

Dynamic Decision Making Laboratory

Carnegie Mellon University

A Note from Coty

2014-2015 was ... busy. Our research activities involved several projects related to the different basic research grants from the Army Research Laboratories, Army Research Office, Qatar National Research Fund, and the National Science Foundation. We really appreciate the trust that our sponsors give to our work and their commitment to research and good science.

During this year we advanced our understanding of dynamic decision making in many fronts. One of the principal themes that emerged this year more clearly was that of “scaling up” of experiments, models, and theories. We have been working on expanding from *individual* experiential choice to pairs, groups, society, and crowds; and from simple and out of context binary-choice to contextual and complex decisions. There are several noteworthy examples, some of which you will see described by the lab members in this newsletter. Two additional examples that make me particularly excited carry from past collaborations with my colleagues and friends: Christian Lebiere, Ion Juvina, Jolie Martin, Noam Ben-Asher, Varun Dutt, Tomas Lejarraga, and others.

In a Cognitive Science paper (Gonzalez, Ben-Asher, Martin, & Dutt, 2015), we investigated the dynamics involved in repeated interactions between two players in a Prisoner's Dilemma (PD) under various levels of interdependency information. We developed a set of hypotheses regarding how one player (Player1) accounts for the information (outcomes and actions) of the other player (Player2) under various levels of explicit information about the Player2. A selfish hypothesis suggests behavior in agreement with the homo-economicus model where Player 1 only uses her own information to make a choice. On the other extreme, a complete fairness hypothesis suggests behavior in which Player1 considers Player2's information equally as her own to make a choice. The hypotheses testing process was done by comparing behavior of an Instance-Based Learning Model for the PD (IBL-PD) against observed behavior of humans playing the PD repeatedly. Human pairs showed a gradual and slow emergence of cooperation over the rounds of the PD with increased explicit information about the other player. The IBL-PD model comparisons suggest that: the Selfish hypothesis holds when the information of Player2 is unknown to Player1, but fails completely under full information of Player2. The complete fairness hypothesis also fails completely by suggesting a sharp and quick increase in the proportion of cooperation over the rounds of the PD. The best account of human behavior is captured by a dynamic expectations hypothesis, in which Player1 adjust the consideration of Player2's information according to the gap between the Player1's expectations about Player2's outcomes and the actual outcomes in each interaction. In other words, Player 1 adapts her actions between selfishness and complete fairness according to how Player2 behaves in each round.

In another example, a paper in Memory and Cognition (Lejarraga, Lejarraga & Gonzalez, 2014) investigated whether groups make better judgments and decisions than individuals in dynamic tasks and how groups and individual adapt to changes in the environment. We find that groups (3-people working on a repeated binary choice task) performed better than the average individual while the decision task was stable. However, group performance was worse than individual performance after a change in the decision environment. Group performance was closer to a Bayesian benchmark which model which assumed perfect memory while individuals were closer to an IBL model which assume decay of memory, frequency and recency effects. Groups seem to have coordinated their responses in a way that led them to behave as if they had better memory, and this “stickiness” delayed their subsequent adaptation to changes.

Inside this issue:


From The Director	1-2
News From Our Members	2
New Research	3-6
Recent Publications	7
Conference Proceedings	
Workshops, Presentations, & Invited Talks	

Continuing with good news, we obtained a new NSF grant from the Decision, Risk and Management Sciences program. This grant will allow deeper involvement of Ph.D. students into the research of decisions from experience and description in daily life contexts. Our new research program will help generate new insights regarding choice based on a combination of explicitly described information and experience, and it will help in unifying disparate decision theories. Our research will potentially impact multiple contexts including financial, health, energy, and others; through communication policies, the design of information and decision support systems, marketing strategies, etc. by clarifying what effects different amount, type, and format of information have on risk taking and decision making.

On a personal note, I got a promotion this year: I am able to remove the “Associate” part off of my title. This year I learned to use acrylics in my paintings, traveled to interesting places, and I biked all over Pittsburgh!

Finally, I would like to take the opportunity to thank all of you, collaborators and friends of the DDMLab from all over the world, who in many ways have contributed to our ideas and development. I hope you enjoy reading the contributions of our post-docs, students, international visitors, and other members of the lab in this newsletter. We are looking forward to making 2015-2016 another exciting and productive academic year. Thanks to all of you!

Coty
2015



News From Our Members

Farewell!!!!



Picture from *LSU The Psychology Times*, 09/01/15.

Congratulations to members that departed during this year: **Jason Harman**, started a faculty position at the Department of Psychology, Louisiana State University. Jason is now an assistant professor for the Industrial/Organizational (I/O) psychology at LSU. **Nathan Ashby** moved to new collaborations and activities. Best wishes!!

