Dynamic Decision Making Laboratory

Issue 13 Fall 2021

Carnegie Mellon University

Annual Summary from Coty — Founding Director, DDMLab

Inside this issue:

Dear Friends and Collaborators:

Many greetings from our Dynamic Decision Making Laboratory. We write this newsletter every year, to look back at our research activities from the past academic year and to get started in the new academic year (2021-2022). Thank you for letting us update you!

During last year we continued to be supported by several long-term grants and importantly this past year was a great success with obtaining new research grants on several important programs that will support our work in the coming years.

We are grateful for the support provided by our current research sponsors. This year, we concluded a grant from the **Army Research Office (ARO)**, Network Sciences program. This grant helped us in advancing new theories and cognitive models that expand from individual cognition to networks, teams, and groups. Thank you ARO for the support.

This year we also continued the work sponsored by the **ARO's Multidisciplinary University Research Initiative (MURI)** program on Cyberdeception. Our collaborations with researchers at Harvard University, University of Southern California, and University of Texas at El Paso, have helped advance new deception and masking techniques for cyberdefense. 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 *Adaptive Cyberdefense*. Our vision of using Instance-Based Learning Theory to advance human-autonomy for cyber operations is now conceptually possible, and we will bring additional demonstrations of this in future years.

This year, 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 scaled up our concept of Cognitive Machine Theory of Mind from simple gridworlds to more complex navigation tasks and team/group tasks. Thanks to the challenges brought up by this program, our IBL models have now advanced and scaled up to include mechanisms for credit assignment and to improve complex space representations.

Furthermore this year we obtained three new grants which we are about to start: (1) A grant from the **Air Force Research Laboratory (AFRL)** intends to improve learning science in human-machine team effectiveness; (2) A SBIR proposal from the **Air Force** on Trust Resilience in User-System Team Modeling (TRUST'M); and (3) and new US-Australia International Multidisciplinary University Research Initiative (**MURI-AUSMURI**) to develop the science behind robust Human-AI Cybersecurity Teams. Stay tuned for news and these initiatives in the coming years!

I am very grateful to all the post-docs and collaborators for the many publications that have emerged during this year. Please check all our new publications in our web site!

Invited lectures and plenary talks were all online, done remotely. I was invited to NYU's Neuroeconomics colloquium, to the Experimental Organization Science online seminar, and internationally, to give a talk at the Psychology department in BUAP-Mexico.

Most of our invited talks and all our papers at conferences were presented remotely. In conferences "Gather" or simply "Zoom" made it seem like we were connected and really attending the conferences.

Many of our talks were pre-recorded (see projects and events in this newsletter). Pre-recording has some advantages… but for me, Gather or Zoom are not a good substitute for the real presence. During this time, and now that I am back in "real" office in campus, I realize the power of eye contact, context and gestures in human-to-human communication. That physical presence makes life more enjoyable.

Teaching was all done remotely during Spring 2021 but it all went well. I look forward to teaching in person again in Spring of 2022.

I continue to enjoy biking, now every day to my office, while weather permits.

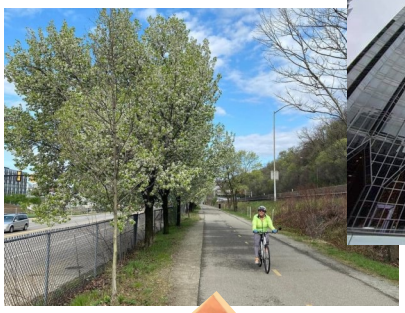

I also continue to paint… but not as often as before. My weekends have been filled with grant proposals, grant reports, papers, editorial reviews and much less with painting. I look forward to seeing all of you in the physical world some time soon.

