PLANNING CHALLENGES IN HUMAN-ROBOT TEAMING

KARTIK TALAMADUPULA

Committee Members

- Dr. Subbarao Kambhampati, Chair
- Dr. Chitta Baral
- Dr. Huan Liu
- Dr. Matthias Scheutz
- Dr. David E. Smith

Source: Robot Comics

Planning for Human-Robot Teaming

- **Human-Robot Teaming (HRT)** is becoming an important problem
- Requires a lot of different technologies
  - Perception (Vision), Actuation, Dialogue, **Planning** …
- Most current **robots** are **glorified remote-operated sensors**
- Autonomous Planning is an important capability
  - Supporting ***flexible*** **HRT** with **constant changes**
- The broad aims of this **thesis** are to
  1. **Engineer** an effective **integration** of planning techniques into a Human-Robot Teaming system
  2. Analyze the design tradeoffs involved in doing so

Contributions

1. **Engineering Approach**
   - Planners have not been used extensively in HRT scenarios
   - Introduce planner into an architecture for HRT
   - Use/extend automated planning methods
     1. **QUANTIFIED GOALS** in an open world
     2. **REPLANNING** for a changing, open world
     3. Handling **MODEL CHANGE** during planning
     4. **PLAN RECOGNITION** to enhance planning
2. **Analysis of Solution Methods**

Contributions

1. **Engineering Approach**
   - Planners have not been used extensively in HRT scenarios
   - Introduce planner into an architecture for HRT
   - Use/extend automated planning methods
     1. **QUANTIFIED GOALS** in an open world
     2. **REPLANNING** for a changing, open world
     3. Handling **MODEL CHANGE** during planning
     4. **PLAN RECOGNITION** to enhance planning
2. Analysis of Solution Methods

USAR Human Factors Case Study

Joint work with C. Bartlett, N. Cooke, Y. Zhang, S. Kambhampati

Planning Challenges in Human-Robot Teaming

1. **OPEN WORLD GOALS**
   - Provide a way to specify quantified goals on unknown objects
   - Consider a more principled way of handling uncertainty in facts
2. **REPLANNING**
   - Handle state and goal updates from a changing world while executing
   - Present a unified theory of replanning, to analyze tradeoffs
3. **MODEL UPDATES**
   - Accept changes to planner's domain model via natural language
4. **PLAN RECOGNITION**
   - Use belief models of other agents to enhance planning

Urban Search and Report (USAR)

