

Pranav Rajbhandari

Research Statement

My research focuses on developing ML algorithms for strategic scenarios. By exploring complexities in both adversarial and cooperative systems, I plan to advance AI's ability to function effectively in these environments.

Strategic Agents:

Many of my research projects focus on developing AI agents for competitive settings, providing a foundation for more intricate multiagent interactions.

In the Reliable Autonomous Systems Lab (RASL) at Carnegie Mellon University (CMU), I designed AI agents to play Jenga, a non-deterministic, adversarial game. By adapting existing methods like AlphaZero and Monte Carlo Tree Search, I created agents capable of strategically choosing which blocks to remove, keeping the tower stable while making the opponent's moves difficult.

I continue this work with my current AlephZero project, which extends AlphaZero's algorithm to a wider class of 'board games' whose board size varies throughout a game. I approach solving these games by using a seq-to-seq model to create a context-dependent encoding of each board position, trained with AlphaZero's algorithm. This allows AlephZero to learn games such as 'Jenga' and '5D Chess with Multiverse Time Travel', which motivated my research.

These projects focus on single agent strategic settings where the agent must perform well against various different opponents. They have potential applications in fields like human-robot interaction.

Multiagent Systems:

While at the Naval Research Laboratory (NRL), I worked on several projects exploring multiagent and swarm control. One project involved utilizing a neuroevolution algorithm to produce emergent swarm behaviors from individual agent policies. Inspired by natural insect swarms exhibiting complex behaviors, I adapted an existing neuroevolution algorithm for swarm control to imitate this.

Later, I collaborated on a multiagent path planning project, where we developed an algorithm to maximize information gain from a set of targets. This project focused on efficiently calculating routes and dwell times to best optimize a metric, using methods inspired by approximate solutions to the Traveling Salesperson Problem.

In my most recent project, I developed BERTeam, a framework designed to select strong teams of agents to play multiagent adversarial games. Inspired by Natural Language Processing sequence generation techniques, BERTeam uses samples of effective team configurations to select a set of agents for a cooperative team. Combined with coevolutionary reinforcement learning, BERTeam selects from a diverse population of trained agents, to form teams robust to a variety of opponents. This project combines the cooperative task of creating a coherent team with the strategic task of constructing teams to best respond to opponent behaviors.

Overall, my work in this area advances the coordination of groups of agents in complex environments, and has applications in domains such as disaster response, search and rescue, and team games like robotic capture-the-flag (which inspired BERTeam).

Reinforcement Learning and Robotics:

My research also concerns robotic control through Reinforcement Learning (RL). At RASL, I focused on training a robotic grasping task in simulation. I used CycleGAN as a 'real-life' filter, facilitating transfer to a physical robot by training on more realistic images. In a similar project at the NRL, I explored transfer learning in the context of robotic manipulation. Through this project, I compared strategies for transferring knowledge across RL agents trained on related tasks.

Currently, I am collaborating with the University of New South Wales to optimize the locomotion of a simulated swimming robot. After obtaining initial parametrized locomotion by approximating the

swimming of mosquito larvae, we optimized the parameters through local search. I guided the search with RL, taking advantage of the episodic nature of evaluations to efficiently approximate a local gradient.

Overall, my work in this area is focused on improving efficient training methods to develop robotic agents more adept at complex tasks, with applications in domains like manufacturing and human-robot interaction.

Future Research:

My research will focus on enhancing strategic AI's utility in multiagent systems. Broad domains that interest me include adversarial game playing and multiagent cooperation.

Strategic agents in adversarial settings: Adversarial settings (particularly board games) are commonly used to measure improvements in AI, as they directly compare strategic ability of competitors. I plan to further explore my work in this area, expanding the classes of games that agents are able to master. I am specifically interested in imperfect information games (e.g. poker and other card games), where a strategic agent must reason about the knowledge other players have, and behave accordingly. This creates a more challenging scenario, requiring complex and innovative approaches. Additionally, solutions can be more applicable to real-world scenarios, where complete information is often not given.

I also plan to explore social/diplomacy games based on communication (e.g. Settlers of Catan, Werewolf). Recent developments in language models allow for stronger understanding of language, as well as more realistic text generation. Developing a strategic planner that interacts with players through language models could allow an agent to better understand and navigate the social dynamics in these games. This work could also contribute to making agents behave more naturally in human-AI interactions.

ML and game theory: Grounding problems in a mathematical framework allows easier theoretical analyses of solutions. I plan to combine my work with my background in discrete mathematics to formulate tasks as game theoretic scenarios, then develop and analyze algorithms to solve them.

One particular topic I plan to explore is the team selection problem in my BERTeam project. This scenario can be rephrased as a normal form game, where coaches each choose a team to compete in a multiagent environment. I plan to update BERTeam's model, loss, and training method to construct an algorithm that approaches a Nash Equilibrium in this game.

Cooperation in complex tasks: As autonomous agents become more common in everyday life, coordination between them will become more useful in solving complex tasks. Continuing my work with multiagent/swarm control, I plan to develop frameworks for cooperation in challenging environments.

One specific direction I would like to explore is learning robotic tasks that must be solved by teamwork between multiple agents. This combines my past work in multiagent cooperation and robotics, and can be applied to areas like autonomous search and rescue and disaster relief as well as manipulation tasks like multiple robot assembly. Improving coordination in these scenarios will greatly improve our ability to automate difficult or dangerous tasks, reducing the risk for human workers.

Broader Impact/Significance:

My research seeks to improve strategic decision making in multiagent systems, creating well-coordinated and adaptable agents. This will lead to better understanding of the abilities and limitations of AI in decision making, as well as potentially providing insight to mechanisms of decision making in humans. Additionally, my work has practical applications in areas that would benefit from strategic collaborative agents, such as healthcare, industry, and autonomous exploration. Overall, I hope to contribute to a future where autonomous agents can act strategically and cooperatively in complex real-world scenarios.