



CSE 591: Human-aware Robotics

Instructor: Dr. Yu (“Tony”) Zhang

Location & Times: CAVC 359, Tue/Thu, 9:00--10:15 AM

Office Hours: BYENG 558, Tue/Thu, 10:30--11:30AM

Nov 10, 2016

Slides adapted from Subbarao Kambhampati, Heni Ben Amor

This set of slides borrows from various online sources; it is used for educational purposes only.

Challenges in human-aware robotics

- **Perception of humans**
Human recognition, human tracking, and activity recognition ✓

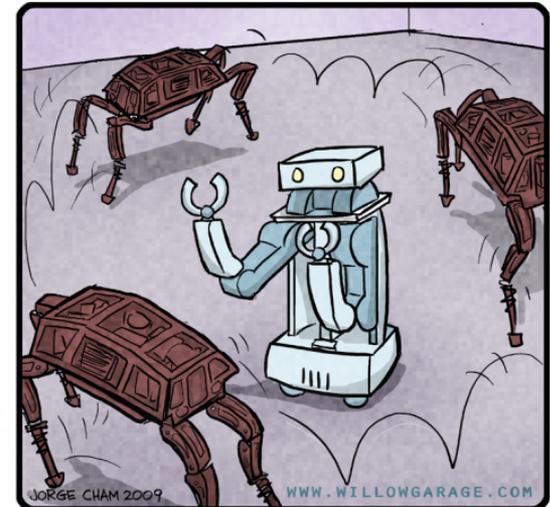


- **Modeling of humans**
Goal and intent recognition, human decision and behavioral models, expectation, model learning ✓



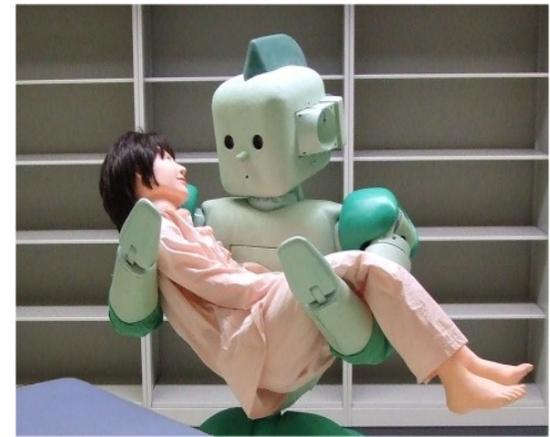
- **Human-robot interface**
Command recognition, gesture recognition ✓

R.O.B.O.T. Comics



"SIT, BOY, SIT! SIT, I SAY,
SI... OH, FORGET IT."

- **Human-aware decision making**
Human-aware planning, reinforcement learning and inverse reinforcement learning.



Dimensions of HIL Planning

Human aware robot planning

	Cooperation Modality	Communication Modality	What is Communicated	Knowledge Level
Crowdsourcing	Interaction (Advice from planner to humans)	Custom Interface	Critiques, subgoals	Incomplete Preferences Incomplete Dynamics
Human-Robot Teaming	Teaming/ Collaboration	Natural Language Speech implicit	Goals, Tasks, Model information	Incomplete Preferences Incomplete Dynamics (Open World)
"Grandpa Hates Robots"	Awareness (pre-specified constraints)	Prespecified (Safety / Interaction Constraints)	No explicit communication Interaction constraints	Incomplete Preferences Complete Dynamics
MAPGEN	Interaction (Planner takes binding advice from human)	Direct Modification of Plans	Direct modifications, decision alternatives	Incomplete Preferences Complete Dynamics

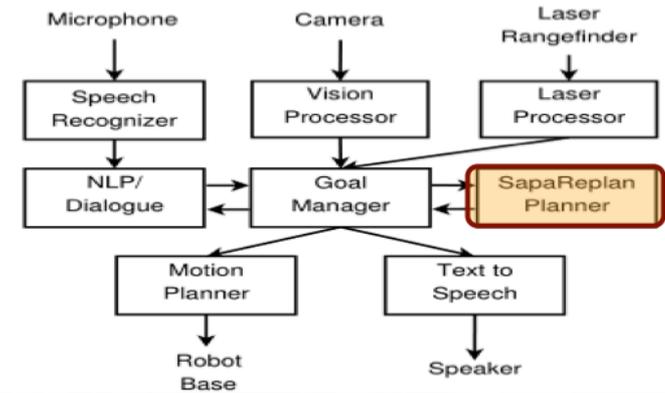
Decision theoretic Assistance,
Resource conflicts and plan for serendipity
Grandpa hates robots



Model Updates (via natural language)

- › “To go into a room when you are at a closed door, push it one meter.”
 - › Precondition: “you are at a closed door”
 - › Action definition: “push it one meter”
 - › Effect: “go into a room”

- › NLP Module
 - i. Reference resolution
 - ii. Parsing
 - iii. Background knowledge
 - iv. Action submission (to planner)



[Cantrell, Talamadupula et al., HRI 2012]

[In collaboration with hrilab, Tufts University]

Dimensions of HIL Planning

Human aware robot planning

	Cooperation Modality	Communication Modality	What is Communicated	Knowledge Level
Crowdsourcing	Interaction (Advice from planner to humans)	Custom Interface	Critiques, subgoals	Incomplete Preferences Incomplete Dynamics
Human-Robot Teaming	Teaming/ Collaboration	Natural Language Speech implicit	Goals, Tasks, Model information	Incomplete Preferences Incomplete Dynamics (Open World)
"Grandpa Hates Robots"	Awareness (pre- specified constraints)	Prespecified (Safety / Interaction Constraints)	No explicit communication Interaction constraints	Incomplete Preferences Complete Dynamics
MAPGEN	Interaction (Planner takes binding advice from human)	Direct Modification of Plans	Direct modifications, decision alternatives	Incomplete Preferences Complete Dynamics

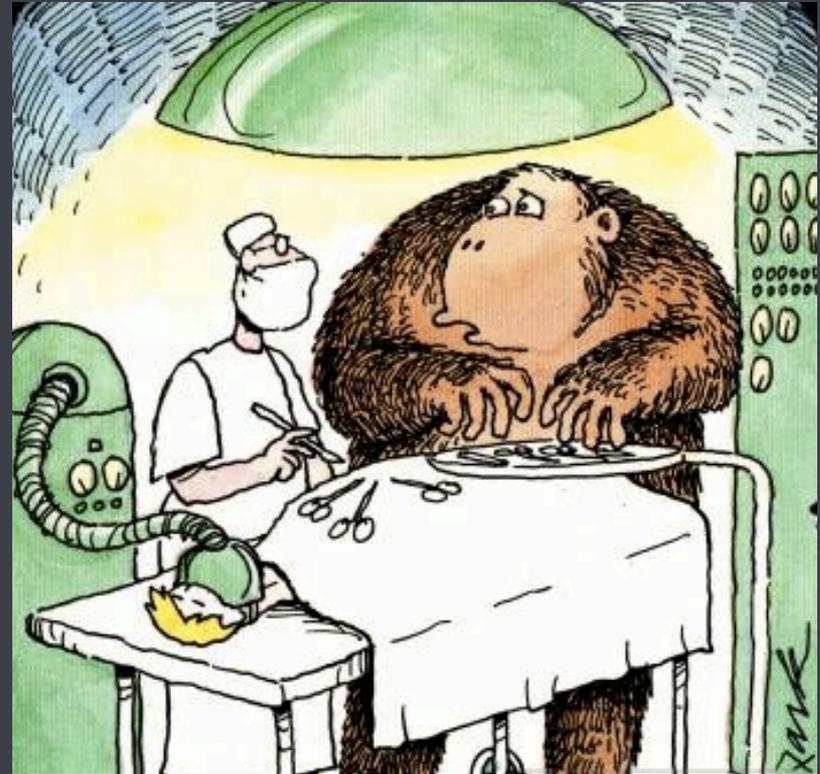
Decision theoretic Assistance,
Resource conflicts and plan for serendipity
Grandpa hates robots + ...