Joint work with C. Bartlett, N. Cooke, Y. Zhang, S. Kambhampati

An Integrated System for USAR

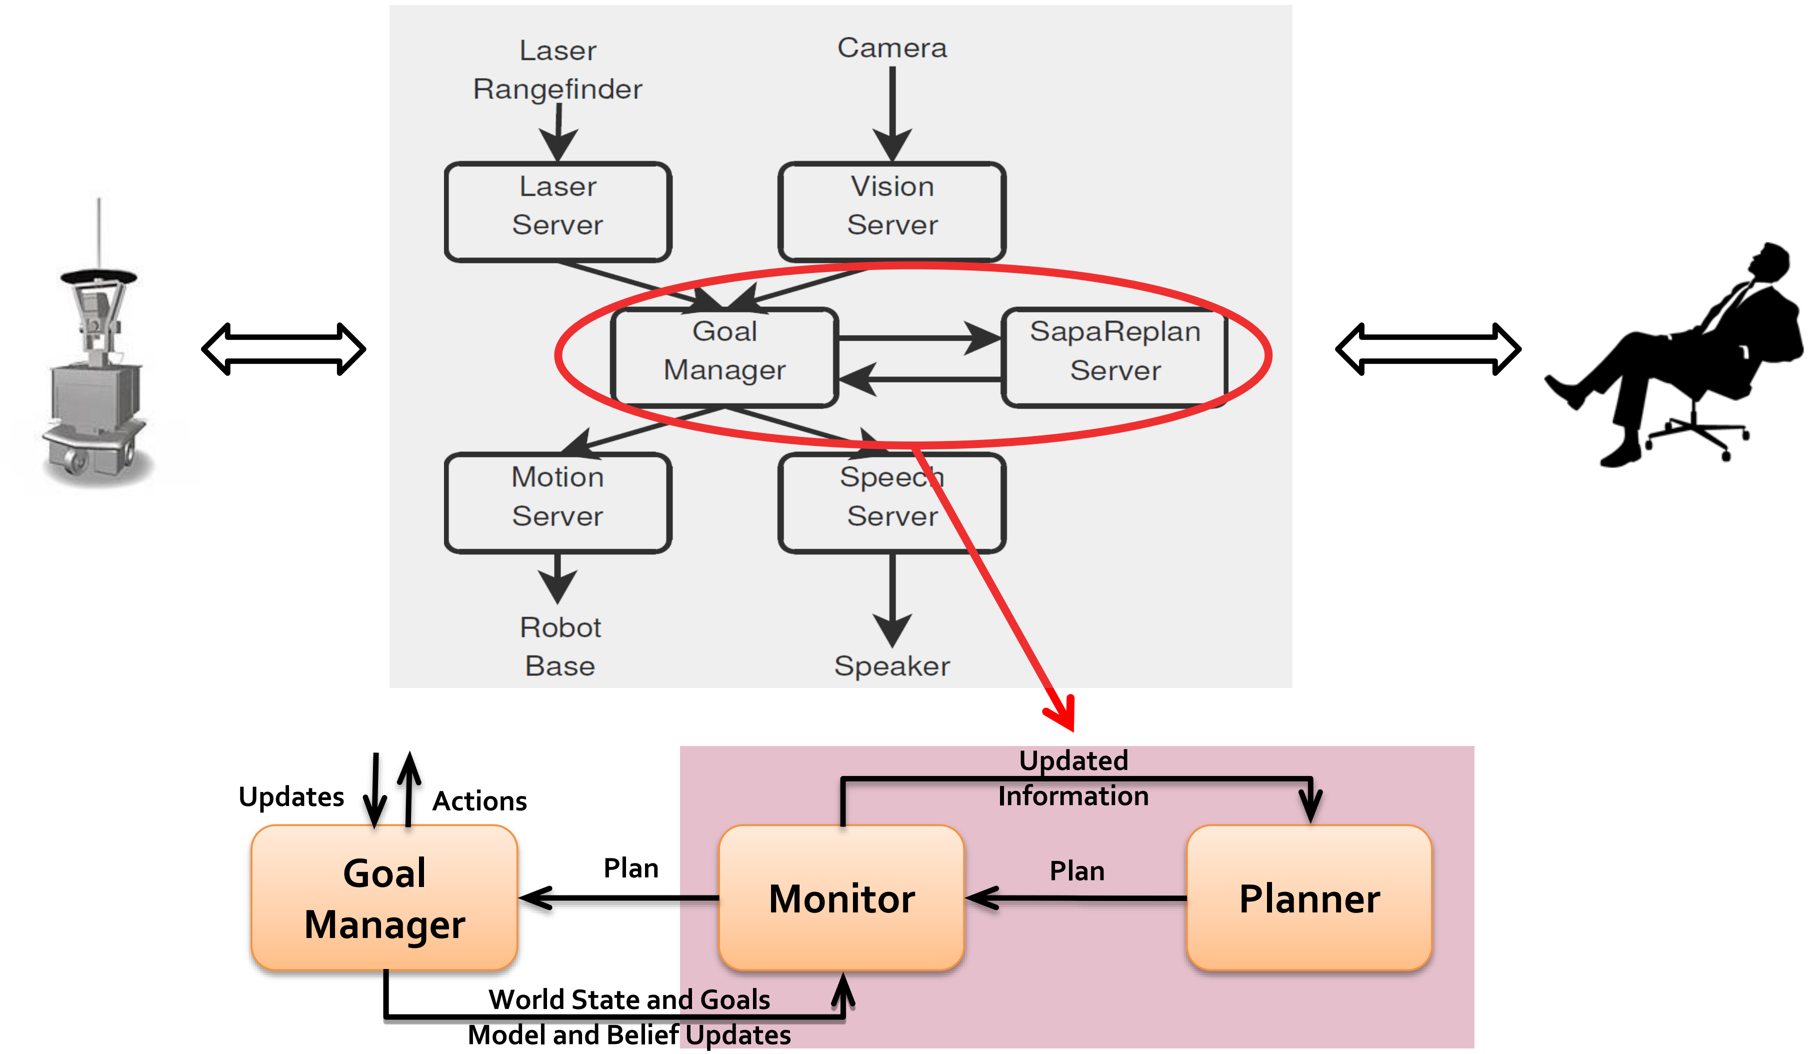

Planner's Role

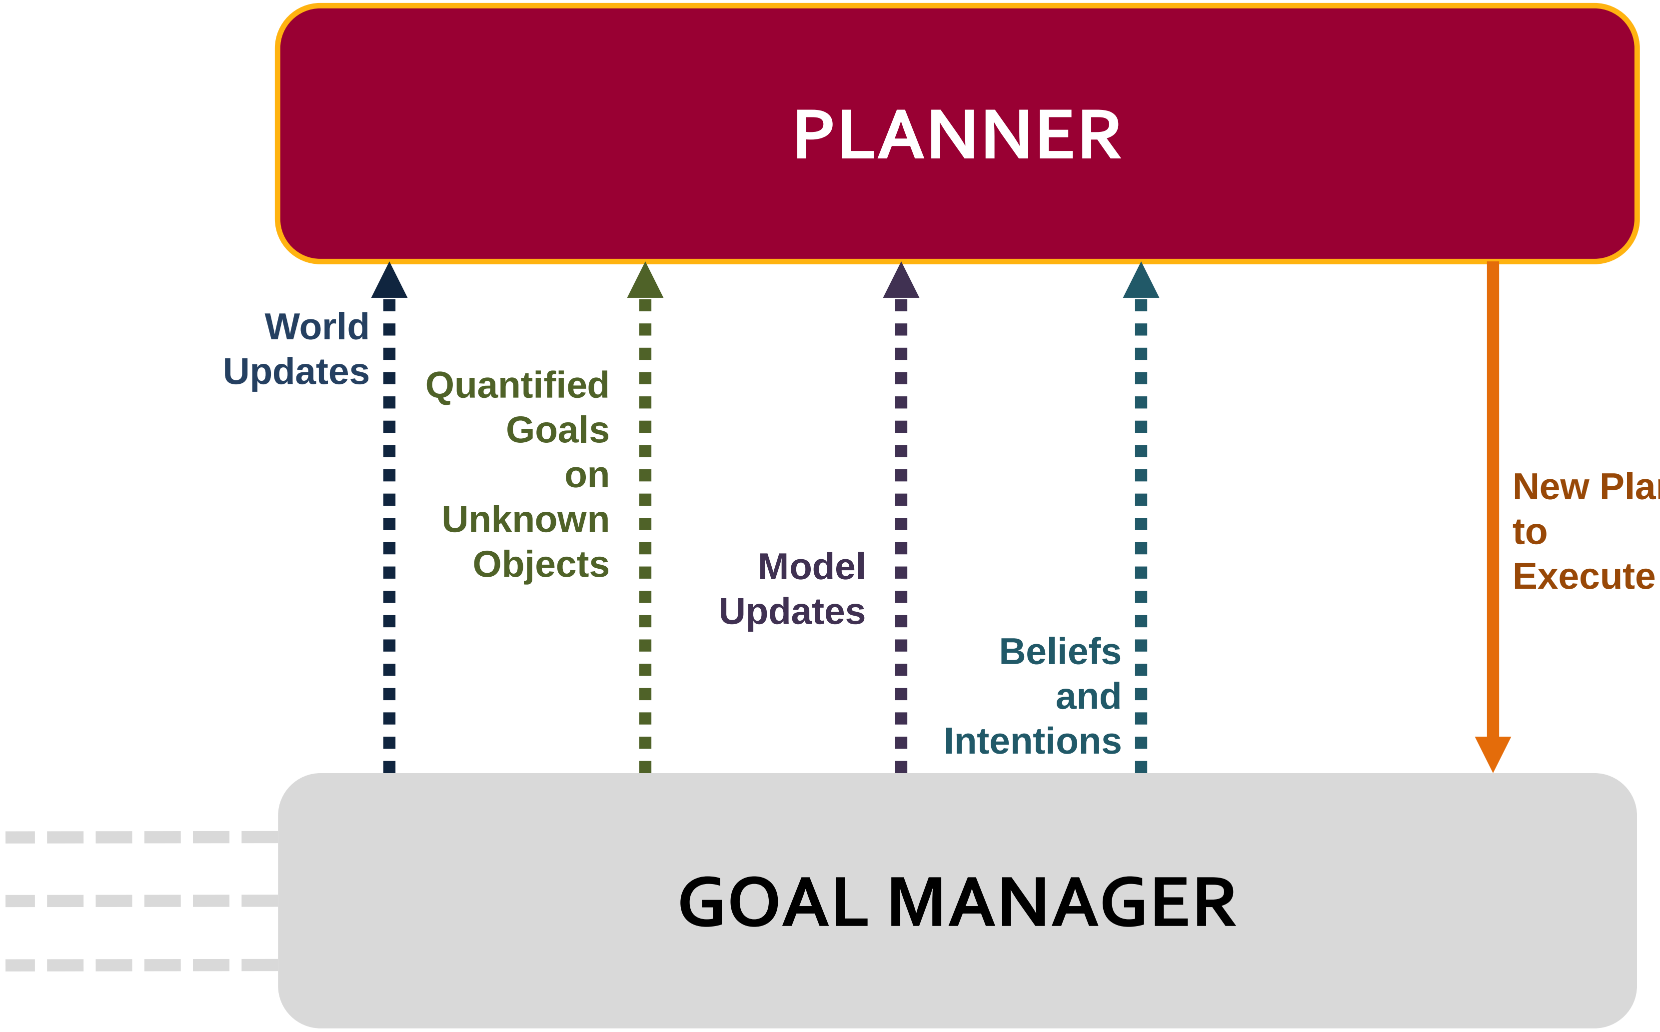

THE TRADITIONAL PLANNER

AN ISLAND UNTO ITSELF

Full Problem Specification

PLANNER

Fully Specified Action Model

Fully Specified Goals

Completely Known (Initial) World State

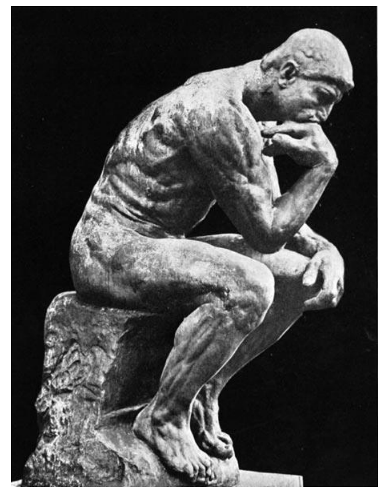

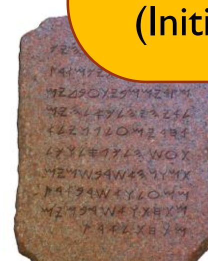

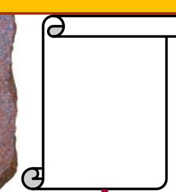

Plan (Handed off for Execution)

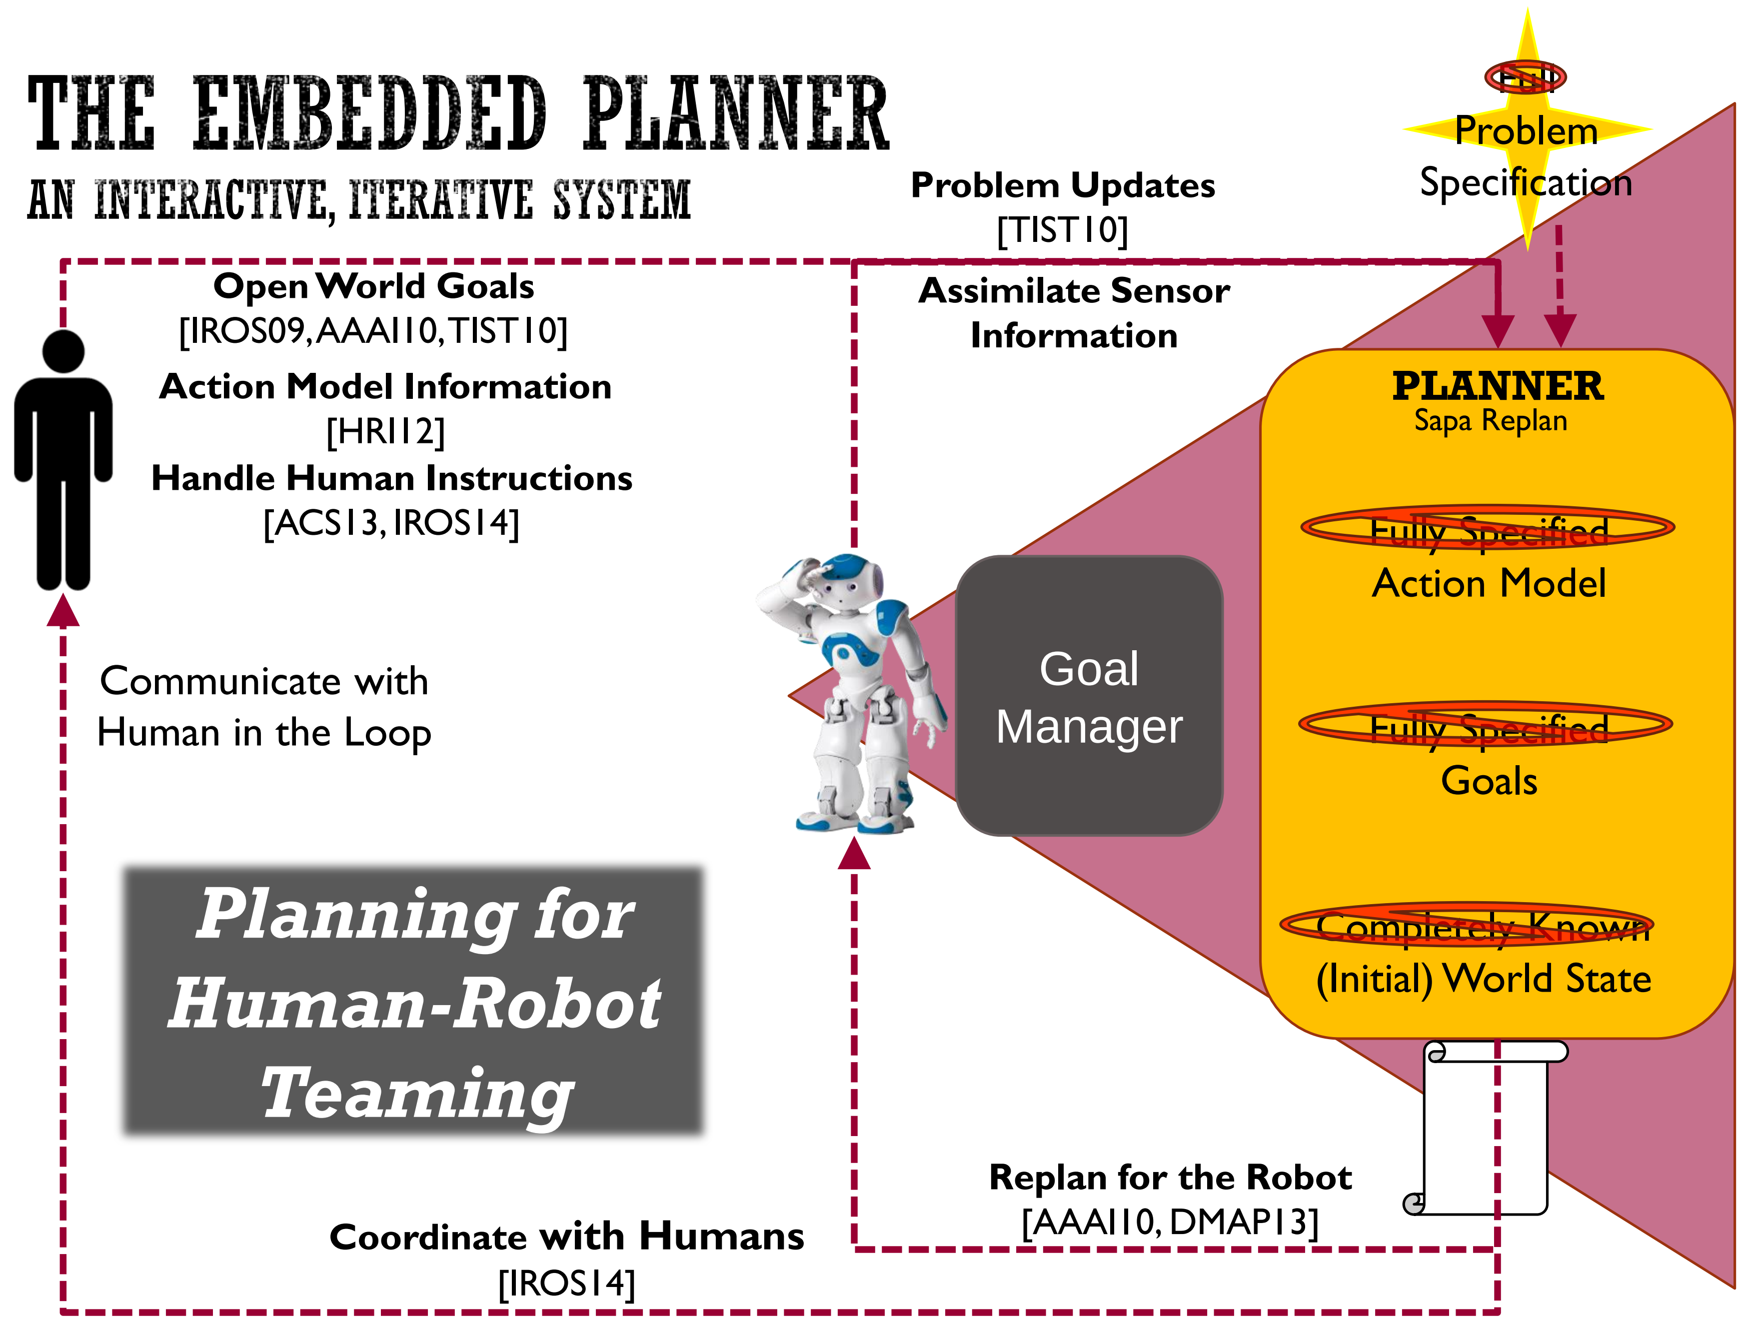
THE EMBEDDED PLANNER
AN INTERACTIVE, ITERATIVE SYSTEM
Full
Problem
Specification
Problem Updates
[TIST10]
Assimilate Sensor
Information
Open World Goals
[IROS09, AAAI10, TIST10]
Action Model Information
[HRI12]
Handle Human Instructions
[ACS13, IROS14]
PLANNER
Sapa Replan
Fully Specified
Action Model
Fully Specified
Goals
Completely Known
(Initial) World State
Goal
Manager
Communicate with
Human in the Loop
Planning for
Human-Robot
Teaming
Replan for the Robot
[AAAI10, DMAP13]
Coordinate with Humans
[IROS14]