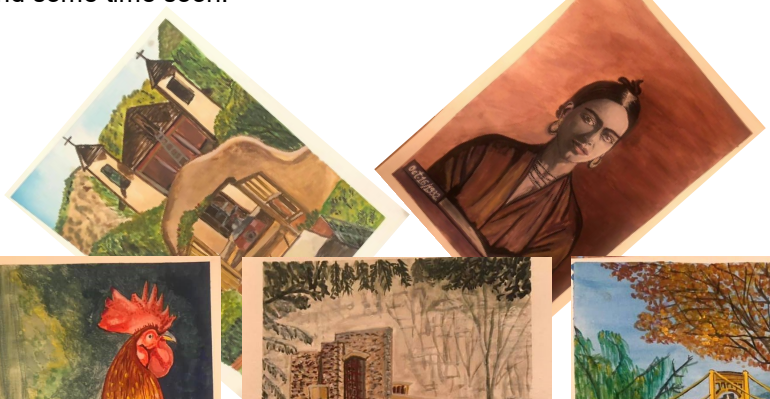

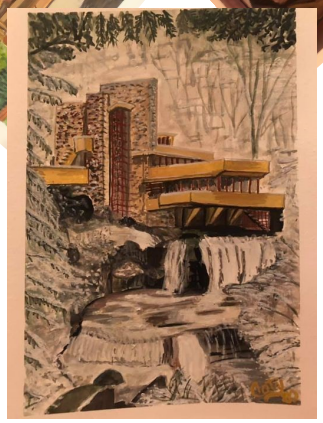

I hope you enjoy this newsletter. Thank you so much for reading our update, and thank you so much for contributing to our efforts in so many ways!! We look forward to a new academic year full of interesting research, surprises and new adventures in 2021-2022. ***Let's all work to make this a better world!!***

Coty

DDMLab comes back to Main Campus!!

One of the most exciting news this year is that, after 14 years of being "away" from campus (really just a close walking distance), our lab is now coming back to main campus and re-joining our colleagues at the Social and Decision Sciences Department in Porter Hall!

Our newly renovated suite is great and it has been finished right as the university is ready to re-start "in-person" activities. Our gratitude goes to Dean Richard Scheines, Linda Babcock, Gretchen Chapman, and to everyone "behind the scenes" that made our new lab space possible.

Here are a few pictures of our new space under construction and now that is finished. All of you are welcome to visit us! Our new address:

Dynamic Decision Making Laboratory
Carnegie Mellon University
5000 Forbes Ave.
Porter Hall 2223-G

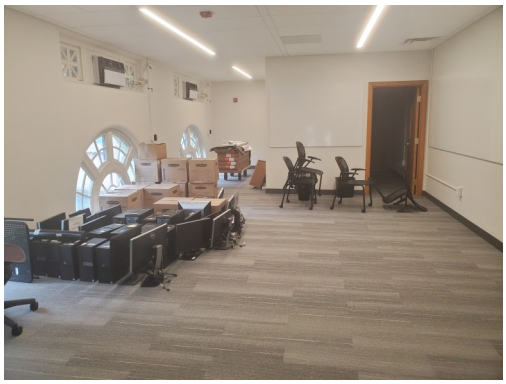

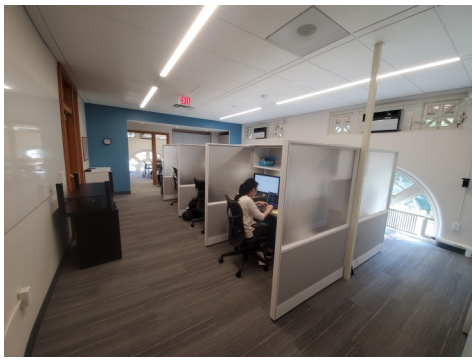

Farewell to our lab space in Winthrop Street Lab. We enjoyed our time there but we look forward to having more interactions with our colleagues in Social Decision Sciences and the Dietrich College.

News About Our Members

Farewells

We have to say good bye to several of our post-docs this year, including Palvi Aggarwal, Kuldeep Singh, and Farnaz Tehranchi. But it is so exciting that all of them moved on to pursue an academic career as Assistant Professors. Our loss will be their new institution's gain!

We are very proud and happy for all of you!! Good luck in your future endeavors!

