



Dynamic Decision Making Laboratory

ISSUE 12

FALL 2020

Carnegie Mellon University

Annual Summary from Coty — Founding Director, DDMLab

Dear Friends and Collaborators:

First, and foremost, I hope all of you and all your loved ones are doing well. We are living unprecedented times, unexpected situations, complex dynamic decision making challenges. I hope all of you are handling that well. We are also doing our best.

We write this newsletter every year, to reflect on the research activities of the DDMLab during the past academic year and as we start a new academic year (2020-2021). Thank you for letting us update you!

We are grateful for the support provided by our current research sponsors. This year, we continued the work sponsored by the **Army Research Office (ARO)**, Network Sciences program. This grant has helped us in developing new algorithms and cognitive models that expand our understanding of how network behavior emerges from individual learning, and how networks form from individual interactions. We also have 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 which we are testing in complex and naturalistic settings. We also continue to be part of an exciting long-term research program with the **Army Research Laboratories (ARL)**. This program on cybersecurity, is a Collaborative Research Alliance (CRA) that allows us to interact with many other universities and organizations. We continue to maintain an active research program on socio-cognitive aspects of cybersecurity including: decision making on defenders, attackers, and end-users, and developing cognitive models that can capture those processes. Also the **Defense Advanced Research Projects Agency (DARPA)**'s **SocialSim** program helped us to collaborate with Christian Lebiere in the development of cognitive models that capture phenomena in large scale data sets such as Twitter. Finally, this year has been exciting as we became part of the new **DARPA** program: **Artificial Social Intelligence for Successful Teams (ASIST)**. We have started our work of developing a Cognitive Machine Theory of Mind framework, using Gridworld navigation tasks. Towards understanding Human-machine collaboration, we have created models that are able to predict the state of mind and the intentions of players in Gridworlds. This research has already produced two publications and will continue to be a very productive program, given the large potential and amazing collaborations with the DARPA's group.

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, for obvious reasons, were limited this year. Early in the Spring semester of 2020, I was part of a small but exciting workshop on *Learning in Networks*, at NYU, where I saw many good friends with whom we were able to discuss ideas related to learning of networks as emerging from individuals.

In the Fall of 2019 there were also two very interesting and exciting invited meetings.

Inside this issue:


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

On September 9-13, 2019 I attended the *Collective Intelligence: Theories and Applications* summer school in Bari, Italy, where I gave a talk about our work regarding the *Emergence of Collective Behavior in Networks from Individual Choices*. This was an amazing trip where I made new connections and enjoyed the talks of many good people. I have to thank my hosts, Ilaria Giannoccaro and Giuseppe Carbone, at Politecnico di Bari. Both were amazing hosts and I will never forget the old town of Bari and the food!! The best food I have tasted in my life. This is a place I would like to visit again.




I also gave a keynote talk at the workshop on Computational Cybersecurity in Compromised Environments (C3E) workshop on September 15-17, 2019, at Stanford, Menlo Park, CA. We presented the work on developing dynamic and personalized algorithms for cyberdeception and signaling.

Teaching during Spring 2020, was a great experience. But, as you all know, the regular face-to-face classes had to be shifted to online teaching in early March. One good thing about the whole situation is that I was able to introduce the Epidemics S-I-R System Dynamics model at the end of the course, and then, there was no better practical example for that model than the Corona Virus example made available by the System Dynamics Community. Thus, making the best of the situation, the students were able to link the concepts of system dynamics modeling with a very real, practical example.



Ending this newsletter with a personal note is usually fun, but unfortunately not this year. My mother passed away in May. A woman that lived for her children and gave all her time, love, life, to us. I miss her, but I am trying to follow her example of positivism and strength.

I continue to enjoy biking and painting on my weekends. These pictures show some of my favorite paintings from this year and some of my favorite biking adventures.




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 2020-2021. **Let's all work to make this a better world!!**

2020 *Coty*

News About Our Members

Farewells

Many good people left our lab during this year. Hanshu Zhang and Korosh Mahmoodi ended their positions as a Post-doctoral fellows this summer. Their positivity, creativity, and thoughtfulness will be greatly missed. We wish them the best in all their upcoming adventures.