Fielded Prototype

- Planning Artifact: **Sapa Replan**
  - Extension of **Sapa** metric temporal planner
- **Partial Satisfaction Planning**
  - Builds on Sapa^{PS} planner
- **Replanning**
  - Uses an execution monitor to support scenarios with real-time execution

[Benton et al., AIJ07]
[Talamadupula, Benton, et al., TIST10]

Planning Challenges in Human-Robot Teaming

1. OPEN WORLD GOALS

- Provide a way to specify quantified goals on unknown objects
- Consider a more principled way of handling uncertainty in facts

2. REPLANNING

- Handle state and goal updates from a changing world while executing
- Present a unified theory of replanning, to analyze tradeoffs

3. MODEL UPDATES

- Accept changes to planner's domain model via natural language

4. PLAN RECOGNITION

- Use belief models of other agents to enhance planning

Open World Goals

- When to start sensing?
  - Indicator to start sensing

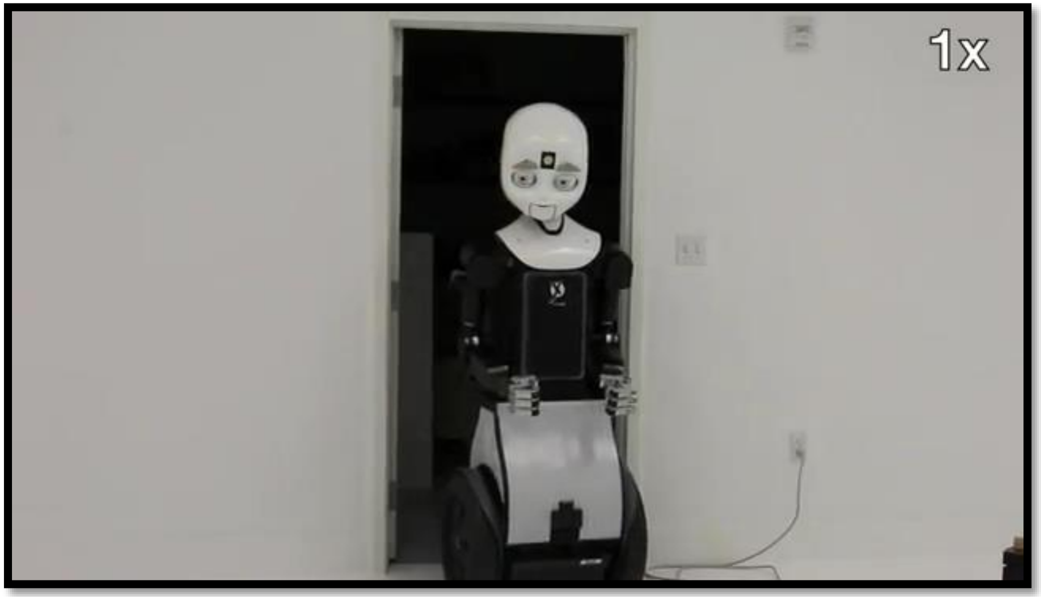

- What to look for?
  - Object type
  - Object properties
- When to stop sensing?
  - When does the planner know the world is closed?
- Why should the robot sense?
  - Does the object fulfill a goal?
  - What is the reward? Is it a bonus?

[Talamadupula, Benton et al., ACM TIST 2010]

Open World Quantified Goals (OWQGs)

1. When to sense
2. What to sense
3. When to stop
4. Why sense

"Wounded persons may be in rooms. Report the locations of as many wounded people as possible."

```
(:open  (forall ?r - room           Quantified Object(s) [1]
   (sense ?p - person               Sensed Object [2]
      (looked_for ?p ?r)            Closure Condition [3]
     (and (has_property ?p wounded)
          (in ?p ?r))               Quantified Facts [2]
(:goal
   (and (reported ?p wounded ?r)
   [100] - soft))))                 Quantified Goal [4]
```

[Talamadupula, Benton et al., ACM TIST 2010]

Solution Approach

Tricking the Robot for Profit

1. OWQG is provided to the planner
2. Planner uses an **optimistic determinization**
   - Given an OWQG, assume the presence of object
     - Create a runtime object (may exist only in planner)
     - E.g.: For every room, assume wounded person
3. **Replan**
   - Make a new plan that uses runtime object to achieve the open world goal; (assumed) profit from reward
4. **Execute**
   - Up to the sensing action (closure condition)
   - Delete runtime object
   - Real object either exists, or doesn't

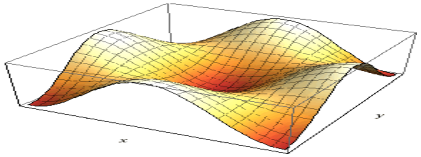

Replanning for Changing Worlds

- New Information
  - Sensors
  - Human teammate
- New Goals
  - Orders: Humans
  - Requests

- Requirement
  - New plan that works in new world (state)
  - Achieves the changed goals

[Talamadupula et al. AAAI10]

How to Replan

The Engineering Solution

- Problem changes from [I, G] to [I`, G`]

- Solution:
  1. Stop execution of old plan π
  2. Assimilate state changes I → I`
  3. Assimilate goal changes G → G`
  4. Give the new instance [I`, G`] to planner
  5. Execute the new plan π`

- (Re)Planning System: Sapa Replan

Sapa Replan: Execution Monitor

- Implement **rational choice** over possible courses of action
  - Two possible choices
    - **Continue** currently executing plan
    - **Deliberate** (replan)
- **Objective Selection**
  - Two possibilities
    - Update goal description: **Replan**
    - Update goal description: **Replan + Restart search**
  - Net Benefit
    - Partial Satisfaction Planning

[ASU-TR08, IROS09, TIST10, AAAI10, SPARK11]

Specifying Changes

- Use an update syntax

$$U = \langle O, E, G_n, T \rangle$$

O: Set of objects (constants)
E: Set of new events (predicates)
G_n: Set of new goals
T: Current time point

- Example

```
(:update
 :objects
          room3 - room
 :events
          (at 125.0 (not (at room2)))
          (at room3)
          (visited room3)
 :goal    (visited room4) [500] - hard
 :now     207.0)
```

[Talamadupula et al. AAAI10]

Replanning + Open World Goals
USAR Example

Original Plan

```
(move-hallway hall_start hall1)
(move-hallway hall1 hall2)
(move-hallway hall2 hall3)
(move-hallway hall3 hall_end)
(deliver medkit1)
```

```
(:open  (forall ?r - room
   (sense ?p - person
      (looked_for ?p ?r)
     (and (has_property ?p wounded)
          (in ?p ?r))
(:goal
   (and (reported ?p wounded ?r)
   [100] - soft))))
```

New Plan

```
(move-hallway hall2 hall3)
(enter room1 hall3)
(sense-for !person1 room1)
(report !person1 room1)
(exit room1 hall3)
(move-hallway hall3 hall_end)
(deliver medkit1)
```

```
(:update
 :objects
          room1 - room
 :events
          (at 90.0 (not (at hall1)))
          (at hall2)
          (connected hall3 room1)
 :goal
 :now     103.0)
```

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”

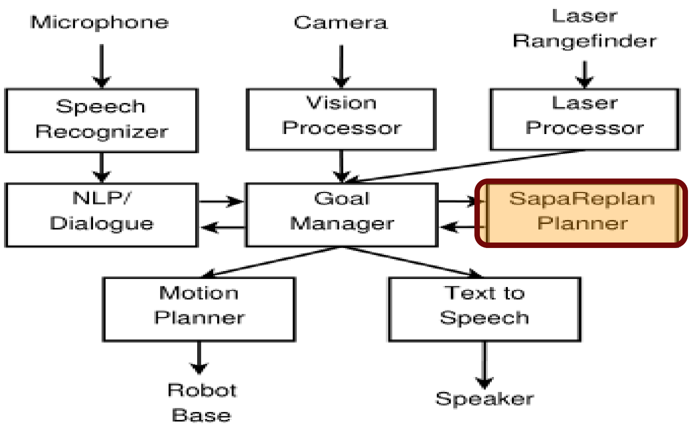

- NLP Module
  i. Reference resolution
  ii. Parsing
  iii. Background knowledge
  iv. Action submission (to planner)

[Cantrell, Talamadupula et al., HRI 2012]

[In collaboration with hrilab, Tufts University]

Example: Action Addition

New Action: "push"

"To go into a room when you are at a closed door, push it one meter."

```
(:durative-action push
 :parameters (?door - doorway ?cur_loc - hallway ?to_loc - zone)
 :duration (= ?duration (dur push))
 :condition (and  (at start (at ?cur_loc))
                  (at start (door_connected ?door ?cur_loc ?to_loc))
                  (over all (door_connected ?door ?cur_loc ?to_loc)))
 :effect (and (at start (not (at ?cur_loc)))
              (at end (open ?doorway))
              (at end (at ?to_loc))))
```

From natural language

Architecture

Background knowledge

Why Support Model Updates?