Prof. Palvi Aggarwal

Prof. Kuldeep Singh

Palvi Aggarwal joined UTEP as Tenure Track Assistant Professor in Computer Science starting in Fall 2021. Kuldeep Singh also joined the Computer Science department at UTEP as an Assistant Teaching Professor. Both of them will continue and advance the work from the DDMLab on Cyberdeception and modeling human behavior in the context of cybersecurity.

Erin McCormick, PhD

Prof. Farnaz Tehranchi

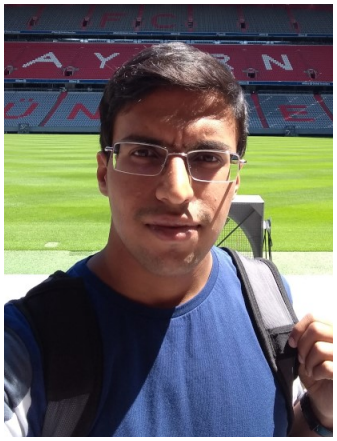

Aditya Gulati

Disha Das

Erin McCormick defended their dissertation successfully this year, but will remain with us for the upcoming year as a visiting scholar. She is in the final stages of formalizing her new position of Research Psychologist with the Air Force Research Laboratory.

Farnaz Tehranchi has started a Tenure Track Assistant Professor position in the School of Engineering Design, Technology and Professional Programs at Penn State. Congratulations Farnaz!

Aditya Gulati will be a new PhD student at the Institute of Human(ity)-centric AI which is a part of the University of Alicante, Spain.

Disha Das is currently senior completing her undergraduate degree at Carnegie Mellon University.

News About Our Members

Welcome New Members

Yinuo Du joined the lab last year as a master student and will be continuing this year as a Ph.D. student this semester in Societal Computing program. She will be involved with behavioral cybersecurity projects and will be developing Instance-Based-Learning models to capture the cooperative decisions made by groups of defenders in the Multi-Attacker-Defender game.

Baptiste Prebot is joining the lab as a post-doctoral fellow. He received his Ph.D. in Cognitive Engineering from the University of Bordeaux, France. As a former UX designer, his research interests are oriented towards the teaming of human and autonomous AI's in complex systems and how to facilitate the use of technology to optimize the performance in the Human-Autonomous AI teams.

Nhat Phan is a new post-doctoral researcher with the DDMLab. He received his Ph.D. in Applied Mathematics in 2016 at the University of Lorraine, France. His research interests include issues related to optimization applied in machine learning and decision making. Nhat performs research studies on developing Instance-based Learning models enhanced with reinforcement learning techniques in goal-seeking and collaboration tasks.

Ed Matlack will be a visiting scholar beginning in October. His research last year in the San Jose State Intelligent Conversational Agents and Neural Networks Lab focused on multi-category sentiment analysis of Covid-19 Twitter feeds and is currently involved in projects for the ICANN Lab for updating economic models in turn taking, intelligent agent based conversations.

Miso Demko is an undergraduate research student at Carnegie Mellon University. He will be assisting social media and research communication for our findings as well as giving support to our Human/AI projects.

Tony Xi is a is a first-year master student at the Information Technology Institute (INI) at Carnegie Mellon University. At the DDMLab, he will be working on the CybORG challenge, an autonomous attacker-defender game, using an Instance Based Learning implementation.

Research Updates from Lab Members During 2020-2021

From Palvi Aggarwal

It has been a wonderful year at Dynamic Decision Making Lab. I continued working on projects focused on cybersecurity and decision making with internal and external collaborators. With the support of Coty, other lab members and our collaborators, I also got success in getting a new a tenure track Assistant Professor position at the University of Texas, El Paso.

Last year, in collaboration with the Harvard University and University of Southern California, we developed a Risk-based Cyber Camouflage Games by assuming a risk-averse attacker and creating Prospect Theory based masking strategy of defense. This work was published in Gamesec-2020 conference and won the best paper award. Despite a lot of challenges due to covid, we had success in evaluating the masking algorithms using the CyberVAN testbed against human participants. I have also developed Instance-Based Learning Models to capture the human decisions in such complex scenarios. Currently, this work is under reviewed in a journal. I presented this work in the annual CyLab Partners Conference in September-2020 and in a workshop on Adversarial Cybersecurity, IIT Mandi-India. In another project, we conducted experiments using the HackIT tool to evaluate the effectiveness of 2-sided deception. Our work has been published in the Adaptive Cyber Deception workshop at IJCAI conference this year. Our work on signaling strategies has been accepted in Cognitive Science journal.

