What Research Do We Do?

We study learning and decisions from experience in situations involving repeated, interdependent decisions in dynamic environments.

- How does experience influence our decisions?
- What kinds of experiences would produce better decisions and better adaptation?
- How does experience transfer to new situations?

We also study humans making decisions in a wide range of decision contexts that we bring to the laboratory in the form of dynamic simulations (MicroWorlds or DMGames).

- How do operators of complex industrial plants make dynamic allocations of limited resources?
- How can forensic examiners be more accurate at identifying matches between physical samples?
- How might cybersecurity analysts improve their detection of cyberattacks?

Our driving theory is the Instance-Based Learning Theory (IBLT), which in essence proposes that people make choices by retrieving the best outcomes from their past experience. The process involves:

- Retrieving memories (*instances*) that resemble the current situation (instances are triplets: situation-decision-utility)
- Filtering memories according to their maximum experienced expected value (utility or *blended value*)
- Evaluating and storing new instances reflecting each possible option in the decision situation
- Selecting the option with the maximum blended value

Carnegie Mellon University

People of the DDMLab

Cleotilde (Coty) Gonzalez, PhD
Research Professor and Director of DDMLab

Maria Ferreira, PhD
Post-Doctoral Fellow

Xiaohong Cai, PhD
Post-Doctoral Fellow

Grace Roessling, PhD
Post-Doctoral Fellow

Marko Morrison
PhD Student

Anu Aggarwal
Research Programmer

Don Morrison
Senior Research Programmer

Jeffrey Flagg
Research Associate, Lab Manager

Vlad Miloserdov
Undergraduate Research Assistant

Dynamic Decision Making Laboratory
Carnegie Mellon University
Department of Social and Decision Sciences
Porter Hall 223-G
Pittsburgh, PA 15213
Phone: (412) 268-9547
Email: ddmlab@andrew.cmu.edu
URL: http://www.cmu.edu/ddmlab

Last Updated October 2025

Dynamic Decision Making Laboratory

The Dynamic Decision Making Laboratory (DDMLab) was founded in 2002 by Dr. Cleotilde Gonzalez to investigate decision making in complex dynamic environments. Such environments are characterized by the need for people to make multiple, interdependent, real-time decisions in reaction to both external changes, as well as the effects of their past decisions.

At the DDMLab, we seek to build models and methods that help explain, predict, and draw recommendations for improving decision making in dynamic environments. We use multiple research methods, but notably, we rely on laboratory experiments where we collect human behavior using dynamic simulations (Decision Making Games) and on computational cognitive models based on Instance-Based Learning Theory (IBLT) and the ACT-R cognitive architecture to understand and predict such behavior.

Practical applications of our research extend from front-end system design activities to back-end training and decision-support. On the front-end, we can provide principled guidance and empirical support for the design of systems that exploit DDM strengths. On the back-end, we can help decision makers exploit system strengths. In this respect, DDM theory and methods are particularly suited for the design of training interventions. But the closely related activity of decision support design is no stretch for the skill set our multidisciplinary team provides.

The laboratory consists of post-doctoral fellows, research-programmers, doctoral students and research assistants. Lab members come from different fields, including Behavioral Decision Research, Psychology, and Computer Science.

The DDMLab is part of the Social and Decision Sciences Department at Carnegie Mellon University.

http://www.cmu.edu/ddmlab

DDMLab

How Do We Do Research?

Our research approach includes laboratory experiments and cognitive models, which form a learning cycle that compares human data from experiments against theory-informed data from computational models.

In our research, we address questions such as:

- How does experience influence our decisions?
- What kinds of and how much experience produces better performance and better adaptation to novel environments?

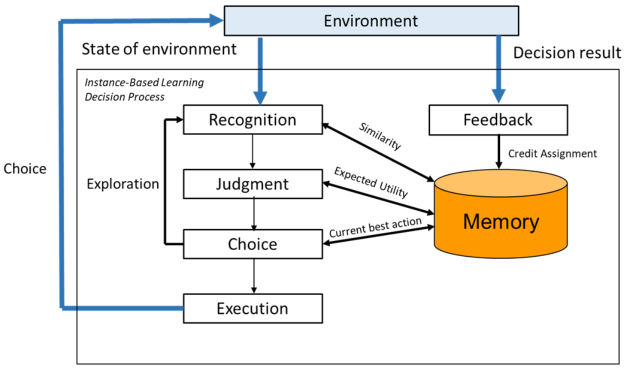

Our research approach includes laboratory experiments and cognitive models, which form a learning cycle that compares human data from experiments against theory-informed data from computational models.

- We collect behavioral data using complex, dynamic simulations.
- Our experiments often involve extended practice to help us understand how experience develops, changes, and transfers to new situations.
- We create cognitive models that rely on IBLT and the mechanisms proposed in the ACT-R cognitive architecture to represent and predict human behavior in decision making tasks.
- Our theory and methods are applied to many domains that deal with the prediction of behavior in complex systems.

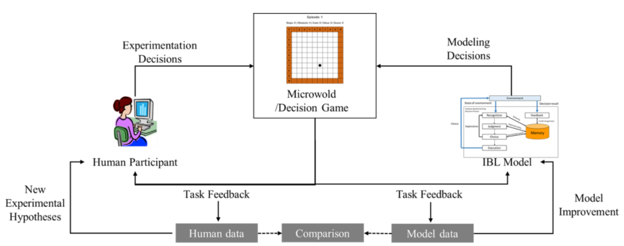

The DDMLab is constantly upgrading its systems for computational modeling. A fantastic new addition is SpeedyIBL, a Python library that allows to create single or multiple IBL agents with fast processing and response time without compromising the performance.