- **One ground truth** model of the world
  - Neither human nor robot have this
  - Human may know more though …

- **Impossible to specify everything up-front**
  - But during execution …

1. **Addition**
   - Human sees a closed door, but knows robot can push it
2. **Deletion**
   - Taking a picture might ignite vapors
3. **Modification**
   - No power, so robot must needs light for taking a picture

Model Revision

- Model represented in **PDDL**
- PDDL domain model

$$\mathbb{M} = \langle \mathbb{C}, \mathbb{P}, \mathbb{F}, \mathbb{A} \rangle$$

  - $\mathbb{C}$: set of constants (objects)
  - $\mathbb{P}$: set of predicates
  - $\mathbb{F}$: set of functions
  - $\mathbb{A}$: set of actions (operators)
- Revision should support **modification** of any of these **on the fly**

How to Update a Model

(The Engineering Solution)

1. **Pause execution** of the current plan

2. Provide a way of **updating an existing model**
   - (Currently restricted to only actions)
   - Planner API for architecture can access and edit various action constituents
     i. Cost
     ii. Duration
     iii. Variables (Parameters)
     iv. Preconditions
     v. Effects

3. **Replan with new model**, generate new plan
   - Discard old plan

4. **Execute new plan**

Plan & Intent Recognition

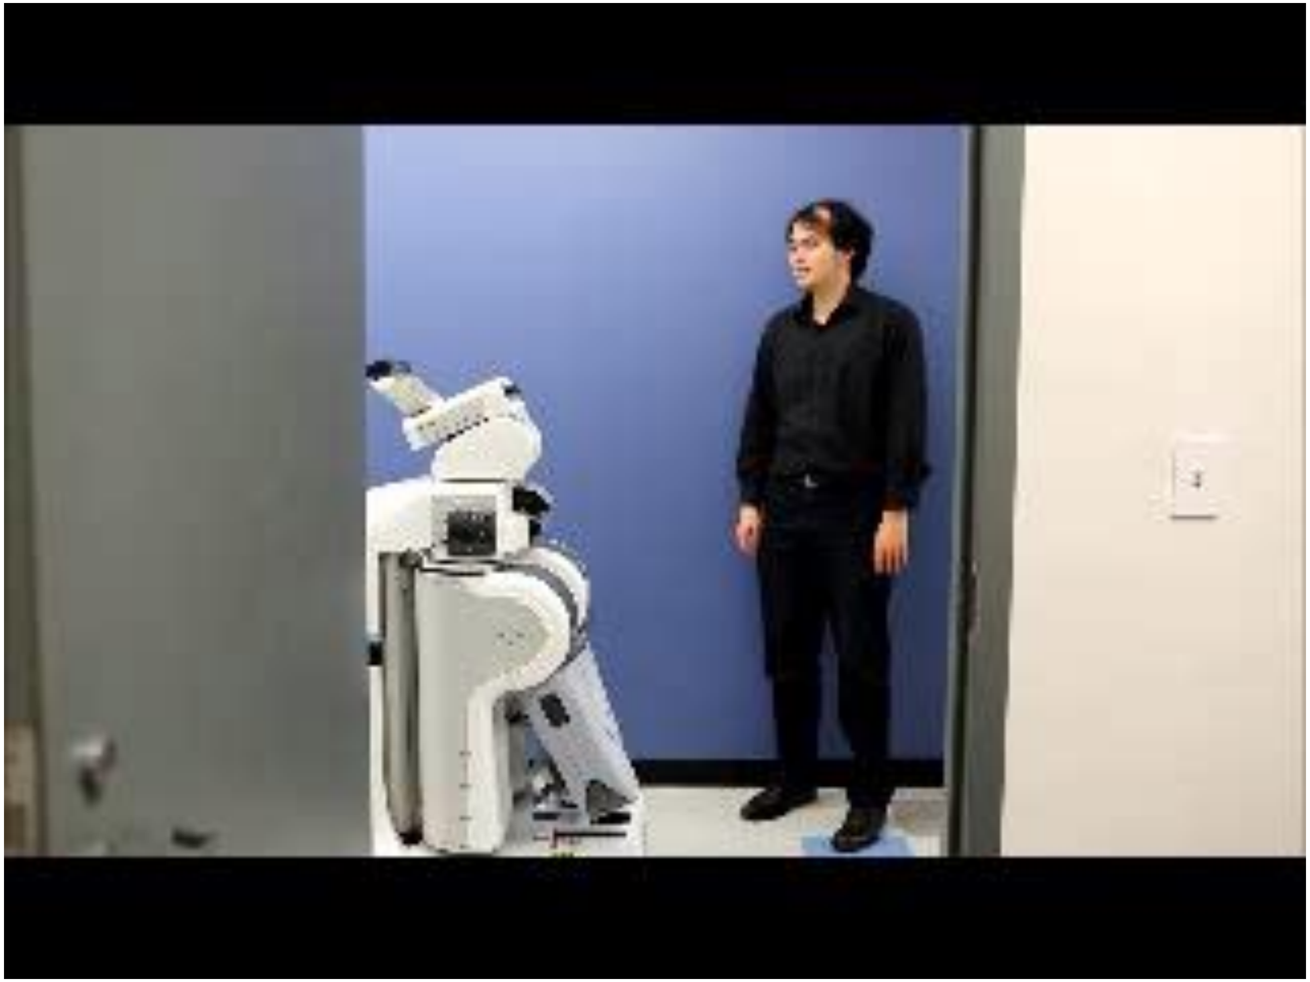

[In collaboration with hrilab, Tufts University]

[Talamadupula, Briggs et al., IROS14]

Proposed Approach

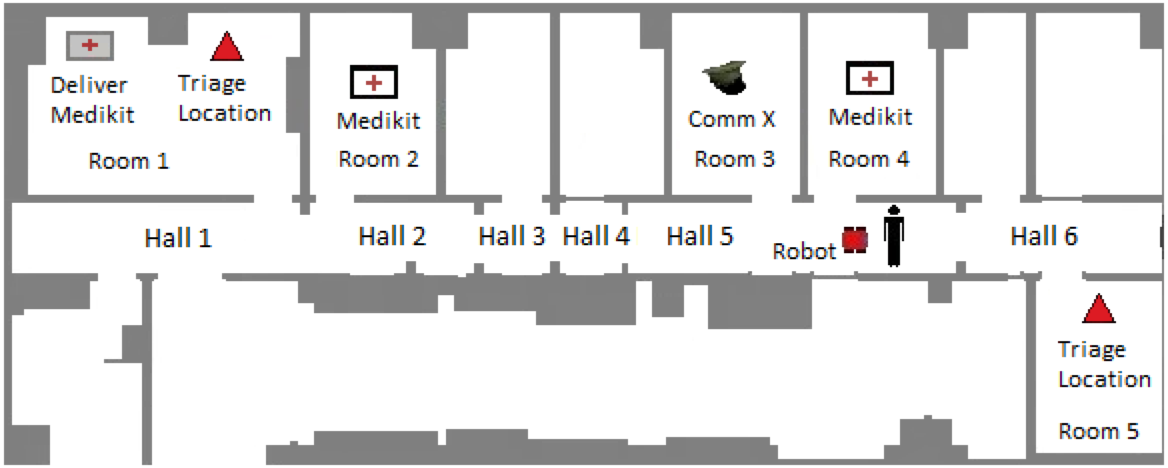

1. Map the robot's beliefs and knowledge about CommX into a new planning instance
2. Generate a plan for this instance – prediction of CommX's plan
3. Extract relevant information from the predicted plan
   - Which medkit will CommX pick up?
4. Use the extracted information to deconflict robot's plan

Solution

Comm X's Goal

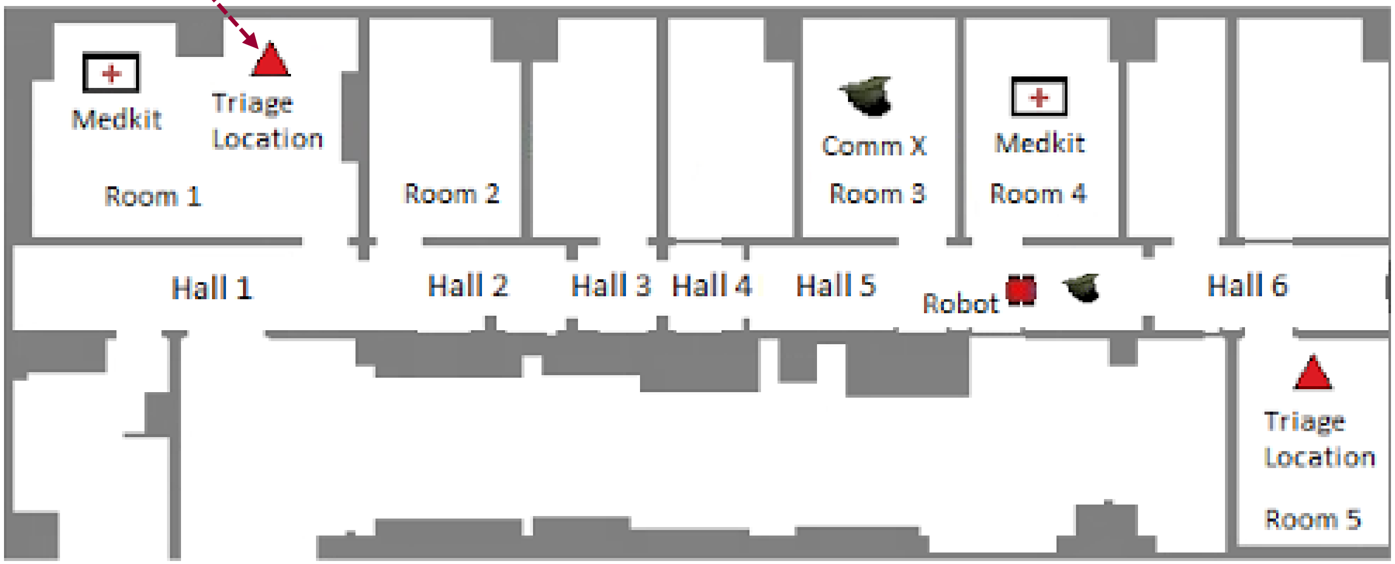

PREDICTED PLAN FOR COMMX

move commx room3 hall5
move_reverse commx hall5 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

