Issue 9

Fall 2017

Dynamic Decision Making Laboratory

Carnegie Mellon University

Annual Summary from Coty—Founding Director, DDMLab

Dear Friends and Collaborators:

This newsletter will give you an update on DDMLab activities in the past academic year, 2016-2017. We have news to share with you regarding new and continuing grants, new manuscripts, new ideas and, most importantly, the people who make up this lab. Thank you for letting us update you!

I can easily say that this past year was extremely busy, perhaps the busiest so far. (But every year seems that way.) The extreme action we have experienced over the past year comes from many research opportunities and proposal submissions to several granting agencies, but also from many other activities.

Inside this issue:

Starting with grants—old, new, and pending—we continue to have the fortune of being supported by many organizations, including **The Army Research Laboratory (ARL), Army Research Office (ARO), National Science Foundation (NSF), Defense Advanced Research Projects Agency (DARPA), and The National Institute of Standards and Technology (NIST). ARL** continues to fund our work on the socio-cognitive aspects of cybersecurity under a large cybersecurity Collaborative Research Alliance (CRA). I continue to be the leader of the socio-cognitive work done in this program. During this past year, we advanced our work on cognitive models and the understanding of attackers, defenders, and end-users' behavior through a number of experiments. We addressed the basic psychological issues of attack detection and agility using behavioral game theory to study attacker-defender interactions, risk, and decisions from experience. **NSF** funded our work on decisions from experience in individuals and groups, as well as the connections and interactions between experience and description-based decisions. We also advanced instance-based learning (IBL) theory and computational models. We created a new Bayesian implementation of IBL (IBL-Bayes); a new model that processes description-based choice through lexicographic rules and experience-based decisions through the same IBL memory mechanisms (Lexicographic-IBL, LIBL); and a tool for visualizing the results from IBL models of choice (Shiny-IBL). **NIST** supports our work on forensic science through the Center for Excellence in Forensic Science (CSAFE) in the Statistics department at Carnegie Mellon University and Iowa State University. Last year, we concentrated on investigating human's identification decisions based on physical evidence. We have designed and conducted pilot studies that involve breech face images from the *NIST Ballistic Imaging Database Evaluation (NBIDE)*. We have also worked on related experiments regarding eyewitness identification decisions and the effects of face similarity in different lineup procedures on the accuracy of identification. Starting this year, **DARPA** and **ARO** will provide new funds to support our work on group and network learning. **ARO**'s network science grant through the Social and Cognitive Networks program will support our study of group learning and decisions from experience in groups. **DARPA**'s Computational Simulation of Online Social Behavior (SocSim) program will support new, exciting collaborations that aim at integrating cognitive architectures (ACT-R) into traditional network approaches in computer science. **ARO** will also provide us with new funds through a multi-university research initiative (MURI) program to advance psychological theories of deception for the purposes of understanding cyber-deception and applying these theories in practical cybersecurity problems. Several other proposals are currently under review, and we hope to share more good news regarding research grants soon.

Our research publications this year encompassed various themes at the DDMLab. First, for **dynamics of control**, we published a paper with Eric Qi, Nalyn Sriwattanakomen, and Jeff Chrabascz in *System Dynamic Review* (Gonzalez, Qi, Sriwattanakomen, & Chrabascz, 2017). The manuscript presents results from an experiment that tested how the

graphical elements of flow graphs influence the judgment of the accumulation of a stock over time. An example of this concept is how the size of areas under the curve (i.e., accumulation) of a graph of CO_2 emissions (inflow) and absorptions (outflow) would influence people's understanding of CO_2 accumulation. The results from this experiment provide conclusive evidence that the design of a graph influences the salience of the areas between the flows but does not eliminate stock-flow failure. Second, for **dynamics of choice**, we published an article in the *Journal of Behavioral Decision Research* (Ashby & Gonzalez, 2017) on learning to make time-dependent decisions from experience. The paradigm used in that research has the potential to answer broad questions about the value of time, rewards, and people's ability to adapt their response time to received rewards. Related to our work on cyber security and **attacker-defender interactions**, we published one article in *Frontiers in Psychology* on behavioral game theory. In Moisan & Gonzalez (2017), we present an experimental study that manipulates attack strategies and the level of uncertainty confronted by a defender. We find that adaptive strategies are the most challenging for human defenders to identify and learn to recognize, even more so than random attack strategies. For a complete list of journal articles and book chapters, see our lab web page: www.cmu.edu/ddmlab).