Farewell and best wishes to **Manos Konstantinidis**! He will start as a Post-Doctoral Fellow in the School of Psychology, University of New South Wales in Australia. **Fred Moisan** will be starting as a Research Associate in the department of Economics, University of Cambridge in the UK. **Hassen Gharbi** has concluded a visit to our lab as a Fullbright scholar. Finally, congratulations to **Hau-yu Wong**, our lab manager of many years! She started this year in the very prestigious Masters program in Human-Computer Interaction at **Carnegie Mellon University**. Please come and visit! Our best wishes to all of them.

Welcome!!!!

Pantelis Analytis

and



Prashanth Rajivan,

will join our lab as Post-Doctoral Fellows this year. Pantelis is currently a postdoctoral scholar at the Center for Adaptive Behavior and Cognition (ABC) at the **Max Planck Institute for Human Development** in Berlin, and he will be joining us the beginning of next year. Prashanth is currently a postdoctoral scholar at **Indiana University**, Bloomington and he will join us in November.

Fei Lu, our new lab manager joined our lab in July. Welcome here and all new lab members.





Group photo, July 2015

Farewell Party for Jason, Hau-yu and Hassen

New Research

Below is a summary of research highlights during 2014-2015:

From Cristóbal De La Maza *Ph.D. Student*

Working in the lab the past year and collaborating with Coty has been exciting. We are investigating the effects of communication of probabilistic information and laypeople difficulty understanding cumulative risks in the flood context. For example, suppose that there is a 1% chance of flooding in a specific location each year. When asked to assess the probability of at least one flood occurring in thirty years, people commonly use variants of two heuristics for computing the cumulative risk deviating from the normative response (26%): 1) They take the average probability (1%), or 2) They sum the probabilities for all periods (30%). Following work from Doyle 1997, Juslin et al. 2015 we will ask participants to provide direct assessments of the cumulative risk, measuring risk from their judgments. Further, we will compare these judgments to the perceived risk derived from their choices between risky options, allowing us to compare perceived risk across two elicitation modes, each with their own heuristics and biases.

From David Hagmann *Ph.D. Student*

I am a fourth year PhD student at the Department of Social and Decision Sciences at Carnegie Mellon University. My current research topics include information preferences (and anomalies), decisions from experience, and behavioral interventions (Nudges). In the past year, I have started to co-author with Coty and former lab member Jason (who is now a faculty at Louisiana State University). Our work examined how the timing of feedback affects choices in repeated decisions from experience. I am looking forward to extending our research on experiential choice this year, thanks to a generous grant to Coty from the National Science Foundation.

New papers this year:

Loewenstein, G., Bryce, C., Hagmann, D., & Rajpal, S. (2015). Warning: You Are About to Be Nudged. *Behavioral Science & Policy*, 1(1), 35–42.

Golman, R., Hagmann, D., & Miller, J. H. (2015). Polya's Bees: A Model of Decentralized Decision Making. *Science Advances*.

Golman, R., Hagmann, D., & Loewenstein, G. (2015). *Information Avoidance* (SSRN Scholarly Paper No. ID 2633226). Rochester, NY: Social Science Research Network. (Under review at the Journal of Economic Literature)

New Research (Cont'd)

From Jason Harman

Previous Post-Doctoral Fellow

Assistant Professor at Louisiana State University

The past year has been immensely exciting and productive. Beyond job interviews (which resulted in my current position at Louisiana State University), having a baby, and getting married, we managed to produce some fantastic research and a number of older projects from my first and second year in the lab finally went to press. Work on the Allais paradox, stocks and flows, and recommendation systems were published in the Journal of Behavioral Decision Making, The Journal of Operations Management, and the proceedings of ACM RecSys respectively. Separate collaborations with current and former lab members Katja Mehlhorn, Manos Konstantinidis, and David Hagmann are all in manuscript form and should be in press in the near future. Possibly the most exciting project of the past year emerged from the lab's work building models to enter into the choice prediction competition. One of these models, Lexicographic Instance Based Learning (LIBL), is an extension of IBL to decisions from description and managed to account for 14 known decision making phenomena while retaining psychological plausibility and parsimony. I presented this model at the Summer Institute on Bounded Rationality in Berlin this summer and received very enthusiastic and positive feedback. I'm very glad that LIBL and other projects will insure my continued collaboration with the DDMLab far into the future. My three years in Pittsburgh have been invaluable both personally and professionally and being a DDMLab member is a title I will always be grateful for and proud of.