Contributions

1. Engineering Approach

   › Planners have not been used extensively in HRT scenarios

   › Introduce planner into an architecture for HRT

   › Use/extend automated planning methods

      1. **QUANTIFIED GOALS** in an open world
      2. **REPLANNING** for a changing, open world
      3. Handling **MODEL CHANGE** during planning
      4. **PLAN RECOGNITION** to enhance planning

2. Analysis of Solution Methods

Relevant Publications

1. **Coordination in Human-Robot Teams Using Mental Modeling and Plan Recognition.**
   Talamadupula, K.; Briggs, G.; Chakraborti, T.; Scheutz, M.; and Kambhampati, S.
   Proceedings of the IEEE/RSJ Inter...
   Ro...
2. ***Th...***
   of...
   Tal...
   arX...
3. **Arc...**
   Ope...
   Tala...
   Proc...
   Cogn...
4. **On t...**
   Tala...
   ICAP...
5. **A The...**
   Talan...
   ICAPS...
   (DMA...
6. ***Tell m...***
   Updat...
   Cantre...
   Kambh...
   Procee...
   Confer... ...teraction (HRI), 471--478, 2012.
7. **Planning for Agents with Changing Goals.**
   Talamadupula, K.; Schermerhorn, P.; Benton, J.; Kambhampati, S.; and Scheutz, M.
   ICAPS 2011 Systems Demos and Exhibits
   Placed 3rd for Best Demo
8. **Planning for Human-Robot ...**
   ...enton, J.;
   ...hop
   ...horn, P.;
   ...(TIST),
   ...d Robot:
   ...pati, S.;
   ...ificial
   ...Robot:
   ...ati, S.;
   ...d
   ...ring
   ...Schermerhorn, P.; Benton, J.; Scheutz, M.; **Talamadupula, K.**; and Kambhampati, S.
   Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 3912--3917, 2009.

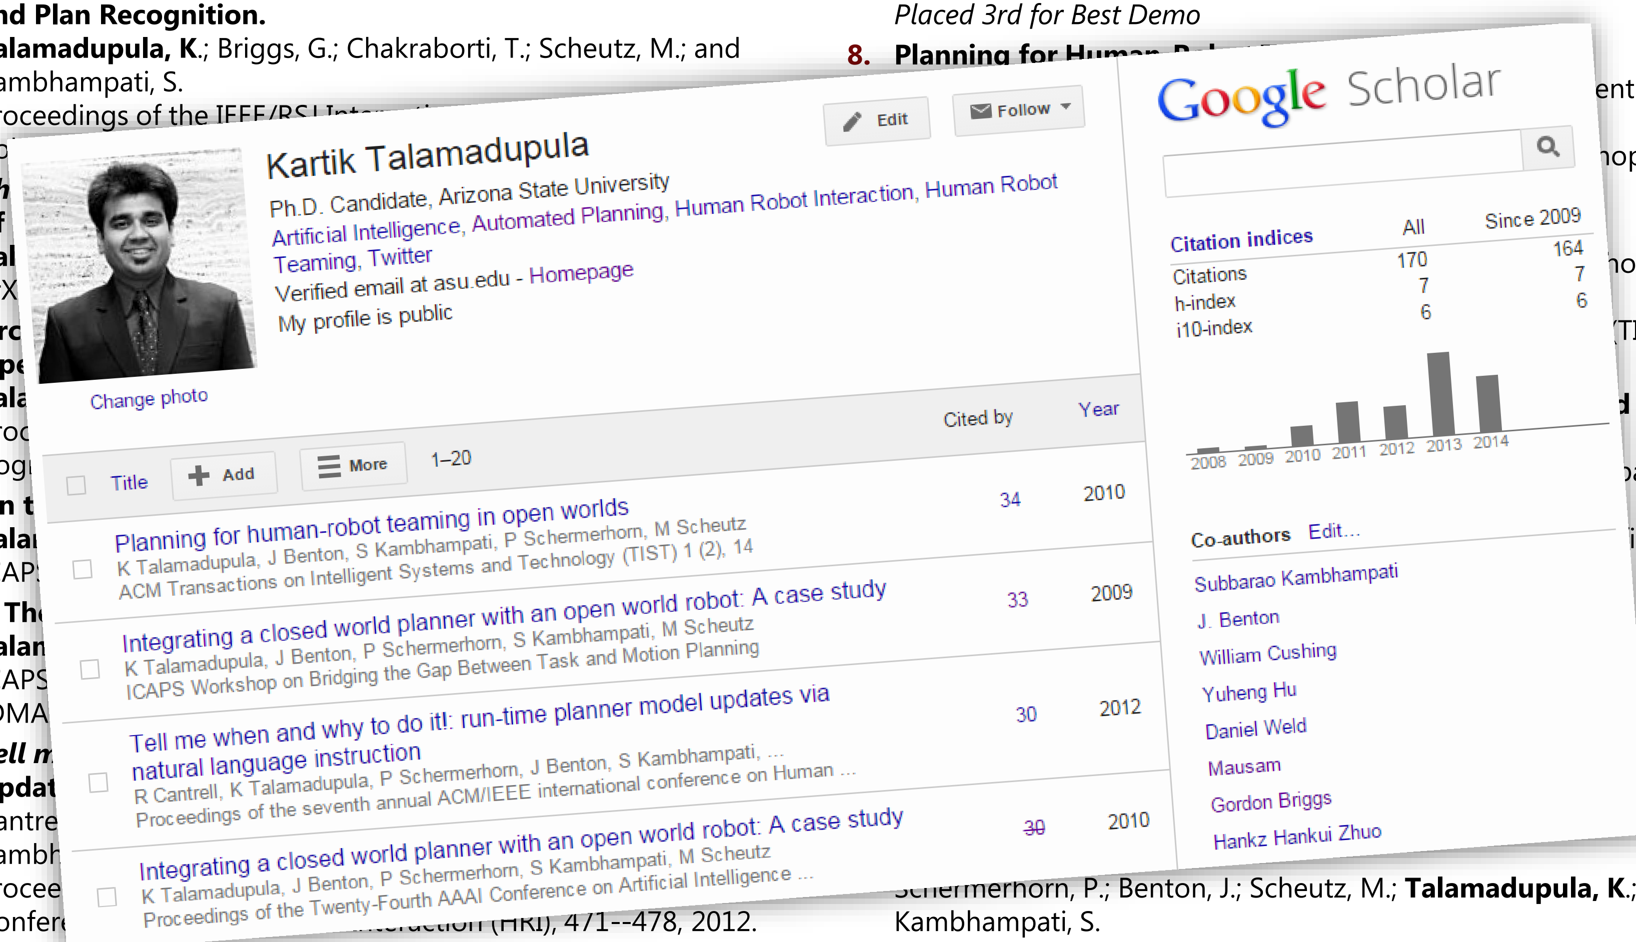

Related Work

Human-Robot Teaming

Symbiotic Autonomy [Rosenthal et al. 2010]

Seeking Human Help [Rosenthal & Veloso 2012]

Replanning with Dynamic Information [Coltin & Veloso 2013]

Generalized Architectures for Distributed Human-Robot Teams [Scerri et al. 2003] [Schurr et al. 2005]

Mixed-Initiative Planning [Bagchi et al. 1996]

Advisable Planning [Myers 1996]

Continuous Planning & Execution [Myers 1998]

TRAINS-95 [Ferguson et al. 1996]

(Open World) Goals

Local Closed Worlds [Etzioni et al. 1997]

Sensing Goals [Scherl & Levesque 1993] [Golden & Weld 1996]

Temporal Goals [Baral et al. 2001] [Bacchus & Kabanza 1996]

Trajectory Constraints (Preferences) [Gerevini et al. 2009]

Replanning & Execution Monitoring

Contingent Planning [Albore et al. 2009] [Meauleau & Smith 2003]

CASPER [Knight et al. 2001]

IxTeT-eXeC [Lemai & Ingrand 2003]

STRIPS [Fikes et al. 1972]

Plan Stability & Repair [Fox et al. 2006] [Van Der Krogt & De Weerdt 2006]

Minimal Perturbation Planning [Kambhampati 1990]

Plan Re-Use [Nebel & Koehler 1995]

Plan Validity [Fritz & McIlraith 2007]

Multi-Agent Systems

Inter and Intra Agent Commiments [Wagner et al. 1999]

Inter-Agent Commitments [Meneguzzi et al. 2013] [Komenda et al. 2012] [Komenda et al. 2008] [Bartold & Durfee 2003] [Wooldridge 2000]

Coordination Using Mental Models

Joint Human Behavior [Klein et al. 2005]

Common Ground [Clark & Brennan 1991]

Coordinated Assembly Tasks [Kwon & Suh 2012]

Object Hand-overs [Strabala 2013]

PLAN & INTENT RECOGNITION

- **Modeling human agent key to teaming**
  - Can augment robot's planning capabilities
  - Information can be used for inter-plan coordination
- **Required information**
  - Action/capability model of the human agent
  - Goal(s) of the human agent
  - Current state of the human agent
- **Planner *simulates* human's mental process**
  - Produces a *predicted plan* that can be used by robot for coordination purposes

[Talamadupula, Briggs et al., IROS14]

Can Belief Models Enhance Planning?

