



Dynamic Decision Making Laboratory

ISSUE 14

FALL 2022

Carnegie Mellon University

Annual Summary from Coty — Founding Director, DDMLab

Inside this issue:

Annual Summary from Coty	1-2
Farewells and Welcome	3
Research Updates From Our Members	4-8
Summary of Publications, Conference Proceedings and Presentations	9
Adventures and Travels	10
Thank You and Contact Information	11

Dear Friends and Collaborators:

This is an annual update on the activities in the Dynamic Decision Making Laboratory, and we hope you enjoy reading about us from the past academic year (2021-2022) and our news to get started the 2022-2023 academic year. Thank you for letting us update you!

During last year we obtained a new research grant from the **Air Force Office of Scientific Research (AFOSR)** aimed at investigating trust establishment between human and machine recommenders in intelligence analyses. We also obtained a grant to organize a workshop on *Cyber Deception: Modeling, Computation and Adaptive Learning* from the **Army Research Office (ARO)**.

We also continued to be supported by our new grants on Cybersecurity and by our long-term grants, enjoying our work with many collaborators. We continued the work sponsored by the **ARO's Multidisciplinary University Research Initiative (MURI)** program on Cyber deception. Our collaborations with researchers at Harvard University have helped advance the algorithms on deceptive signaling in cybersecurity by adding new uncertainty constraints. The MURI project together with our long-term grant from the **Army Research Laboratories (ARL)'s Collaborative Research Alliance (CRA)** on Cybersecurity, allowed us to make significant advances on the concepts of collaboration among cyber defenders. Using Instance-Based Learning Theory to advance models of human-autonomy teaming for cyber operations. Our grant with the US-Australia International Multidisciplinary University Research Initiative (**MURI-AUSMURI**) has advanced significantly in its first year with development of better understanding of human-AI Cybersecurity Teams.

We also continued our work on cognitive theory of mind for human-machine teaming in the **DARPA's Artificial Social Intelligence for Successful Teams (ASIST)** program. This year we made significant advances to the development of metrics of team collective intelligence (CI), and advanced the understanding of models that can be used in collaboration with humans in teams to advance their CI.

This year we also concluded the work with the Network Science Program of the **Army Research Office (ARO)**, and with the **Air Force Research Laboratory (AFRL)** regarding learning science in human-machine team effectiveness.

Last but not least, we worked very hard this year to submit a large proposal for an Institute for AI-enabled Societal Decision Making. Send good vibes our way, we are looking forward to engage in this effort.





Many thanks to my current and past post-docs and all the great collaborators that have contributed to our success this year! We look forward to continue working with you.

Other relevant development notes of importance include:


I started a new course in Spring of 2023. *Decision Models and Games* is a new course that was proposed as part of the Dietrich College General Education Curriculum.

I was granted a Fulbright scholarship and will spend my Spring Semester of 2023 at the University of Malaga, Spain. My expectation is to be able to slow down a bit during the Spring semester but I would still be directly connected with all my projects remotely during that time from Spain.

And... Yes, I continue cycling every day to my office and every weekend for longer rides. I think that soon I will need to get an electric bike. But, while I can, I still pedal with all my force through the hills of this wonderful city.



Painting, continues to be my other passion... mostly working on improving my watercolor skills now... Here are a few of the paintings from this past year:



I hope you enjoy this newsletter. Thank you so much for reading our update, and thank you so much for contributing to our efforts in so many ways!! We look forward to a new academic year full of interesting research, surprises and new adventures in 2022-2023. **Let's all work to make this a better world!!**

Coty

Farewell and Welcome New Members!

Farewells

We have to say good bye to one of our post-docs this year, Nhat Phan. But it is so exciting that Nhat moved on to the University of Massachusetts to work on applied mathematics and teaching. Our loss will be their new institution's gain!

To everyone on this list, we are very proud and happy for all of you!! Good luck in your future endeavors!



Prof. Nhat Phan



Erin McCormick, PhD



Jiajia Hu



Ed Matlack



Miso Demko

Erin McCormick remained with us this year as a visiting scholar. She is now in her new position of Research Psychologist with the Air Force Research Laboratory.

Jiajia Hu will continue as a Ph.D. student in the HCII, now moved to a lab that is expected to be a better alignment with her computational data-driven science interests.

