Social Hierarchical Learning: Enabling Human-Robot Teaming

Bradley Hayes
Dept. of Computer Science, Yale University

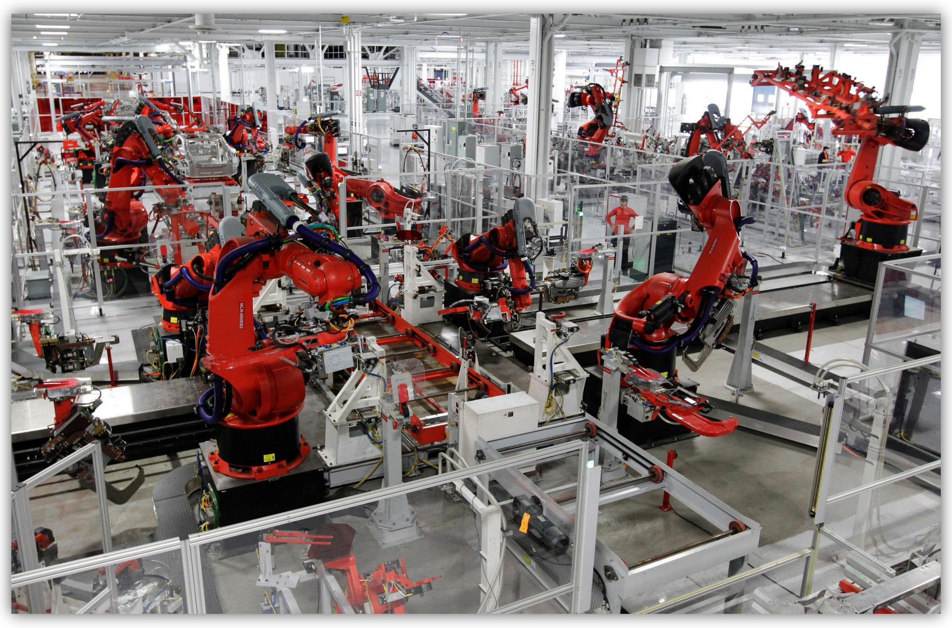

Human-robot teaming has the potential to enable robots to perform well beyond their **current limited and isolated roles**. Many modern robotics advances remain inapplicable in domains where tasks are either too complex to properly encode, beyond modern hardware limitations, too sensitive for non-human completion, or too flexible for static automation practices.

In these situations human-robot teaming can be leveraged to **improve the efficiency, quality-of-life, and safety of human workers**. We desire to create collaborative robots that can provide assistance when useful, remove dull or undesirable responsibilities when possible, and assist with dangerous tasks when feasible.

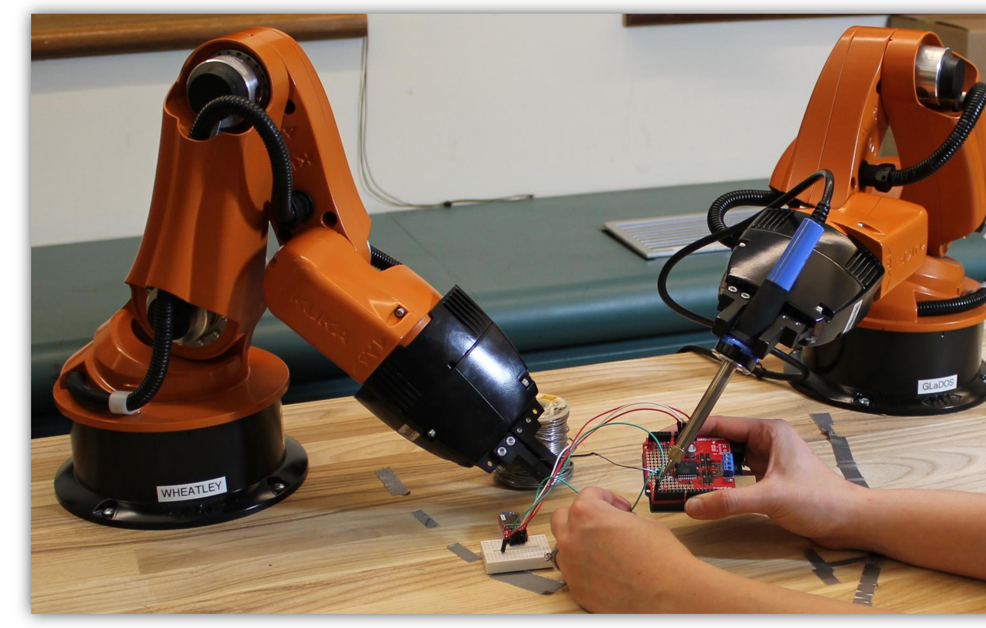

Task Comprehension

Tasks are learned by observing action sequences and building SMDPs from recorded environment states and their associated action-based transitions.