- **Communication Bandwidth**
  - Even with good NLP, there are still bandwidth issues between humans and robots
  - Humans are not always fully explicit about what they are going to do, or what they want
- **Natural Teaming**
  - Agents have good models of each other
  - Enables them to
    - **Anticipate**: Actions of other teammates
    - **Recognize**: The intentions of other teammates
- Can **affect the robot's planning** in turn

Beliefs, Intentions & Teaming

- Agents have **beliefs** and **intentions**
  - An agent can model its *team members'* beliefs and intentions

$$\{ \phi \mid bel(\alpha, \phi) \in Bel_{self} \}$$

$$\{goal(\alpha, \phi, P) \mid goal(\alpha, \phi, P) \in Bel_{self}\}$$

- This information can be used to **predict the plans** of team members

[Briggs & Scheutz, SIGDIAL11]

Proposed Approach

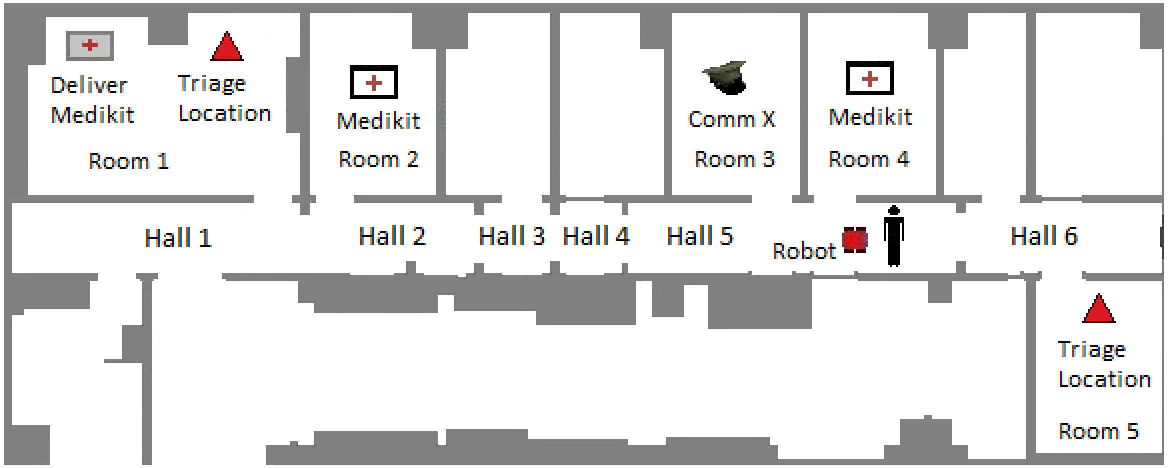

1. Map the robot's beliefs and knowledge about CommX into a new planning instance
2. Generate a plan for this instance – prediction of CommX's plan
3. Extract relevant information from the predicted plan
   - Which medkit will CommX pick up?
4. Use the extracted information to deconflict robot's plan

[Talamadupula, Briggs et al., IROS14]

Mapping to Planning

- Used for high-level plan synthesis

- Can be used to **simulate** the agent's plan
  - Based on known beliefs and intentions
  - Some information about agent's capabilities

- Automated Planning Instance:
  - **Initial State**: All known beliefs of that agent
  - **Goal Formula**: All known goals of that agent
  - **Action Model**: Precondition/Effect description

[Talamadupula, Briggs et al., IROS14]

Mapping to Planning

- Beliefs of another agent α

$$bel_{\alpha} = \{ \phi \mid bel(\alpha,\phi) \in bel_{self} \}$$

- Intentions of another agent α

$$goals_{\alpha} = \{ goal(\alpha,\phi,P) \mid goal(\alpha,\phi,P) \in bel_{self} \}$$

where P is a goal priority

- Mapping to a planning problem

$$I = \{ \phi \mid bel(\alpha,\phi) \in bel_{robot} \}$$

$$G = \{ \phi \mid goal(\alpha,\phi,P) \in bel_{robot} \}$$

$$O = \{ o \mid o \in (\phi \mid \phi \in (I \cup G) \}$$

[Briggs & Scheutz, SIGDIAL11]

[Talamadupula, Briggs et al., IROS14]

Use Case Scenario

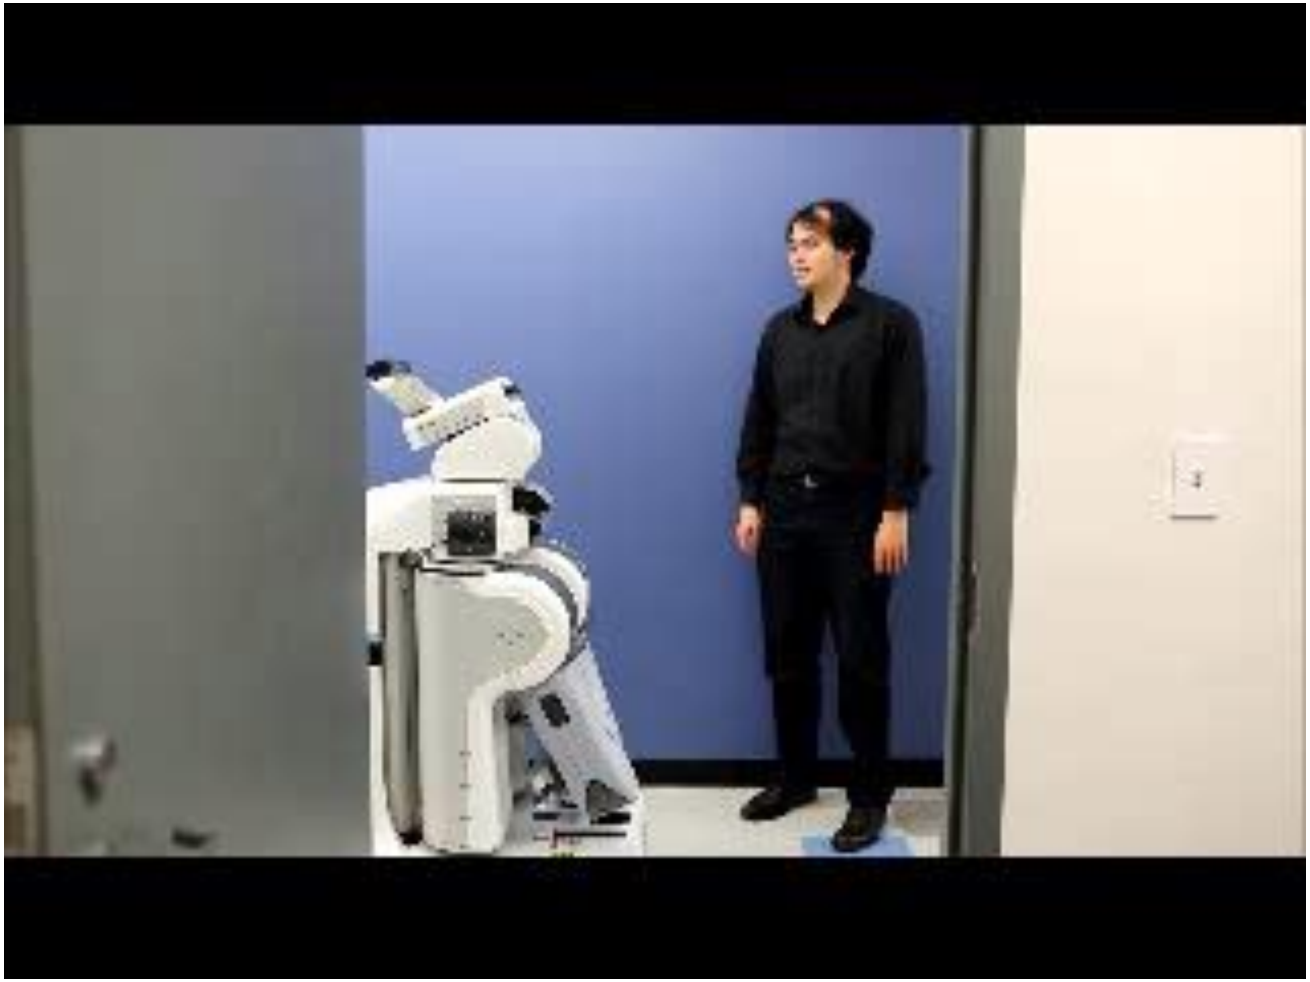

[In collaboration with hrilab, Tufts University]

[Talamadupula, Briggs et al., IROS14]

Use Case Scenario

Comm X's Goal

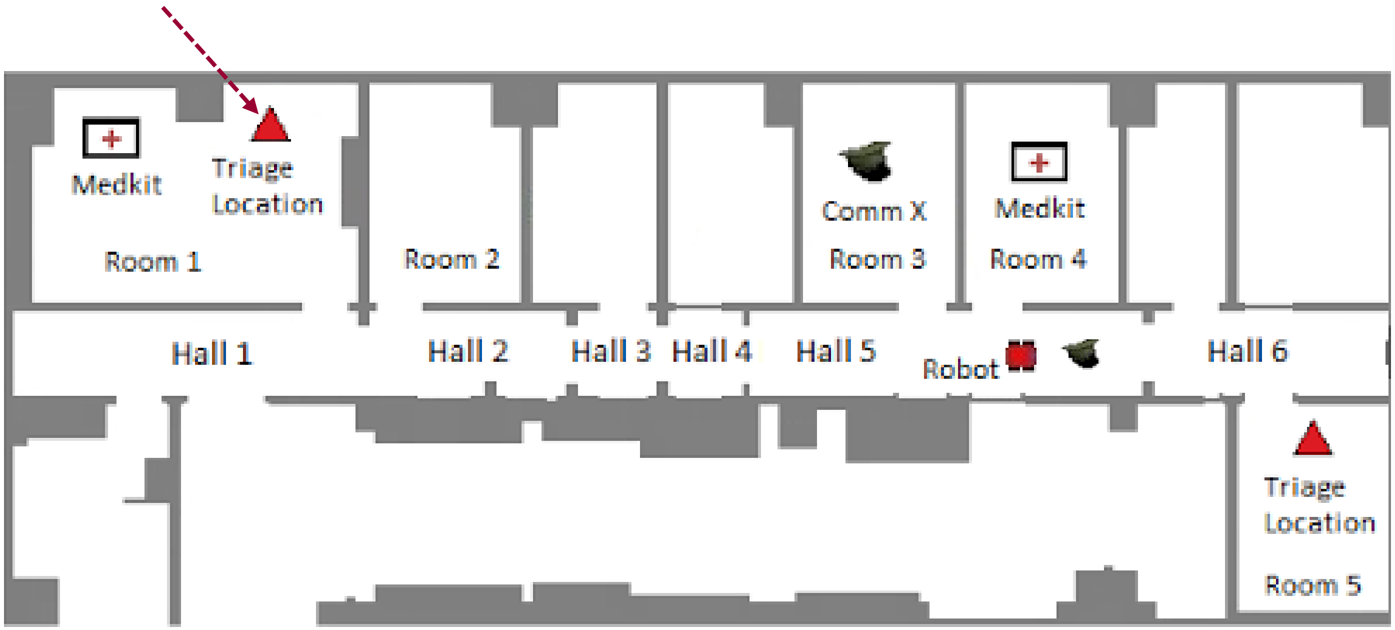

CommY:	"CommX is going to perform triage at Room 1."
Robot:	"Okay."
CommY:	"I need you to take a medkit to Room 5."
Robot:	"Okay…"
	"I am picking up the medkit at Room 4."