Ed Matlack is continuing his independent research in the measurement of meaning in decision making and memory retention.

Miso Demko is currently finishing his undergraduate studies at CMU.

Welcome!



Tyler Malloy (left) is a new post-doctoral researcher with the DDMLab. He is expected to receive his PhD in Cognitive Science in November from Rensselaer Polytechnic Institute in Troy, New York. After this, he will join our lab. His research interests include cognitive modeling of human perception, learning, and decision making.

Research Updates from Lab Members During 2021-2022

FROM NGOC NGUYEN

This past year, I developed a multiplayer version of Minimap, aka Team Minimap, in which multiple people can join as a rescuer team with different roles to navigate through a dynamic environment and triage victims. The development lays the groundwork for investigating not only humans' behavior in a teamwork setting, but also how to build artificial agents that are socially intelligent and beneficial for human collaboration when making joint decisions. In particular, I have been working on a project that studies the relationship of Collective Intelligence (CI) measures and the performance of human players at the team level as well as at their individual level in the Team Minimap environment. The goal of this study is to analyze the extent to which the CI measures can predict the team performance and whether the awareness of CI measures will help boost the team performance. The study is part of the collaboration with Anita Woolley's Collective Intelligence-Lab at CMU, and we expect that the results will inform the design of interventions. Furthermore, I am setting up a Human-AI teaming experiment that examines the effect of different strategies of agents who serve as teammates of human partners. In this research work, we are not only interested in the performance of human-AI teams, in terms of score, but also in the collaborative fluency measures. The results of this study are anticipated to inform the design of collaborative agents in human AI collaborations as well as guide approaches in human behavior modeling.

Relatedly, I have been involved in a project that investigates the utility of cognitive IBL models in designing Reinforcement Learning (RL) agents with Chase, Coty, and Henny Admoni's HARP Lab at CMU. Specifically, the goal of this project is to evaluate the influence of cognitive IBL models on training of RL agents by playing with in a collaborative cooking task, Overcooked. Regarding publications, we had success in publishing our work on SpeedyIBL, an advanced implementation of the full set of Instance-based Learning Theory mechanisms to the Behavior Research Methods journal. Together with Nhat, we also prepared and gave a SpeedyIBL tutorial to our lab members. The collaborative work with Anita's and Henny's groups on studying the effect of anger on performance and teamwork in Minimap has been accepted in the INGROUP 2022 conference.

FROM NHAT PHAN


During the past year, I have worked on several projects. In the first project on multi-agent systems, we developed different IBL models, combined with rein-

forcement learning (RL) techniques, for coordinated multi-agent systems in stochastic environments. Through numerical experiments, we demonstrated that our models outperform existing related work in RL. We submitted this work to a journal and it currently under review.

In the second project, we have developed a new Python library version for creating IBL agents called SpeedyIBL. SpeedyIBL makes significant improvements in the computation time of running IBL models. This work has been published in the Behavior Research Methods journal. We also prepared and gave the SpeedyIBL tutorial to our lab members.

In addition, I have been involved in a defensive cyber-deception project where we used the IBL for training attackers to develop adaptive defense strategies. In another cyber project, we develop adaptive training algorithms to help humans distinguish phishing emails from ham emails. We have collected human data for these projects. Recently, I have worked on TRUST™ project where we develop IBL models that are aimed to help improve the relevance of suggested articles.

With one and half wonderful years with DDM Lab, I am now moving to UMASS Lowell to work on applied mathematics and teaching. I am grateful to the lab for proving me a great exposure and opportunities during my postdoc. I hope to continue to collaborate with this wonderful team in future.



For more about SpeedyIBL see our lab GitHub page!

<https://github.com/DDM-Lab/SpeedyIBL>

Research Updates Continued

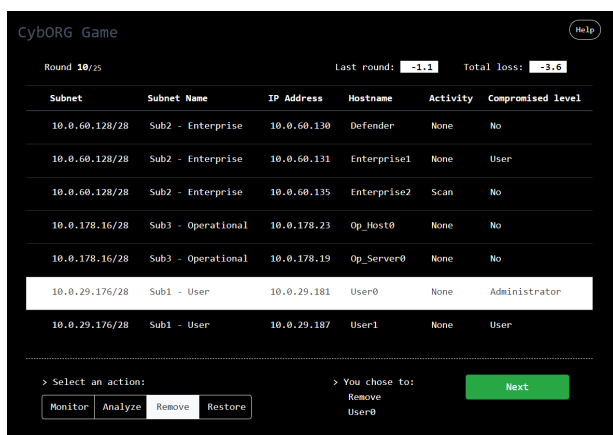
FROM BAPTISTE PREBOT

Since I joined the DDMLab as a postdoc last October, I have mainly been working with Yinuo and Coty on a cyber defense task, as part of the CATCH MURI project studying shared human-machine models for cybersecurity.