Undergraduate students Max Gamerman, Peijie He, Andrew Ye, and Max Yeh joined the DDMLab for the Fall and some also for the Spring semester and were important contributors to our lab. They will no doubt continue to have a significant impact at CMU and beyond.



Welcome New Members



Erin Bugbee (left) is a new doctoral student interested in creating cognitive models of human learning and decision making to improve human decisions and artificial intelligence.

Chase McDonald (middle) is a new doctoral student interested in including both empirical and simulated multi-agent systems to social decision-making, cognition, and intelligence.

Farnaz Tehranchi (right) is beginning post-doc position and will be working on issues involving deception/signaling and human-machine teaming.

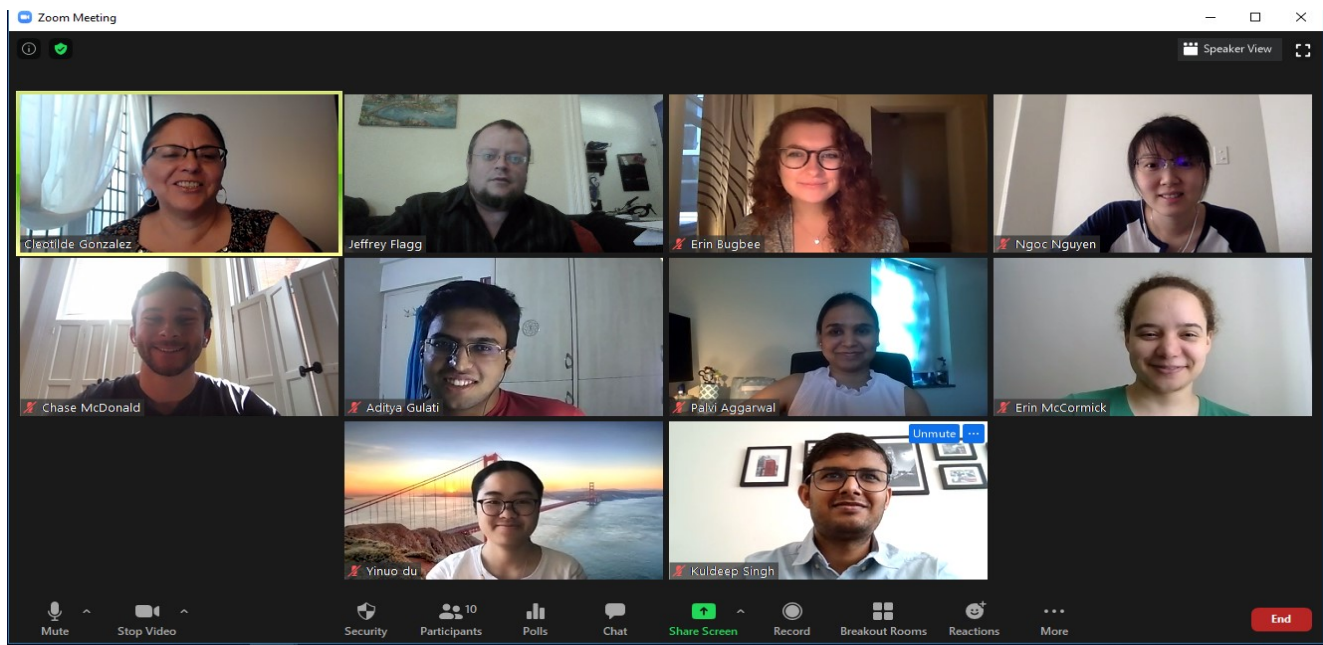
Aditya Gulati (left) is a final year masters student and is currently working on human-machine teaming. He is interested in decision making and data science.

Yinuo Du (middle) is a first year master student, she is interested in the area of Behavioral Cybersecurity and will work to explore the effectiveness of honeypots.

Disha Das (right) is an undergraduate research student in computer ad decision science and will be supporting our cyber-deception projects.



