

Artificial Intelligence (AI) research at the University of Michigan is comprised of a multidisciplinary group of researchers conducting theoretical, experimental, and applied investigations of intelligent systems. Current projects include research in cognitive architectures, distributed systems of multiple agents, machine learning, data mining, computer vision, natural language processing, robotics, computational healthcare, human computing, and computational social science, and others.

Research in the Michigan Artificial Intelligence laboratory tends to be highly interdisciplinary, building on ideas from computer science, linguistics, psychology, economics, biology, controls, statistics, and philosophy. In pursuing this approach, laboratory faculty and students work closely with colleagues throughout the University. This collaborative environment, coupled with our diverse perspectives, leads to a valuable interchange of ideas within and across research groups.

COGNITIVE ARCHITECTURES & COMPUTATIONAL COGNITIVE SCIENCE

AI Faculty: Nikola Banovic, Satinder Singh Baveja, John Laird, Emily Mower Provost

Affiliated Faculty: David Kieras (CSE), Richard Lewis (Psychology)

At Michigan AI, we are exploring architectural structures to support human-level AI systems, as well as computational models of human behavior and the structure of the human brain. All these efforts draw from AI, psychology, and neuroscience, so that our research is inherently interdisciplinary. Michigan AI is unique in the breadth of cognitive architecture research it supports, including active groups in Soar, EPIC, Act-R, Cognitive Constraint Modeling, and neurally-inspired architectures.

MULTIAGENT & ECONOMIC SYSTEMS

AI Faculty: Satinder Singh Baveja, Edmund Durfee, Sindhu Kutty, Arunesh Sinha, Michael Wellman

Affiliated Faculty: Grant Schoenebeck (CSE)

Environments with multiple autonomous agents present special opportunities and pose distinct challenges for design and analysis of AI systems. Agents may take actions, reason strategically, and coordinate with other agents. A multiagent environment is effectively a social system, and thus analyzing multiagent behavior can often be informed by social science. Multiagent systems research at Michigan AI considers all perspectives, from individual agent to social designer. We design planning and learning algorithms suitable for multiagent contexts, and methods for analyzing networks of agents as organizations, economies, and societies. Our work also spans the range from theory to practice. We conduct fundamental research in distributed coordination, algorithmic game theory, and social computing, and apply our techniques to real-world problems in areas such as healthcare, electronic commerce, and finance.



MACHINE LEARNING

AI Faculty: Satinder Singh Baveja, David Fouhey, Danai Koutra, Benjamin Kuipers, Sindhu Kutty, John Laird, Honglak Lee, Rada Mihalcea, Emily Mower Provost, Edwin Olson, Jenna Wiens

Affiliated Faculty: Long Nguyen (Statistics), Ambuj Tewari (Statistics), Ji Zhu (Statistics)

Research in machine learning at Michigan AI encompasses unsupervised and supervised learning, as well as deep learning and reinforcement learning. A number of research projects derive from an interest in building long-lived and flexibly-competent agents rather than the more usual agents that perform one complex task repeatedly. We work on deep learning and representation learning, where we learn abstract representations of the data by exploiting its hierarchical and composition structure. In reinforcement learning, we focus on building autonomous agents that can learn to act in complex, sequential, and uncertain environments. We are also interested in the application of these algorithms to problems in computer vision and language processing, and more generally to sensory information processing and perception, to analyses of large-scale interconnected data (including social, communication, brain networks), and downstream tasks in medical domains or social sciences.

COMPUTATIONAL HEALTHCARE & ASSISTIVE TECHNOLOGIES

AI Faculty: Nikola Banovic, Satinder Singh Baveja, Edmund Durfee, Danai Koutra, Walter Lasecki, Honglak Lee, Rada Mihalcea, Emily Mower Provost, Veronica Perez-Rosas, Jenna Wiens

Affiliated Faculty: Kayvan Najarian (CMB)

The recent increase in the availability of clinically relevant datasets provides unique opportunities for improving patient care. Such data arise from diverse sources including, but not limited to, individual patient health records, genomic data, data from wearable sensors, clinical notes, and medical imaging. Michigan AI research in this area encompasses computational methods from machine learning, computer vision, and natural language processing, and spans the entire learning pipeline from problem formulation, to model learning and output. Working in close collaboration with clinicians, we address important healthcare problems ranging from patient risk stratification and clinical time-series prediction to models of medication adherence and methods for AI-enhanced counseling. We also create assistive technologies by employing automated

planning, constraint-based and temporal reasoning, and probabilistic inference. We build systems that help cognitively-challenged people with complex healthcare needs to schedule and remember daily activities; that monitor the behavior of a person at home; or that adapt computer interfaces to meet the needs of persons with visual or fine-motor impairment. With its strong ties to Michigan Medicine and unparalleled computational and data resources, Michigan AI provides a collaborative environment for carrying out high-impact research at the cutting edge of machine learning, data mining, and medicine.

