AI collaboration. Through many rounds of feedback and revision, including rebuilding parts of the system from the ground up, CORA became more polished and better suited for research and training applications. I am especially grateful to Jeffrey and Marko, whose thoughtful feedback helped improve the project throughout development.



[Learn more about CORA in this video!](#)

One exciting milestone was seeing CORA selected as a finalist for the Games for Change Student Challenge at the GEE! Awards. This recognition reflected the contributions of everyone involved and encouraged me to continue exploring the potential of AI-powered game-based learning.

I am fortunate to have collaborated with Xiaohong, Connor, Morgan, and Samsara, who each brought valuable ideas and expertise to the project. I am also deeply grateful to Coty for her guidance, encouragement, and support. I am currently on a gap year before beginning my PhD at CMU, where I look forward to continuing work on CORA and exploring the research questions it has inspired.

From Kibo Schaffer

I recently joined our lab to maintain the lab server and provide technical assistance on our projects.



Since then, I made some updates to our backup and security policies. I also added some niceties: using a copy-on-write filesystem called ZFS, snapshots are taken every 15 minutes in-case of accidental file deletion or modification. Lab members now log into isolated environments so that they can install and configure their own services (like Apache) for testing and development. CUDA is available for local LLM inference. I also consolidated all of the existing experiments into containers so that they can be more easily migrated and reproduced.

Before the summer ends, I plan to ask lab members about the pain-points in keeping projects synchronized between our github, server and local copies during development, to figure out tools that might help mitigate version skew. Soon I will be contributing to PyIBL and related projects.

From Don Morrison

The latest version of PyIBL is now 5.2.2, with some additions to its plotting capabilities and improved documentation.



As usual it can be downloaded from PyPi with `pip install pyibl`, and its documentation is at <https://ddm-lab.github.io/pyibl-documentation/>. While not publicly released there is now also WebPyIBL, an interactive demonstration of IBL operating on a spectrum of discrete choice tasks learning from experience, with a variety of plots available and easy GUI adjustment of parameters. As well as having PyIBL at its heart, and running on many simultaneous threads on our Janus server so it is quite speed, it also extends what can be done in ShinyIBL.


While I am now gone I hope someone will pick it up and work through the remaining minor bugs so it can be publicly released.

What is PyIBL?

PyIBL is a Python implementation of a subset of Instance Based Learning Theory. It is made for making computational cognitive models supporting research in how people make decisions in dynamic environments. PyIBL is a module of Python code, useful for creating Python programs.

Typically PyIBL is used by creating an experimental framework in the Python programming language, which uses one or more PyIBL Agent objects. The framework then asks these agents to make decisions, and informs the agents of the results of those decisions. The framework, for example, may be strictly algorithmic, may interact with human subjects, or may be embedded in a web site.

PyIBL is a library, or module, of Python code, useful for creating Python programs; it is not a stand alone application. Some knowledge of Python programming is essential for using it.



[Learn more here!](https://ddm-lab.github.io/pyibl-documentation/)

From Jeffrey Flagg

Over the past year, I have continued to support our members and projects. In 2025 this involved a lot of work with Anu and Maria on various ReSCIND projects and in 2026 a major thrust has been with the AISDM and the American Red Cross. I continue to serve on CMU's OSF Advisory Board which helps promote best practices in the greater CMU community. I have also recently taken the CMU Center of Excellence's generative AI training course. I am eager to begin applying the lessons we learned.



Artwork at the 2026 [CMU Community Picnic](#). That is four years in a row! The piece is based off CMU's brand standard colors.

I led the relaunch of the DDMLab website using the Cascade CMS to align with university standards. We believe it is a significant improvement!

In the upcoming year, I hope to continue developing binary choice and Theory of Mind studies. I also plan on developing AI assisted training tools for onboarding new DDMLab members.

From Vlad Miloserdov

At the DDMLab, I worked to develop a network simulation platform designed to study how AI agents interact with human cybersecurity experts in dynamic, high-stakes environments. The platform

allows researchers to plug in different AI agents and evaluate their performance across a range of metrics, including dependability and predictability, to help assess how much we can trust AI systems when they're operating alongside humans under real-world conditions.



