



Dynamic Decision Making Laboratory

ISSUE 16

FALL 2024

Carnegie Mellon University

Annual Summary from Coty — Founding Director, DDMLab

Dear Friends and Collaborators:

The Dynamic Decision Making Laboratory (DDMLab) is proud to update you on our activities during the 2023-2024 academic year. Thanks to all of you who are interested in reading this update!

Starting with updates on new research grants, we have some good news and continued updates on the work on our recent grants. Top on the list is our collaborative efforts with University of Texas, El Paso, University of Washington, and Peraton Laboratories in a **IARPA** program **ReSCIND**, intended to use psychologically-inspired insights regarding cognitive biases of cyber attackers to develop effective cyber defenses. As it will be clear through the research summaries of the DDMLab members in this newsletter, this grant which started in January 2024 has been the focus of much of our activities. It has been a very involved project, keeping us very busy since it started.

The next new grant in the lab is a collaboration with Katherine Flanigan, Mario E. Berges and Christopher McComb, my CMU colleagues at Civil and Environmental Engineering, with whom I will be developing Instance-Based Learning models to uncover relevant relationships between built environment features and the resultant human behaviors or actions. Our work aims to construct an understanding of how physical space characteristics influence human decision making, actions, and interactions within those spaces. This work which is about to start was funded by the **National Science Foundation (NSF)**.

We also continued to work on our current grants. Notably, the National Institute for AI in Societal Decision Making ([AI-SDM](#)), for which I am a co-Director for research. The institute, supported by the **National Science Foundation (NSF)**, has been a major activity for the DDMLab during its first year. We have been involved in multiple projects related to disaster management and public health, while developing basic research on dynamic decision making and aiming to advance our modeling framework and approach. In addition to the AI-SDM, we continue to be supported by the **Army Research Office (ARO)** Multi-University Research Initiative (MURI) program on cyber deception in collaboration with Harvard University and the **ARO US-Australia** International MURI ([CATCH](#)) on building cyber security teams of humans and bots (AI, machines).

Finally, during last year, we concluded a long-term program with the **Army Research Laboratories (ARL)** on collaborative research alliances on cybersecurity as well as the **Defense Advanced Research projects Agency (DARPA)** program on Artificial Social Intelligence for Successful Teams (**ASIST**), focused on collective intelligence and creating new models that can be used to collaborate with humans in teams to advance collective intelligence and adaptation.


As it will be clear through the research summaries of the DDMLab members in this newsletter, all these funding agencies are providing resources into advancing foundational research on dynamic decision making that has implications for a number of applied significant problems, notably: cybersecurity and human-AI interactions and collaborations.

Inside this issue:

Annual Summary from Coty	1-2
Welcome new members	3
Research Updates From Our Members	4-7
Projects and Events	8
Summary of Publications, Conference papers	9
Adventures and Travels	10
Thank You and Contact Information	11

Beyond the thrill of getting new research grants, managing exciting research projects, and aiming at catching up with my post-docs and students productivity, this year was generally less full of other activities, compared to 2023 where Spain and related activities brought up a lot of excitement to the lab.

On a personal note, I have continued to enjoy cycling and painting. Although I have been significantly less productive this year, I have paused my painting for a number of months and then took it again just for a brief time. Here are some of the paintings produced this year. Some of them I painted over someone else's drawing ... so I sort of cheated.



On the cycling side, have you heard of the **El Camino De Santiago**? The “Way of St. James” is a pilgrimage in the northern Spain, ending in the Cathedral of Santiago de Compostela. This year, I did the Camino starting on the coast of Portugal, by bicycle. Most people do this journey by foot.

It is not only a spiritual or cultural journey, but a great physical activity!! About 245 miles in 8 days. Starting in Lisbon, we moved on to Porto, Esposende, Caminha, Baiona, Pontevedra, Caldas dos Reis, and finally, Santiago de Compostela!. What an experience!!



El Camino- day 4

Distance	Elev Gain	Time	Achievements
41.96 mi	2,510 ft	5h 22m	2


Congrats! You just logged your 2nd biggest climb.



