

The efficiency and optimality of robot decision making is often dictated by the fidelity and complexity of models for how a robot can interact with its environment. It is common for researchers to engineer these models a priori to achieve particular levels of performance for specific tasks in a restricted set of environments and initial conditions. As we progress towards more intelligent systems that perform a wider range of objectives in a greater variety of domains, the models for how robots make decisions must adapt to achieve, if not exceed, engineered levels of performance. In this talk I will discuss progress towards model adaptation for robot intelligence, including recent efforts in natural language understanding for human-robot interaction.