This project led us to develop a new microworld environment for cybersecurity studies. The Interactive Defense Game - or IDG - is a simplified, game-like version of a computer network security task in which a human plays the role of a cyber analyst trying to prevent an attacker from progressing through a corporate network. It relies on a fully reconfigurable network simulation, allowing manipulation of experimental parameters such as network size, type and attacker strategies. My work then focused on setting up an experiment and collecting data from human participants facing different types of attack strategies. Main results showed consistency with the predictions of IBL models and that more aggressive but consistent attacks are more easily learned by defenders than slower but more insidious attacks.

As a next step in this project, I developed a team version of the task, to pair Human with IBL and other AI models and study their collaboration, shared mental models and trust dynamics.

In addition, I worked with Nhat and Coty on creating an IBL model to support a natural language-based item recommendation system for intelligence analysts. The goal is to provide the recommendation system with an estimate of the expected human decision about the recommended item in order to trigger adaptive actions or user interface adaptations.



FROM ERIN BUGBEE

My research in the DDMLab has focused on building computational cognitive models of how humans learn

and make sequential decisions from experience. A major milestone this year was completing and defending my second year research paper, “Deciding When to Stop: Cognitive Models of Sequential Decisions in Optimal Stopping Tasks.” I investigated how humans behaved in two optimal stopping tasks, in which participants were presented with sequences of alternatives and had to decide when to make a final selection. I developed an Instance-Based Learning model of both tasks and showed that the IBL model learns when to stop in alignment with the human participants. The manuscript is currently under review. Relatedly, I have been working on how to model decision making across different sequential decision tasks. One such task is the Balloon Analog Risk Task (BART), in which participants pump up a balloon. The value of the balloon increases with each pump, but there is some probability that the balloon will burst. The key decision is when to stop pumping and collect the current value of the balloon, measuring risk tendency. I developed an IBL model of the BART and displayed how this model inductively predicts pumping behavior, an advancement over existing models that rely on fitting participant data. I presented this work at CogSci 2022, which I attended in Toronto, and this was the first in-person conference I have attended in my PhD. This was also my first accepted first-author paper.

Leveraging the models I’ve mentioned, Chase McDonald, Coty, and I have been pursuing a line of research on applying wisdom of crowds techniques to model predictions. We propose that aggregating over individual model agents at each decision in a sequential decision task can lead to better performance than most individuals alone. We showed this in our 2022 International Conference on Cognitive Modeling paper, which I presented in July at Virtual MathPsych/ICCM. I look forward to pursuing this interesting line of work further.

Additional projects I have been involved with include running experiments using the boxgame task and improving ShinyIBL, an interactive web application built using the Shiny R package. In ShinyIBL, the user can directly adjust various simulation settings and model parameters and then observe how these changes affect the agent’s choices, helping to build intuition for how the model works. Finally, this past summer I was an Applied Scientist Intern at Amazon Science, where I taught machine learning courses and developed interactive articles to explain machine learning concepts. You can find one of the articles I wrote and developed here: <https://mlu-explain.github.io/logistic-regression/>. I look forward to another incredible year in the DDMLab as I work towards proposing my dissertation.


Research Updates Continued

FROM CHASE McDONALD

In the last year, I've worked on and presented several projects that have covered a variety of topics. On the Science of Understanding project with the Air Force Research Lab, we carried out human experiments that tested the hypotheses from our previous simulation experiments. Our prior work suggested that experiential diversity improved the ability of agents to adapt to novel situations, and this result held true in our empirical experiments. In addition to this extension, I presented joint work with Coty and Ngoc on multi-agent specialization and coordination at the 2021 AAI Fall Symposium Series. This work had been previously presented at the 2021 ACM Collective Intelligence conference.