Overall, this academic year was enjoyable, unique, full of work and full of surprises. We are ready to start the new academic year 2024-2025. I hope you will continue to join us in making all our work count.

Let's all work to make this a better world!

Welcome New Members!



Anu Aggarwal

Anu is a new Research Associate and programmer. She has 6 years of industrial experience in software development. She is a Master degree holder in Cyber Law & Information Security from the Indian Institute of Information & Technology Allahabad. She completed her Bachelor's in Computer Science from University Institute of Engineering & Technology (UIET), Panjab University, Chandigarh. She is interested in exploring cognitive biases and their applications in the field of cyber security.

Arya Mane

Arya is an undergrad student studying Statistics & Machine Learning and Business at CMU. She is interested in exploring how we can take principles of how humans make decisions and its implications in the tech sphere, specifically in cyber security. On campus she is part of Design For America, and outside of campus she enjoys spending time in the outdoors.



Research Updates from Lab Members During 2023-2024

FROM TYLER MALLOY

This was my second year as a postdoctoral researcher in the DDMLab. In the beginning of this year, I created and ran an undergraduate course called 'Human and Machine Decisions from Experience' which taught students using the python programming language and the PyIBL package for Instance Based Learning. This was a great opportunity to share the PyIBL library with students interested in research in decision science. The final assignment of this course had students build IBL models to predict the learning and decision making of participants in a cybersecurity task, identifying emails as being either phishing or dangerous. I hope to continue this class this fall during my third and final year in the lab.

Our research into integrations of cognitive models and artificial intelligence methods has increased thanks to a grant from the Microsoft Accelerating Foundation Models Research Program, which included \$20,000 of azure credits that we could use to prompt GPT-4, Dall-E, or run azure servers and services. We used these resources to generate phishing emails to test people's ability to tell if an email generated by a large language model is safe or dangerous. This research gave us valuable insight into the potential risks associated with the misuse of AI methods, and how best to train the public to be aware of these risks. This grant has already been applied onto a publication in *Frontiers in Psychology*, and several submissions currently under review elsewhere.

In May I attended the Multi University Research Initiative meeting at the CMU Silicon Valley campus where I presented some of our recent work integrating AI methods and cognitive models, as well as our work in cybersecurity more broadly. There I met with other researchers investigating various social and technological applications of cognitive modeling and artificial intelligence. I received valuable feedback on our planned study of human participant abilities in identifying phishing emails generated by LLMs, as well as our work in the CybORG environment with the Individual and Teaming Defense Game tasks. I plan on continuing this research into real world applications of cognitive modeling and AI in cybersecurity applications.

FROM MARIA FERREIRA

Over the past year, I've had the privilege of working in the DDMLab, where my focus has been projects covering Binary Choice, Phishing, and Cybersecurity. Notably, my work on IBL modeling with Model Tracing and Model Fitting using PyBL in a Binary Choice task has paved the way for a significant achievement—the acceptance of a six-page paper at the ICCM 2024 conference. This success underscores the potential of our research to make a meaningful impact in various domains.

My collaboration with Palvi Aggarwal and Arunesh Sinha has been instrumental in applying our findings of Model Tracing and Model Fitting to the Phishing and Insider Attack Game Tasks. This collective effort has led to significant insights, such as the enhanced performance of the IBL model through parameter calibration and the interpretability of our Model, which revealed to be similar to that of neural network models. With Palvi's team, we were able to have an abstract and a full paper accepted at the HFES 2024 conference, and we are eagerly awaiting the results of a submission with Arunesh's team.

In collaboration with my postdoc colleague Tyler Malloy, we've delved into incorporating large language models (LLMs) into the Phishing task, aiming to discern differences in detecting human-generated and AI-generated emails. We are currently awaiting the results of various submissions with our main findings.