A key contribution of this work was introducing a novel framework for measuring these metrics, which we termed dependencies. Conceptually, if you imagine the human experts and AI agents as nodes in a graph, dependencies are the edges and they allow us to define and constrain how these agents are allowed to interact with one another. This framework gave us a structured way to isolate specific strengths and weaknesses in AI agent behavior, rather than evaluating performance in the abstract.

One notable finding from our beta testing: when AI agents were not required to seek approval before taking costly or high-stakes actions, they consistently made sub-optimal decisions in non-trivial, ambiguous situations. This result underscored the importance of human-in-the-loop oversight for maintaining reliability as these systems are deployed in increasingly dynamic contexts.

Come see our newly resigned website!

In 2026, the DDMLab launched a complete redesign of our website to better showcase the lab's research, people, and impact. The new site introduces a modern visual identity, streamlined navigation, and expanded content that reflects the breadth of our work. The redesign makes it easier for visitors to explore ongoing projects, publications, software, educational resources, and opportunities for collaboration while providing a more accessible and engaging experience across desktop and mobile devices. [Click here to take a look!](#)



From Jet Jacobsen


I am a new visiting research associate in the DDM Lab. I received my undergraduate degrees in Economics, Philosophy and Business



Analytics in May of 2025 from Duquesne University. As an incoming research associate I have begun my work in the lab by brainstorming new directions and planning future experiments for the Optimal Stopping Task problem space. I am excited to be working on this project with the guidance of Coty and Erin Bugbee. As a visiting research associate I aim to deepen my knowledge of decision making through working on Optimal Stopping Tasks and other DDMLab projects.

From Raghav Sharma

I joined the DDMLab this summer as a research assistant working with Coty and Grace Roessling on human-AI teaming. My contribution is the Brigade Kitchen: a four-agent, Overcooked-style cooking game built on the lab's cogrid platform as a testbed for studying how humans and AI agents coordinate under pressure. Four cooks share ten interdependent subtasks, from taking orders and prepping ingredients to cooking, plating, and putting out kitchen fires, with limited tools and plates forcing genuine teamwork.



We have a long history of using games to study human behavior. [See a recent example here!](#)

Undergraduate Research

The DDMLab welcomes the participation of undergraduate students who might be interested in gaining research experience in the area of dynamic decision making. All our undergraduate student commit to registering for an [independent study](#) or research for credit for a full semester or more. The number of units vary between 6 and 12, and correspond to the number of hours per week that the student is willing to commit to research within the lab. Working with our talented team is a great first step for budding researchers!



Any seat can be played by a human, a rule-based agent, or, in future work, an LLM, and every session is fully logged and exactly replayable for analysis. I graduate from CMU in December, and I am excited to see the experiments on team roles, visibility, and leadership that this platform enables.



Fun times at the Dietrich College holiday party!

Project Highlight: C.O.R.A.

Collaborative Operations & Resource Management with Agentic AI C.O.R.A.

CORA is a real-time disaster response simulation inspired by American Red Cross tabletop exercises, where players open shelters and kitchens, manage staff, dispatch vehicles, and make high-stakes decisions to support displaced communities under tight budgets and time pressure. What sets CORA apart is that players do not work alone: AI teammates with distinct personalities propose actions, manage problems, and challenge decisions in real time, making human-AI collaboration a core part of gameplay rather than an add-on.



Learning by consequences, not instructions.

Rather than teaching disaster response through direct explanation, CORA relies on cascading consequences: a staffing shortfall strains logistics, a delayed shelter opening ripples into community satisfaction.

Players build systems thinking, prioritization under uncertainty, and — through negotiating with AI teammates — judgment about when to trust automation and when to override it. A configurable admin system lets facilitators customize maps and disaster conditions, so scenarios stay dynamic and replayable across skill levels and training contexts.

Built for research as much as training.

Beyond its use as a training tool, CORA is designed to function as a research instrument for studying human decision-making. The game logs structured data on every player and AI-agent decision — choices made, reasoning shown, outcomes reached — creating a dataset for studying how people allocate scarce resources under pressure and how they collaborate with AI recommendations. This year, the team began drafting a paper that frames CORA explicitly in these terms, situating it within the literature on human-AI teaming and decision-making under uncertainty.