These SMDPs are converted into goal-centric Hierarchical Task Networks, where vertices indicate intermediate goals to be achieved during the task's execution.

Each goal state is representative of a collection of possible actions that may be taken from a valid previous task state to progress the activity.

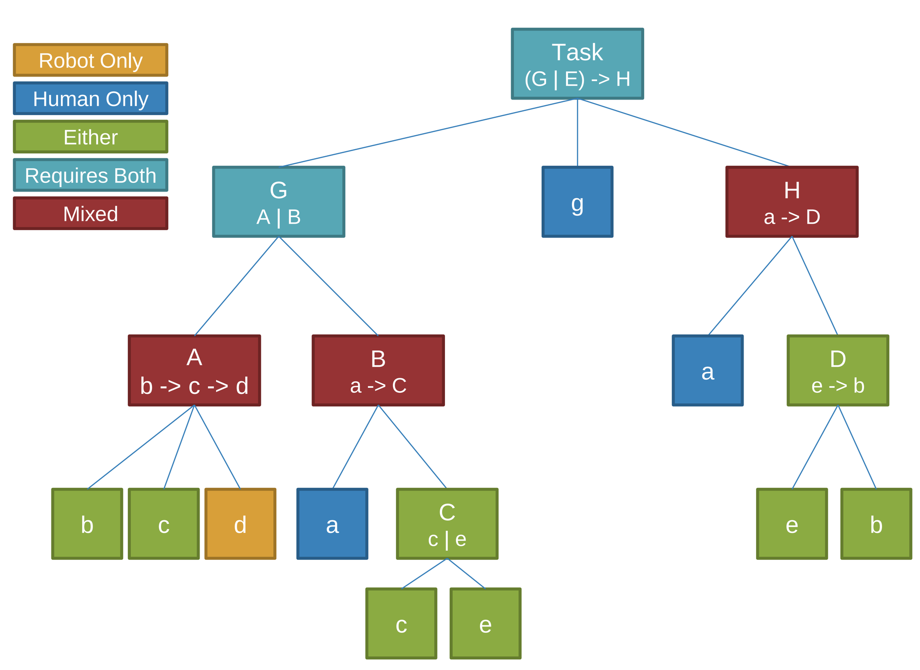

Assistive Behaviors

Joint Object Manipulation	Knowledge Transfer
Task Progression Guidance	Collaborative Tool Use
Materials Retrieval	Materials Stabilization

A taxonomy of assistive behavior types for human-robot collaborative activities.

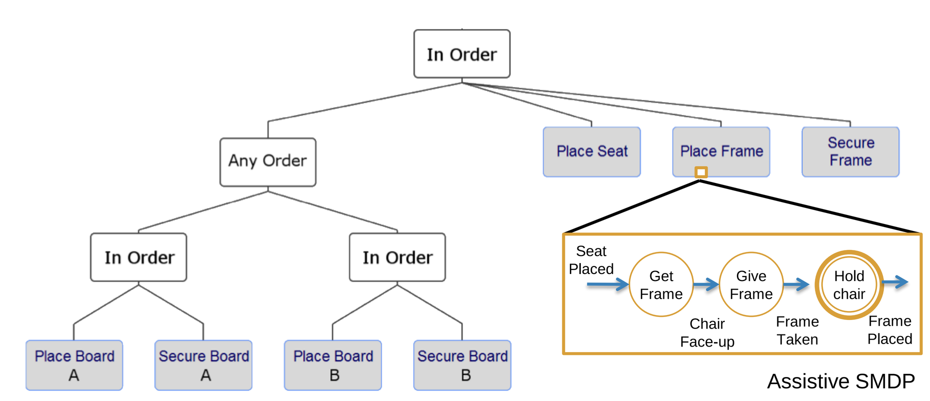

A goal-centric Hierarchical Task Network describing the goal steps for the assembly of an IKEA Chair, with attached assistive action SMDP

Similar to learning tasks, once an HTN is known we learn SMDPs by demonstration for many explicit types of assistive behaviors, associating them with the HTN goal state active at the time of training.