In addition to our published journal articles, we continue to publish manuscripts and abstracts that appear in conference proceedings. We presented our work in many well-renowned international conferences. For example, this past summer, our lab members attended a number of conferences. In the 17th International Conference on Social Dilemmas in Taormina, Sicily (June 20-23, 2017), Efrat and I presented basic work on attacker-defender interactions and the effects of uncertainty and variability on risk taking over time. The pictures below show a group photo from a trip to the Sicilian countryside, Efrat and I enjoying the reception, and picture of the conference hotel.

The 50th Annual Meeting of the Society for Mathematical Psychology and the 15th Annual Meeting of the International Conference on Cognitive Modeling (MathPsych/ICCM 2017) in Warwick was a great experience. Jeff and I presented work regarding IBL-Bayes and Shiny-IBL. The picture to the right shows Jerker Denrell (University of Warwick) and me trapped at Warwick Castle. The Conference in Cognitive Science (CogSci 2017) in London was an interesting conference where I attended editorial board meetings.

Teaching was again a great experience this year. In the spring of 2017, I taught a course titled "Dynamic Decisions" that I created last year. This course was designed for junior and senior students in Social and Decision Sciences, but it is also open to other students. The course includes sections on system dynamics modeling, decisions from experience, and learning and experiential choice with dynamic simulations. On one of the assignments, students were able to experiment with Shiny-IBL. We received great feedback and we plan to continue improving this course. Erin McCormick, a Ph.D. student in the lab, was the Teaching Assistant for this course, and the students loved her!

On a personal note, I was a productive painter this year! I used mostly acrylics and watercolors. The pictures on the right are some of my favorites from this year. I also got to appreciate the corners of Pittsburgh on my bike. Pittsburgh is the city of surprises. You think you have seen everything, but you can always discover something new.

In closing, I want to thank all of you, DDMLab collaborators and friends from all over the world!!! Thank you to all of you who spent time with us this year, talking and discussing ideas. Thank you for contributing to our efforts, and thank you for sharing your ideas and lending us your support.

We look forward to many of you visiting our lab this year!

Enjoy the newsletter!!!

Coty.
2017

Sunset Through the Waves, watercolors

El Gallo, acrylics

A Woman's Face, acrylics

News From Our Members

Farewell

Dr. Jeffrey Chrabaszcz, August 2016-September 2017. Ph.D. in Neuroscience and Cognitive Science, University of Maryland, College Park, MD. Dr. Chrabaszcz took a Data Scientist position at the Software Engineering Institute (SEI) at Carnegie Mellon University.

Dr. Michael Yu, May 2014-2017. Ph.D. Social and Decision Sciences, Carnegie Mellon University. Dr. Yu took a Postdoctoral Research position at the Biomedical Ethics Unit at McGill University in Montreal, Québec.

Fei Lu, June 2015-2017. M.A. in Developmental Psychology, University of Pittsburgh, Pittsburgh, PA. Fei enrolled in a Ph.D. program in Industrial/Organizational Psychology at Seattle Pacific University.

Welcome New Members

Sabina Sloman completed her B.A. in Economics at McGill University in Montreal, Québec. She joined the lab in August and is currently enrolled in a Ph.D. program in Social and Decision Sciences at Carnegie Mellon University.

Visitors

Melisa E. Chávez Guerrero visited us from February to May of this year. She is a Ph.D. student at the Experimental Analysis of Behavior program at the National Autonomous University of Mexico (UNAM).