Out of the cybersecurity projects, I also worked with Coty on a project that focuses on an important question of when to stop sampling. This work is accepted as a book chapter for publication this year. I also worked with Don Morrison and Coty on conducting team human experiments and developing cognitive models.

With three wonderful years with DDM Lab, I am now moving to UTEP to continue working on behavioral cybersecurity. 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.

From Ngoc Nguyen

This past year, we published work on Machine Theory of Mind in *Topics in Cognitive Science*. Also, I have continued to work on the human machine teaming project that is focused on addressing the key question of how machines can team up with humans to perform a task in an effective way. To gain an understanding of how humans players perform sequential decision tasks where feedback is delayed, I have developed different interactive browser-based applications motivated by goal-seeking scenario called Gridworld and One-step Gridworld. Based on Instance-Based Learning Theory (IBLT), we built IBL models of a player with different mechanisms of credit assignment, including a variant of anIBL model enhanced with temporal difference (IBL-TD). When comparing these models against collected human data in a goal-seeking task, we found that there is no single model that can accurately represent humans' performance with respect to all considered metrics. The analyses of humans' and models' data in terms of process measures shed light on the differences between the behavior of humans and models, and how to create AI systems that emulate humans. A manuscript is currently in preparation.

During the past year, I also developed an interactive application based on search and rescue task to study humans' behavior in more complex and richer environment, called Minimap. We have run a few human experiments using Minimap. First, we manipulated levels of visibility and environmental complexity to understand how these variables affect humans' learning to make sequential decisions. Also with our DARPA-ASIST team, I have extended the Minimap and ran a study to gain insight into the role that anger in shaping individual performance and collaborative behavior.

We have also recently started to work on multi-agent systems. We used IBL-TD in a coordinated multi-agent object transportation problem (CMOTP), a fully cooperative multi-agent learning environment. Experimental results show IBL-TD outperforms state-of-the-art Deep Reinforcement Learning (DRL) models. This work was presented at the *ACM Collective Intelligence* conference 2021. Currently, we are working on an experimental design for team-level experiments to test the insights obtained from our multi-agent models.

Research Updates Continued

FROM KULDEEP SINGH

My research in the past year mostly focused on the designing of an anti-phishing training program. We ran multiple human experiments by manipulating cognitive parameters, trying to find the impact of these cognitive factors on human decision while encountering phishing emails. We analysed collected data and found that the frequency of phishing emails and detailed feedback have impact on human participants performance. Whereas, penalizing for incorrect classification does not help in participant's performance. We also did the email level analysis to find out why human are good in detecting some phishing emails compared to others. The results were published at the HFES 2020 conference. Furthermore, we are now moving in direction of adaptive training, currently we have designed the new experiment and ready to run the study with human participants. We have also worked with our collaborators, Christian and Drew Cranford, to develop a cognitive model for phishing. The work is presented at ICCM 2021 conference.

I also worked on a multi defender game. We ran a human study in a group information sharing game, where multiple defenders collaborate and share threat information with each other. The major objective of this study is to find out the cognitive factor, which can motivate humans to share threat information in a cyber-security closed group. We analysed the collected data and found that defenders (participants) majorly follow the reciprocal strategy and share information only when they receive information from other defenders in the group. We also found that security investment has an effect on information sharing in the group. We saw that defenders tried to build a trust during the game and share more with one other defender in the group. These results were published in the *ACM Collective Intelligence* 2021 conference.

FROM ERIN MCCORMICK

This past year, I defended and deposited my dissertation, in addition to continuing work on multiple projects with the Lab. My dissertation covered both a manuscript with Coty and Sam Cheyette, former DDMLab member and current graduate student at UC Berkeley, and several experiments on how time constraints affect decision making, with Coty and Stephen Broomell.