Embarking on the ReSCIND project earlier this year has been a captivating experience. The goals of this project are multifaceted. First, we aim to investigate whether attackers demonstrate bias, drawing from established literature on the subject. Secondly, we intend to utilize any identified biases to thwart attackers and prevent their success. Collaborating closely with Anu Aggarwal, we've been dedicated to developing cybersecurity scenarios aligned with traditional bias methods (cyber isomorphs). Our recent involvement in a two-day site visit to Peraton Labs, New Jersey, in May, where IARPA and several esteemed partners (ARLIS, LLNL, MITRE, and MITLL) evaluated our work, marked a significant milestone in this project. Lately, we have started collecting data to validate our cyber isomorphs, and we are currently analyzing it.

As we look ahead, the upcoming year in the DDMLab holds great potential for learning, networking, and advancing groundbreaking research. I am very excited about the opportunities and challenges that lie ahead.

Research Updates Continued

FROM RODERICK SEOW

My first year at the DDMLab has been an absolutely eye-opening experience. Together with collaborators from both academic and non-academic organizations, I was involved in multiple exciting projects as part of the AI Institute for Societal Decision Making, and even got to visit training sites and sessions to learn more about real-world disaster and emergency management. My projects were mostly centered around cognitive modeling - both in terms of applying cognitive models to improve human-AI outcomes and designing models to test theories about the cognitive mechanisms that underlie human learning and decision making.

In collaboration with Hoda Heidari, we investigated the appropriate level of adaptation and anticipation that an AI partner should possess in a human-AI coordination task. Using our novel multi-attribute, multi-agent task, we compared team performance between types of human-IBL teams. IBL partners were designed at different k-levels, which reflects the number of recursive steps taken when making a prediction about the human partner's choices. Preliminary results will be presented at the 2024 meeting of the Psychonomic Society, and further studies are currently under development.

Together with Yunfan Zhao, Milind Tambe, and Duncan Wood, we developed personalized time-series forecasters for a repeated intervention allocation problem using Instance-Based Learning models. For context, we used real-world data from ARMMAN - a NGO that serves expectant mothers in India - and modeled how mothers' engagement with the program changed over time and in response to receiving voice calls (interventions) from ARMMAN. Unlike the previous time-series forecasting approaches, each IBL model has individually-fitted parameters

that specifically capture each beneficiary's engagement consistency over time, and how responsive they were to an intervention. We show that the personalized IBL models better predicted future engagement levels than existing approaches. This work has been submitted as a poster to the 2024 meeting of the Society for Judgment and Decision Making and will be submitted as a paper to the 2024 workshop on Health Recommender Systems.

Tyler Malloy and I have also started collaborating to investigate the role of dimensional attention within the Instance-Based Learning framework. Preliminary simulations with IBL models reveal that while increased attention to stimuli dimensions that are predictive of reward (or risk) during training does not appear to result in noticeable differences in training performance, the differential allocation of attention is necessary to account for the empirical advantage that humans demonstrate on intradimensional vs. extradimensional transfer tasks. We are also investigating how these attentional weights are learned through experience, and are currently exploring information-based vs. error-based learning mechanisms.

As part of an ongoing partnership with the American Red Cross, we are developing an emergency management game based on the Disaster Resource Management Simulation exercise. This game will be both used as a testbed for research and a digital tool for training Red Cross personnel. Guided by the core decision-making dilemmas and considerations faced by the DRMS trainees and our research interests in multi-agent resource allocation problems, we wrote and submitted a game design proposal to the Entertainment Technology Center at CMU in order to recruit game design students to help develop a digital version of the game.



Lab members at the inaugural AI-SDM Annual Review.



Research Updates Continued

FROM ERIN BUGBEE

This year was another wonderful year as a PhD student in the DDMLab, shaped by a wide array of academic and professional experiences. I have continued to investigate human decision making through behavioral experiments and cognitive models, gaining insights into how people balance exploring options and collecting information with exploiting knowledge gained and stopping search. A significant achievement was proposing my dissertation in March, titled "Balancing Exploration and Exploitation: Sequential Decision Making in Humans and Machines." This milestone was the culmination of years of research in the lab.