Research Updates from Lab Members During 2019-2020



FROM PALVI AGGARWAL

Last year, I started deploying the feature masking algorithms on a real-time testbed for cybersecurity research called CyberVAN. We conducted an exploratory experiment to test the effectiveness of a masking strategy to reduce the utility of attackers. I am happy to announce that our initial results are accepted for publication in the HFES-2020. I also presented this work in the annual CyLab Partners Conference in September-2019. We found a generalized tendency in human attackers to act according to the certainty bias (or risk aversion). Motivated from these findings, with the help of our collaborators, we developed new algorithms of masking and started the data collection. I have also been developing Instance Based Learning Models to capture the human decisions in such complex scenarios. Using the concept of cyber deception, one of our projects was evaluating signaling to create uncertainty in attacker's decision making. We continued our research in this direction and published papers in conferences such as HICSS-2020 (Best Paper Award), HFES-2020, and Topics in Cognitive Science Journal.

We also had success publishing a book chapter summarizing our work on the effectiveness of various algorithms for defensive cyber-deception in an adversarial decision-making task using human experiments. In this chapter, we propose six algo-

gorithms with different degrees of determinism, adaptivity, and customization to the human adversary's actions. We find that the adaptivity of the strategy is more important than the expected optimality of the algorithm. Humans learned and took advantage of defense algorithms that are deterministic, non-adaptive, and not customized. We plan to continue this work in the direction of developing adaptive and customizable defense algorithms.

In another project, we conducted experiments using the HackIT tool to understand the scanning strategies of attackers during the reconnaissance phase. Furthermore, we are moving towards two-sided deception where we identified various features on real and honeypot nodes that could be used for effective deception. We designed experiments in HackIT to test the effectiveness of deception and we look forward to data collection in coming months. I am also happy to mention that we were managed to get 1-year funding to study how to induce and exacerbate biases for cybersecurity using the HackIT tool.

In addition, I worked with Coty and Don on Python IBL model (pyIBL) for developing ideas on network IBL model. Along with the research work at DDM Lab, I got opportunity to serve as track chair for Cognitive Security track in Computational Security in Compromised Environments (C3E) meeting in Menlo Park, CA. I represented the DDM Lab at the Human Factors and Ergonomics Society (HFES 2019) Meeting in Seattle, WA.

Research Updates Continued

FROM KULDEEP SINGH

This past year I continued working on cybersecurity projects. In one of the projects, we focused on training of end-user and prepare them against phishing emails. We investigated how frequency of phishing emails, type of feedback and incentive structure impact the learning of end-users in phishing task. Our research findings suggest that the detailed feedback using high frequency of phishing emails help end-users identify phishing emails. Along with our collaborators, we are also developing cognitive models to predict end-user's actions on various emails. In this project, we also started working in the direction of finding that why some emails are more difficult to detect compared to others. We use our previous experiments dataset of phishing and ham emails and expand that data set by collecting annotations of the features that make the emails phishing. Our results indicate that the similarity of the features of phishing emails to benign emails, play a critical role in the accuracy of detection. The phishing emails that are most similar to ham emails had the lowest accuracy while the phishing emails that were most dissimilar to the ham emails were detected more accurately. This work is accepted for publication in HFES-2020. We are looking forward to publish the extended version of our papers this year.

In another project, we have designed and developed multi attacker and defender game (MAD game). We are investigating, how to motivate multiple defenders to share information and take advantage of information sharing against cyber-attacks by learning the behavioral aspect of an attacker in the network. The main goal of this task is to develop game-theoretic models and cognitive models to analyze the impact of collaboration among a set of defenders. We will be conducting human team experiments to test various factors in multi attacker defender game scenarios. I look forward to make progress in upcoming year and suggest some interesting implications of our work to help secure the cyber-world.

FROM NGOC NGUYEN

I joined the DDMLab last October 2019, and have been working on the human machine teaming project that is focused on addressing the question of how machines can work with humans to perform a task in an effective way. We specifically have developed a cognitive model of an observer which is

based on Instance-Based Learning Theory (IBLT). Results from our experiments demonstrate that the IBL model of an observer is able to predict next action, intentions, and false beliefs of agents in new gridworld environments given past observations of the agents' actions. This research work is accepted for publication in the 42nd Annual Meeting of the Cognitive Science Society (CogSci 2020). As a next step, we have been conducting human experiments to test the predictions of our IBL observer against those of humans.