The manuscript received encouraging and constructive reviews, and we have been invited to resubmit. 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.

The investigation of how time constraints affect decision making consists of two experiments, which together test whether specific changes in the decision process should and do allow decision makers to adapt to time constraints. Previous literature suggested the weighted additive decision strategy is too complicated and effortful for decision makers to use under severe time constraint, despite its standing as the normatively optimal decision strategy for multiple-alternative, multiple-attribute choice tasks. However, we found evidence for decision making as if participants were using this strategy, even under severe time constraints. While the assessment of these results is ongoing, this evidence suggests that previous studies may have used an experimental paradigm that constrained decision making strategies, and that the contextual and internal cues that lead to systematic deviations from normal decision making under time constraints merit further research, to understand the true strengths and limitations of human decision making.

Additionally, I continued working on the Science of Understanding projects with Coty and Leslie Blaha. We conducted several experiments to test human adaptation to change within-subjects and for different types of change, and to create data to use Recurrence Quantification Analysis to assess variation in human decision making and compare it to the Instance-Based Learning model.

Research Updates Continued

From Erin Bugbee

This past year was my first year as a PhD student in the DDMLab. My research has centered around learning from experience and sequential decision making. In a collaboration between CMU and the Air Force Research Laboratory, I have been working on developing an Instance-Based Learning model for the optimal stopping problem. The optimal stopping problem is a sequential decision making problem involving choosing when to take an action to maximize expected reward, and some examples include deciding when to select an apartment or choosing when to purchase a flight ticket when the price changes dynamically. I gave a talk on this work at the Society for Mathematical Psychology conference and am currently drafting a paper. Moving forward, I plan to continue studying how people make these types of sequential decisions, how they may deviate from what is theoretically optimal, and how they learn to approximate their choices to what is theoretically optimal with experience.

In addition, I have been working on the treasure hunting game, or the boxgame, with Coty, Palvi, and Nhat. In this game, participants choose between two boxes that may be protected by a computer bot. Participants receive information about the protection status of the boxes, but these signals may be truthful or deceptive. We have studied how people make decisions from experience when receiving potentially deceptive information and have developed an IBL model that makes choices resembling those of humans in this task, as well as an IBL model that can dynamically provide signals to influence the choices of the decision maker. To extend this work, I am developing an interactive version of the boxgame to be used in future experiments. Finally, Coty and I are writing a review of how decision biases emerge from experience.

From Chase McDonald

After joining the lab last year, I've been involved in a number of projects in the DDMLab. Ngoc, Coty, and I have pursued several lines of inquiry, the first of which was a part of the ACM Collective Intelligence conference earlier this year. In that work, we explored the environmental factors that facilitate coordinated behavior in a multi-agent goal-seeking task. In particular, we showed how inter-agent observability, reward distribution, and relative spawn positions impacted the emergence of coordinated behavior among IBLT agents. Additionally, we have an ongoing project, which we are currently preparing for journal submission, that explores mechanisms for addressing the temporal credit assignment problem in IBLT; that is, how to assign credit to actions with delayed rewards. I have also been a part of the Science of Understanding project with AFRL. We presented our work exploring the diversity hypothesis—that diverse experience leads to improved ability to adapt to novel tasks—at ICCM. In this work, we demonstrated how exposing IBLT agents to varying levels of decision complexity affected their ability to adapt to a novel level of decision complexity.

There are several projects that I am working on now. Our work on the diversity hypothesis will be extended to test whether the predictions from our simulation extend to human participants. We will ask if varying exposure to decision complexity impacts humans' ability to adapt to distinct levels of decision complexity. Additionally, we will continue to work on multi-agent systems and investigate the credit assignment problem from a structural perspective. Specifically, we will explore mechanisms that allow agents to determine who was responsible for an outcome and attribute credit accordingly. These mechanisms for making responsibility judgments will then be used to both improve learning capabilities in multi-agent settings and facilitate cooperative behavior.