In my work this past year, I found that people learn from feedback in an optimal stopping task I developed, and that the type of feedback affects learning and performance. Additionally, I discovered an interesting result: having information about the distribution from which the values of options are sampled can hinder necessary exploration. This result is both novel and impactful, as it implies that having knowledge about the environment may discourage exploration, so perhaps exploration may be encouraged by making people aware of the potential information to be gained through further search. I am grateful to have received the Tata Consultancy Services Presidential Fellowship, which will continue to support me this fall. I look forward to building upon this work in the remainder of my dissertation.

This past year has been full of a variety of adventures. In the fall, I attended `posit::conf(2023)` in Chicago as an Opportunity Scholar, where I connected with the R and Python communities, discovered the latest tools, and learned techniques I started integrating into my work, such as using Quarto for reproducible data analysis and for updating my personal website. In October, I traveled to Puebla, Mexico with Coty and Ngoc Nguyen, where we presented a workshop on SpeedyIBL at the XI Latin American Conference on Human Computer Interaction (CLIHIC 2023). Being in Puebla for the Mexican holiday *Día de los Muertos* and seeing Coty's university and where she lived for many years was extremely special.

I presented a poster at SJDM in San Francisco and had a paper accepted to CogSci, which were both fantastic opportunities to share my work. This past summer, I interned as a Machine Learning Engineer at Apple in Austin, Texas where I gained experience building machine learning models for fraud detection. I wrapped up the summer by attending `posit::conf`

(2024) in Seattle, where I was selected to give a talk on how I use R and Python in my research, titled "To Explore or To Exploit: Decoding Human Decision Making with R and Python."

Exciting work lies ahead as I complete the remaining experiments in my proposal and prepare for my dissertation defense in the spring. I look forward to another year of learning and discovery in the DDMLab.

FROM CHASE McDONALD

A major focus of the last year has been on my dissertation, which I successfully proposed this past July. My dissertation proposal, On Human-AI Complementarity in Shared Workspaces, was comprised of three chapters. The first contains completed work, the second is in progress, and the third is entirely proposed. In the first chapter, we developed and open-sourced a framework for human-AI experimentation comprised of two libraries: CoGrid and Interactive Gym. The former is a library for the simple development of multi-agent grid-based simulation environments. The latter, Interactive Gym, allows for multi-agent simulation environments to be easily ported into the browser to run human experiments with any number of humans or AI agents acting together. The second and third chapters put these tools to use. In the second, we propose a method for developing controllable reinforcement learning agents called Interpretable Behavior Conditioning (IBC). We first demonstrated that this method allows us to build more robust collaborative AI partners, then proposed experiments to investigate human preferences for control and predictability over their AI partners. The final chapter takes an alternative approach to human-AI interaction and proposes an investigation into how we can use empowerment-maximizing AI partners to induce the zone of proximal development in humans to improve human learning.

In addition to the work in my dissertation, we've continued work on several long-running projects. Our paper on the effects of collective intelligence awareness displays, presented last November at the Collective Intelligence Conference, is now under review at the Collective Intelligence journal. We also presented our work on mechanisms for human-like credit assignment at the AAAI Spring Symposium earlier this year.

Over the next year, I'm looking forward to continuing my work in the DDM Lab and wrapping up my dissertation as I finish my final year in the PhD program.

Research Updates Continued

FROM YINUO DU

In the past year, I continued my work on human and AI decision making in the context of cybersecurity. The milestone is my dissertation proposal, titled “Human and AI Decision-making in Cybersecurity: A Multiagent Perspective”, which integrated my four lines of work: cognitive modeling of the attackers, reinforcement learning for network defense, human-AI team defense, and cross-organization collaboration. Part I is under R&R to ACM-Transaction of Social Computing. Part III is under review at Computers in Human Behavior. My co-authors helped immensely in revising and improving the manuscripts. I’m grateful for their guidance during my first time submitting to both journals.