[Talamadupula, Briggs et al., IROS14]

Solution

Comm X's Goal

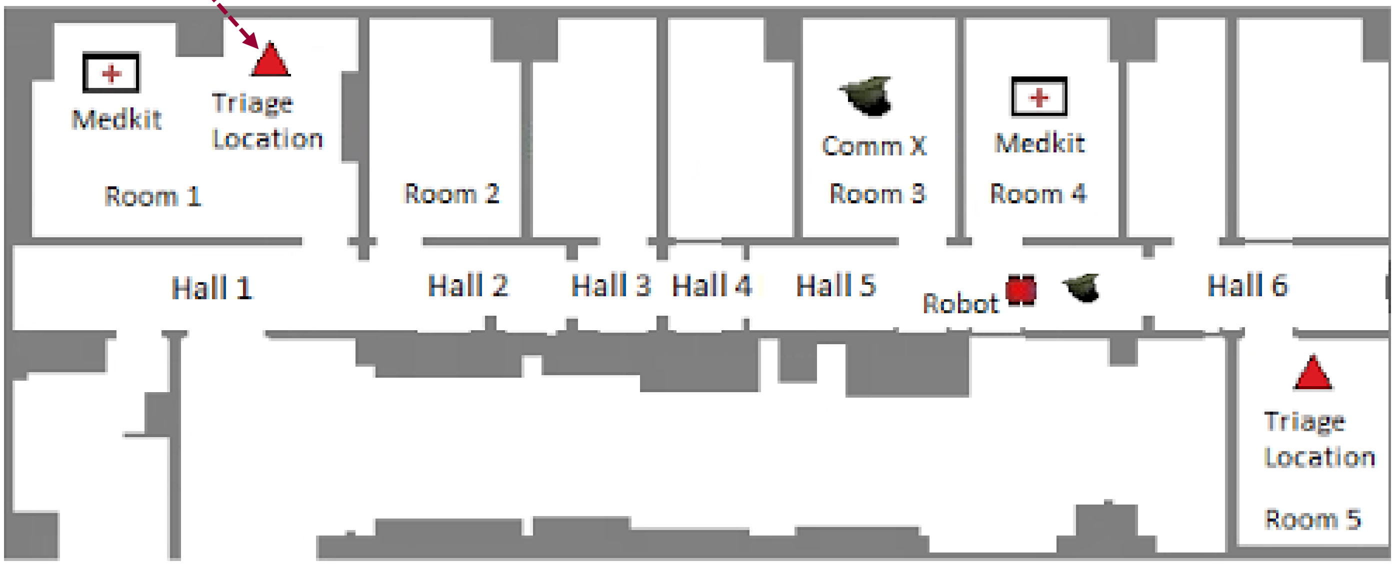

PREDICTED PLAN FOR COMMX

move commx room3 hall5
move_reverse commx hall5 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

Preliminary Evaluation

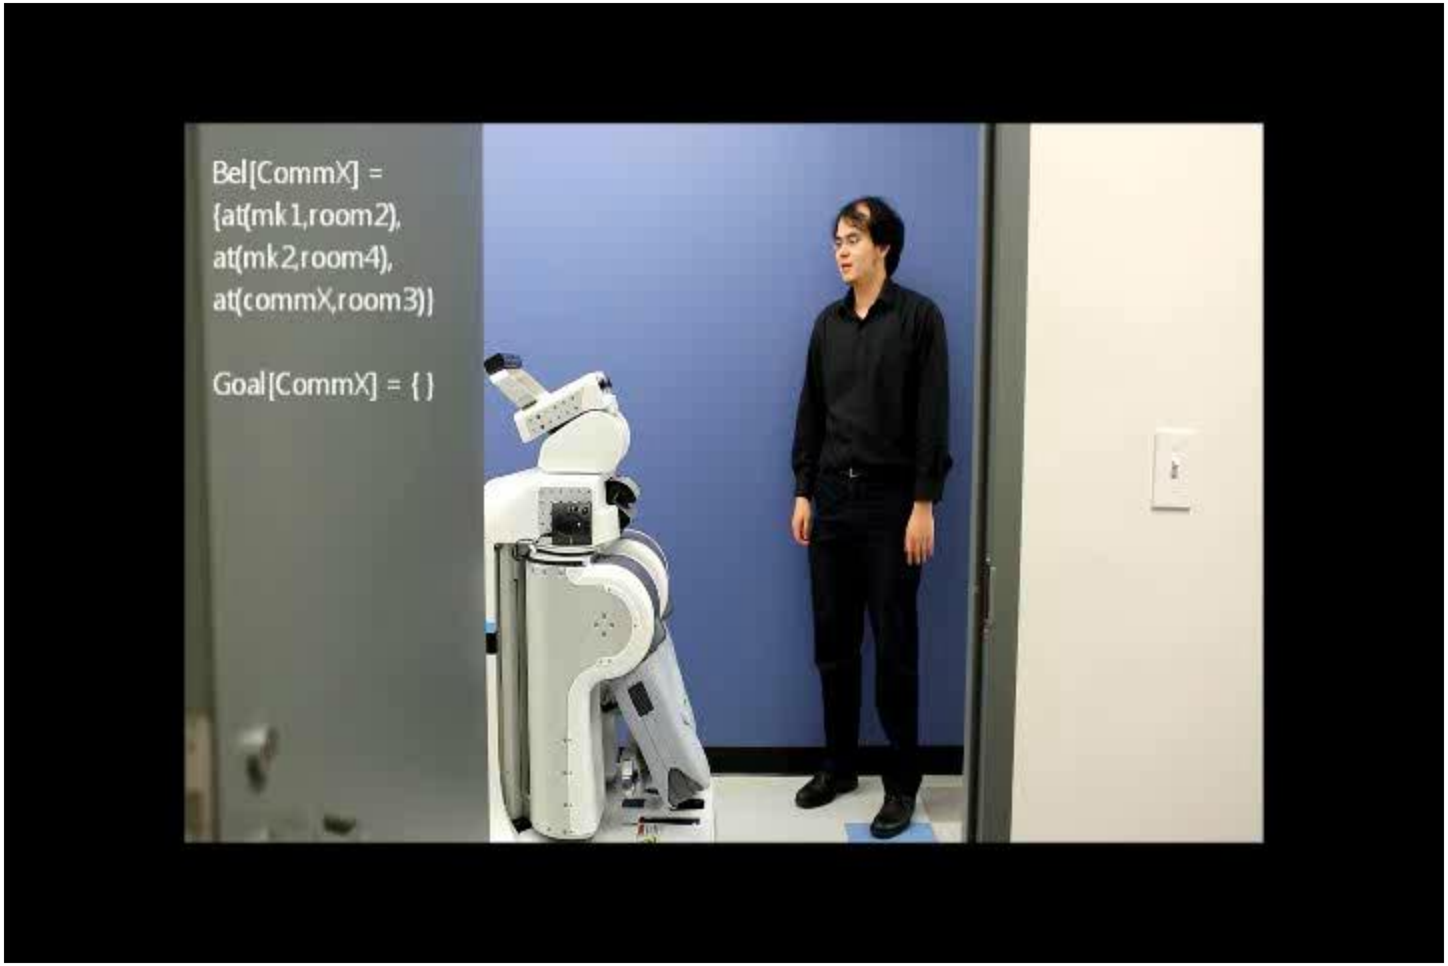

[In collaboration with hrilab, Tufts University]

[Talamadupula, Briggs et al., IROS14]

But what if we don't have full knowledge regarding the team member's goal(s)?

Intent Recognition

- Extend the goal set to a *hypothesized goal set*
  - Contains all possible goals of CommX
- Given a sequence of observations of CommX's actions, recompute the probability distribution over the hypothesized goal set
  - Plan recognition as planning [Ramirez & Geffner 2010]
  - Compiles plan recognition problem into a classical planning problem
- Given more observations, the distribution converges towards the most likely goal
  - (assuming correct observations and rational agency)
- *Incremental* Plan Recognition
  - Can accept a *stream* of observations
  - Incremental re-recognition: Replanning when compiled to classical planning

[Talamadupula, Briggs et al., IROS14]

Evaluation: Intent Recognition I

[Talamadupula, Briggs, Chakrabarti et al., IROS14]

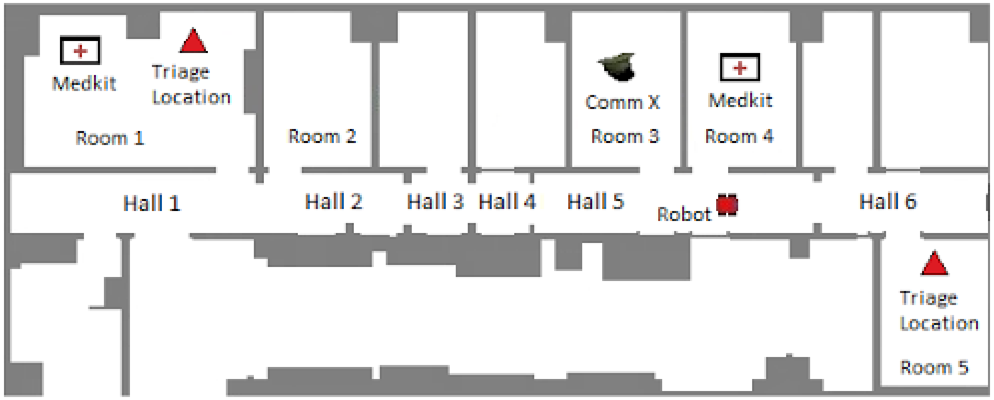

BELIEF IN GOAL

(conducted_triage commX room1)
(conducted_triage commX room5)

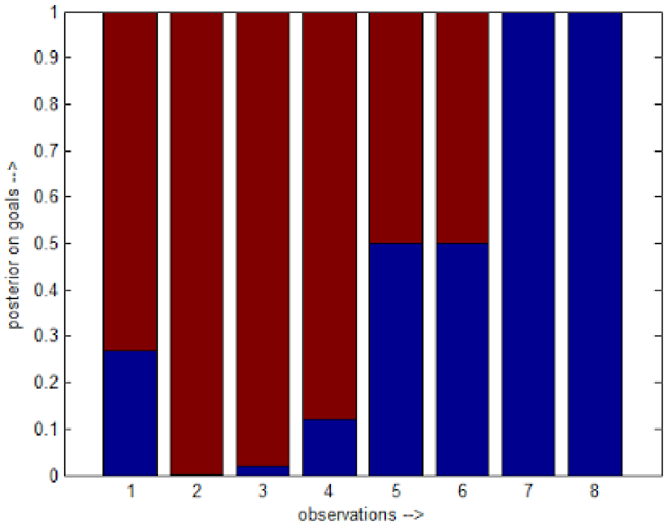

observations -
move commx room3 hall5
move_reverse commx hall5 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