From Aditya Gulati

Over the past year, I have been working with Coty and Ngoc on understanding the impact of group size and task complexity on groups where members cannot communicate with each other. For this, we have been using the minimap task - a search and rescue task where groups need to rescue victims scattered around the map. We built an IBL agent for this task and analysed the data collected from simulations run with these agents. We looked at the importance of structural complexity on predicting performance and how different levels of complexity impact a group's ability to find victims on the map. In addition, we have also studied the impact of groups on the performance of different members in the group. Our findings were presented as a part of my master's thesis at IIIT Bangalore which I successfully defended earlier this year.

Research Updates Continued

FROM FARNAZ TEHRANCHI

This past year, I worked on two projects. The first one was the Rock, Paper, and Scissors study. Hanshu initially conducted this study, and she investigated individual behaviors. I studied pairs' performance and behaviors to investigate exploiting and exploitable behaviors in the Rock, Paper, and Scissors Repeated Choice game. We want to answer research questions such as how players may learn to exploit others' behavior or learn to be exploited by others. Also, do players make choices according to the probabilities of the opponent's past moves, or do they rely on their past moves? We first look at the paired results and compare them with random and Nash models to answer these questions. Then, we develop IBL models using PyIBL with different attributes.

The second project is the Capture-the-Flag study. I am using GridWorld that Nogc and Chase were also working on. In my experiment, I have two agents and two goals or flags. An agent tries to find their own goal or flag. To develop this scenario, I used the multi-agent version of GridWorld with only two agents. As soon as one of the agents reaches a goal, the trial stopped. This stop condition simulates competition between agents similar to the Capture-the-Flag game scenario.

FROM DON MORRISON

PyIBL continues to evolve, the most recently released version being 4.2. Work this past year has focused on extensions to the PyIBL API to enable its use in ways outside it's "sweet spot," most notably support for allowing alternative decision procedures still within the IBL paradigm. There have also been improvements to its runtime performance.

In exploring PyIBL's use in several networked experiments we stumbled over a simple binary choice task that, when modeled in a network of users, produced a counter-intuitive result. We believe we understand why and have built and run a networked human experiment with nodeGame and using participants from Amazon Mechanical Turk, to explore whether or not humans make the same initially counter-intuitive choices.

We have also built an online version of the moments@work game, now dubbed moments@cmu, for use as a tool in fostering more sensitive and inclusive work environment.

FROM YINUO DU

In the past year, I've been collaborating with Dr. Palvi Aggarwal and Dr. Kuldeep Singh on a project to research on the effectiveness of two-sided deception. After identifying various features on real and honeypot nodes that could be used for effective deception, we designed and conducted human-subject experiments on HackIT. The initial results were recently published in the IJCAI-ACD workshop. In the future, I'll expand my exploration on cyber simulation platforms to CyberVAN and CybOrg. This summer, I've been developing Instance-Based-Learning models to capture the cooperative decisions made by groups of defenders in the Multi-Attacker-Defender game. I'll be moving forward to explore more methods to help promote cooperation among defenders.

FROM JEFFREY FLAGG

Recently I have been very active getting the DDMLab ready to move from Winthop to our new space in Porter Hall. With any move, you don't really start to realize how much stuff you have until you have to start putting it in boxes! There is still a lot to do, but we are almost there, and by the time this newsletter is published, we should be fully back! 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 attended several seminars conducted by the CMU Library team on various professional development topics. I have also worked to help the DDMLab fully utilize the Open Science Framework. We are now preregistering all projects and should only continue to improve our documentation in the future. I am also serving on CMU's OSF Advisory Board and am hoping to use the knowledge gained from our lab's use of OSF to assist the greater CMU community.

FROM MYCHAJLO DEMKO

I am a returning fifth-year undergraduate in the Dietrich College of Humanities, joining the DDM Lab with prior undergraduate course experience in dynamic decisions and disaster analysis. I will be contributing to a project involving Human-AI teams by reviewing the pertinent Human Factors literature in assessment of emergency response scenarios. I am interested in supporting learning in order to enhance networks of collaboration, co-design, and innovation.

