Understanding Human Intelligence in the Era of Artificial Intelligence

Ruohan Zhang
The University of Texas at Austin

Humans accomplish complex sensory-motor and cognitive tasks with ease and grace. How does the brain accomplish this? For example, playing video games requires rapidly moving the eyes to perceive relevant information in a dynamic environment, memorizing important object locations, predicting the future state of the moving targets, and making decisions that maximize the expected reward. These cognitive abilities are fundamental for artificial intelligence (AI), but only recently have AI programs been able to play simple Atari games with the skill of a human player. Beyond these limited applications, AI is still far from handling many mundane tasks that we as humans can easily perform daily.

My research lies at the intersection of computer science and cognitive psychology. I explore how human intelligence and machine intelligence are related to each other intending to build human-centered intelligent machines. Although biological brain and machine hardware are inherently different, the underlying mathematical principles, i.e., the *software* that defines general intelligence, may be similar in nature. Biologically inspired AI designs, through reverse-engineering the human brain, have given birth to at least three major breakthroughs in the field of AI: artificial neural networks, imitation learning, and reinforcement learning. They belong to a research field called machine learning, in which statistical learning algorithms allow machines to learn from past experiences.

My research follows this trend and aims to discover the mechanism of intelligence by studying human brains and behaviors and applying these insights to the field of AI. Models that seek to explain human intelligence inspire me to build intelligent machines. Through building AI systems, we can test these models and attempt to reproduce human-level intelligent behaviors, therefore machines can better assist humans in daily life. More importantly, designing machines to be human-like facilitate communication between humans and machines hence they can better assist humans in daily life.

Modeling Human Visual Attention and Decisions

One outstanding intelligent mechanism that is unique to humans is *visual attention* – the ability to allocate cognitive resources on important things, which can be revealed by human eye movements. Human eyes have high-resolution vision only in the central 1-2 visual degrees of the visual field, known as the fovea, covering only the width of a finger at arm's length. To compensate for the limited area of high visual acuity, humans learn to move their eyes to direct the foveae to the correct place at the right time to process important task-relevant visual information. A wealth of information is encoded in this type of human gaze behavior which reveals the momentary attentional priorities of the observer.

Unlike human eyes, machines use cameras that have a full resolution at every pixel. Despite this uniform visual acuity, for a given task, not all visual information is relevant at the current time and place. Redundant information poses a heavy burden on computational resources for AI systems. We conjecture that machines could learn an attention mechanism from humans to discard irrelevant information and prioritize resources for the current task, and that integrating this mechanism into the machines will make them a lot more efficient.

Learning visual attention from human gaze I am particularly interested in visual attention for reward-seeking visuo-motor tasks such as game playing and navigation. Games are ideal as a starting point since they are simple but capture many features of human cognition. With high-speed eye trackers, we have collected a large dataset of human subjects' eye movements together with their keyboard actions while playing Atari video games [19]. Our lab has also collected a dataset for human subjects walking on a rough rocky terrain with eye and body movements recorded. Utilizing these datasets, we designed a deep neural network model[1] to learn the mapping between the visual stimulus and the corresponding human gaze position [13, 15]. The result has shown that, given a novel game or outdoor image, the model was

[1]Inspired by biological neurons, computer scientists have implemented artificial neurons in software and thereby constructed what is called an artificial neural network. It can theoretically learn any function that maps inputs to outputs given sufficient example data.

able to predict where the human player or walker would allocate gaze with very high accuracy. This indicates that the machine has learned to distinguish important vs. unimportant visual information in these images.