There were several lines of new work this year. In particular, I was part of a project led by Erin Bugbee to investigate the use of cognitive models for wisdom of crowds aggregation in sequential decision-making tasks. We demonstrated the efficacy of modeling crowd decisions with Instance-Based Learning Theory, and presented this work at the 2022 International Conference on Cognitive Modeling. A further milestone last year was the completion of my second-year research paper, a requirement for the doctoral program. My paper, Sufficient Social Reasoning to Learn Productive Cycles of Cooperation, investigated how cooperative cycles can emerge in social groups with diverse needs and capabilities. Specifically, this involved manipulations to the social reasoning capabilities and preferences of agents in a simulated trading game. The results emphasized the importance of other-regarding preferences and causal attribution in learning jointly cooperative strategies. Lastly, I spent the summer as an intern at Riot Games, where I conducted applied research in deep reinforcement learning for games.

I will continue work started with Coty and Ngoc, alongside collaborators in the Robotics Institute, Henny Admoni and Michelle Zhao. Our ongoing project has investigated human-AI interaction in cooperative tasks. It asks how we can design an AI agent to collaborate effectively with human partners by varying intrinsic motivations, training regimes, or cognitive reasoning abilities, as well as how we can quantify an agent's collaborative capacity. This is aligned with my planned dissertation work, where I will investigate mechanisms for developing autonomous partners and the kinds of objectives and training regimes that allow them to support human learning and complement human action in reaching known or unknown human goals.



FROM YINUO DU

This past year was my first year as a PhD student in the DDMLab. My research has centered around cognitive modeling in the context of cyber security.

Continuing the collaboration with Palvi and Kuldeep at UTEP, I have been working on developing an Instance-Based Learning model for a cyber threat information sharing problem.

The problem is framed as a multi-player prisoner's dilemma to capture the benefit of information sharing and risk of leaking private information resulting in a loss. I gave a talk on this work at the Society for Mathematical Psychology conference and am currently drafting a journal submission. Moving forward, I'm interested in exploring how a human weighs the prediction and outcome of himself and other cooperative or competitive players in the game.

In addition, I have been working on the CybORG game with Baptiste and Coty. In this game, participants take the role of attacker or defender and interact with a simulated cyber scenario through abstract tactical actions. We have studied how people make defensive decisions and have built an IBL model of defenders and attackers. The modeling work has been accepted by HFES (defender) and HICSS (attacker, conditionally). Baptiste and I are preparing for human experiments. To extend this work, we are designing and developing a team-version of this task to enable the interaction between a human-ai defender and an automated attacker.

Finally, in a collaboration with Cylab, I have been working on cyber deception evaluation. The aim is to design and develop a platform that can empirically evaluate the effectiveness of cyber deceptive techniques and strategies informed by various methods for example game-theoretic or reinforcement-learning models.

Research Updates Continued

FROM TYLOR MALLOY

As a member of the Dynamic Decision Making lab I hope to explore how humans quickly learn new tasks that require decision making under risk and uncertainty. My previous research in human-like artificial intelligence drew inspiration from limitations in human cognition that impact generalization and robustness in areas like multiplayer games. The other main area of my research experience is in cognitive modeling, the focus of my PhD thesis, which explored how cognitive limitations in visual perception impact behavior. I hope to apply similar techniques onto contexts like cybersecurity where humans have biases from previous experiences that shape how they represent tasks, assess risks, and resolve uncertainty.

FROM DON MORRISON

We've acquired a new, powerful server, Janus, for the lab. This is being used to host online demos, run online experiments, and as a powerful compute server for big simulations. All our existing experiments demos have now been moved here, and some shared infrastructure for running online experiments and demos have been enabled. Several new demos and experiments have been created and deployed on Janus already.

Work is currently ongoing on PyIBL to speed it up, using lessons learned from SpeedyIBL. In addition work is also ongoing to make the PyIBL API more flexible to allow a greater variety of tasks to be modeled in PyIBL.

FROM JEFFREY FLAGG

As last year's newsletter was being published, we were in the final stages of completing our move from Winthrop to Porter Hall. We are now well established in our new space and are happy to be fully back on campus. Another large transition occurred with CMU's official migration to Google's suite of tools. The migration to Gmail and GDrive was largely smooth, it did cause some lingering issues with our YouTube channel (please click and subscribe!). However, these issues are now resolved and we are enjoying the advantages of this migration.