For my final year of PhD, I will facilitate the real world deployment and empirical evaluation of the reinforcement learning defense strategies I developed in part II. My focus will be on part IV, which examines human decision making in multiplayer prisoner’s dilemma games. I’ll analyze the human data we collected in online games, design cognitive models that capture human learning in groups, and extrapolate from small groups of three to large social networks. On the side, I will facilitate the HackIT project led by previous DDMLab Palvi Aggarwal. The project will be relevant to Part I of my dissertation, with a specific focus on the cognitive biases of the attackers.

In the past year, I explored the exciting rising domain of large language models and leveraged this new tool to investigate collective problem solving in groups. This work is accepted to CogSci 2024. I was also fortunately involved in Tyler’s Accelerate project. I learned a lot from collaborating with him and other members in DDMLab. I also had the opportunities to attend the AAAI Human-like Learning Symposium and the WiCyS conference, where I connected with the community of human learning and cybersecurity researchers. In the past summer, I attended the LearnLab summer school, where I learned to design cognitive tutors and connected with the community of cognitive modelers, learning scientists, and teachers.

In terms of teaching, I mentored two master students on their course project in AI for Social Good. In the past summer, I participated in the REUSE program hosted by S3D and mentored an undergraduate student. I’m also participating in the Future Faculty Program and practicing on course syllabus design. Helping others on their research and designing courses

helped me to reflect on my research, organize my knowledge structure, and improve my interpersonal skills. In terms of community service, I joined the AI-SDM student leadership council, the S3D PhD admission committee, the REUSE admission committee, and the SCS Teaching Award committee. I am glad to contribute to the community. Conversations with fellow members in the committee are a great source of inspiration.

I’m looking forward to wrapping up my dissertation work and exploring new opportunities.

FROM DUNCAN WOOD

In my first year of the Ph.D. program, I have assisted Roderick on the project on public health for AI-SDM. We have continued our analysis of the IBL-TARI model. This allows the model to learn how “interested” the mother is in the “engage” action, contextualized by the periods since her last intervention. This approach produces greater prediction accuracy than the TARI model trained on 300 mothers. Through our analysis, Roderick and I realized that the IBL-TARI model is characterized by two parameters, not one. The parameters describe how much weight the mother places on the similarity of the time since the last intervention and the similarity of the previous engagement. When both of these are zero, the mother is modeled as only accounting for recency and frequency. I presented our findings at the AI-SDM annual review.

FROM DON MORRISON

In the past year there have been several releases of PyIBL, the most recent being version 5.1.5. Changes since version 5.0 a year ago include: A new plot() method that can display a variety of plots describing the evolution of a model’s internal state over time. A new aggregate_details property that provides a pandas.DataFrame containing detailed information about a model’s history. A new noise-distribution property that allows providing an alternative to the default logistic noise distribution that is expected be useful for some esoteric experiments and models. A variety of small improvements to the model development and debugging utilities trace and instances(). A revised web site design, which is in the process of being populated with more complex example of PyIBL’s use and improved documentation. For the latest see the PyIBL website at <http://pyibl.ddmlab.com/>!

Recent Publications

FROM ANU AGGARWAL

As It has been 6 months since I started working in DDMLab as a Research Associate. In these past months, my major focus of research has been in Cyber Psychology. This study funded by IARPA aims on improving cyber security defenses by developing defenses that leverage the limitations in the decision making attributes of attackers based on cognitive vulnerabilities (cogVuln). I have been working on this project in collaboration with Maria and Coty from DDMLab and also with UTEP, UW and Peraton Labs. We have designed our first experiment that validates the existence of cogVulns in the established methods and in the cyber domain. We have manipulated the scenarios i.e. bias triggers to identify the cogVulns and defined different metrics for measuring each cogVuln. Our next steps are designing the scenarios in a realistic cyber environment which includes various bias sensors and triggers in order to measure the cogVulns in the attacker.