Learning decisions from humans Learning attention from humans is only the first step towards having an AI agent that can interact with the world by taking action. Using a machine learning technique called imitation learning, AI algorithms can learn how to perform a task from human-demonstrated actions by imitating the human teacher. However, this type of learning can only capture *what* the human teacher did, without knowing the reason *why* the decision was made. Naturally, visual attention is a good indicator of why a particular decision was made. An example from our data shows a human subject directing gaze on an obstacle, which is a strong indication to justify the subject's next action to turn right to avoid it. We conjecture that attention learning and decision learning should be combined, and have proposed an Attention-Guided Imitation Learning (AGIL) framework [15, 14]. We have shown that incorporating attention into a basic decision-learning algorithm, called behavioral cloning, led to a performance increase of 115% (in terms of game scores) for game AIs. A followed-up work has shown that attention can also improve the performance of more advanced decision-learning algorithms, such as inverse reinforcement learning, behavioral cloning from observation [9], and potentially deep reinforcement learning [11]. Thus a fundamental aspect of human cognition, visual attention, is an important aspect of efficient information processing. Visual attention can be treated as a type of human guidance information provided to AI learning systems that make AIs more human-like and efficient [18].

Modeling Human Intrinsic Reward

Reward is a fundamental mechanism of intelligent human behavior that drives attention and decisions. Gaze and other actions are a means to accomplish certain behavioral goals that are defined through rewards. Reinforcement learning is the subfield of AI that studies reward-seeking behaviors. Given the observed behavior, an algorithm called inverse reinforcement learning (IRL) can be used to guess the internal reward function that explains that behavior. Therefore, revealing the underlying reward mechanism that resides in the brain can help us as well as machines better understand and predict human actions.

In [21] we used virtual reality and motion capture technologies to record human navigation behaviors. The participants needed to approach targets, avoid obstacles, and stay close to a path. We assumed that complex, multitasking human behaviors can be decomposed into simple behavior *modules* such as obstacle avoidance and path following [17, 8]. Each module represents a simple goal with unique rewards. Human attention, which is limited in capacity, plays the role of scheduling and coordinating these modules to collectively make decisions to gain rewards. This is known as the hierarchical module hypothesis and has been studied extensively by our lab [2, 3]. By estimating the reward for each module, we can understand why the human makes a certain navigation decision. We formalized and extended a previous IRL algorithm called modular inverse reinforcement learning [8]. Given observed human data, the algorithm can estimate rewards and discount factors (the degree of impulsiveness in seeking future reward) for each module. To verify the results, we showed that an AI avatar using estimated rewards and discount factors to make decisions was able to navigate the virtual rooms in a human-like manner. The importance of this work is that we are now able to model human unconstrained behavior in a naturalistic environment.

Enabling machines to infer human decisions is particularly useful in the scenario of human-robot interaction. As follow-up research, work with colleagues [6] used the modular IRL model as the pedestrian's cognitive model to predict their walking directions. With this information, robots can anticipate human movements and safely navigate around pedestrians.

Neural Communication Models

An important problem in understanding the brain is that we have only a limited understanding of how a particular cognitive operation, such as attention, is implemented neurally. We look to mammalian nervous systems to study how attentional mechanisms are biologically implemented and realized. Recently, we have proposed a novel neural coding theory which explains how neurons communicate with each other to receive relevant signals and filter out irrelevant ones [1]. Intuitively, the theory hypothesizes that neurons communicate through 40-80 hertz gamma frequency in a manner similar to radio stations. A group of neurons designated to a particular computational process could tune to a particular frequency in that range. It is, therefore, possible to form multiple separate networks that might constitute neural trains of thought that can be kept from crosstalk. Initial analysis of awake mouse neural recording data tends to support this theory. This model suggests that the attention mechanism may reflect the use of additional neurons by a computational process by switching these neurons to the frequency band associated with that process. If this theory is borne out by additional experiments, it would constitute a significant advance in both neural coding and artificial neural networks.