Recently, we have continued to investigate the construction of IBL model of agents (i.e. players) in a goal seeking task in gridworld environments, and then provide a comparison of its performance with that of Reinforcement Learning (RL). In particular, we have evaluated the accuracy and efficiency of the RL and IBL agents in goal seeking gridworlds under various levels of decision complexity. This investigation can pave the way for understanding the characteristics of each player model in the navigation task under uncertainty and under increasing complexity decision conditions. This work is accepted for publication in the 18th International Conference on Cognitive Modeling (ICCM 2020). As a follow-up study, we have implemented an interactive, single-player version of the gridworld task in order to verify the findings reported in the ICCM 2020 paper with human participants.

From Korosh Mahmoodi

My research in the past year has focused on analyzing the complex trust patterns that emerged between the human pairs in MTurk, who we hired to play the Prisoner's Dilemma game. This game added one extra feature: the choice of trusting the other player's decision rather than their own. We analyze the temporal complexity of the short time series of trust decisions (200 trials) and we developed a new form of Diffusion Entropy Analysis (DEA). We found that the formation of trust decisions between the participants was the mechanism that generated mutual cooperation. We supported the experimental results with an agent-based model (Selfish Algorithm, SA), and a cognitive model (Instance-Based Learning, IBL). Although I recently left the lab, I do continue to collaborate with the wonderful DDMLab members.

Research Updates Continued

From Hanshu Zhang

I worked on several projects in my previous year at the DDMLab. In one project with Fred Moisan, Palvi Aggarwal, and Coty, we explored how people behaved in the sender-receiver deception game with the joint effect of social preference and cost of truth-telling. We found that though nice people lied less, they still sent deceptive messages when they realized that telling the truth was costly to themselves. In another project with Fred Moisan and Coty, we examined the players' strategy in the Rock-Paper-Scissors game, we found that in addition to the "win-stay/lose-shift" heuristic strategy that previous research has proposed, there were also a small group of players who behaved the "win-shift/lose-stay" strategy. Our work on the Rock-Paper-Scissors project will be presented at the HFES annual conference this year in October.

Recently, we collected data (thanks for Kuldeep and Palvi's help with programming) on these two projects. Because of the restriction in analyzing the pair-level strategies, we instruct participants in the sender-receiver game and Rock-Paper-Scissors game play with designed computer algorithms and expected that they would learn the specific strategy to maximize their gains through the game. We are working on finalizing the analysis and hope that we would submit soon.

FROM ERIN MCCORMICK

This past year, I have been involved in multiple active research projects. I submitted a manuscript with Coty and Sam Cheyette, former DDMLab member and current graduate student at UC Berkeley. We investigated 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 deci-

sions from experience are made in dynamic environments.

For my dissertation, I have been collecting data for a project on time constraints and decision making, advised by Coty and Stephen Broomell, another faculty member in the Social and Decision Sciences. I test whether specific changes in the decision process should and do allow decision makers to adapt to time constraints, and will contribute to our understanding of how decision processes change under time constraints.

Additionally, last October I received a year-long appointment to the ORISE Postgraduate Research Participation Program via the Air Force Research Laboratory (AFRL) and Department of Energy, and started several new projects with Coty and Leslie Blaha, a research psychologist for AFRL who is leading a growing set of CMU-AFRL collaborations. For our projects, we applied a novel visual analytics technique from Physics, called recurrence quantification analysis (RQA), to sequences of decisions--a common type of data in dynamic decision making research. For this type of data, RQA augments existing statistical methods by giving us a much clearer, closer look at individual choice sequences. We've applied RQA to data from my own past work, and also to data courtesy of two DDMLab alumni, Emmanouil Konstantinidis and Jason Harman. So far on this topic, we've written a CogSci proceedings paper, given a talk at the annual meeting of the Society for Mathematical Psychology, and developed multiple working papers.