Another project that I have been working on is "Interactive Adversary Game (IAG)" in collaboration with an undergrad student Arya Mane and Coty at DDMLab. This idea is motivated from some of the past work in DDMLab on Interactive Defense Game (IDG) and Team Defense Game (TDG). IAG focuses on developing an environment where human participants acting as attackers can perform the cyber attack on a defined network and go through various cyber kill chain steps while doing so. We are developing a generic game where the network or any modifications in the scenario is possible. IAG will also act as a simulator to test the various strategies of attacker and defender. This will work not only for humans as attackers however this will also use attacker cognitive agents. In the future our goal is to develop the cognitive agents, Attacker and Defender, following different attack and defense strategies, running the simulations and also running the study for human participants.

Over the past six months, I've had valuable opportunities to learn and grow. I'm thoroughly enjoying my research area and look forward to continuing this work in the coming year!

FROM JEFFREY FLAGG

In the Fall 2023, I had the pleasure of helping organize the ARO Workshop: The Future of Cyber Deception. The goal of this workshop was to bring together leading researchers from cyber security,

AI researchers, as well as cognitive scientists, who work on the research area of cyber deception against adaptive and intelligent adversaries. This interdisciplinary gathering of experts worked to improve our understanding of characteristics and effectiveness of deception techniques and algorithms, the dynamic nature of cyber and cyber-physical attacks, and the adaptive nature of cyber defense.

I have also continued to contribute to various experiments during the year. I have helped test materials, review papers, maintain our website and social media materials, and integrate new members. I have had the pleasure to work with Coty and former lab member Ngoc Nguyen on exploring aspects of Theory of Mind. We hope to publish our work in the near future! I also have continued to serve on CMU's OSF Advisory Board which helps promote best practices in the greater CMU community.

FROM ARYA MANE

This Over the past few months at DDMLab, I have been deeply immersed in an exciting cybersecurity project, focusing on leveraging instance-based learning to develop an engaging attack simulation game. This game is designed to model and analyze various attack scenarios, providing insights into defensive strategies and vulnerabilities. Working alongside Coty and Anu has been incredibly rewarding; I have learned a lot and gained valuable experience. Together, we have delved into the intricacies of instance-based learning to create a dynamic and educational tool for understanding cybersecurity threats. I'm eager to continue exploring and innovating in this space with Coty, Anu, and the rest of the team!



Viewing the April eclipse with the DDMLab!

Highlighted Projects and Events

Army Research Office Workshop: The Future of Cyber Deception October 4-6, 2023

Cyber networks are ubiquitous but also vulnerable to attacks on multiple dimensions. Attackers no longer consist of individuals or small groups, but have expanded to include active engagement by state and other highly organized actors with significant resources, who employ increasingly sophisticated military intelligence techniques in their approaches and strategies for launching cyberattacks. It is thus extremely important to design effective defense algorithms to protect cyber systems from such intelligent and well-organized attackers. Among different defense mechanisms, cyber deception is an economical technique, which can work effectively by providing misleading information to deceive attackers.



The goal of this workshop was to bring together leading researchers from cyber security, AI researchers as well as cognitive scientists, who work on the research area of cyber deception against adaptive and intelligent adversaries, as well as practitioners and other interested parties. This expert interdisciplinary gathering of experts in the area of cyber deception worked to improve our understanding of characteristics and effectiveness of deception techniques and algorithms, the dynamic nature of cyber and cyber-physical attacks, and the adaptive nature of cyber defense. This workshop shed light on key challenges and interdisciplinary research opportunities, with the ultimate goal of improving cyber defense.

IARPA: ReSCIND

This year, the DDMLab was a part of a major new grant! The Intelligence Advanced Research Projects Activity (IARPA) Reimaging Security with Cyberpsychology-Informed Network Defenses (ReSCIND) grant will help fund multi-site work at CMU, University of Texas at El Paso, University of Washington, and Peraton Labs. Coty is the principle investigator on this project.