Challenges in Human-Aware Planning & Decision Making

- Interpret what humans are doing
 - Plan/goal/intent/preference/capability recognition ([Expectation](#))
 - Plan with incomplete domain models
 - Robust planning with “lite” models
 - (Learn to improve domain models)
-
- Continual planning/Replanning
 - Commitment sensitive to ensure coherent interaction
 - Explanations/Excuses
 - Excuse generation can be modeled as the (conjugate of) planning problem
 - Asking for help/elaboration
 - Reason about the information value

Explicability among Humans

- We, as humans, while interacting make an effort (behave explicably) to be understood.
- Explicability is crucial for good team performance.



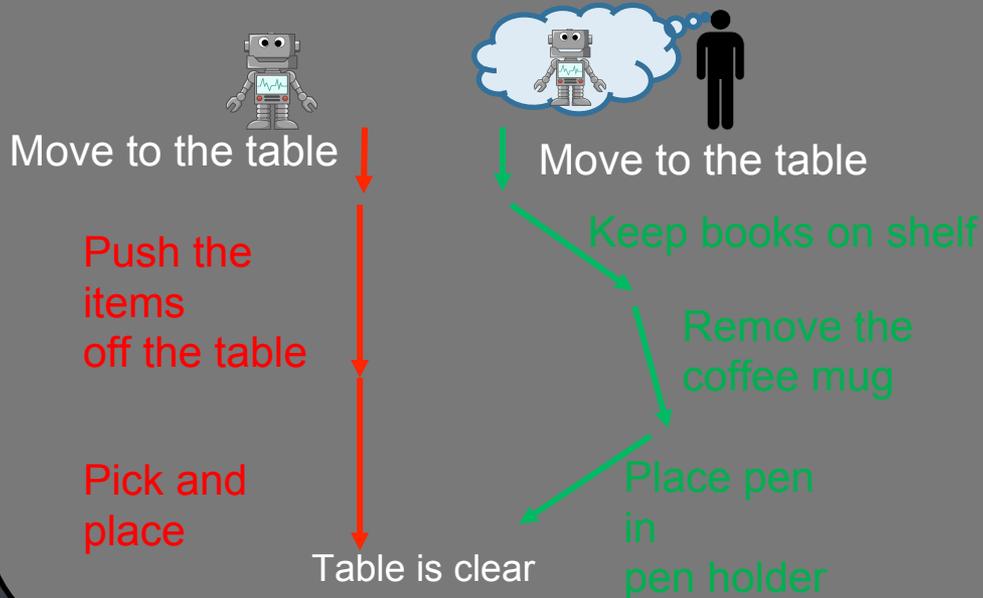
"You clumsy ape! I ask for a hemostat and you hand me a banana. Where'd you go to med school, the Bronx Zoo?"

Explicability

Explicability is not only important for human-human teams but also for human-robot teams.



Goal: Clear a cluttered table

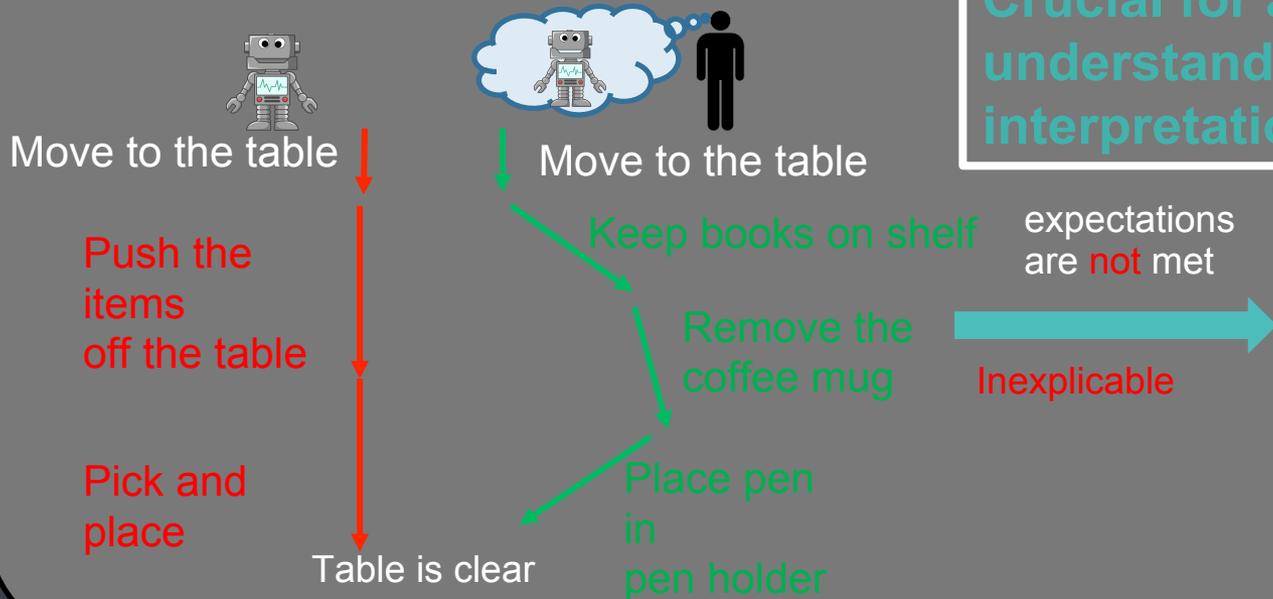


Explicability

Explicability is not only important for human-human teams but also for human-robot teams.



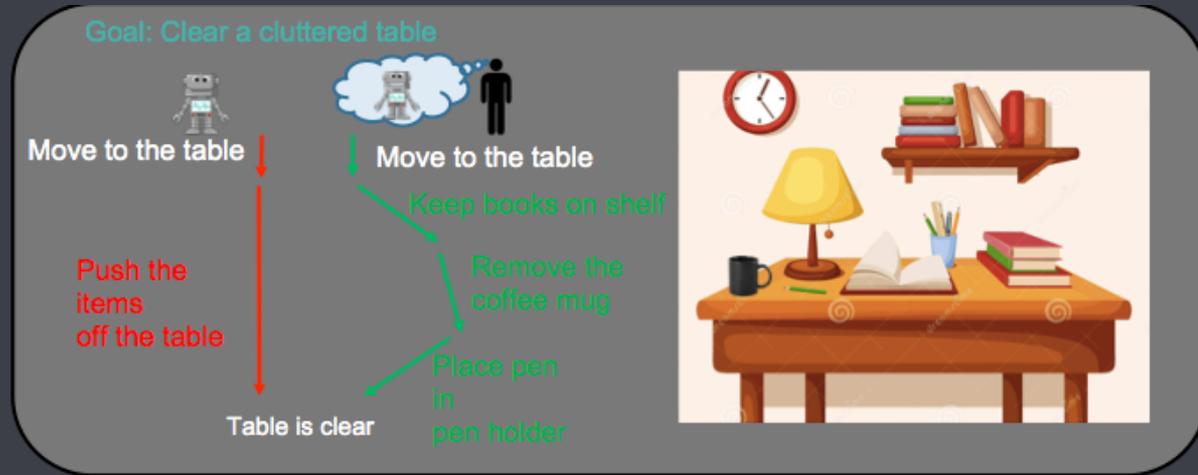
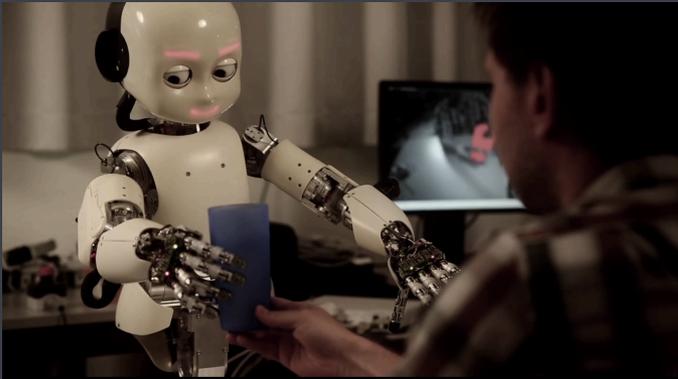
Goal: Clear a cluttered table



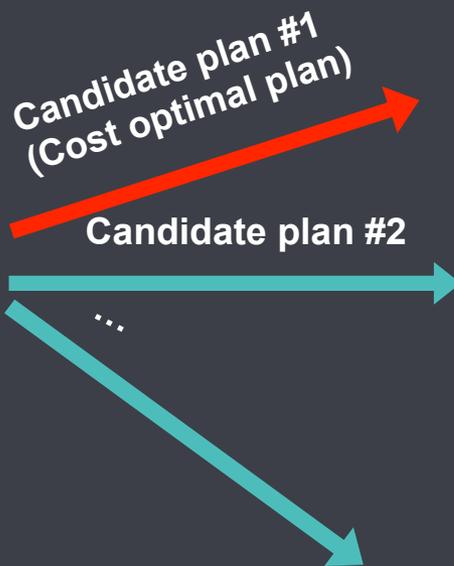
Crucial for an agent to understand other agent's interpretation of itself.