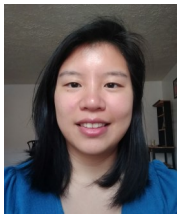

Emily Ho joined us this August as a visiting Ph.D. student from Fordham University's Psychometrics and Quantitative Psychology program

Research Updates from Lab Members During 2016-2017

From Jeffrey Chrabaszcz

My focal project over the last year has been to understand how people combine information from experience and descriptive information. To this end, I created a study in oTree to elicit choices in response to description and sampling periods for shared gamble environments. This has given us a rich dataset to begin modeling with Lexicographic IBL and a modified version of Ido Erev's BEAST model. These stimuli were also designed to allow us to test whether classical certain/risky choices differ from risky/risky choices, which follows on a paper by Glöckner *et al.* from last year.

I have also developed a Bayesian implementation of IBL in collaboration with Coty, Cristobal, and Manos Konstantinidis. This model introduces a probabilistic choice rule into IBL and allows us to fit multilevel models that simultaneously vary the noise and decay rates by subject, gamble, and trial. We presented this model and a re-analysis of some previously-published data from the lab at the Society for Mathematical Psychology this year. One of the most surprising findings was that the model suggested negative decay values to be plausible. While this this finding is *a priori* unlikely, it suggests some ways to model choice autocorrelation and primacy in later computational work.

From Cristobal De La Maza Guzman

For the past year, I have been working on three projects. Here I present a short description for each:

1) Information search and the reversed description-experience gap. New evidence has shown that the description-experience gap can be reversed for certain tasks in decisions under risk, namely when the choice is between two risky lotteries (non-reduced task) instead of one risky and one safe lottery (reduced task). In the paper for this project, we reviewed the different factors affecting the presence of the description experience gap and proposed a theoretical model that explains how changes in environment structure can affect the exploration process, which leads to different information search strategies and thereby biases probability judgments in decisions from experience in different directions. We used the instance based learning model to further understand choices in both reduced and nonreduced tasks. In so doing, we recovered the pattern exposed by the reversed description-experience gap and improved upon results obtained by other existing models such as prospect theory.

2) Discovering preference structure with graph theory. In this project, we proposed a new way to discover non-traditional (and traditional) preference structures empirically using graph theory. We explored the approach with simulated and stated preference data. Using simulations, we varied the proportion of decision makers in each of the most common classes of economic and psychological models of choice. We also applied this classical data approach to observed (in)transitive preferences as well as newer studies. We discovered important structures across a variety of dimensions, such as number of decision makers, alternatives, and attributes.

3) High dimensional parametric preference learning. We developed and tested a method for high-dimensional parametric preference modeling that is able to automatically capture variation in both the content and structure of individual preferences. Specifically, we

Research Updates (Cont'd)

proposed a sparse, grouped. mixed multinomial logit discrete choice model to estimate both the content and structure of preferences across a variety of datasets and tasks in economics and psychology. The approach is flexible enough to model a wide variety of phenomena while retaining a parametric form that is readily interpretable, making it useful for aiding individual and societal decisions. We illustrated the approach with applications to Generalized Extreme Value (GEV) models and context dependence.

From David Hagmann

Coty and I have explored the description-experience (D-E) gap across choices involving two risky gambles. We generated three different lottery profiles that have the same expected value but vary whether outcomes are equally likely or dissimilar (i.e., either the high or low outcome is more likely). We then extended those gambles to three outcomes and looked at choices in the sampling and repeated choice paradigms. We observed a consistent D-E gap in the sampling paradigm, with participants preferring the maximin option.

For my dissertation work, I have conducted experiments on persuasion where people are motivated to defend their beliefs, and I am working on additional experiments in this new paradigm. My other work includes the development of a scale to measure information preferences and a project that shows that introducing an "easy way out" via nudges can undermine support for traditional economic policies (e.g., taxes). Another one of my projects, which is nearing revision, looks at the desire to explore new options in the presence of losses.

I am spending this year as a visiting scholar at Wharton, where I will study whether people actively avoid information related to their taxes. Such avoidance could be costly when it leads to missed opportunities to take advantage of tax incentives and, more generally, when it undermines the effectiveness of tax incentives.