The ReSCIND program aims to improve cybersecurity by understanding how human cognition impacts cyber behavior and affect cyber actors' success in network attack activities. Our teams will investigate how well-known cognitive patterns, such as loss aversion and the representativeness bias, may mitigate factors in the efficacy of cyber-attack behavior. This research will contribute to the broader goals of improving cyber defense practices by delaying and thwarting attacks.



Recent Publications

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

Gonzalez, C. (2023). Building Human-Like Artificial Agents: A General Cognitive Algorithm for Emulating Human Decision Making in Dynamic Environments. *Perspectives on Psychological Science*. <https://journals.sagepub.com/doi/epub/10.1177/17456916231196766>

Aggarwal, P., Cranford, E. A., Tambe, M., Lebiere, C. & Gonzalez, C. (2023). Deceptive Signaling: Understanding Human Behavior Against Signaling Algorithms. In T. Bao et al., (eds.). *Cyber Deception, Techniques, Strategies, and Human Aspects*. Vol. 89, pp. 83-96, Springer. https://doi.org/10.1007/978-3-031-16613-6_5

Aggarwal, P., Nowmi, S. R., Du, Y. & Gonzalez, C. (2024). Evidence of Cognitive Biases in Cyber Attackers from An Empirical Study. In *Proceedings of the 57th Hawaii International Conference on System Sciences, HICSS 2024*, January 3-6, 2024, Waikiki, HI. (pp. 934-943). <https://hdl.handle.net/10125/106491>

Du, Y., Prebot, B., & Gonzalez, C. (2024). Turing-like Experiment in a Cyber Defense Game. In *Proceedings of the AAI Symposium Series*, Vol. 3, No. 1, pp. 547-550. <https://ojs.aaai.org/index.php/AAAI-SS/article/view/31271>

Nguyen, N. T., Phan, N. D., & Gonzalez, C. (2023). Learning in Cooperative Multiagent Systems Using Cognitive and Machine Models. *ACM Transactions on Autonomous and Adaptive Systems*. Vol. 18, Issue 4, Article No: 15, pp. 1-22. <https://doi.org/10.1145/3617835>

Malloy, T. & Gonzalez, C. (2024). Applying Generative Artificial Intelligence to cognitive models of decision making. *Frontiers in Psychology*, 15, 1387948. <https://doi.org/10.3389/fpsyg.2024.1387948>

McDonald, C., Malloy, T., Nguyen, T. N., & Gonzalez, C. (2023). Exploring the Path from Instructions to Rewards with Large Language Models in Instance-Based Learning. In *Proceedings of the AAI Symposium Series*, (Vol. 2, No. 1, pp. 334-339). <https://ojs.aaai.org/index.php/AAAI-SS/article/view/27697/27470>

Park, S., Puranam, P., & Gonzalez, C. (2024). Decision Centralization and Learning from Experience in Groups: Separating Context from Aggregation Effects. *Management Science*.

Prebot, B., Du, Y., & Gonzalez, C. (2023). Learning about simulated adversaries from human defenders using interactive cyber-defense games. *Journal of Cybersecurity*. Volume 9, Issue 1. <https://doi.org/10.1093/cybsec/tyad022>



Dietrich College
Holiday Party
2023!

Adventures During 2023 - 2024



Three lab members defended dissertation proposals in 2024. Congratulations to Yinuo Du, Chase McDonald, and Erin Bugbee!



Lab members: singing, biking, and painting!





Dynamic Decision Making Laboratory

DDMLab-Porter Hall

Carnegie Mellon University
5000 Forbes Avenue
Social and Decision Making Sciences
Porter Hall 223-G
Pittsburgh, PA 15213
e-mail: [ddmlab\(at\)andrew.cmu.edu](mailto:ddmlab@andrew.cmu.edu)
Phone: 412-268-9547
412-268-8439

THANK YOU FOR YOUR INTEREST IN OUR
2024 NEWSLETTER!

KEEP UP WITH US ON OUR [WEBPAGE](#) OR FOLLOW US ON

[FACEBOOK](#)



[LINKEDIN](#)



[TWITTER](#)



OR [YOUTUBE](#)