A year of growth.

This year's development focused on turning CORA from a working prototype into something ready for real playtesting and deployment: standardizing terminology to match actual Red Cross operational language, rebuilding core gameplay systems for stability, adding accessibility features, and expanding the instructor tools, including an in-game map editor that lets facilitators visually design new scenarios without touching code. The team also deepened the integration of AI teammates into the simulation, giving them a shared structured log format with human players and a redesigned conversation interface as the primary way players interact with them.




In June, CORA was named a finalist for the Games for Change Student Challenge at the GEE! Awards, which is a validation that a game built for research rigor can also stand on its own as genuinely engaging play.

Project Highlight: Phishing Training

Comparing Human and LLM Susceptibility to Phishing Attacks

In collaboration with former DDMLab member Tailia Malloy, we have designed an experimental paradigm to investigate whether LLMs exhibit cognitive biases similar to those observed in humans in the context of phishing attacks. A growing body of research has examined which cognitive bias triggers—including authority, urgency, and scarcity—make humans more susceptible to phishing. At the same time, there is increasing interest in the cognitive biases exhibited by LLMs, yet relatively little is known about how or why these bias triggers influence LLM decision-making. To address this gap, we have developed a controlled experimental paradigm that directly compares how humans and LLMs respond to malicious messages. Specifically, we are interested in determining whether LLMs exhibit the same underlying cognitive biases as humans and, by comparing human and LLM responses, whether we can identify similarities and differences in how the two systems represent and process malicious messages.



During the study, participants work to determine if an email is phishing and decide what the best response to the message would be. Should you open a link? Download attachments? Maybe reply or archive it? Or delete it or mark as spam? After making their decisions, participants identify the components of the message that influenced their choice and explain what led them to make that choice. Would you get them all correct?

Motivation

The central motivation for this research stems from the increasingly blurred behavioral boundary between humans and LLMs as these systems continue to advance. LLMs are beginning to exhibit emergent patterns of behavior that resemble human cognitive tendencies, despite these behaviors not being explicitly programmed. From a scientific perspective, this convergence raises important questions about the nature of representations and decision-making in artificial systems and the extent to which these mechanisms parallel - or diverge from - those observed in humans. A controlled, experimental comparison of human and LLM behavior provides a concrete framework for investigating these questions.



Future Directions

We are interested in extending this work to human training interventions. In particular, IBL could be used to identify the cognitive bias triggers that influence an individual's susceptibility to phishing and, based on their previous responses, predict whether they are likely to fall for future phishing attacks. This approach could ultimately support more personalized and adaptive cybersecurity training.

Building Better from the Beginning



How will the next generation of people work with AI? In a June 2026 article, Coty and the DDMLab were profiled on how we go about answering this question. We believe the answer lies in harnessing the unique, different abilities that humans and AI have to create a truly collaborative effort.

[Read more here!](#)

PUBLICATIONS & CONFERENCE PAPERS

In the past year, DDMLab members and our collaborators authored several journal articles and chapters. For a full list, please see the publications page on the lab's website at <https://www.cmu.edu/dietrich/sds/ddmlab/publications.html>

Gonzalez, C. (2026). Reply to 'Human-AI complementarity needs augmentation, not emulation'. *Nature Reviews Psychology*. <https://doi.org/10.1038/s44159-026-00537-2>.

Gonzalez, C., & Heidari, H. (2025). A cognitive approach to human-AI complementarity in dynamic decision-making. *Nature Reviews Psychology*. <https://doi.org/10.1038/s44159-025-00499-x>.

Gonzalez, C., & Malloy, T. (2026). Toward Complementary Intelligence: Integrating Cognitive and Machine AI. *Current Directions in Psychological Science*, 0. <https://doi.org/10.1177/09637214251407571>.