Research Updates Continued

FROM NHAT PHAN

I joined the DDMLab in March this year as a visiting scholar. During the past 6 months, I have worked on several projects. In the first project is on multi-agent systems. We developed different IBL models combined with reinforcement learning (RL) techniques for coordinated multi-agent object transportation problems (CMOTPs). Through numerical experiments, we demonstrated that our models outperform existing related work in RL. The results were published in the ACM Collective Intelligence 2021. We are extending this work to submit it to a journal. In another project, we have developed a new version of Python library for creating IBL agents called SpeedyIBL that makes significant improvements in computation time of running an IBL model. We demonstrated the efficiency of SpeedyIBL on a wide range of tasks from single-agent single-state, single-agent multi-state to multi-agent multi-state environments. We are writing up a manuscript reporting the implementation and benefits of SpeedyIBL. In addition, I have been involved in defensive cyber-deception project where I implemented IBL observer as a defender. Recently, I have worked on a phishing project that is focused on training human participants to help them learn to distinguish phishing emails from ham emails. In this project, I have implemented the adaptive and non-adaptive training. We will be conducting human experiments to test our models.

FROM BAPTISTE PREBOT

As a new post-doctoral fellow in the Dynamic Decision Making Lab, I will be studying human behavior and the integration of human machine teams in cyber security contexts. I previously received my Ph.D. in Cognitive Engineering from the University of Bordeaux, France, working on the real-time evaluation of Team Cognition in military Command and Control (C2) using psychophysiological and behavioral monitoring. Thus, the intended work at the DDMLab connects with my previous research and more general research interests on Cognitive Warfare, Team Cognition and Human consideration in building AI teammates. I look forward to work with DDM Lab tools and theories to develop cognitive computational models of human and teams decision making in this context. I am quite excited by the opportunity to participate in the scientific life of the lab.

FROM ED MATLACK

My research last year in the San Jose State Intelligent Conversational Agents and Neural Networks Lab (ICANN) focused on multi-category sentiment analysis of Covid-19 Twitter feeds, where we worked on optimizing conversational agent response by matching sentiment in tone. I am currently involved in projects for the ICANN Lab for updating economic models in turn taking, intelligent agent based conversations. My first project in the DDMLab will be working on understanding how end-users learn to detect phishing emails. This is my first practical research in the area of behavioral cyber-security, which I hope to expand into other areas of research including how information availability, quantity and diversity affect decision making over time. Coming from an Information Science and History background, I am interested in the intersection of technology driven information growth and decision making over time. Specifically I'd like to use IBLT to model how diminishing information half-life makes decision making potentially more exploitable.

FROM TONY XI

I am first-year master student at the Information Technology Institute (INI) in the Information Networking program. My research interests involve cognitive science and its applications in the area of cybersecurity. At the DDMLab, I will be working on the CybORG challenge, an autonomous attacker-defender game, using an Instance Based Learning implementation.

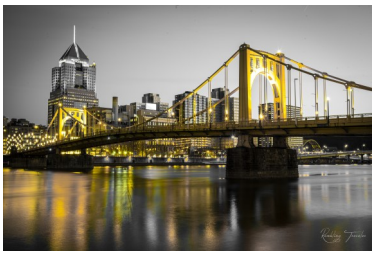

Photos by Kuldeep Singh.

Recent Publications

In the past year, we published several journal articles and articles in conference proceedings 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:

Cranford, E. A., Gonzalez, C., Aggarwal, P., Tambe, M., Cooney, S., & Lebiere, C. (2021). Towards a Cognitive Theory of Cyber Deception. *Cognitive Science*. 45(7), e13013. https://doi.org/10.1111/cogs.13013.

Harman, J. L., Yu, M., Konstantinidis, E., & Gonzalez, C. (2021, March 20). How to Use a Multi-Criteria Comparison Procedure to Improve Modeling Competitions. *Psychological Review*. In press. https://doi.org/10.1037/rev0000274.

Nguyen, T. N., & Gonzalez, C. (2021). Theory of Mind From Observation in Cognitive Models and Humans. *Topics in Cognitive Science*. In press. https://doi-org.cmu.idm.oclc.org/10.1111/tops.12553.