- **Confusion and surprise**
- **Cognitive load**
- **Safety concerns**

Plan Explicability for Autonomous Robots



Autonomously
construct task plans

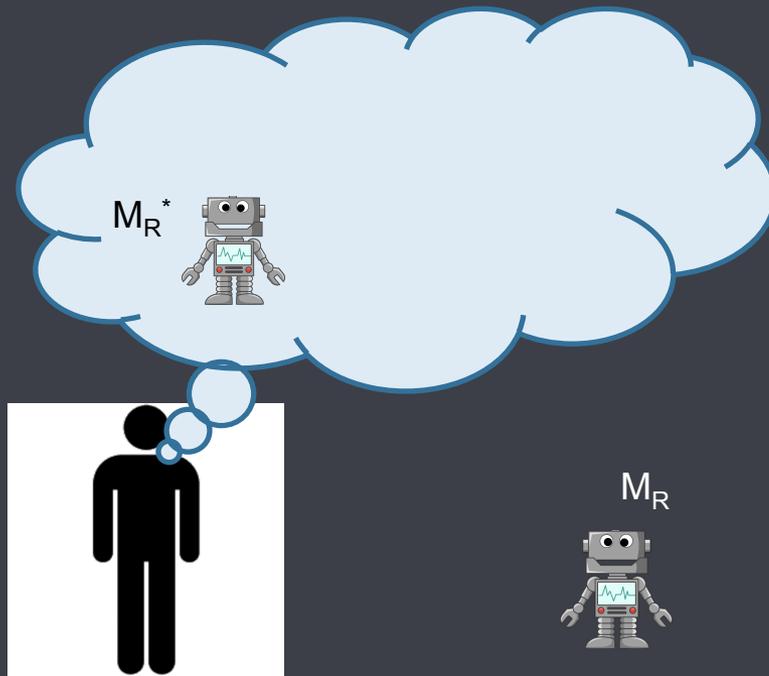


Filtered by plan
explicability measure

Explicable Plan

How do we compute the
plan explicability measure?

Plan Explicability



How do we compute the plan explicability measure?

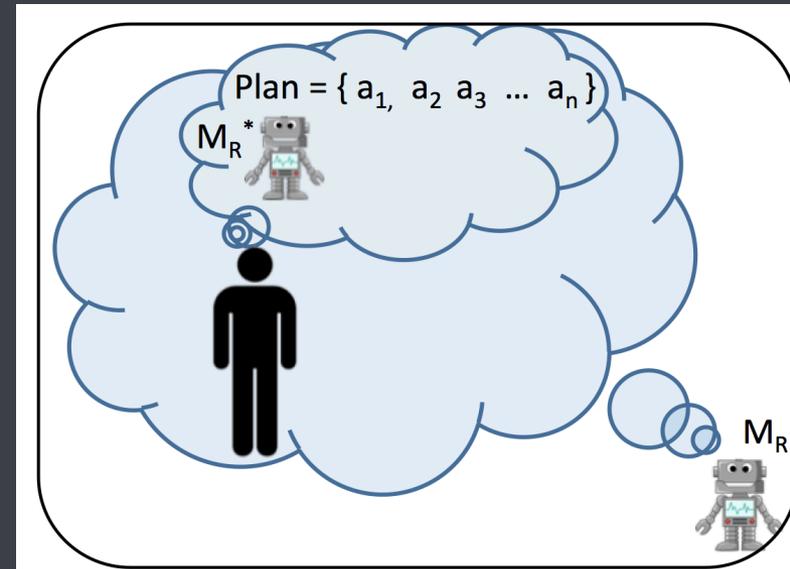
- Human's expectation of an agent is associated with another model of the agent in **human's mind**
- Plan explicability can be interpreted as a "distance" between the plans generated by M_R and M_R^*

Problem Formulation

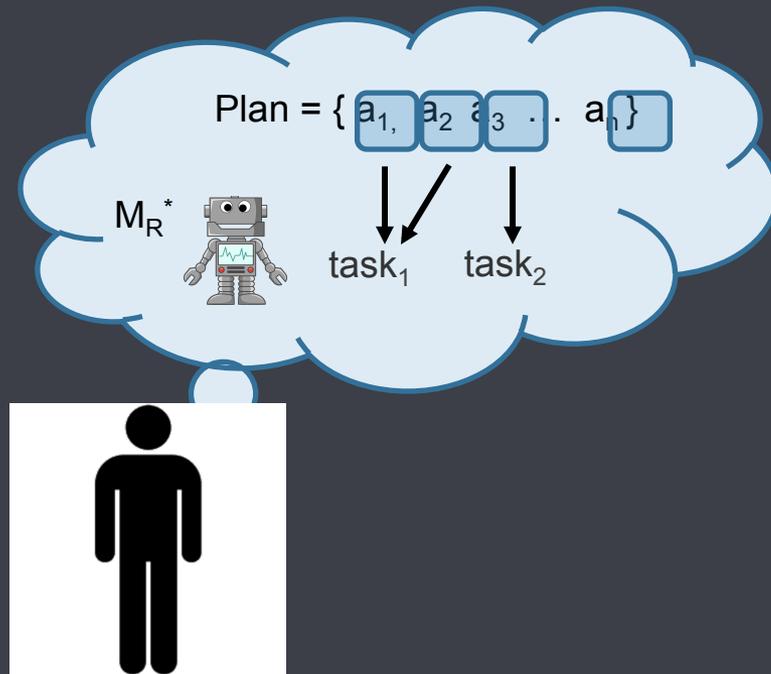
Given a goal, the objective is to find a robot plan that minimizes a weighted sum of cost of robot plan and differences between the robot plan and human's expectation of the robot plan):

$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot \operatorname{dist}(\pi_{M_R}, \pi_{M_R^*})$$

How do we obtain M_R^* ?



Problem Formulation



How do we obtain M_R^* ?

- Human's understanding of the other agents' behavior is related to how we associate it with tasks

The easier it is for humans to associate tasks/sub-goals to a plan, the more explicable the plan is.

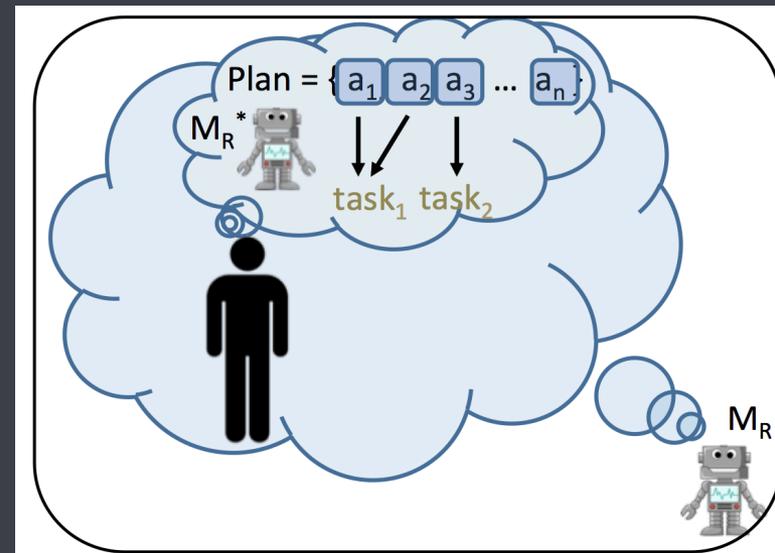
Psychological Studies: G. Csibra and G. Gergely 2007, R. Vallacher and D. Wegner, 1987.