Aggarwal, A., Ferreira, M., Aggarwal, P., Rajivan, P., & Gonzalez, C. (2025). Cognitive Biases in Cyber Attacker Decision Making: Translating Behavioral Insights into Cybersecurity. 2025 IEEE EuroS&P Active Defense & Deception Workshop. July 4th, 2025, Venice Italy. <https://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=11129591>.

Carreira, C., Aggarwal, A., Cuevas, A., Ferreira, M. J., Hibshi, H., & Gonzalez, C. (2025). Evidence of Cognitive Biases in Capture-the-Flag Cybersecurity Competitions. In *Proceedings of The 59th Hawaii International Conference on System Sciences (HICSS-59)*. Page 1219-1228. January 6-9, 2026, Maui, HI, USA. <https://doi.org/10.48550/arXiv.2510.05771>.

Chen, V., Gonzalez, C., Woolley, A. W., Lee, M., Wu, T., Conitzer, V., & Singh, A. (2026). Align AI to Dynamic Human-AI Workflows. *arXiv preprint arXiv:2607.14240*. <https://arxiv.org/pdf/2607.14240>.

Du, Y., Miloserdov, V., Ferreira, M. J., Prebot, B., Malloy, T., & Gonzalez, C. (2026). CHART: A Configurable Testbed for Human-AI Teaming Research in Cybersecurity Operations. In Chiou, E. K., Lange, D. S., Wong, J. H., & Marble, J. (Eds.) *Advancements in Human Agent Teaming Research Infrastructure: Testbeds, Metrics, and Concepts*. CRC Press, Taylor & Francis. Accepted.

Lou J., Liang, Z., Liu, P. Zhang, Y., Gonzalez, C. & Tang, P. (2026). Computational Modeling of Human Adaptation in Urban Infrastructure Management under Extreme Conditions: A Case Study of Subway Flood Scenarios. *Journal of Computing in Civil Engineering*. Accepted. <https://arxiv.org/abs/2606.06429>.

Malloy, T., Fang F., & Gonzalez, C. (2025). Improving Online Anti-Phishing Training Using Cognitive Large Language Models. *Computers in Human Behavior*. 108760, ISSN 0747-5632, <https://doi.org/10.1016/j.chb.2025.108760>.


McDonald, C., Nguyen, T. N., Botelho, C., Dishop, C., Woolley, A. W., & Gonzalez, C. (2026). Exploring the effects of real-time feedback on collaborative processes to enhance collective intelligence in teams. *Collective Intelligence*, 5(1), 26339137261431725. <https://doi.org/10.1177/26339137261431725>.

McDonald, C. & Gonzalez, C. (2026). Co-Grid & the Multi-user Gymnasium: A Framework for Multi-Agent Experimentation. *Behavior Research Methods*. <https://arxiv.org/abs/2604.15044>.


Project Highlight: AI-SDM

NSF AI Institute for Societal Decision Making

The National Science Foundation supports the AI Institute for Societal Decision Making, which brings together AI and social sciences researchers to develop human-centric AI for societal good. DDMLab members work with the AI-SDM to explore disaster management scenarios that require complex, often life-saving, decisions to be made under uncertain, dynamic, and resource-constrained circumstances. On [September 24-25](#), the AI-SDM will host an academic workshop focusing on human-AI complementarity for decision making.



The AI-SDM is advancing AI techniques for autonomous decision-making while considering human factors that govern the acceptance of decisions. These AI advancements are not only inspired by human decision-making but are also actively applied to enhance decision-making processes in critical areas.



Coty and Xiaohong at the 2026 AI-SDM Workshop.



Using science to promote positive social change: Participants at the 2026 NSF AI-SDM Annual Retreat.



Dynamic Decision Making Laboratory

Carnegie Mellon University
5000 Forbes Avenue
Social and Decision Making Sciences
Porter Hall 223G
Pittsburgh, PA 15213
e-mail: ddmlab@andrew.cmu.edu

Stay Connected All Year!



<http://www.ddmlab.com>



<https://www.facebook.com/ddmlab.cmu>




<https://bsky.app/profile/ddmlab.bsky.social>



<https://www.linkedin.com/showcase/27154823>



<https://www.youtube.com/@dynamicdecisionmakinglab>



Carnegie Mellon University