FROM DON MORRISON

Version 4.0 of PyIBL (<http://pyibl.ddmlab.com/>) was released this year. It is the first publicly available version to include similarities and partial matching. It also includes a flexible mechanism for implementing delayed feedback, as well as a simplified and more Pythonic API. Work is currently underway to expose more of the underlying mechanisms for researchers wishing to experiment with more flexible uses of Instance Based Learning.

Research Updates Continued

FROM ERIN BUGBEE

I am a new doctoral student in Behavioral Decision Research through the Department of Social and Decision Sciences. I received my Bachelor of Science in Statistics and Bachelor of Arts in Behavioral Decision Sciences from Brown University in 2020. My previous research spans the explore-exploit dilemma, trust in humans versus machines, computational neuroscience, and music information retrieval.

As an incoming doctoral student in the Dynamic Decision Making Lab, I look forward to building computational cognitive models of learning and decision making founded in Instance-Based Learning Theory. I am particularly interested in the intersection of human and machine learning, as well as network science and applications to behavioral game theory. Some topics that I hope to investigate include the explore-exploit dilemma, human-machine interaction, and collective intelligence.

FROM CHASE McDONALD

I am a new doctoral student in the Social and Decision Sciences department. My research interests include both empirical and simulated multi-agent systems and their connections to social decision-making, cognition, and intelligence. I received my Bachelor's degree in Computational Social Science from the University of California, Los Angeles, and wrote my undergraduate thesis on the emergence of inter- and intra-group norms under discrimination in multi-agent reinforcement learning systems.

In the DDMLab, I'll begin by working on networked Instance-Based Learning models, which will include the development of novel hypotheses about the interactivity mechanisms among individuals that support pro-social behavior and improve group outcomes. This ties into my previous research, as I have been involved in several projects that use ideas from sociology and social psychology to facilitate cooperation in groups of reinforcement learning agents. I am looking forward to continuing this line of work using IBLT, and to develop computational models of both human behavior and cognition, and the systems that can improve empirical outcomes.

FROM YINUO DU

My work in DDMLab is to explore methods to improve the effectiveness of honeypots. I got the opportunity to work with Prof. Gonzalez, Dr. Aggarwal and Dr. Singh on HackIT project. The study aims to help decide the balance point between the degree of deception and the (fake) vulnerabilities present on (real node) honeypot. Specifically, we would like to test how changing different numbers of features on honeypots, and 2-sided deception where both honeypots and real nodes features are subject to manipulation would help improve deception. We already identified modifiable features on honeypots, designed and implemented experiments in HackIT to investigate the best ways to deceive hackers. The project is going on well and we are about to launch the experiments.

FROM ADITYA GULATI

I am currently doing my integrated masters in computer science at IIIT Bangalore. I started working at the DDMLab in July 2020 as a visiting scholar during my summer break. I spent this time working with Coty and Ngoc on studying human-machine teaming in environments of varying complexity. We looked at designing IBL models to run in a search and rescue scenario designed in Minecraft. We also looked at various metrics to measure the complexity of a decision and how well these metrics map to the complexity of the environment. From August, I also started work on my masters thesis in collaboration with the DDMLab. As part of this work, I will be looking at how multiple agents work together in a multi-agent grid-world. The aim is to understand the impact of different kinds of interactions on aggregate behavior.

FROM JEFFREY FLAGG

Over the last year I have continued preparing IRB protocols, testing materials, running experiments, and helping to integrate new lab members. In addition, I have been maintaining the DDMLab website and procedural documents, preparing our Newsletter, and running monthly lab meetings. I have had the pleasure to help review several paper submissions for DDMLab members. I also helped to research individual differences in Theory of Mind. As many of us know, over the last several years there has been a growing trend for openness in research. I have worked within the Open Science Framework develop preregistration processes for future DDMLab studies.

Research Updates Continued

FROM ALISON BUTLER

I have been working with Coty on a project related to end-of-life decision making. We are interested in the question of how people make decisions when they cannot necessarily obtain personal experience with the choices involved -- such as decisions about what type of medical care people would want when they were near the end of life, or what type of interpersonal treatment they would desire if they were dying. This work will shed light on how clinicians and other healthcare professionals can identify emotional and social difficulties involved in completing a living will or advanced directive, so that they are better equipped to help patients take the initial steps in completing documents like the Five Wishes well in advance of actual end-of-life settings.