From Efrat Aharonov-Majar

I joined the lab exactly a year ago. During this year, I worked mainly on two research themes. The first one addresses repeated decisions to take protective actions over time. In collaboration with Coty and Noam Ben-Asher, I recently submitted a paper on the effect of patterns of attacks over time on individuals' ability to predict attacks and carry out timely backups. For the second stage of this research, we examined how training with different patterns in the first stage affects backup decisions in a novel environment with rare attacks. In a second ongoing project, we are examining the effect of the dynamic features of updates on the timing of security update implementation.

My second research theme addresses group decision-making. In one study, we explore the effects of feedback (full or none) and the order of decisions on decisions to contribute to or make monetary requests from a joint pool when group members vary in initial wealth. The order of decisions is either random or based on members' initial wealth. In another study, I examined the ability of groups to maximize outcomes in a repeated choice task. The outcome of coordination (being in the majority) was certain or noisy. We discovered that learning to coordinate was easier when the outcomes were certain than when they were noisy. This finding replicates previous findings on the behavior of individuals. Overall, it has been a busy and productive year!

From Erin McCormick

This year, I have continued to investigate human adaptation to continuous, exogenous changes in the probabilities and payoff values of choice outcomes in decisions from experience (a project with Coty and Sam Cheyette, former lab member and current graduate student at University of Rochester). The project has become a part of my dissertation, alongside a line of investigation into how decision makers adapt to time pressure (with Coty and Stephen Broomell, another faculty member in the Social and Decision Sciences department). Adaptation to certain types of time pressure can be conceptualized as a form of adaptation to change. This conceptualization creates a link between the two areas that we hope to leverage for a better understanding of both topics. For example, a decision maker running errands just as stores are about to close for the night must adapt to a choice environment with a continuous change (a decrease) in the time available to complete the desired tasks.

From Nalyn Sriwattanakomen

In the past few months, I've transitioned to the role of lab manager, but I've also had the chance to engage in exciting research over the past year. With Coty, Jeff, and Eric Qi, I published a brief article in *System Dynamics Review* about the relationship between the graphical features of stock-flow graphs and a phenomenon known as stock-flow (SF) failure, in which people are unable to make accurate judgments about the accumulation of a stock over time. We found a relationship between graphical features and the visual discriminability of the areas between the flow curves, which are integral to making accurate judgments about the stock. However, we found no significant relationship between visual discriminability and SF failure, which suggests that graphical features are not a substantial contributor to SF failure as previously surmised.

Research Updates (Cont'd)

My current focus is a study on partisan blindness and monetary incentives that I performed with Coty, Matthew Cronin, and John Sterman in February. My role in this project has revolved around data collection, analyses, and interpretation. Using photos from President Obama's and President Trump's inaugurations, we tested the effect of a $1 incentive on Trump and Trump non-supporters' statements of fact under varying degrees of political content. We found that in the most politicized context, in which participants were asked to identify which photo belonged to which president's inauguration, Trump supporters were less accurate than Trump non-supporters, and the $1 incentive increased Trump supporters' accuracy—that is, statement of fact.

From Don Morrison

My recent work has focused on supporting experiments (predominately with Efrat) involving the interactions between multiple participants recruited from MTurk. The platform we are using is oTree <http://www.otree.org/>. However, multi-participant experiments in MTurk have a problem wherein participants sometimes drop out before a full group has formed without warning or telling us, which tends to confuse oTree. We have therefore augmented it with a lightweight, purpose-built, Python-implemented waiting room that gathers groups together, notices when waiting participants have become unresponsive, and collects still-active participants into correct-sized groups to be sent off to oTree simultaneously. With appropriate arguments in the URL used to bind it to an experiment, we can configure various properties of this waiting room and use it for multiple, simultaneous experiments with possibly different parameters. In addition, PyIBL <http://pyibl.ddmlab.com/>, our library of reusable Python code for building IBL models, continues to advance and has been used for a variety of projects.