I have continued assisting on various experiments, helping test materials, reviewing papers, maintaining our website and social media materials, and integrating new members. I have had the pleasure to help review several paper submissions for DDMLab members. I also have continued to serve on CMU's OSF Advisory

Board which helps promote best practices in the greater CMU community. Recently, we have restarted our attempt to study individual differences in Theory of Mind processing. In the coming months, we will run studies to identify important ToM components and develop better assessment tools for future research.

FROM ERIN MCCORMICK

This past year, I officially started as a Research Psychologist for the Air Force Research Laboratory (AFRL), while continuing to conduct multiple projects with members of the DDMLab, both from my dissertation work and for the cooperative research agreement between CMU and AFRL.

Coty, Sam Cheyette (former DDMLab member and current graduate student at UC Berkeley) and I had a manuscript accepted for publication in *Memory and Cognition*. This project investigates human adaptation to changes in the probabilities and payoff values of choice outcomes, in a repeated binary choice from experience task. We found that successful adaptation in this task is influenced by an understudied factor in dynamic choice environments: the direction, or trend, of the change. Specifically, a decreasing trend--where an initially favorable option worsens over time--lead to better adaptation than an increasing trend--where an initially unfavorable option improves over time--regardless of feedback. This empirical pattern raises the issue of attention's influence on choice, and this research project contributes to our understanding of how decisions from experience are made in dynamic environments. This line of work has also lead to multiple additional experiments, currently in use for the Science of Understanding project with Coty and Dr. Leslie Blaha

The Science of Understanding project includes several manuscripts with Coty and Dr. Leslie Blaha, a Senior Research Psychologist at the Air Force Research Laboratory, currently embedded at Carnegie Mellon University. These manuscripts continue our work on using Recurrence Quantification Analysis to characterize human decision making in dynamic decision contexts, both from a behavioral perspective and from a cognitive modeling perspective using the Instance-Based Learning model. DDMLab also kindly hosted a recurrence quantification analysis tutorial Leslie and I had developed, to teach researchers about the practical interpretation of recurrence plots and the general use of recurrence quantification analysis for decision making research. The tutorial was well-attended by both lab members and other research collaborators.

Research Updates Continued

FROM ED MATLACK

I worked on developing datasets related to our spear phishing research. This included close collaboration with former lab member Prashanth Rajivan and his team at the University of Washington. This work focused on understanding qualitative predictors within spear phishing messages including error categorization and weighting of personal data. I plan to continue my collaboration with Prashanth while also pursuing new independent research in the measurement of meaning in decision making and memory retention.

FROM MISO DEMKO

I joined the DDM Lab during fall 2021 tasked with leading the lab's social accounts and with an opportunity to learn more about dynamic decision theory under Prof. Gonzalez's guidance, by studying and replicating an existing experiment. I first conducted a social media audit and identified a number of priorities to improve on - frequency of engagement, bridging scientific language barriers to reach different groups of our online community, and communicating life of the lab more holistically. During this process, I was exploring science communication techniques, strategies, and theories.

Prof. Gonzalez shared an opportunity with me to get involved with a workshop on science communication, which has led to a collaboration and findings beyond my time at the DDM Lab. I am deeply appreciative of Prof. Gonzalez's support in letting me pursue an exploratory approach to science communication beyond the immediate needs of the lab, and for Jeffrey's and a number of the lab members' efforts in helping me to experiment with creating different forms of science communication.

FROM TONY XI

During the past year, I've worked with Yinuo and Baptiste on the CybORG project. We examined the CybORG challenge and developed an IBL model that acts as an autonomous defender in the challenge. I helped to develop an interactive web-based game that incorporated our IBL defender model. I also implemented an autonomous attacker model based on IBL and performed experiments on the performance against heuristic and other IBL agents. This year I'll continue to be affiliated with the project and help on different tasks.



Still haven't seen our new lab space on campus? Visit us any time at 223-G Porter Hall!

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>

Some Recent Publications Highlights:

Gonzalez, C. (2022). Learning and Dynamic Decision Making. *Topics in Cognitive Science. TopiCS*. Vol. 14, Issue 1, pp. 14-30. <https://doi.org/10.1111/tops.12581>.