FROM MAX GAMERMAN

We implemented a research project on stock flow failure and the correlation heuristic, examining why both well-educated experts in various STEM and humanistic fields as well as individuals unaccustomed to system dynamics behavior fail to understand supply chain management applications and accumulations in systems over time. Our experiment aimed to manipulate randomly assigned conditions of financial and social motivation to assess how incentive structures influence reasoning and analytical thinking patterns in the context of stocks and flows. Our insights and findings were presented at the 2020 International Conference of the System Dynamics Society.

FROM PEIJIE HE

Last year I joined the DDMLab as an undergraduate research assistant and was working on an individual project under the guidance of Coty and help of many lab members (Thank you guys!). My project is about adults' linear thinking behaviors. This is a very interesting topic to study since although the system dynamics societies know that adults do have problem responding to dynamic problems, they never use simple non-linear problems tested

on children to see how adults would respond to those questions. I presented the research as an interactive poster on the System Dynamics Conference in July 2020. Although I didn't make it to Bergen this year due to the COVID-19, I still enjoyed it a lot since it was the first time I experienced a virtual conference with all the cute avatars.

FROM ANDREW YE

During the former half of the year, I worked with Kuldeep and Palvi to create a browser-based game that modeled information sharing between multiple agents. The game functioned in multiplayer, where each person would enter as a defender, and would be randomly selected to be a victim of a cyberattack. Each defender would then choose whether to share information about their current status, and with whom this information would go to. The game would go on for several rounds, and aimed to model the benefits and drawbacks of defenders working together or separately against cyberattacks. I spent my time creating a bare-bones version of the game using the NodeGame Javascript package, as well as cost matrices determined by Kuldeep, Palvi, and Coty. Although I worked remotely in the spring semester, I've really enjoyed my time at the DDMLab so far, and look forward to working with Coty and the rest of the lab members!

FROM MAX YEH

At the DDMLab, I had the opportunity to work with Coty on modifying and improving the *ShinyIBL* interface. I would like to thank Erin McCormick for helping me learn the basics of Shiny and for walking me through the ShinyIBL code. Along with modifying the IBL interface, Coty and I also began developing a study to test the effectiveness of Shiny IBL compared to other visualization tools such as CogTool and IBLTool. Everyone at the lab was so fantastic and welcoming. I feel so lucky to have been part of such a great team. I will be finishing up my Senior year and am excited for what lies ahead for me.

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:

Gonzalez, C., Aggarwal, P., Cranford, E., Lebiere, C. (2020). Design of Dynamic and Personalized Deception: A Research Framework and New Insights. *Proceedings of the 53rd Hawaii International Conference on System Sciences HICSS 2020*, January 7-10, 2020. pp. 1825-1834. doi.org/10.24251/HICSS.2020.226.

Aggarwal P., Gonzalez C., Dutt V. (2020) HackIt: A Real-Time Simulation Tool for Studying Real-World Cyberattacks in the Laboratory. In: Gupta B., Perez G., Agrawal D., Gupta D. (eds). *Handbook of Computer Networks and Cyber Security*, Springer, Cham.

Cranford, E. Gonzalez, C., Aggarwal, P., Cooney, S., Tambe, M., Lebiere, C. (2020). Adaptive Cyber Deception: Cognitively Informed Signaling for Cyber Defense. *Proceedings of the 53rd Hawaii International Conference on System Sciences HICSS 2020. January 7-10, 2020*, pp. 1886-1894. **Conference Best Paper Award**. doi.org/10.24251/HICSS.2020.232

Cranford, E. Gonzalez, C., Aggarwal, P., Cooney, S., Tambe, M., Lebiere, C. (2020). Toward personalized deceptive signaling for cyber defense using cognitive models. *Topics in Cognitive Science.*, 12, pp. 992-1011.

Mahmoodi, K., West, B. J., & Gonzalez, C. (2020). Selfish Algorithm and Emergence of Collective Intelligence. *Journal of Complex Networks*, In Press.