From Manos Konstantinidis

Post-Doctoral Fellow

I joined the DDMLab in November 2014 and since then I have been working on three interconnected lines of research: First, the development and application of computational cognitive models in

experiential learning and decision-making from experience, with an emphasis on IBL and reinforcement-learning approaches. As part of this work, we submitted a model in the recent Choice Prediction Competition (CPC 2015); our submission finished third among 25 other models. Other modeling endeavors include the use of sequential-sampling models in decisions-from-experience. Second, we seek to understand and explore important topics in decision-making research, such as the exploration-exploitation dilemma, adaptation to change in uncertain environments, and how people behave in environments with increasing structural complexity (i.e., more options to choose from). To this end, we have conducted several behavioral experiments and analyzed the data with the aid of computational models. Our third line of research concerns investigating category learning and classification in uncertain and noisy environments. This project is particularly interesting as it attempts to connect basic cognitive research with applications into the Cyber Security domain (e.g., how cyber security analysts-defenders categorize events as threats or non-threats).

From Fred Moisan

Post-Doctoral Fellow

A recent project consists in studying the emergence of cooperation among two individuals repeatedly playing the well-known prisoner's dilemma (PD). In each round of this game, each individual is asked to make a choice between an individualistic option that maximizes their own payoff and a cooperative option that can benefit all participants equally. While previous studies have found that the level of cooperation is directly related to the payoff structure of the PD, the role of social preferences in this dynamic context remains largely unexplored. In an experimental setting, we have therefore measured people's social preferences and determine how such individual differences interact with the payoff structure of the PD to stimulate cooperation. Our

New Research (Cont'd)

results indicate that other-regarding preferences can indeed explain choices to cooperate in games with a sufficiently strong incentive to act selfishly. Moreover, our analysis suggests that prosociality in the PD is related to a distaste of advantageous inequality, and that the game partner's type of social preferences (in addition one's own) matters to explain one's level of cooperative behavior. Finally, we find that social preferences are stable across repeated interactions, which leads us to discuss specific conditions on payoffs under which being prosocial can be best supported.

In another line of research, we have been investigating how recommender systems can effectively help people behave more optimally in repeated strategic interactions. More specifically, we have conducted experiments to better understand how human players learn to trust or distrust such systems depending on the accuracy of the recommendations. This problem is particularly relevant to the context of cybersecurity where imperfect *intrusion detection systems* are regularly used by computer analysts to help them decide whether an individual's online behavior is to be considered a threat. More recently, we have also started exploring the role of such recommenders to help two human players dynamically coordinate their behavior to improve their long term payoffs. In this case, we are particularly interested in the influence of the fairness of the recommendations provided to both players: are people willing to follow recommendations that clearly favor one player over the other? Or do they prefer ignoring such unfair recommendations, even though doing so would then make everyone worse off?

From Don Morrison
Senior Research Programmer

Development of PyIBL, our library of reusable Python code to quickly and easily create Instance-Based Learning models, continues. Version 2 is available

at <http://pyibl.ddmlab.com>.

The not yet released version 3 adds support for similarity and partial matching, delayed feedback, and enhanced logging capabilities. Pre-releases of version 3 have already been used on several projects, including rapid prototyping of a variety of IBL models for the 2015 Choice Prediction Competition for Decisions under Risk and Ambiguity (CPC2015), supporting two different entries in that competition; and the CyberWar game, work jointly pursued with Noam Ben-Asher at the Army Research Laboratories.

From Mike Yu
Post-Doctoral Fellow

We're hitting a bit of a transition point right now between some older and newer projects. We're wrapping up our writing on our research that investigated the impact of early trust on developing trust over time. In particular, we've found reasonable support that an information-gathering induction can promote early trust-like behaviors and faster trust development, without influencing early judgments of trustworthiness. We've also finished the majority of our simulations regarding cooperation and cognitive agents and have found evidence that a "fuzzy" assignment of current outcomes to past decisions might lead to the emergence of cooperation. A newer R implementation of IBL was also developed to support efforts in the choice prediction competition.

Our current efforts have been focused on developing a study that investigates how the way in which recommendations are provided by a system can influence how likely one is to adhere to a system's recommendations -- both initially and over time.

New Research (Cont'd)

From Hau-yu Wong
Previous lab manager

This past year wrapped up a segment of the NEXCEL project, done in collaboration with Iliano Cervesato at Carnegie Mellon Qatar, with a focus on human recursive reasoning. Through a series of experiments, we tested the extent to which people of various technical backgrounds could reason through recursive processes in plain English and symbolic language, and their ability to construct recursive algorithms in symbolic language. On the most part, people demonstrate some level of difficulty when working with symbolic language. While programming skills are being promoted and increasing among the general populace, our results highlight how important it is to supply decision support to those with and without that background.