Problem Formulation

Given a goal, the objective is to find a robot plan that minimizes a weighted sum of cost of robot plan and differences between the robot plan and human's expectation of the robot plan):

$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot \operatorname{dist}(\pi_{M_R}, \pi_{M_R^*})$$

Hidden model



Problem Formulation

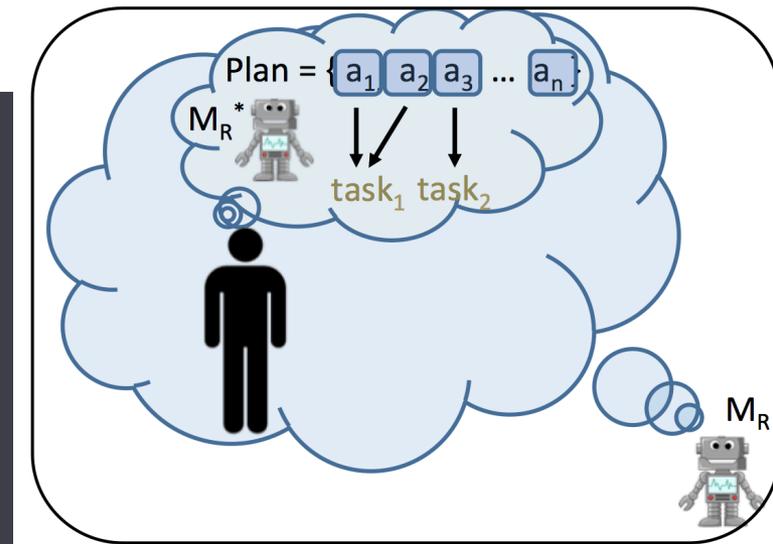
Given a goal, the objective is to find a robot plan that minimizes a weighted sum of cost of robot plan and differences between the robot plan and human's expectation of the robot plan):

$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot \operatorname{dist}(\pi_{M_R}, \pi_{M_R^*})$$

$$\operatorname{dist}(\pi_{M_R}, \pi_{M_R^*}) = F \circ \mathcal{L}^*(\pi_{M_R})$$

Function that takes plan labels as input

Human's labeling scheme for robot plans



Problem Formulation

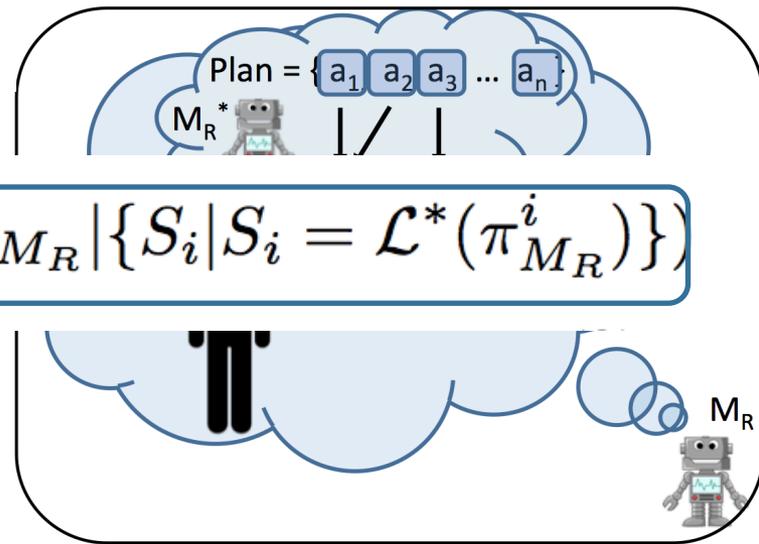
Given a goal, the objective is to find a robot plan that minimizes a weighted sum of cost of robot plan and differences between the robot plan and human's expectation of the robot plan):

$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot \operatorname{dist}(\pi_{M_R}, \pi_{M_R^*})$$

$$\operatorname{dist}(\pi_{M_R}, \pi_{M_R^*}) = F \circ \mathcal{L}^*(\pi_{M_R})$$

$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot F \circ \mathcal{L}_{CRF}^*(\pi_{M_R} | \{S_i | S_i = \mathcal{L}^*(\pi_{M_R}^i)\})$$

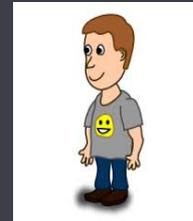
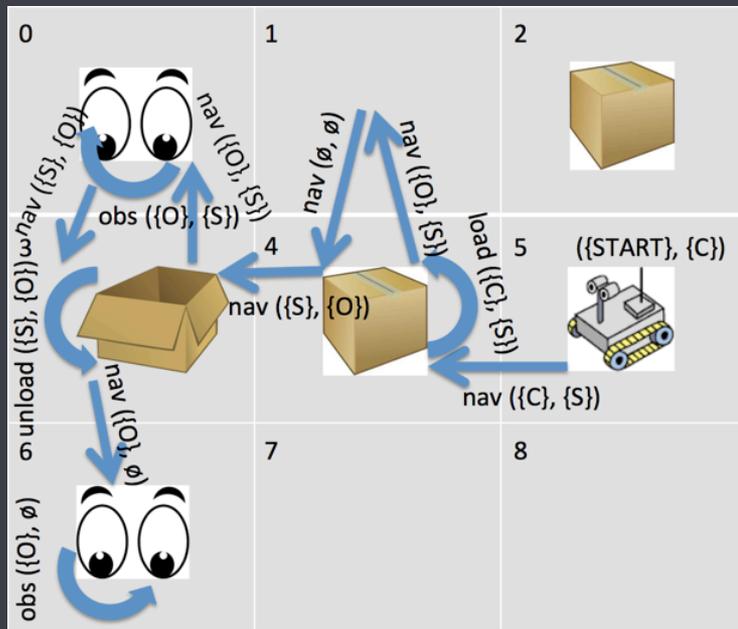
Learned labeling scheme function using linear chain
Conditional Random Fields



Explicability labeling for obtaining training samples

Task labels for actions. For example:

- Collect
- Store
- Observe



$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot F \circ \mathcal{L}_{CRF}^*(\pi_{M_R} | \{S_i | S_i = \mathcal{L}^*(\pi_{M_R}^i)\})$$

Learning using CRF

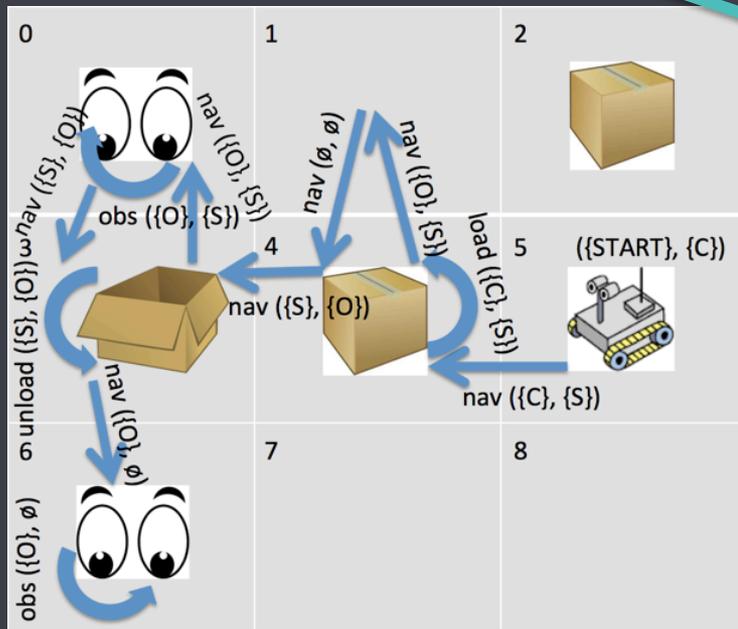
Model:

