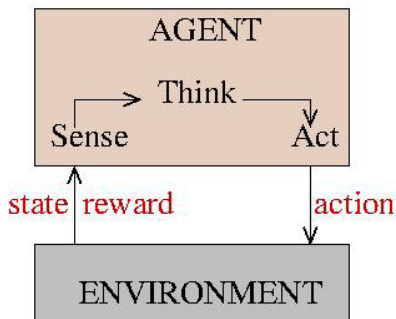


Our group's research is motivated by the goal of creating intelligent agents, especially ones that can *learn*. In pursuit of this goal, we consider questions from a wide variety of topics. Central to our investigation is reinforcement learning (RL), which is a general paradigm for an agent, through trial and error, to discover actions that maximise its long-term gain. RL finds application in a variety of domains, including game-playing, stock-trading, medical decision-making, and environmental preservation. Indeed RL was the key element in training AlphaGo, the program that beat the human champion in the game of Go last year.

As an illustration of RL, consider the task of getting a robot to play soccer. The traditional approach to do so would be for a programmer to directly specify the robot's behaviour in the form of rules such as: 'If farther than 21.5 m from the goal, and within 5.6 m of a teammate who makes an angle of more than 23° with any opponent, then pass the ball to that teammate.' The attraction of RL is that rather than specifying behaviour in this manner, designers may merely specify what is *desired* of the behaviour. Thus, if the soccer-playing robot is rewarded +1 if the ball gets into the opponent's goal, and 0 for every other situation, then by applying an RL algorithm, the robot can eventually learn to play soccer! The technical problem that is solved by RL is determining the credit to be assigned to individual actions for the success or failure of a *sequence* of actions. In general, the outcomes of actions can be stochastic.



Our group contributes to both the theory and the practice of RL. On the theoretical side, we continue to improve the design and analysis of algorithms for formalisms such as Multi-armed Bandits and Markov Decision Problems, which are closely related to RL. On the practical side, we engineer solutions to scale RL to meet the demands of real-world applications. We validate our solutions in domains such as humanoid robotics, on-line advertising, computer games, and robot soccer.