Application-Oriented Interdisciplinary Research

Interdisciplinary collaboration has played a major role in my work since artificial intelligence is a powerful tool that could potentially provide solutions for problems in many fields. The aforementioned work could not be accomplished without the expert advice provided by neuroscientists and psychologists. Additionally, collaborating with researchers in statistics and engineering, we have advanced the tools of AI from both a theoretical perspective, e.g., optimization theory [10, 5], as well as an application-oriented perspective, such as robot soccer [4, 7]. Through these interdisciplinary collaborations, I gradually developed an understanding of how AI could have a positive impact on other research fields. Here I highlight some of my projects that aim at using AI technologies to tackle critical social issues in various domains using large-scale computer simulation.

Traffic management In [20], We studied vehicle routing problem for traffic congestion reduction. Assuming each vehicle knows 1) the capacity of roads and intersections and 2) the starting positions and destinations of other vehicles on the road, we developed a route-planning algorithm based on the reinforcement learning model. The algorithm allows vehicles to navigate towards roads and intersections that have low usage. The core idea is to model the traffic as a density control problem. The route planning algorithm calculates how densities of vehicles at different locations would evolve over time given each vehicles planned route. It then inversely induces routes for each vehicle to control the collective traffic density under road capacity. Experiments in simulation with up to 24,000 vehicles showed that the proposed algorithm reduced traffic congestion compared to ordinary route planning algorithms by a factor of 12 in terms of travel time. This solution is also scalable; given fixed road capacity, the advantage of the algorithm increases as the number of vehicles increases.

Sustainable harvesting The increasing level of autonomy in the foraging industry (e.g., fishing and logging) presents new challenges to environmental sustainability. While efficient, autonomous foraging machines could damage natural resources by overharvesting. These machines require adaptive, intelligent control algorithms to monitor and respond to environmental changes that could indicate overharvesting while working without human supervision. In a simulation, assuming a Verhulsts logistic growth model of natural resources, we designed a control pipeline for a team of foraging robots to harvest regenerative natural resources [16]. The algorithm enables the robots to gather information about the resource being harvested, monitor its status, and adjust harvesting rate to approximate a theoretically optimal point (the maximum sustainable yield) to attain maximum economical profits while ensuring sustainability. Compared to previous studies, this work focuses on the uncertainty in the environment by making robots adaptive and intelligent.

Public opinion about artwork The Participatory Art Museum movement encourages public engagement of non-expert visitors to break the tradition in which visitors are perceived as passive recipients. [12] aims to develop scalable computational tools that help non-profit organizations with limited resources to effectively gather and interpret public opinions and encourage two-way communications. We crowdsourced non-expert interpretations of 21 artworks by artists with distinct period styles. Art historians provided interpretations from books, museum catalogs, and Wikipedia. All these opinions are projected as opinion clouds in a reduced mathematical semantic space, using standard semantic embedding and dimension reduction algorithms in natural language processing research. This technique allows one to visualize with ease that how far or close an opinion is from another in terms of semantics. The hope is that, for the connoisseurs, crowd judgments help to correct homogeneity-generated biases by introducing diverse perspectives and innovative interpretations. For the non-expert viewers, the crowd interpretation presents an interesting alternative that invites an open and dynamic dialogue in which one can see where one's opinion is among others'.

References

[1] Dana Ballard and Ruohan Zhang. Cortical spike multiplexing using gamma frequency latencies. *bioRxiv*, page 313320, 2018.

[2] Dana H Ballard. *Brain computation as hierarchical abstraction*. MIT Press, 2015.

[3] Dana H Ballard and Ruohan Zhang. The hierarchical evolution in human vision modeling. *To appear in Trends in Cognitive Sciences*, 2020.

[4] Katie Genter, Patrick MacAlpine, Jacob Menashe, Josiah Hannah, Elad Liebman, Sanmit Narvekar, Ruohan Zhang, and Peter Stone. Ut austin villa: Project-driven research in ai and robotics. *IEEE Intelligent Systems*, 31(2):94–101, 2016.

[5] Xiangru Huang, Ian En-Hsu Yen, Ruohan Zhang, Qixing Huang, Pradeep Ravikumar, and Inderjit Dhillon. Greedy direction method of multiplier for map inference of large output domain. In *Artificial Intelligence and Statistics*, pages 1550–1559, 2017.