- Conditional Random Fields (CRF)

$$p(\mathbf{x}, \mathbf{y}) = \frac{1}{Z} \prod_A \Phi(\mathbf{x}_A, \mathbf{y}_A)$$

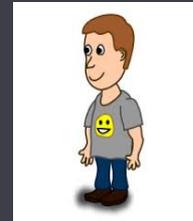
Task labels for actions. For example:

- Collect
- Store
- Observe



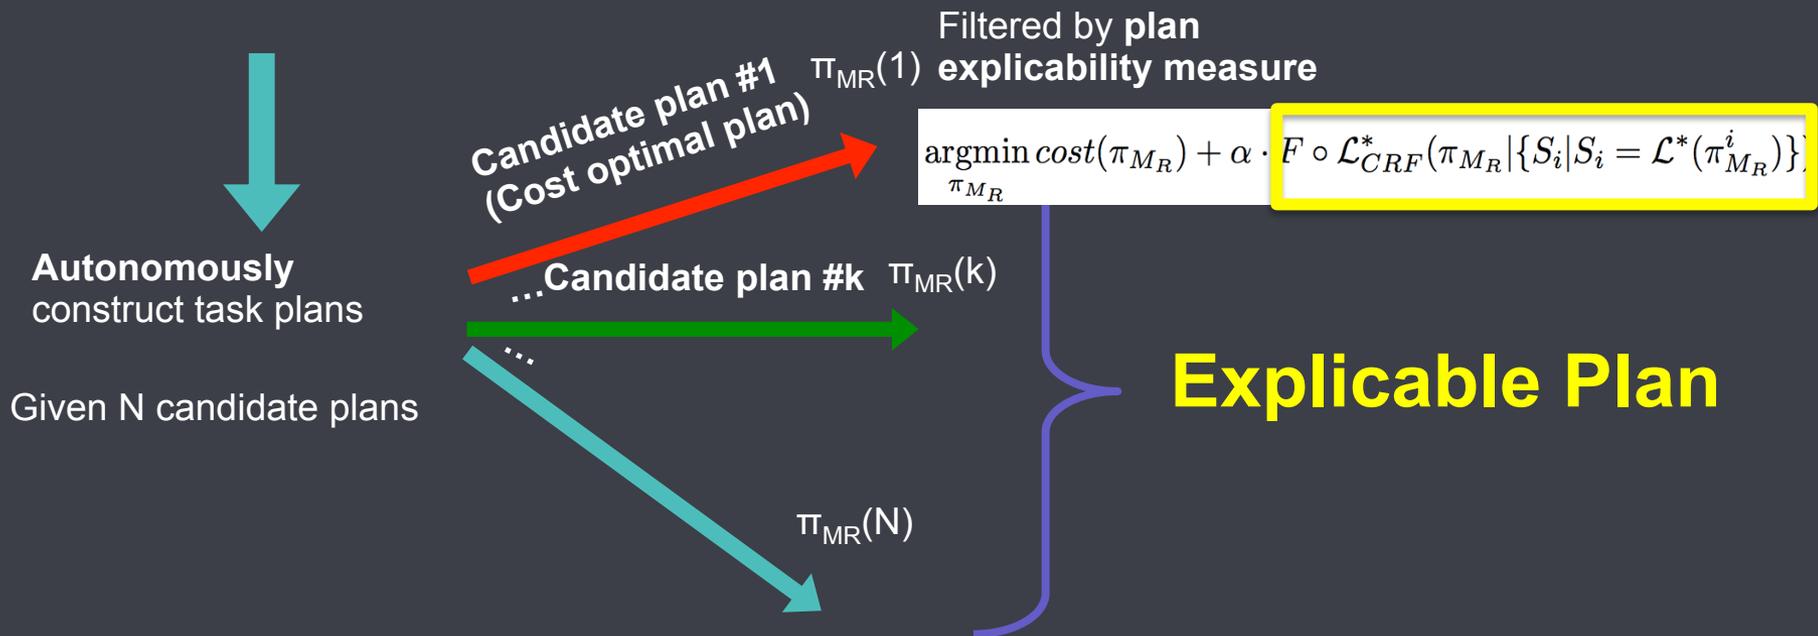
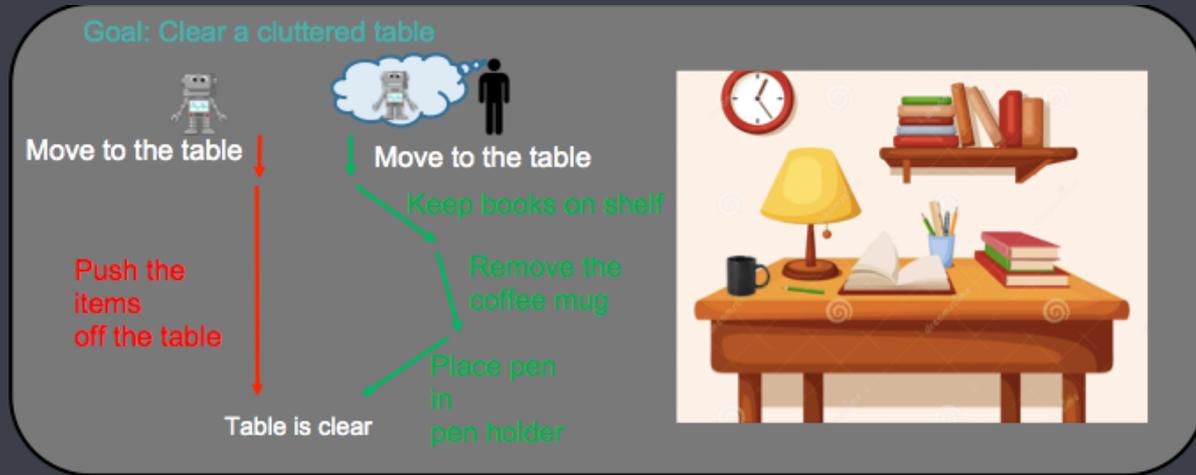
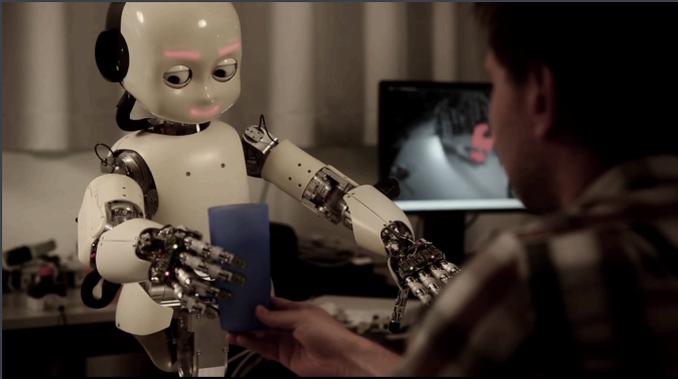
Features:

- Plan features: state, e.g., at rover 5; action name, e.g., rover



$$\operatorname{argmin}_{\pi_{M_R}} \operatorname{cost}(\pi_{M_R}) + \alpha \cdot F \circ \mathcal{L}_{CRF}^*(\pi_{M_R} | \{S_i | S_i = \mathcal{L}^*(\pi_{M_R}^i)\})$$

Plan Explicability for Autonomous Robots



Implemented Scenario

**Goal: build a
tower of
height 3**

**-
Heavy block
is on the left**



Robot in action



Related work

- There exists work on generating legible robot motions
- Two different task plans may map to the exact same motions
- **Plan explicability is focusing on task planning**

- Model learning via learning from demonstration, inverse reinforcement learning, cross-training, and tutoring systems
- Previous approaches are about how one agent teaches the other agent in terms of its plan preference
- Note that a preferred plan may not always be an expected plan; similarly, an expected plan may not always be a preferred plan
- **We are addressing the question of what humans would expect the robot to do rather than what humans themselves would do**