Zhang, H., Moisan, F., & Gonzalez, C. (2021). Rock-Paper-Scissors Play: Beyond the Win-Stay/Lose-Change Strategy. *Games*. 12(3):52. https://doi.org/10.3390/g12030052.

Aggarwal, P., Thakoor, O., Mate, A., Tambe, M., Cranford, E. A., Lebiere, C. & Gonzalez, C. (2020). An Exploratory Study of a Masking Strategy of Cyberdeception Using CyberVAN. 64th International Annual Meeting of the Human Factors and Ergonomics Society (HFES 2020). October 5-9, 2020. Chicago, Il.

Highlighted Projects and Events

This year, the DDMLab launched it's 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!

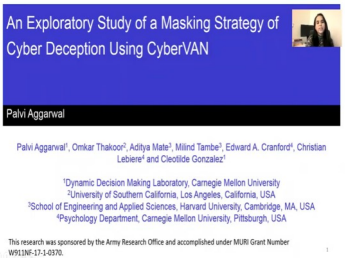

https://youtu.be/_gOfInJAaT4

Palvi Aggarwal and Kuldeep Singh presented work at the 2020 Human Factors and Ergonomics Society.

https://youtu.be/ClL6m0TvdMU

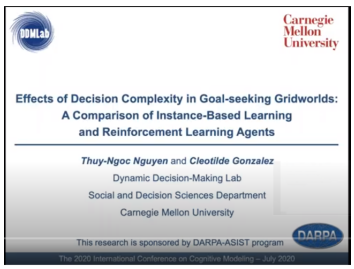

https://youtu.be/CWhjwlEz_hw

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

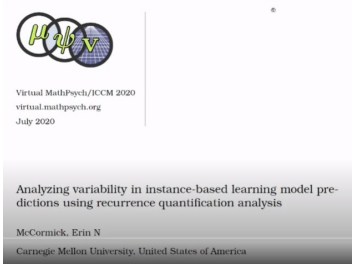

https://youtu.be/YEQn7l2T30Q

Best Paper Award

DDMLab members and collaborators including Palvi Aggarwal and Coty Gonzalez of Carnegie Mellon University, Omkar Thakoor and Phebe Vayanos of the University of Southern California, and Shahin Jabbari and Milind Tambe of Harvard University were awarded the Best Paper Award at GameSec 2020!

Adventures During 2020 - 2021

Left: Finally able to meet up after a long year apart!

Below: Welcoming Thu Phan to the world. Congratulations Ngoc and Nhat!

Above: DDMLab bike ride from the South Side to Point State Park. We got to meet Thu for the first time.

Right: Making gingerbread houses during the 2020 holiday season!

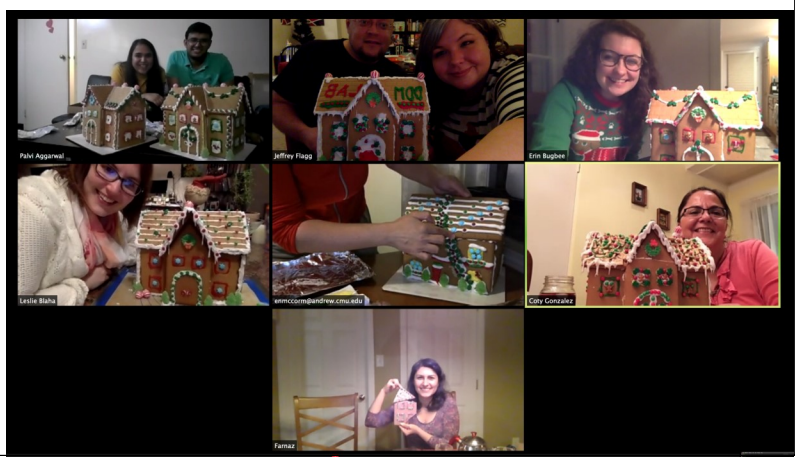

Dynamic Decision Making Laboratory

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 2021 NEWSLETTER!

KEEP UP WITH US ON OUR WEBPAGE OR FOLLOW US ON FACEBOOK LINKEDIN TWITTER OR YOUTUBE