FROM OUR RESEARCH ASSISTANT AND VISITING STUDENTS:

Sam Cheyette
Carnegie Mellon University

I just began research at the end of this summer, and am still developing my first project. The overall goal will be to try to give an account of change detection and responses to changing conditions in humans. IBL has performed well as a predictor of aggregate decision making in binary choice tasks, where the risk and reward probabilities of the two options increase or decrease, linearly and monotonically, or remain constant. It will be interesting to see whether IBL can predict human behavior when risk and reward functions do not follow simple patterns – what if they are quadratic or non-monotonic? – or in different contexts. If IBL does indeed perform well as an aggregate model, we might also see to what extent it can capture individual variation in decisions of this type.

Nalyn Sriwattanakomen
Washington & Jefferson College

This summer, I had the privilege of working with Dr. Gonzalez and Dr. Cervesato from the Qatar CMU campus on a project testing the usability of NEXCEL, an automated spreadsheet tool prototype that helps users make logical inferences from tabular data. This particular project explored how people reason about recursively-defined relations and the degree to which people are able to identify and express recursive rules in a Datalog-based language. We found that graphical displays of recursive relations were more helpful than tabular ones, that people best process recursive rules when they are presented in plain English, that the average person has a poor grasp of recursion, and that recursive reasoning ability and expertise in related fields such as computer programming and logic are not associated. In the fall, I will continue collaborating on the NEXCEL line of research. We hope to investigate how people construct recursive rules when given error feedback and “training wheels” that gradually acclimate users to the Datalog language. Additionally, Dr. Gonzalez and I intend to study how healthy people make end-of-life-care decisions for themselves.

Hui Sun
Tsinghua University

Over the past few months, we've been trying to gain a closer look at human behavior during live interaction of repeated rock-paper-scissors game. Albeit its simplicity, the game elicits complicated behavior in the two coupled players. And consequently, stochasticity of choice embedded in the intransitive dominance of the three moves pose challenge to quantitative analysis. This ongoing project seek to understand the learning process underlying human versus human interactions in the rock-paper-scissors game, with various levels of information.

Recent Publications

In the past year, more than ten articles authored by members of the DDMLab and our collaborators were published in the *Journal of Dynamic Decision Making*, *Journal of Behavioral Decision Making*, *Journal of Operations Management*, *Cognitive Science*, *Journal of Applied Research in Memory and Cognition*, *Memory and Cognition*, *Computers in Human Behavior* and the *Journal of Economic Psychology*. An additional book chapter is currently in press.

For a full list of publications, please see the publications page on the lab's website at <http://www.hss.cmu.edu/departments/sds/ddmlab/papers.html>.

Some recent publications in conference proceedings:

- Ben-Asher, N. & Gonzalez, C. (2015). Training for the unknown: The role of feedback and similarity in detecting zero-day attacks. 6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015). July 26-30. Las Vegas, NV.
- Konstantinidis, E., Ashby, N. J. S., & Gonzalez, C. (2015). Exploring complexity in decisions from experience: Same minds, same strategy. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 1177-1182). Austin, TX: Cognitive Science Society
- Kumar, S., Cervesato, I., Maruthullathil, A., Wong H., & Gonzalez, C. (2014). Role of problem complexity and domain familiarity in relational reasoning – A psychological perspective. *Qatar Foundation Annual Research Conference (ARC'14)*. November 18-19, 2004. Doha, Qatar.
- Harman, J., O'Donovan, J., Abdelzaher, T. & Gonzalez, C. (2014). Dynamics of Human Trust in Recommender Systems. *The ACM Conference Series on Recommender Systems. RECSYS 2014*. October 6th-10th. Foster City, Silicon Valley, USA.

Presentations & Invited Talks

Some invited talks during 2014-2015:

- 2015 June 15. The Role of Behavioral Science in Mission Assurance for Cybersecurity. Workshop on Attack Detection, Forensics and Attribution for Assessment of Mission Impact. June 15-17, 2015, Istanbul, Turkey. The Information Systems Technology (IST) NATO panel.
- 2015 February 26. *Creating Cognitively-Aware Decision Support Technology*. **Key Note lecture**. Co-nielecomp 2015. The 25th International Conference on Electronics Communications, and Computers. Universidad de las Americas, Cholula, Puebla, Mexico. February 25-27, 2015.
- 2015 February 14. *Human Decision Making in CyberSpace*. AAAS 2015 Annual meeting. San Jose CA. February 12-16, 2015.