Dimensions of HIL Planning

Human aware robot planning

	Cooperation Modality	Communication Modality	What is Communicated	Knowledge Level
Crowdsourcing	Interaction (Advice from planner to humans)	Custom Interface	Critiques, subgoals	Incomplete Preferences Incomplete Dynamics
Human-Robot Teaming	Teaming/ Collaboration	Natural Language Speech implicit	Goals, Tasks, Model information	Incomplete Preferences Incomplete Dynamics (Open World)
"Grandpa Hates Robots"	Awareness (pre-specified constraints)	Prespecified (Safety / Interaction Constraints)	No explicit communication Interaction constraints	Incomplete Preferences Complete Dynamics
MAPGEN	Interaction (Planner takes binding advice from human)	Direct Modification of Plans	Direct modifications, decision alternatives	Incomplete Preferences Complete Dynamics



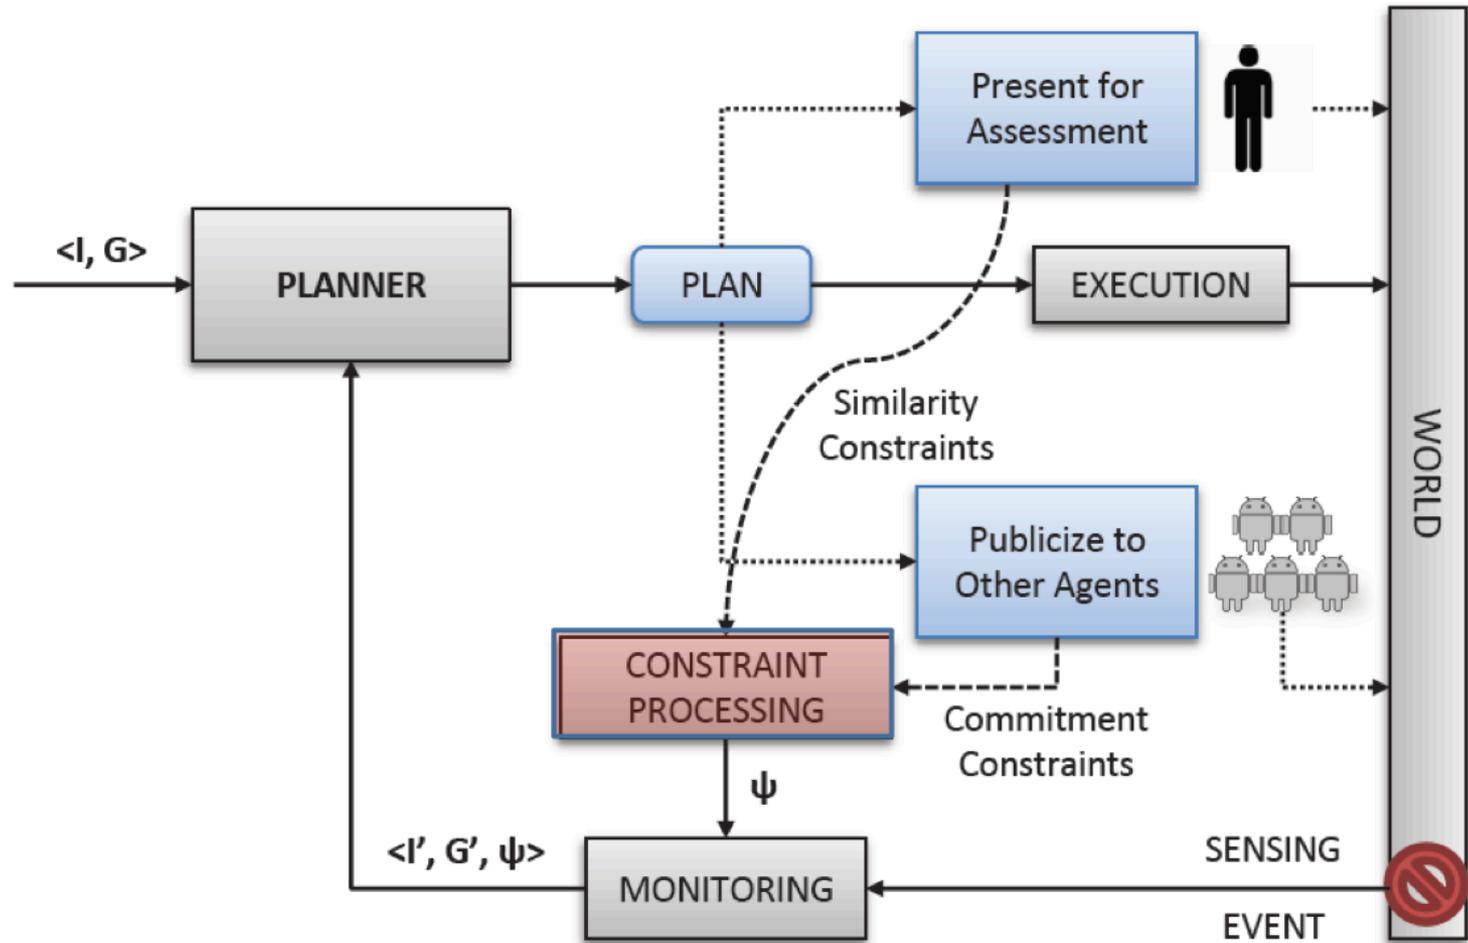
Decision theoretic Assistance,
Resource conflicts and plan for serendipity
Grandpa hates robots + ...

Challenges in Human-Aware Planning & Decision Making

- Interpret what humans are doing
 - Plan/goal/intent/preference/capability recognition ([Expectation](#))
 - Plan with incomplete domain models
 - Robust planning with “lite” models
 - (Learn to improve domain models)
-
- Continual planning/Replanning
 - Commitment sensitive to ensure coherent interaction
 - Explanations/Excuses
 - Excuse generation can be modeled as the (conjugate of) planning problem
 - Asking for help/elaboration
 - Reason about the information value