For implicitly discernible types of assistance, we are able to generate motion primitives and behaviors autonomously from primitives within the task network itself.

Collaborative Task Execution

Task execution is modeled as a multi-agent, goal-directed POMDP

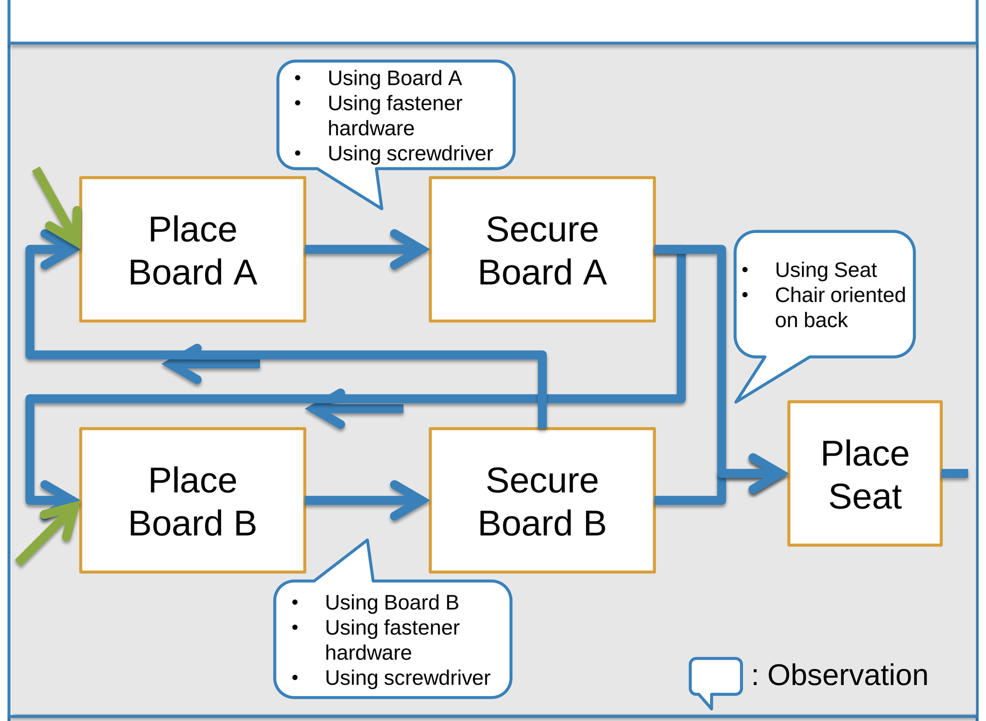

Partial view of goal-directed POMDP for assembling an Ikea Chair. Speech bubbles denote observations from state transitions.

Task progress is measured by determining which task goal the agent(s) are attempting to satisfy. We use active tools, occupied workspace areas, and workpiece features alongside the task network to determine action intent and identify any unexpected deviations.

Teammate preferences are modeled from observations and experiences to determine:

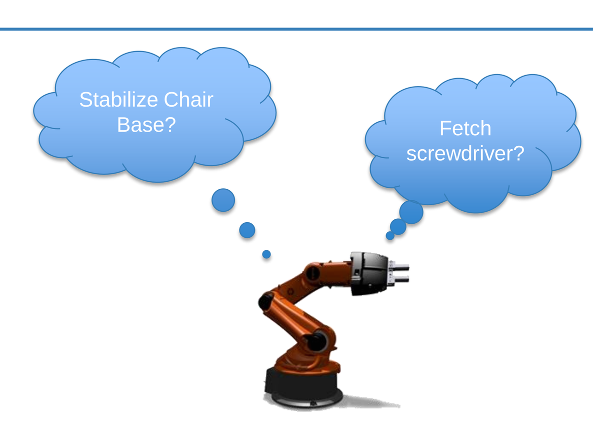

Which assistive actions to perform for a particular user: how does this user prefer to be assisted on each task?