[6] Sebastian Junges, Nils Jansen, Joost-Pieter Katoen, Ufuk Topcu, Ruohan Zhang, and Mary Hayhoe. Model checking for safe navigation among humans. In *International Conference on Quantitative Evaluation of Systems*, pages 207–222. Springer, 2018.

[7] Jacob Menashe, Josh Kelle, Katie Genter, Josiah Hanna, Elad Liebman, Sanmit Narvekar, Ruohan Zhang, and Peter Stone. Fast and precise black and white ball detection for robocup soccer. *Robot Soccer World Cup XXI*, 2017.

[8] Constantin A Rothkopf and Dana H Ballard. Modular inverse reinforcement learning for visuomotor behavior. *Biological cybernetics*, 107(4):477–490, 2013.

[9] Akanksha Saran, Ruohan Zhang, Elaine Schaertl Short, and Scott Niekum. Efficiently guiding imitation learning algorithms with human gaze. *arXiv preprint arXiv:2002.12500*, 2020.

[10] Ian En-Hsu Yen, Xiangru Huang, Kai Zhong, Ruohan Zhang, Pradeep K Ravikumar, and Inderjit S Dhillon. Dual decomposed learning with factorwise oracle for structural svm of large output domain. In *Advances in neural information processing systems*, pages 5030–5038, 2016.

[11] Liu Yuezhang, Ruohan Zhang, and Dana H Ballard. An initial attempt of combining visual selective attention with deep reinforcement learning. *arXiv preprint arXiv:1811.04407*, 2018.

[12] Xiaoyu Zeng and Ruohan Zhang. Participatory art museum: Collecting and modeling crowd opinions. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. AAAI Press, 2017.

[13] Luxin Zhang, Ruohan Zhang, Zhuode Liu, Mary M Hayhoe, and Dana H Ballard. Learning attention model from human for visuomotor tasks. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.

[14] Ruohan Zhang, Zhuode Liu, Mary M Hayhoe, and Dana H Ballard. Attention guided deep imitation learning. *Cognitive Computational Neuroscience*, 2017.

[15] Ruohan Zhang, Zhuode Liu, Luxin Zhang, Jake A Whritner, Karl S Muller, Mary M Hayhoe, and Dana H Ballard. Agil: Learning attention from human for visuomotor tasks. *European Conference on Computer Vision (ECCV)*, 2018.

[16] Ruohan Zhang and Zhao Song. Maximum sustainable yield problem for robot foraging and construction system. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 2725–2731. AAAI Press, 2016.

[17] Ruohan Zhang, Zhao Song, and Dana H Ballard. Global policy construction in modular reinforcement learning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*. AAAI Press, 2015.

[18] Ruohan Zhang, Faraz Torabi, Lin Guan, Dana H Ballard, and Peter Stone. Leveraging human guidance for deep reinforcement learning tasks. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 6339–6346. AAAI Press, 2019.

[19] Ruohan Zhang, Calen Walshe, Zhuode Liu, Lin Guan, Karl S Muller, Jake A Whritner, Luxin Zhang, Mary M Hayhoe, and Dana H Ballard. Atari-head: Atari human eye-tracking and demonstration dataset. In *To appear in Thirty-Fourth AAAI Conference on Artificial Intelligence*. AAAI Press, 2020.

[20] Ruohan Zhang, Yue Yu, Mahmoud El Chamie, Behçet Açikmeşe, and Dana H Ballard. Decision-making policies for heterogeneous autonomous multi-agent systems with safety constraints. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*, pages 546–552. AAAI Press, 2016.

[21] Ruohan Zhang, Shun Zhang, Matthew H Tong, Yuchen Cui, Constantin A Rothkopf, Dana H Ballard, and Mary M Hayhoe. Modeling sensory-motor decisions in natural behavior. *PLoS computational biology*, 14(10):e1006518, 2018.