Gonzalez, C. & Aggarwal, P. (2022). Sequential Decisions from Sampling: Inductive Generation of Stopping Decisions Using Instance-Based Learning Theory. Chapter in *Sampling theories continue to inspire novel judgment and decision research*, Fiedler, K., Juslin, P. & Denrell, J. (Eds.). In Press. <https://doi.org/10.31234/osf.io/t4vmb>.

Aggarwal, P., Thakoor, O., Jabbari, S., Cranford, E. A., Lebiere, C., Tambe, M., & Gonzalez, C. (2022). Designing Effective Masking Strategies for Cyberdefense through Human Experimentation and Cognitive Models. *Computers & Security*. doi: <https://doi.org/10.1016/j.cose.2022.102671>.

Dugarte-Pena, G., Sanchez-Segura, M., Medina-Dominguez, F., Amescua, A. & Gonzalez, C. (2022). An Instance-Based Learning Simulation Model to Predict Knowledge Assets Evolution Involved in Potential Transformation Projects. *Knowledge Management Research and Practice*. <https://doi.org/10.1080/14778238.2022.2064348>.

Konstantinidis, E., Harman, J.L. & Gonzalez, C. (2022). Patterns of choice adaptation in dynamic risky environments. *Memory & Cognition*. <https://doi.org/10.3758/s13421-021-01244-4>.


McCormick, E., Cheyette, S., & Gonzalez, C. (2022). Choice adaptation to changing environments: trends, feedback, and observability of change. *Memory & Cognition*. <https://doi.org/10.3758/s13421-022-01313-2>.

Nguyen, N. T., Phan, N. D., & Gonzalez, C. (2022). SpeedyIBL: A Comprehensive, Precise, and Fast Implementation of Instance-Based Learning Theory. *Behavior Research Methods*. <https://doi.org/10.3758/s13428-022-01848-x>.

Zhang, H., Moisan, F., Aggarwal, P., Gonzalez, C. (2022) Truth-Telling in a Sender–Receiver Game: Social Value Orientation and Incentives. *Symmetry*. 14(8):1561. <https://doi.org/10.3390/sym14081561>.

Highlighted Projects and Events

This year, the DDMLab launched its own [YouTube channel](#). Conferences are often required recorded presentations, even if a speaker will be live as well. As a result, we felt it was the perfect time to begin a YouTube channel for the lab. We hope to continue to expand our content in the following years. Please feel free to watch some of our presentations from recent conferences!



<https://youtu.be/BUw334eQkI4>

What makes a good theory? Listen to Coty discuss this topic during a keynote presentation at the Lorentz Workshop “What makes a good theory? Interdisciplinary perspectives.” You can view other keynote talks [here](#).



Virtual MathPsych/ICCM 2022
mathpsych.org
July 2022

Modeling of multi-defender collaboration in a cybersecurity scenario

Yimin Du - Carnegie Mellon University

<https://youtu.be/reJLvb7H3WA>



Virtual MathPsych/ICCM 2022
mathpsych.org
July 2022

Leveraging cognitive models for the wisdom of crowds in sequential decision tasks

Erin Baghee - Carnegie Mellon University

<https://youtu.be/L74ebUjX0q0>

Ngoc Nguyen and Erin McCormick presented their work at the Virtual MathPsych/ICCM 2020 conference. Great job!

Adventures During 2021 - 2022



Top Left: DDMLab members at a fall festival for apple picking.

Top Right: New DDMLab shirts!




Left: Exploring the Bayernhof Museum of musical instruments.

Bottom: Welcoming Tyler Malloy for his CMU orientation.



Bottom: DDMLab members Baptiste Prebot and Jeffrey Flagg have art displayed at this year's Community Picnic.

Right: Science as art! Recurrence Quantification Analysis by Erin McCormick, Leslie Blaha, and Coty Gonzalez are on display outside of the lab. Not only does science enrich our knowledge but it also beautifies it as well.



Recurrence Quantification Analysis
Condition: Decreasing/Partial McCormick, Blaha, & Gonzalez (2021)



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
Phone: 412-268-9547
412-268-8439

THANK YOU FOR YOUR INTEREST IN OUR
2022 NEWSLETTER!

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

[FACEBOOK](#)



[LINKEDIN](#)



[TWITTER](#)



OR [YOUTUBE](#)