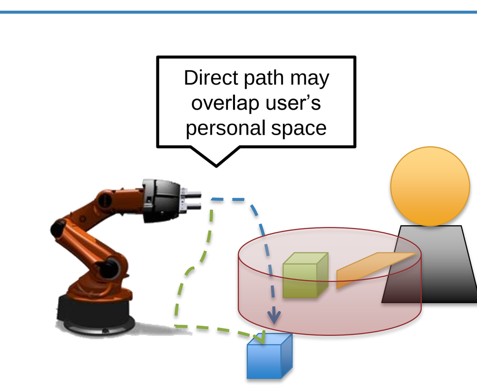

When performing assistive actions: which paths best balance efficiency with respect for the user's safety and comfort tolerances?

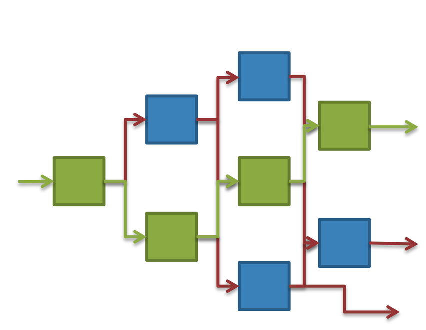

Which action orderings to anticipate from the user: how does this collaborator typically prefer to execute this task?

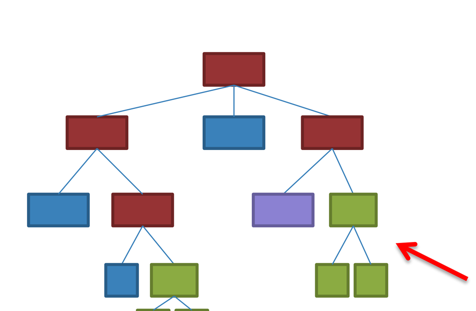

Transferring preferences across similar branches of the hierarchical task network: how can teamwork knowledge be generalized across subtasks?

Actions are safely executed in a shared environment

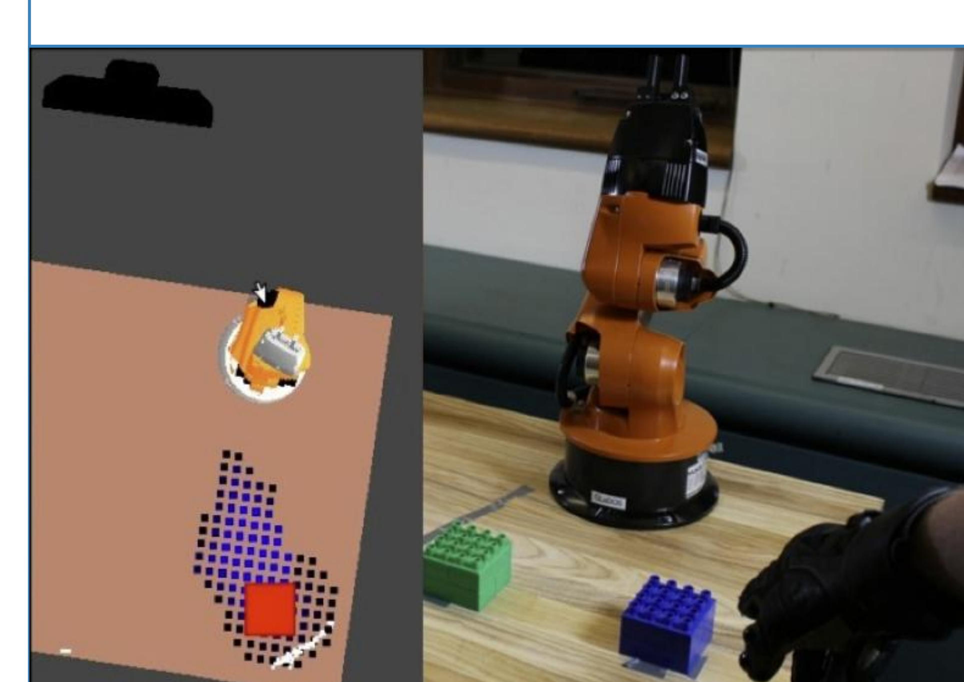

The predicted active workspace of the user is projected into the simulation environment, used to 'trial run' actions before executing them in the real world.

Safe motion primitive execution is accomplished through real-time motion planning with active path correction. Path constraints are integrated into an OMPL-based planner to ensure safe robot operation.

Funded by the Office of Naval Research Grant #N00014-12-1-0822, "Social Hierarchical Learning"