Nguyen, N. T., & Gonzalez, C. (2020). Cognitive Machine Theory of Mind. *42nd Annual Meeting of the Cognitive Science Society (CogSci 2020)*, July 29-August 1, Virtual meeting. pp. 2560-2566.

Rajivan, P., Aharonov-Majar, E., Gonzalez, C. (2020) Update now or later? Effects of experience, cost, and risk preference on update decisions. *Journal of Cybersecurity*, 6(1).

Veksler, V. D., Buchler, N., Lafleur, C. G., Yu, M. S., Lebiere, L., & Gonzalez C. (2020). Cognitive Models in Cybersecurity: Learning from Expert Analysts and Predicting Attacker Behavior. *Frontiers in Psychology*, 11, pp. 1049.

Highlighted Projects and Events



2.8 Million DARPA Grant

Anitia Woolley, Coty Gonzalez, and Henny Admoni received a major grant in 2019 to study team collective intelligence and theory of mind involving human/machine interactions. They aim to develop a machine theory of mind to direct a artificial team teammate that can interact with groups to improve performance.

Best Paper Award

DDMLab members and collaborators including Drew Cranford (CMU, Psychology), Christian Lebiere (CMU, Psychology) and Milind Tambe (Harvard, Computer Science) were awarded the Best Paper Award at the 53rd Hawaii International Conference on Systems Sciences.



Adventures During 2019 - 2020



In early November of 2019, Coty, Palvi, and Kuldeep traveled to Seattle for the 63rd International Annual Meeting of the Human Factors and Ergonomics Society. Coty participated in a panel on human-machine teaming, Kuldeep gave a talk on phishing detection, and Palvi presented a poster.

It also gave an opportunity to catch up with some former DDMLab members. Thanks, we had a great time!



Undergrad RA's Max Gamerman and Peijie He gave virtual talks at the 38th International Conference of the Systems Dynamics Society. We are all very proud of their work!

THE 38th INTERNATIONAL CONFERENCE OF THE SYSTEM DYNAMICS SOCIETY
 Virtually everywhere including Bergen, Norway

Incentives and Correlational Thinking: An Experiment with Piece-Rate Payments

Max Gamerman, Carnegie Mellon University
 Dr. Cleotilde Gonzalez, Carnegie Mellon University
 The Dynamic Decision Making Laboratory @ CMU
www.cmu.edu/ddmlab

<https://youtu.be/vC9MGIV9y4w>

A test of Adults' Linear Thinking in Non-linear Math Problems for Children

Peijie He and Dr. Cleotilde Gonzalez
 Dynamic Decision-Making Laboratory,
 Carnegie Mellon University

<https://youtu.be/o5M8rz8A22M>

DDMLab Carnegie Mellon University

Effects of Decision Complexity in Goal-seeking Gridworlds: A Comparison of Instance-Based Learning and Reinforcement Learning Agents

Thuy-Ngoc Nguyen and Cleotilde Gonzalez
 Dynamic Decision-Making Lab
 Social and Decision Sciences Department
 Carnegie Mellon University

This research is sponsored by DARPA-ASIST program
 The 2020 International Conference on Cognitive Modeling – July 2020

https://youtu.be/CWhjwIEz_hw

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

While not typical “adventures,” we have all tried to make the most of what has been a very unusual year!

μψν

Virtual MathPsych/ICCM 2020
 virtual.mathpsych.org
 July 2020

Analyzing variability in instance-based learning model predictions using recurrence quantification analysis

McCormick, Erin N
 Carnegie Mellon University, United States of America

<https://youtu.be/YEQn7I2T30Q>



Dynamic Decision Making Laboratory

Winthrop Lab

4609 Winthrop Street, 1st Floor

Pittsburgh, PA 15213

e-mail: ddmlab@andrew.cmu.edu

Phone: 412-268-9547

412-268-8439

THANK YOU FOR YOUR INTEREST IN OUR
2020 NEWSLETTER!

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

[FACEBOOK](#)



[LINKEDIN](#)



OR [TWITTER](#)

