Algorithmic Challenges in Interactive Learning with Users

Nakul Gopalan

Arizona State University
ng@asu.edu

Abstract

Current thrusts in Artificial Intelligence (AI) and robotics are focusing on learning systems that learn with significantly large amounts of data while using training methodologies or pipelines developed by experts. There needs to be a larger push towards systems that learn from users in real time with interactions that are possible in a short duration, realistic and accessible to human users. While this is possible with the current trend towards large data driven models, the philosophy around algorithm development and experimentation with learning models needs significant work. In this work we specify some algorithmic challenges in developing AI systems that can learn from people.

Building systems that work with people has been left as the domain of researchers who work on developing user interfaces and improving user experiences. Meanwhile the focus of AI, machine learning (ML) and Robotics have been in developing algorithms that enable agents to learn desired behavior. However, for wide scale adoption of learning systems considerations must be given to expectations of a human users on these learning systems. In this work we list a few algorithmic challenges (of interest robotics specifically) in making a learning system functional for human users.

Real Time Feedback

It is critical to have feedback of a behavior being taught so a user can assess the state of an agent's understanding of the behavior being taught. Humans teaching tasks to each other (or to animals) communicate, and expect to understand where the learner needs help. This is an instantaneous feedback that a teacher expects whether it is check if a dog can sit down or a worker knows how to operate a machine. For example, when teaching a robot to pick an object or teaching a video game agent to solve a level, a user teaching the behavior might expect to test the behavior at any point during the teaching process. Robot Learning or Reinforcement learning is not at a place where users can expect real time, instantaneous feedback. The collected data is processed by an algorithm whose *time complexity* might be too large for instantaneous deployment of the learned behavior. This issue of time complexity lies within the learning algorithm and also the deployment of the agent itself. It might take too long for a robot to update and use new policy weights, which itself hurts the user's experience and might disincentivize the users from teaching or using the robot.

AAAI 2024 Spring Symposium on User-Aligned Assessment of Adaptive AI Systems. Stanford University, Stanford, CA, USA.

Communication and Doubt

Agents need to communicate what they know and what they do not know. This is a challenging question as it presumes some knowledge about unseen tasks that agents are learning. It is a chicken and an egg problem but humans solve it all the time. We ask uncertain questions about tasks that we are learning. We are also taught the art of asking good questions in formal settings with explicit feedback about the quality and the hardness of the questions asked. Moreover, we also communicate when are confident about our outcomes of learning. Agents have no such mechanisms to express doubt. It is unclear how to build these systems in a large scale even with sufficient large scale data. This is a more fundamental question about information gain with explicit actions, which has been modeled previously with active learning, but lacks a modern deep learning counterpart.

Sample Complexity

This is an obvious challenge to machine learning in general. Current ML methods are data hungry and their ability to generalize in low data scenarios is yet to formalized. This problem is important from the perspective of personalization for users and from the perspective of correcting an agent from performing incorrect behaviors. The work in fine-tuning large models towards novel behaviors is critical here. We need to demonstrate that find manipulation tasks or dynamic skills within robots can be learned with fine-tuning with few sample points. Older methods with more structure (based on representations such as fixed point attractors or Gaussian Processes) demonstrated significant savings in sample complexity but lacked the rich input space of neural networks with visual and 3-D scenes. Perhaps a balance in providing structure to networks is required in learning sample efficient behaviors for agents.

Subjective vs Objective metrics

Safety and Trust are widely studied for any system in the human factors community. Unfortunately, these are subjective metrics which are collected from users after they have used a system. There is a challenge in using subjective metrics to optimize a learning algorithm. We know how to use objective metrics to improve performance, but we lack methods to improve upon subjective metrics. Humans provide such corrective feedback with language without explicitly providing an objectively correct solution. We need methods that can consume subjective human feedback and learn to perform the right thing.

Robustness of Algorithms and Systems

Most user studies with learning systems are relatively short in duration. A few minutes or hours is insufficient in observing what preferences users might have with robots when working with them in long term settings. For example, users might prefer teaching a robot novel tasks for the first few days of getting a novel robot. However, this novelty might wear off and the users might completely stop using the robot because the learning or teaching mechanism was unrealistic for a user. We lack robust algorithms that can continue to function over days on robot like systems. This lack of robustness is a significant challenge to the progress of the field of robotics, but also in turn to the progress of the field of AI. Without such robustness it would be difficult to understand user preferences for systems.

Et cetera

While this work lists some challenges that the authors believe are important, this is by no means a complete list. There are other challenges such as user modelling, explainability, mode of interaction and teaching, and quantifying the learning effect within users which are critical in attaining functional agents that can learn interactively from users.