A Generalized Model of Replanning



Failures in Planner-Based Systems

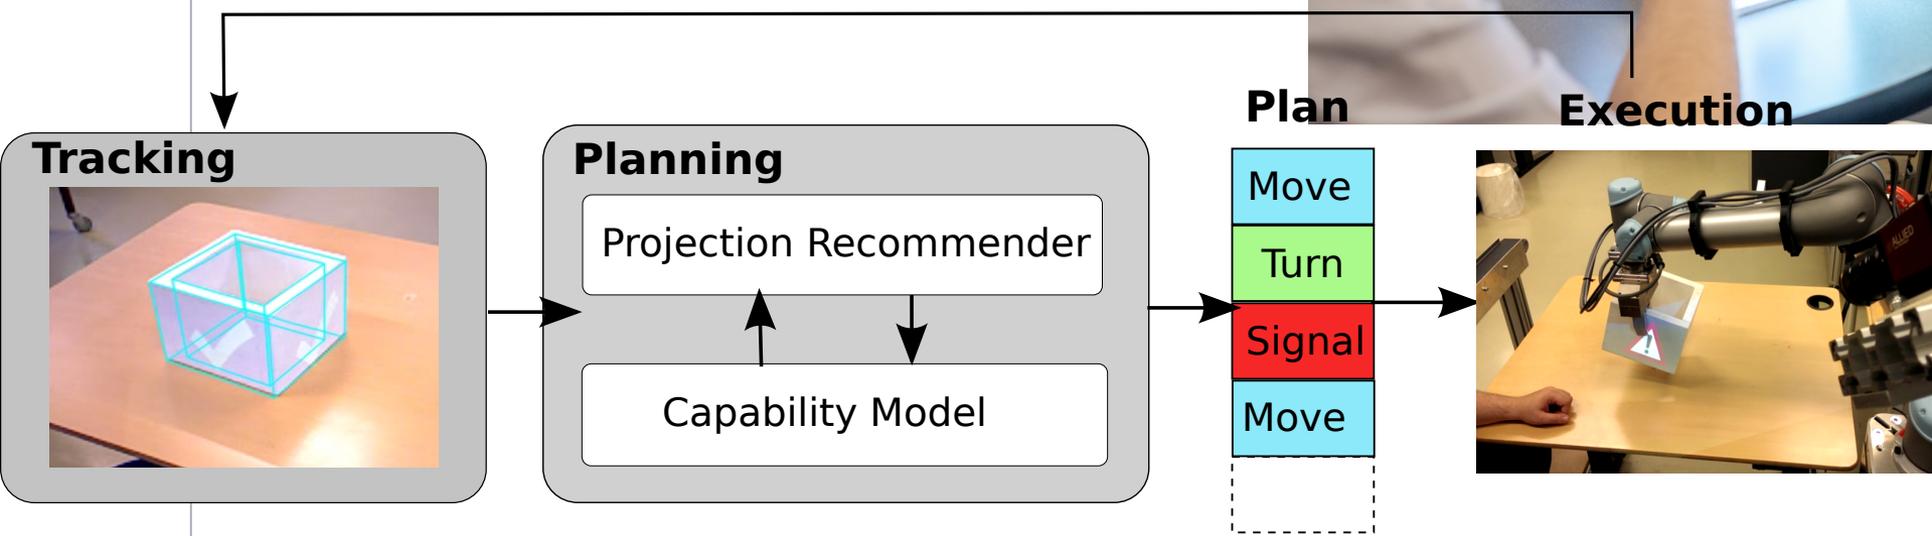
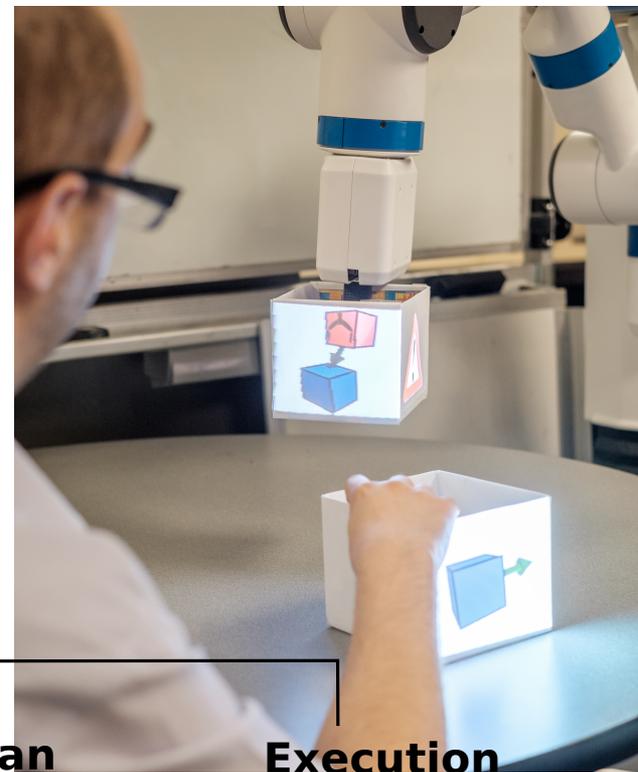
When acting in a uncertain, dynamic environment, things can go wrong:

- Execution failures
 - Error diagnosis
 - Continual Planning
- Planning failures
 - Domain is incorrectly modelled
 - Incomplete world knowledge
 - Missing resources
 - Maybe the task is just unsolvable

Coming up With Good Excuses

Göbelbecker, Keller, Eyerich, Brenner & Nebel (2010)