From Orsi Kovács

In the past year, I was involved in forensic research projects at the lab through the Center for Statistics and Applications in Forensic Evidence's collaboration with the CMU Statistics Department. I started working with Coty and Amanda Luby on the design and implementation of an experiment exploring human factors in eyewitness identification decisions with a special focus on the different effects that simultaneous and sequential lineups have on identification accuracy. With our generalized linear model approach, we expect the results to be consistent with previous research: Sequential representation of lineups should lead to a decrease in false positive and true positive decisions. Our current focus is to investigate in greater depth the interplay between lineup representations and other variables such as foil similarity and lineup size.

From Prashanth Rajivan

In the past year, I have continued to work on multiple projects from previous years and had the opportunity to initiate new projects and collaborations. I am especially excited about the fact that each of my projects is targeting a different human component in cybersecurity. We continued our data collection effort with the Mid-Atlantic Collegiate Cyber Defense Competition (MACCDC) in collaboration with Norbou Buchler at the Army Research Lab. We submitted a paper about this work to the *Journal of Computers & Security*. In this paper, we present results that indicate that leadership and face-to-face interactions are important factors that determine the success of cybersecurity teams. Functional specialization within a team and an adaptive leadership approach to different security defense sub-tasks could be important predictors of overall performance. In coming years, we plan to continue and expand this work to advance the team science of cybersecurity.

I am pleased by the progress we have made on the phishing project and about the collaboration with Dan Jones at University of Texas, El Paso on Dark Triad Personalities. We found that certain kinds of deception strategies are more likely to persuade end-users to perceive phishing emails as important and therefore persuade them to respond. We have gained insights into adversarial personalities that lead to higher divergent thinking when writing phishing emails and insights into end-user personalities that are susceptible to phishing attacks. We are working on two publications about this research that we aim to submit to *Frontiers* and *PsychScience*. We are expecting to publish these articles by the end of 2017. New collaborations with Efrat on decision making surrounding security updates and a collaboration with Jeff on multilevel Bayesian models for predicting threat frequency are coming along great.

In the fall of last year, we published a conference paper at the HFES conference on security event categorization. The paper was in collaboration with Manos and Noam. Earlier this year, Coty and I published a book chapter summarizing work on human factors in cybersecurity for a book titled *Human Factors and Ergonomics for the Gulf Cooperation Council: Processes, Technologies, and Practices*. In the past year, I was also able to publish multiple papers from my dissertation work and work I did in collaboration with Jean Camp at Indiana University. The past year has been a truly multi-cultural experience. I have learnt a lot from my fellow DDMlabers, especially as part of our regular luncheons! In summary, the past year was very busy, educational, exciting, collaborative and productive. Now I'm looking forward to the next one!

Adventures During 2016-2017

Silly DDMLab-ers at farewell party for Fei. June 2017 at Church Brew Works

Visit to Santa Fe Institute. October 2016

Visit from our collaborators at UT El Paso. June 2017

Whitewater rafting at farewell party for Jeff. September 2017 at Ohiopyle, PA.

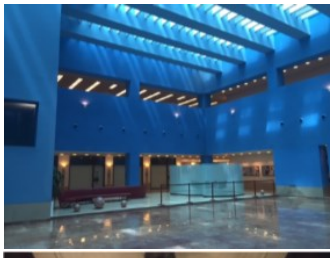

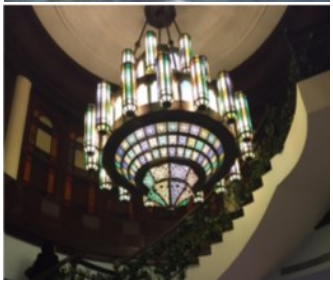

Coty's visit to CMU-Qatar campus. September, 2016

Recent Publications

All of our journal articles and book chapters can be found on our lab web site:

http://www.hss.cmu.edu/departments/sds/ddmlab/papers.html

In the past year, more than ten articles authored by members of the DDMLab and our collaborators were published in various journals such as *Human Factors*, *Decision*, *Journal of Operations Management*, *Cognitive Science*, *Journal of Dynamic Decision Making*, *Frontiers in Psychology*, and *Journal of Cognitive Engineering and Decision Making*.