- SpeedyIBL utilizes fast processing and response time without compromising performance compared to the traditional implementations.
- SpeedyIBL can be used to create IBL agents that can do a wide range of decision games.

SpeedyIBL

Research Program

Human-Machine Collaborations

Funding Source: National Science Foundation AI Institute for Societal Decision (AISDM) & Army Research Office, Australia-US Multidisciplinary University Research Initiative (MURI-AUS-MURI).

Long-term Goal: Long-term goal: *Design synthetic coaches that would have Machine Theory of Mind to support team work and enhance team collaboration.*

- Develop a process of coaching/interventions in Human-Machine teams.
- Use this coach to perceive individual cognitive states and team social states.
- Understand the role of humans and other agents in the context of the task environment.
- Diagnose team success to design interventions to improve teamwork.

Disaster Relief Management Decision Making

Goal: *To test how real professional disaster relief managers make decisions about resource allocation and information gathering during natural disasters.*

- Perceptual Aggregation: How do participants internally weigh different damage indicators when rating a disaster-effected asset?
- Allocation Trade-Offs: After forming damage perceptions, how do people prioritize which locations receive scarce follow-up resources?
- Temporal Drift: Do those weights or priorities shift after the full tabletop simulation, once participants have experienced realistic coordination stress?

Answering these questions supports the development of learned social welfare functions that could inform intelligent decision support tools for future disaster response.

Integrating Theory of Mind Capabilities in AI Partners

Goal: *To facilitate Human/AI coordination by building Theory of Mind capabilities in AI partners.*

- Integrate algorithms for preference inference into the k-level framework.
- Examine coordination by an AI partner relying on predictions of the human agent's actions, which in turn relies on both an appropriate and fine-tuned (Instance-Based Learning) model of human agents.
- Look for individual differences in ToM in areas of cognitive processing, emotional intelligence, and spatial awareness.

Cognitive Models in Sequential Decision Tasks

Goal: *To understand how people make sequential decisions in various tasks involving balancing exploration and exploitation, and to develop cognitive models of their behavior in these tasks.*

- Introduce a novel sequential stopping task to shed light on how, when, and why people decide to stop exploring.
- Systematically examine some of the factors that may influence stopping behavior and validate the predictions of our cognitive model.
- Investigate interventions leveraging wisdom of crowds aggregation techniques to provide personalized, cognitive AI-driven recommendations for when to stop searching.

Cognitively Aware Reinforcement Learning

Goal: *To investigate how cognitive models can be used in tandem with reinforcement learning (RL) agents to learn policies that complement human behavior.*

- We incorporate cognitive models into the RL training and testing pipelines to see how such models can improve performance in cooperative tasks.
- We test these models with human proxies and real humans, and analyze their behavior using collaborative fluency metrics, to see how well they learn collaborative policies.

Research Program

Behavioral Cybersecurity

Funding Source: Intelligence Advanced Research Projects Activity (IARPA) ReSCIND, & MURI-AUS on Human-Bot cyber defense teams.

Long-term goal: *To design effective defense techniques informed directly by dynamics of human behavior, emergent cognitive biases, and psychological deception strategies.*

Impact of Cognition on Cyber Behavior

Goal: *Improve cybersecurity by understanding how human cognition impacts cyber behavior and cyber actors' success in network attack activities.*

- Replicate cognitive biases such as availability, endowment, and recency in cyber domains.
- Identify behavioral signatures linked to cognitive biases in capture-the-flag cyber attack games.
- Insert cyber-isomorphs into an attack kill chain environment.
- Utilize the CyberVAN testbed with penetration testers to investigate triggering cognitive biases as mitigating factors in the efficacy of cyber-attack behavior.

Learning of End-Users in Phishing Training

Goal: *Improve cybersecurity by understanding how human cognition impacts cyber behavior and cyber actors' success in network attack activities.*

- Train end-users with different frequency, recency and content of phishing emails.
- Use an AI chatbot to provide feedback on the accuracy of categorizations and train users to properly identify phishing emails.
- Test their detection capabilities detecting human or AI written phishing emails after training.
- We propose IBL model to perform model tracing during the experiment and select emails to show to participants based on that model

Defense Strategies in Binary Choice Tasks

Goal: To design defense strategies to influence human choices in the box game, a repeated binary choice task.

- Attackers repeatedly attempt to find the highest reward in one of two boxes.
- Construct an Instance Based Learning (IBL) cognitive model that tracks human behavior and makes one-step-ahead predictions of human decisions.
- Introduce an intervention based on predicted choices of the participant and measure post-intervention changes.

Towards Human-AI Collaboration in Autonomous Cyber Operations

Goal: *To study the integration of IBL models for improving Trust and Mental Models sharing in Human-Autonomy teams for cybersecurity.*

- Develop a framework for Human-AI collaboration research for cyber-defense and use cognitive models of Human defenders to predict their decisions.
- Test human defenders in an Interactive Defense Game and compare outcomes to Human-AI (ML & IBL) collaborations in the same scenario.
- In a Team Defense Game, compare the effectiveness of a cognitive (IBL), heuristic, or random AI partner.