Projecting robot intents



Asking for Help Using Inverse Semantics

Tellex, Knepper, Li, Rus & Roy

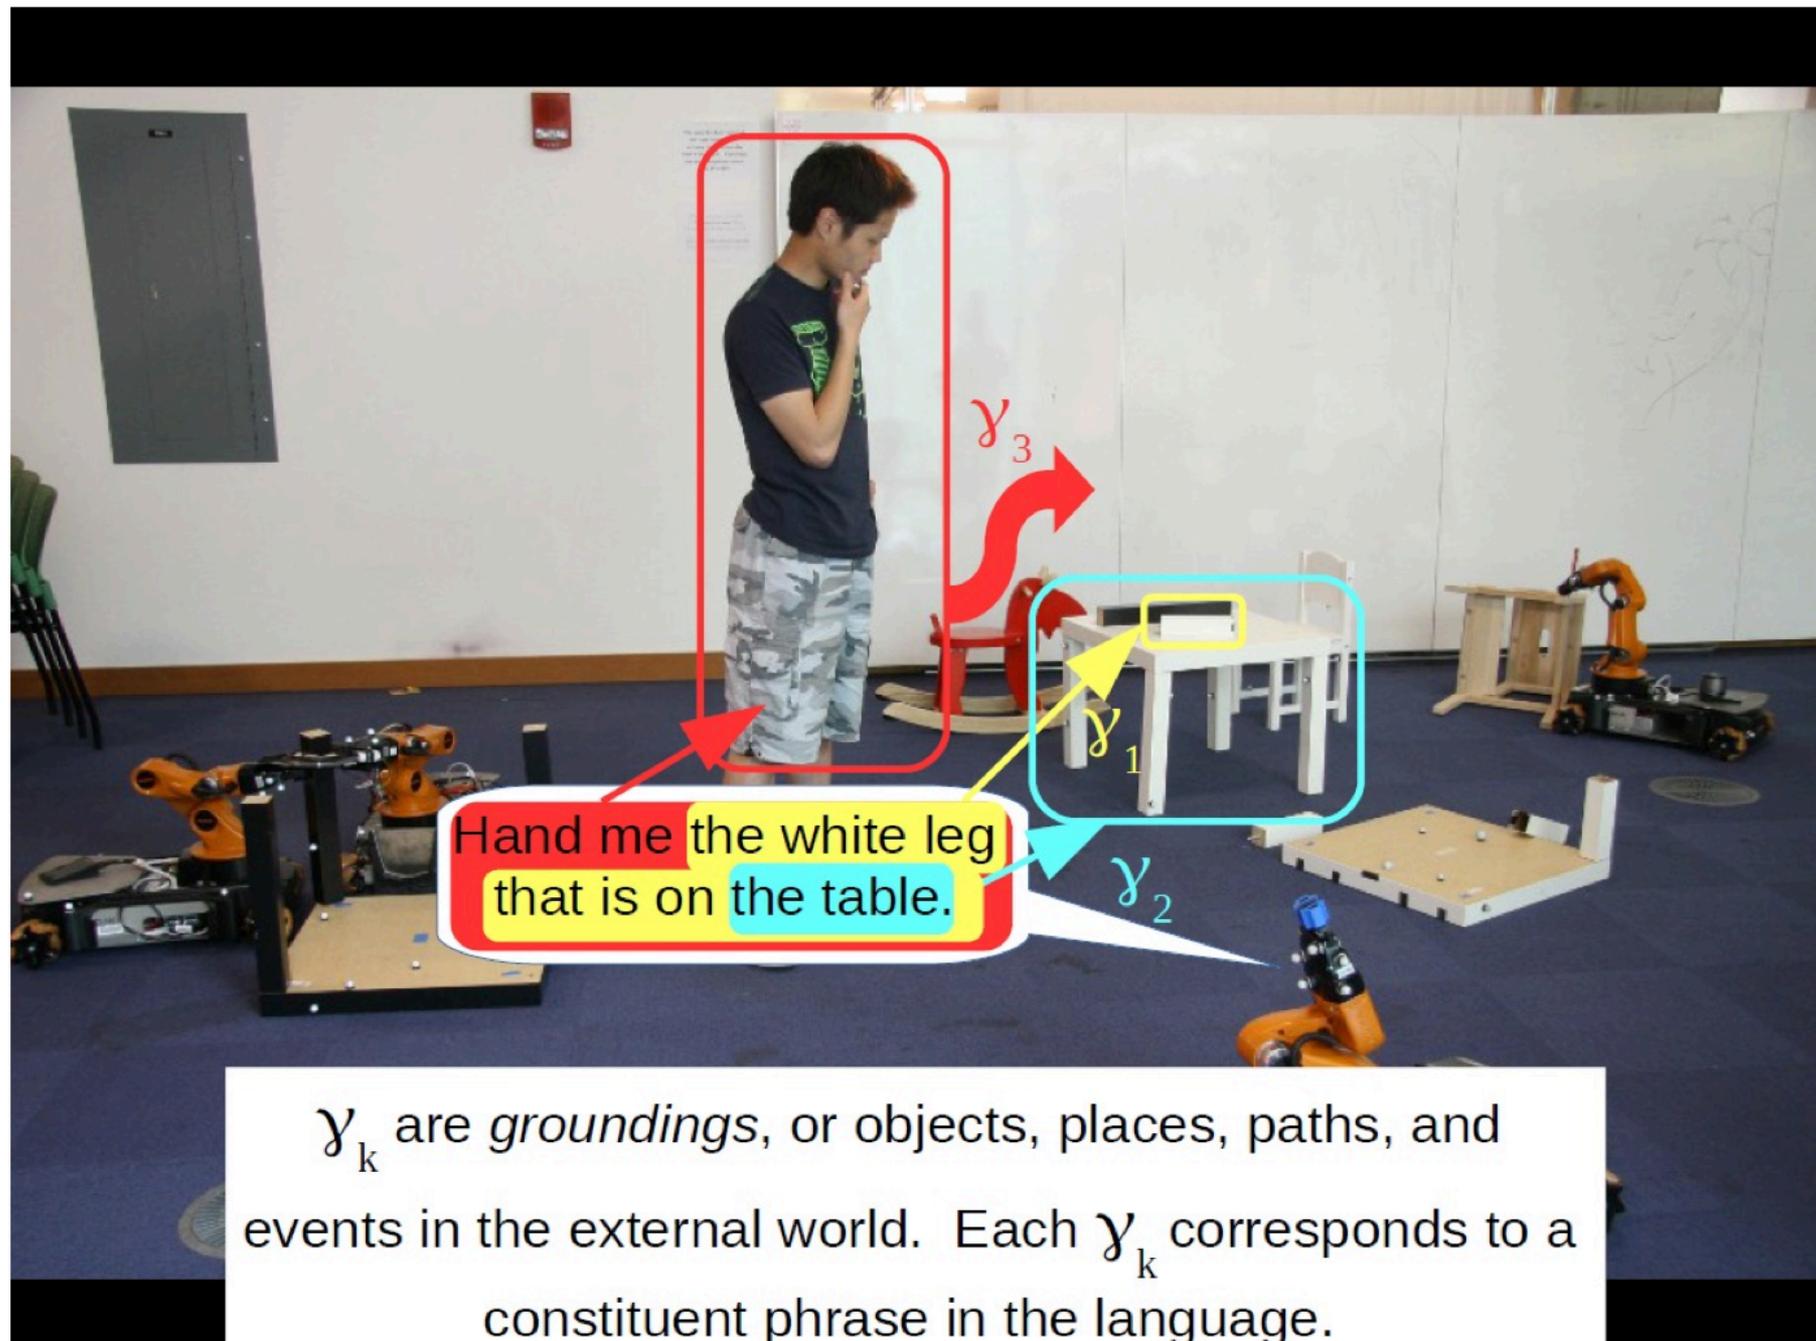


Thursday, Jan 29 2015, 1:55pm – 3:10pm

Session: Science and Systems 2014 (RSS) Presentations 2

Asking for Help Using Inverse Semantics

Stefanie Tellex, Ross Knepper, Adrian Li, Daniela Rus, Nicholas Roy



Other Challenges in Human-aware Robotics

● Trust in human-robot teaming



Trust building

When to trust

Mutual trust

Emergencies

In emergencies, don't trust a robot too much

Published 1 March 2016

[+Share](#) | [m](#) [f](#) [t](#) [in](#)

In emergencies, people may trust robots too much for their own safety, a new study suggests. In a mock building fire, test subjects followed instructions from an “Emergency Guide Robot” even after the machine had proven itself unreliable — and after some participants were told that robot had broken down.



Simulation of rescue robot in operation // Source: gatech.edu

In emergencies, people may trust robots too much for their own safety, a new study suggests. In a mock building fire, test subjects followed instructions from an “Emergency Guide Robot” even after the machine had proven itself unreliable — and after some participants were told that robot had



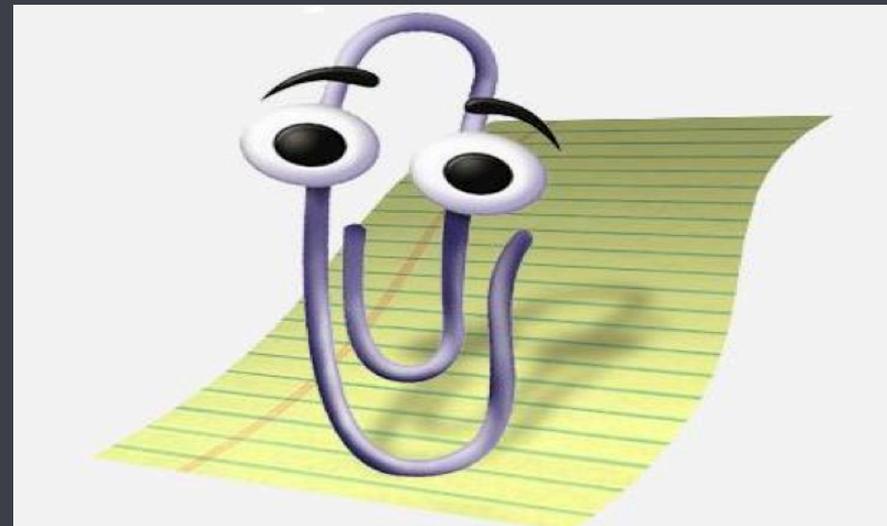
Cyber-physical systems



● Safety in Human-robot Teaming



1. A robot may not injure a human being or, through inaction, allow a human being to come to harm.
2. A robot must obey the orders given it by human beings, except where such orders would conflict with the First Law.
3. A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.



Intelligent Machines: The jobs robots will steal first

By Jane Wakefield
Technology reporter

🕒 14 September 2015 | [Technology](#)

[HOME](#) [ABOUT](#) [LOGIN](#) [SEARCH](#) [CURRENT](#) [ARTICLES](#)
[ANNOUNCEMENTS](#) [AAAI](#)

[Home](#) > [Vol 5, No 2](#) > [Nilsson](#)

Artificial Intelligence, Employment, and Income

Nils J. Nilsson

Abstract

Artificial intelligence (AI) will have profound societal effects. It promises benefits (and may also pose risks) in education, defense, business, and industry. In this article we explore how AI is likely to affect employment and the economy. We argue that AI will indeed reduce drastically the need for human labor. We also note that some people fear the automation of work by machines will lead to widespread unemployment. Yet, since the majority of us probably would rather engage in activities other than our present jobs, we ought thus to greet the consequences of AI enthusiastically. The paper discusses two reasons



Our days at the office could be numbered as an increasing range of jobs are done more efficiently by a machine

If you are sitting at a desk, driving a taxi or carrying a load, stop for a moment and ask: could a robot or machine do this job better?

Challenges in human-aware robotics

- **Perception of humans**
Human recognition, human tracking, and activity recognition



- **Modeling of humans**
Goal and intent recognition, human decision and behavioral models, expectation, model learning

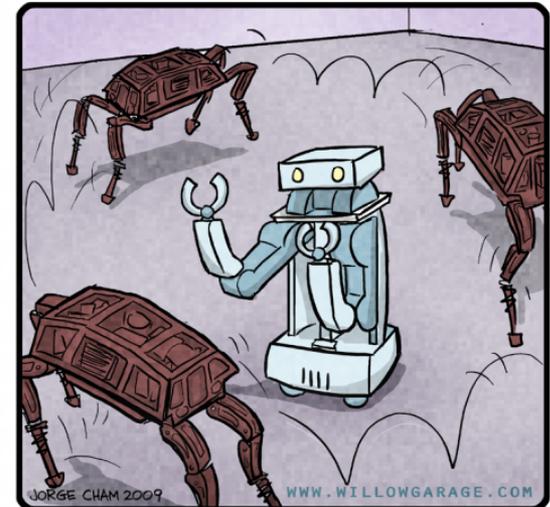


- **Human-aware decision making**
Human-aware planning, reinforcement learning and inverse reinforcement learning.

- **Human-robot interface**
Command recognition, gesture recognition



R.O.B.O.T. Comics



"SIT, BOY, SIT! SIT, I SAY,
SI... OH, FORGET IT."