Some Recent Publications in Conference Proceedings (available upon request):

Ben-Asher, N., Aharonov, E., & Gonzalez, C. (2017). Occasional Downpour or Constant Drip? – The Impact of Adverse Events Distribution on Learning and Risk Taking". Manuscript accepted for presentation at the SPUDM26 conference, Technion, August 20-24, 2017. Haifa, Israel

Gonzalez, C. & Chrabaszcz J. (2017). Some things that glitter are gold: A Shiny app for an Instance-Based Learning Model. Manuscript accepted for presentation at the 50th meeting of the Society of Mathematical Psychology. July 22-25, Coventry, UK.

Chrabaszcz J., Konstantinidis, M., & Gonzalez, C. (2017). A Bayesian Implementation of the Instance-Based Learning Model of Choice. Manuscript accepted for presentation at the 50th meeting of the Society of Mathematical Psychology. July 22-25, Coventry, UK. (pp. TBD).

Aggarwal, P., Gonzalez, C., & Dutt, V. (2017). Modeling the Effects of Amount and Timing of Deception in Simulated Network Scenarios. In *International Conference on Cyber Situational Awareness, Data Analytics And Assessment (CyberSA 2017)*. June 19-20, 2017, London, UK.

Buchler, N., Rajivan, P., Marusich, L., Lightner, L., & Gonzalez, C. (2017). *Wearable Social-Sensors and Structure Observational Assessment of Teaming and Leadership in a Cyber Security Defense Competition*. Track "Human Factors in Cybersecurity". 8TH International Conference on Applied Human Factors and Ergonomics. July 17-21, 2017. AHFE 2017. Los Angeles, CA, USA.

Grabe, J.v. & Gonzalez, C. (2016). Human Decision Making in Energy-Relevant Interaction with Buildings. *Central European Symposium on Building Physics (CESBP 2016).* Dresden, Germany, September 14-16, 2016. pp. 345-352.

Aggarwal, P., Gonzalez, C., & Dutt, V. (2016). Cyber-security: Role of Deception in Cyber-Attack Detection. In D. Nicholson, Advances in Human Factors in Cybersecurity (pp. 85-96). Springer International Publishing. Presented at the *International Annual Meeting of the Human Factors and Ergonomics Society (HFES 2016)*, Washington, DC, September 19-23, 2016.

Dutt, V., Moisan, F., & Gonzalez, C. (2016). Role of Intrusion-Detection Systems in Cyber-Attack Detection. In D. Nicholson, Advances in Human Factors in Cybersecurity (pp. 97-110). Springer International Publishing. Presented at the *International Annual Meeting of the Human Factors and Ergonomics Society (HFES 2016)*, Washington, DC, September 19-23, 2016.

Invited Talks

2017 June 8. *Human Factors in Firearm Analysis*. Center for Statistics and Applications in Forensic Science (CSAFE, 2017 All Hands Meeting), Ames, IA, USA.

2017 March 10. *Human Failure in Stock and Flow Problems: An Updated Review*. Sloan School of Business, Massachusetts Institute of Technology. Boston, MA, USA.

2017 January 26. *Adaptation to Change*. Affective Brain Lab. University College London, London, UK.

2017 January 10. *Psycho-Social aspects of cyber security: why is human (still) the weakest link?*. SaTC PI meeting, NSF, Washington DC, USA.

2016 November 17. *Dynamic Decision Making in Complex Environments*. National League of Cities, City summit. Pittsburgh, PA, USA.

2016 November 9. *Human Factors in Identification Decisions: Cross-Cutting Interdisciplinary Research*. Forensics at NIST. Gaithersburg, MD, USA.

2016 October 11. *Information and Decisions: Description and Experience Come Together in Individual and Social Interactions*. Santa Fe Institute. Santa Fe, NM, USA.