AI IN HUMAN ENVIRONMENTS

AI Faculty: Nikola Banovic, Satinder Singh Baveja, Edmund Durfee, David Fouhey, Chad Jenkins, Benjamin Kuipers, John Laird, Edwin Olson

Michigan AI is innovating technologies that improve how AI systems act with and among people. As examples: driverless vehicles that encounter people (as drivers, pedestrians, etc.) need to anticipate, and sometimes influence their behaviors. A robotic healthcare assistant needs to communicate and learn about a human's evolving needs. AI systems acting autonomously on behalf of users (e.g., financial transactions or telepresence) need to understand bounds on permissible behavior in different social contexts. To address such needs, we conduct research in areas including human-robot interaction (HRI), human behavior recognition, and safety and ethics in AI.

ROBOTICS

AI Faculty: Satinder Singh Baveja, Jason Corso, Edmund Durfee, David Fouhey, Chad Jenkins, Benjamin Kuipers, John Laird, Walter Lasecki, Edwin Olson

Affiliated Faculty: Ella Atkins (Aerospace)

Michigan AI is investigating both theoretical and practical aspects of robots, including aerial, underwater, space, and terrestrial systems. Key areas include: 1) integration of strategic and tactical planning and optimization algorithms to enable robust robot control in the presence of system failures and environmental uncertainties; 2) simultaneous localization and mapping for mobile robots; 3) sensor processing algorithms, including feature matching, object detection, classification, and recognition; 4) AI methods from machine learning, cognitive architectures and multiagent systems to build autonomous robots.



COMPUTER VISION

AI Faculty: Jason Corso, David Fouhey, Chad Jenkins, Benjamin Kuipers, Honglak Lee, Edwin Olson

Michigan AI explores a number of critical problems in the area of computer vision. We focus on the analysis and modeling of visual scenes from static images as well as video sequences. Our research goals include: 1) the semantic understanding of materials, objects, and actions within a scene; 2) modeling the spatial organization and layout of the scene and its behavior in time.

The algorithms developed in this area of research enable the design of machines that can perform real-world visual tasks such as autonomous navigation, visual surveillance, or content-based image and video indexing.

NATURAL LANGUAGE PROCESSING

AI Faculty: Satinder Singh Baveja, John Laird, Walter Lasecki, Honglak Lee, Rada Mihalcea, Veronica Perez-Rosas

Affiliated Faculty: Steve Abney (Linguistics), Kevyn Collins-Thompson (SI), David Jurgens (SI), Qiaozhu Mei (SI), Vinod Vydiswaran (LHS)

Michigan AI is interested in large-scale natural language processing, including word representations, text similarity, information extraction, text summarization, question answering, and applications of these methods to a variety of domain problems. Our techniques range from semantic analysis and parsing to semi-supervised learning and graph-based methods. Some current projects include the automatic identification of similarity relations in text, the development of neural network methods for effective word representations, the development of conversational systems, the extraction of entity relations from text, the automatic detection of deception in text, the identification of attitude and sentiment in online social network discussions.

HUMAN COMPUTATION & CROWDSOURCING

AI Faculty: Sindhu Kutty, Walter Lasecki

Affiliated Faculty: Mark Ackerman (CSE & SI), Ceren Budak (SI), Michael Nebeling (SI), Grant Schoenebeck (CSE)

Human computation incorporates human intelligence into algorithmic processes. While workflow, process management, and organizational theory has been developed in great

depth previously, the introduction of computer science and computational theory has given us a new lens through which to view organized human efforts, and combine it with the efforts of automated systems. We explore crowdsourcing — the practice of making an open call to a group of workers to complete a task, often through social computing channels — to create deployable, large-scale human computation systems that are far more capable than current automated systems are alone, but that are still available on-demand, at a moment's notice. Michigan AI is advancing the frontiers of what can be accomplished with human computation and crowdsourcing: exploring how to create intelligent systems that respond intelligently, in real-time, and with consistency over time. These systems can then be used to train AI and ML systems in real-world environments.

COMPUTATIONAL SOCIAL SCIENCES

AI Faculty: Nikola Banovic, Danai Koutra, Rada Mihalcea, Michael Wellman

Affiliated Faculty: Ceren Budak (SI), David Jurgens (SI), Daniel Romero (SI)

Michigan AI explores computational solutions to central social sciences problems, targeting diverse fields such as economics, psychology, or political science. We develop temporal models for longitudinal data streams, we work on data- and model-driven approaches for financial markets, we track and disentangle the factors behind observed social effects. Example projects use computational linguistics to identify cross-cultural differences in values and behaviors, agent-based models to assess stability of the financial system, and various techniques to study social roles across time, dynamics of political speeches in US congress, and rumor propagation on Twitter.

SYSTEMS FOR ARTIFICIAL INTELLIGENCE

CSE Faculty: Michael Cafarella, Jason Mars, Trevor Mudge, Lingjia Tang, Thomas Wenisch

Software systems and hardware platforms for scalable, accelerated, and distributed execution of artificial intelligence algorithms. Michigan AI works on key research challenges, including: 1) designing frameworks to ease implementation and accelerate execution of inference and training; 2) developing distributed and scalable execution platforms for artificial intelligence; 3) investigating hybrid client + edge + data center deployment models for intelligent applications; 4) developing hardware accelerators, including reprogrammable and ASIC hardware designs, for emerging learning algorithms.



FACULTY IN ARTIFICIAL INTELLIGENCE



Nikola Banovic
Assistant Professor
nbanovic
2630 Beyster



Satinder Singh Baveja
Toyota Professor of Artificial Intelligence
baveja
3749 Beyster



Jason Corso
Associate Professor
jjcorso
4238 EECS



Edmund Durfee
Professor
durfee
3745 Beyster



David Fouhey
Assistant Professor
fouhey
3777 Beyster



Chad Jenkins
Associate Professor
ocj
3644 Beyster



Danai Koutra
Assistant Professor
dkoutra
3633 Beyster



Benjamin Kuipers
Professor
kuipers
3741 Beyster




Sindhu Kutty
Lecturer III
skutty
2640 Beyster



John Laird
John L. Tishman
Professor of Engineering
laird
3753 Beyster



Walter Lasecki
Assistant Professor
wlasecki
2636 Beyster



Honglak Lee
Associate Professor
honglak
3773 Beyster



Rada Mihalcea
Professor
mihalcea
3769 Beyster



Emily Mower Provost
Associate Professor
emilykmp
3629 Beyster



Edwin Olson
Associate Professor
ebolson
3737 Beyster



Veronica Perez-Rosas
Assistant Research Scientist
vrncapr
2709 Beyster



Arunesh Sinha
Assistant Research Scientist
arunesh
2709 Beyster



Michael Wellman
Lynn A. Conway
Professor of CSE
wellman
3757 Beyster



Jenna Wiens
Assistant Professor
wiensj
3765 Beyster



AFFILIATED FACULTY

Steven Abney - Linguistics

Mark Ackerman - Computer Science and Engineering -
Interactive Systems; School of Information

Ella Atkins - Aerospace Engineering

Ceren Budak - School of Information

Michael Cafarella - Computer Science and Engineering -
Software Systems

Kevyn Collins-Thompson - School of Information

Ryan Eustice - Naval and Marine Engineering

Jean-Baptiste Jeannin - Aerospace Engineering

Matthew Johnson-Roberson - Naval and Marine Engineering

David Jurgens - School of Information

David Kieras - Computer Science and Engineering -
Interactive Systems

Richard Lewis - Psychology

Jason Mars - Computer Science and Engineering -
Hardware Systems

Qiaozhu Mei - School of Information

Trevor Mudge - Computer Science and Engineering -
Hardware Systems

Kayvan Najarian - Computational Medicine and
Bioinformatics

Michael Nebeling - School of Information

Mark W. Newman - School of Information

Long Nguyen - Statistics

Daniel Romero - School of Information

Grant Schoenebeck - Computer Science and Engineering -
Theory

Lingjia Tang - Computer Science and Engineering -
Hardware Systems

Ambuj Tewari - Statistics

Richmond H. Thomason - Linguistics and Philosophy

V.G. Vinod Vydiswaran - Learning Health Sciences

Thomas Wenisch - Computer Science and Engineering -
Hardware Systems

Ji Zhu - Statistics