Evaluation: Intent Recognition II

[Talamadupula, Briggs, Chakrabarti et al., IROS14]

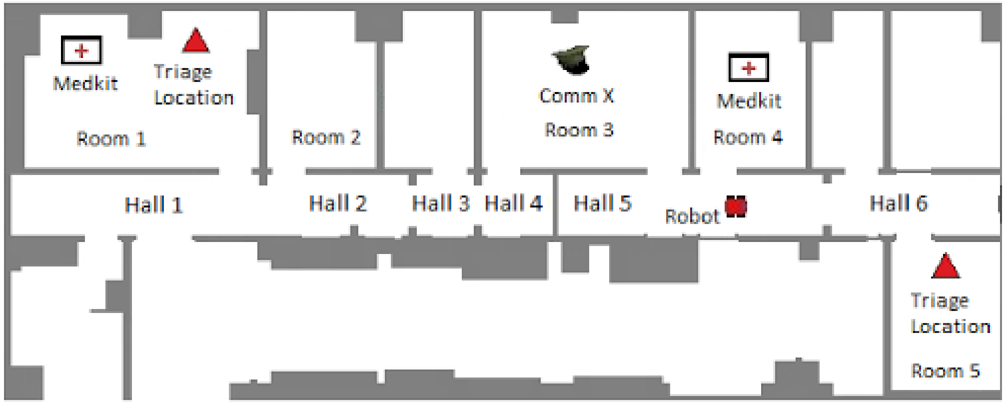

BELIEF IN GOAL

(conducted_triage commX room1)

(conducted_triage commX room5)

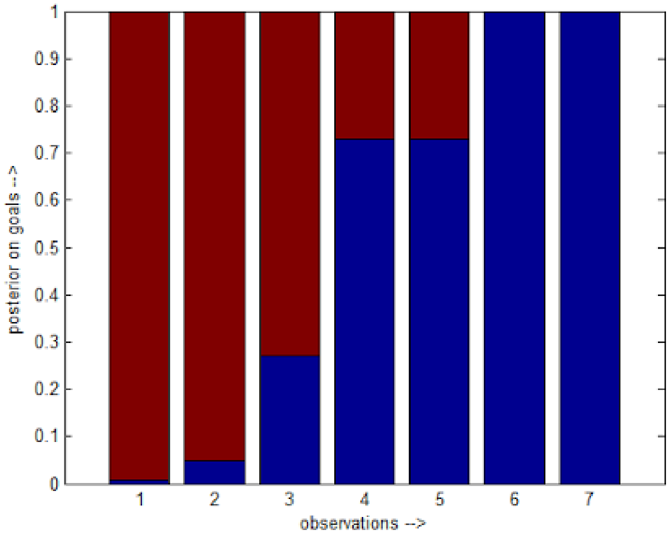

observations -
move commx room3 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

Limitations & Extensions

- **Intentions (and goals) of human** fully known
  - Use observations to determine most likely goals being pursued
- **Model of human** is fully known (and correct)
  - Incomplete models: [Nguyen et al. ICAPS14]
- **High level observations** are **given up-front**
  - Currently given by human (CommY)
  - Going from sensors to observations non-trivial

[Talamadupula, Briggs et al., IROS14]

REPLANNING FOR HUMAN-ROBOT TEAMING

- **Motivating Scenario: Automated Warehouses**
  - Used by Amazon (Kiva Systems) for warehouse management
- **Human: Packager**
  - Only human on the entire floor; remotely located
  - Issues goals to the robotic agents
- **Robot(s): Kiva Robots**
  - Can transport items from shelves to the packager
- **Goals: Order requests; come in dynamically**
  - Goals keep changing as orders pile up
  - World changes as shelves are exhausted; break downs

[IROS09, AAAI10, TIST10, DMAP13, arXiv14]

Replanning Example: Warehouses

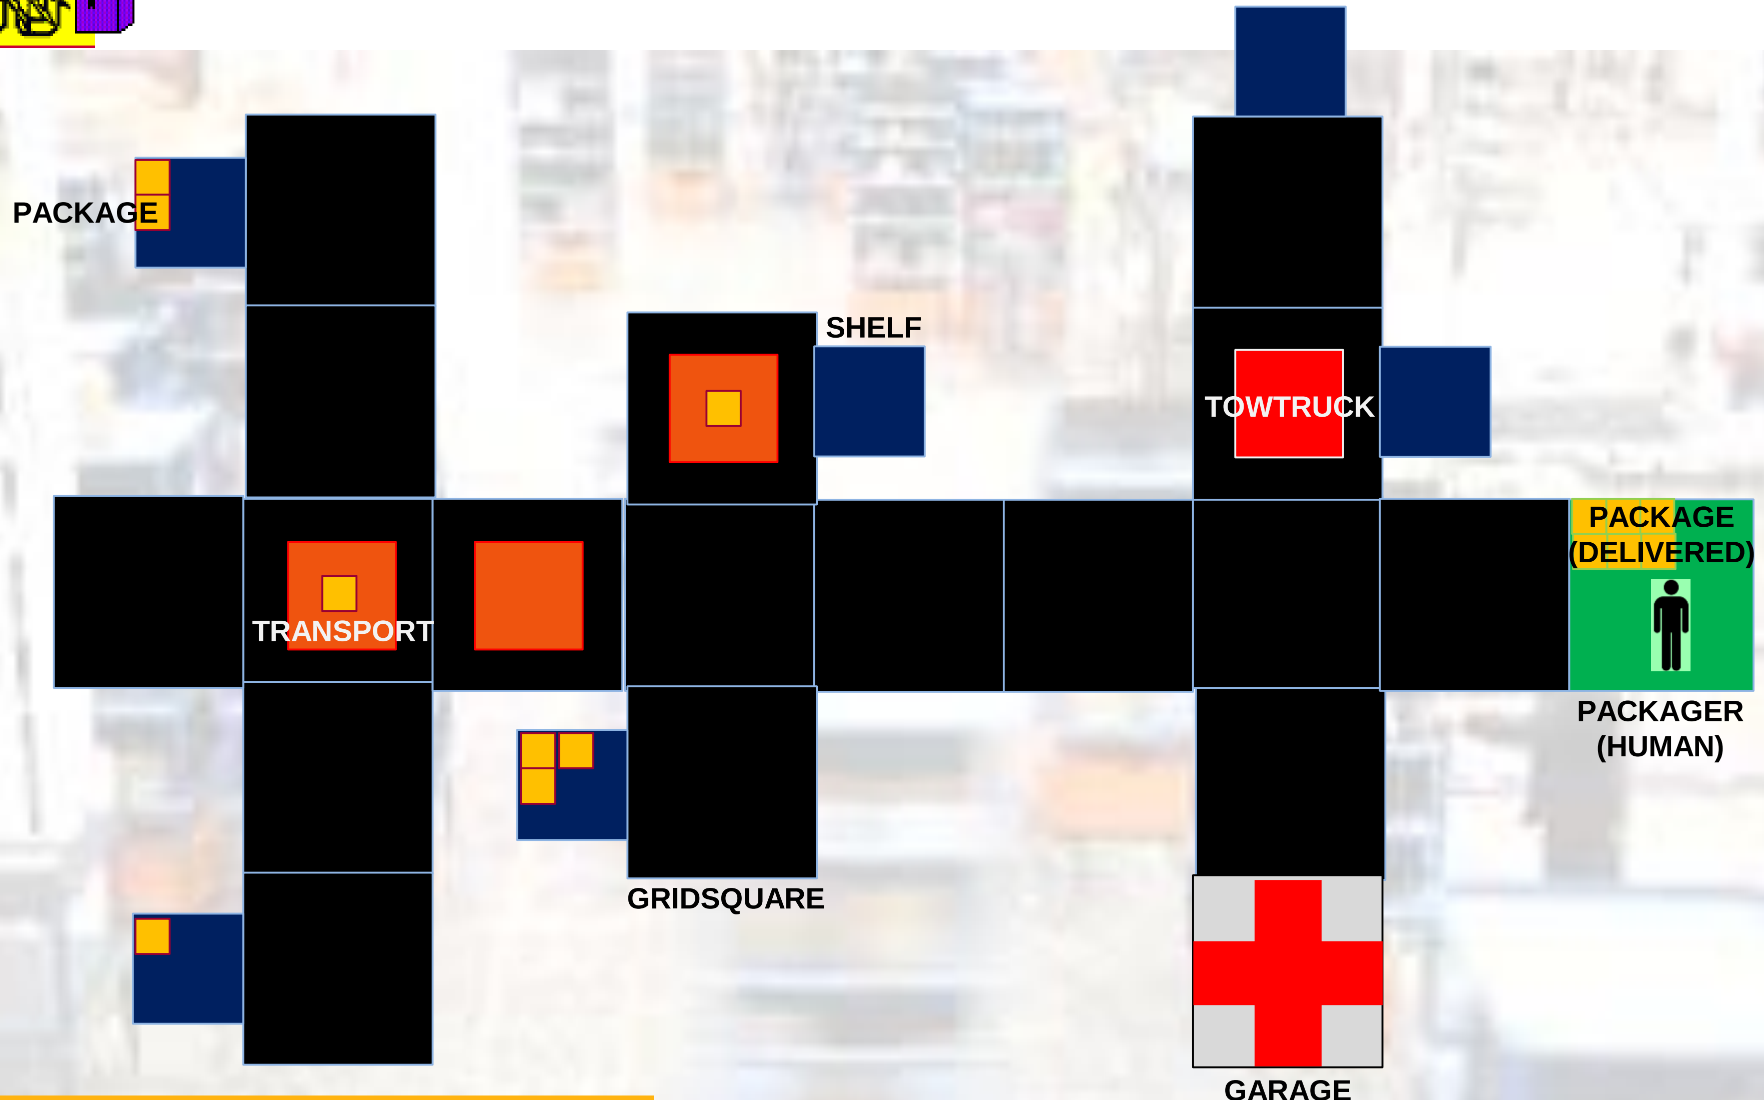

[Talamadupula, Smith et al., Submitted 2014]

Warehouses: Perturbations

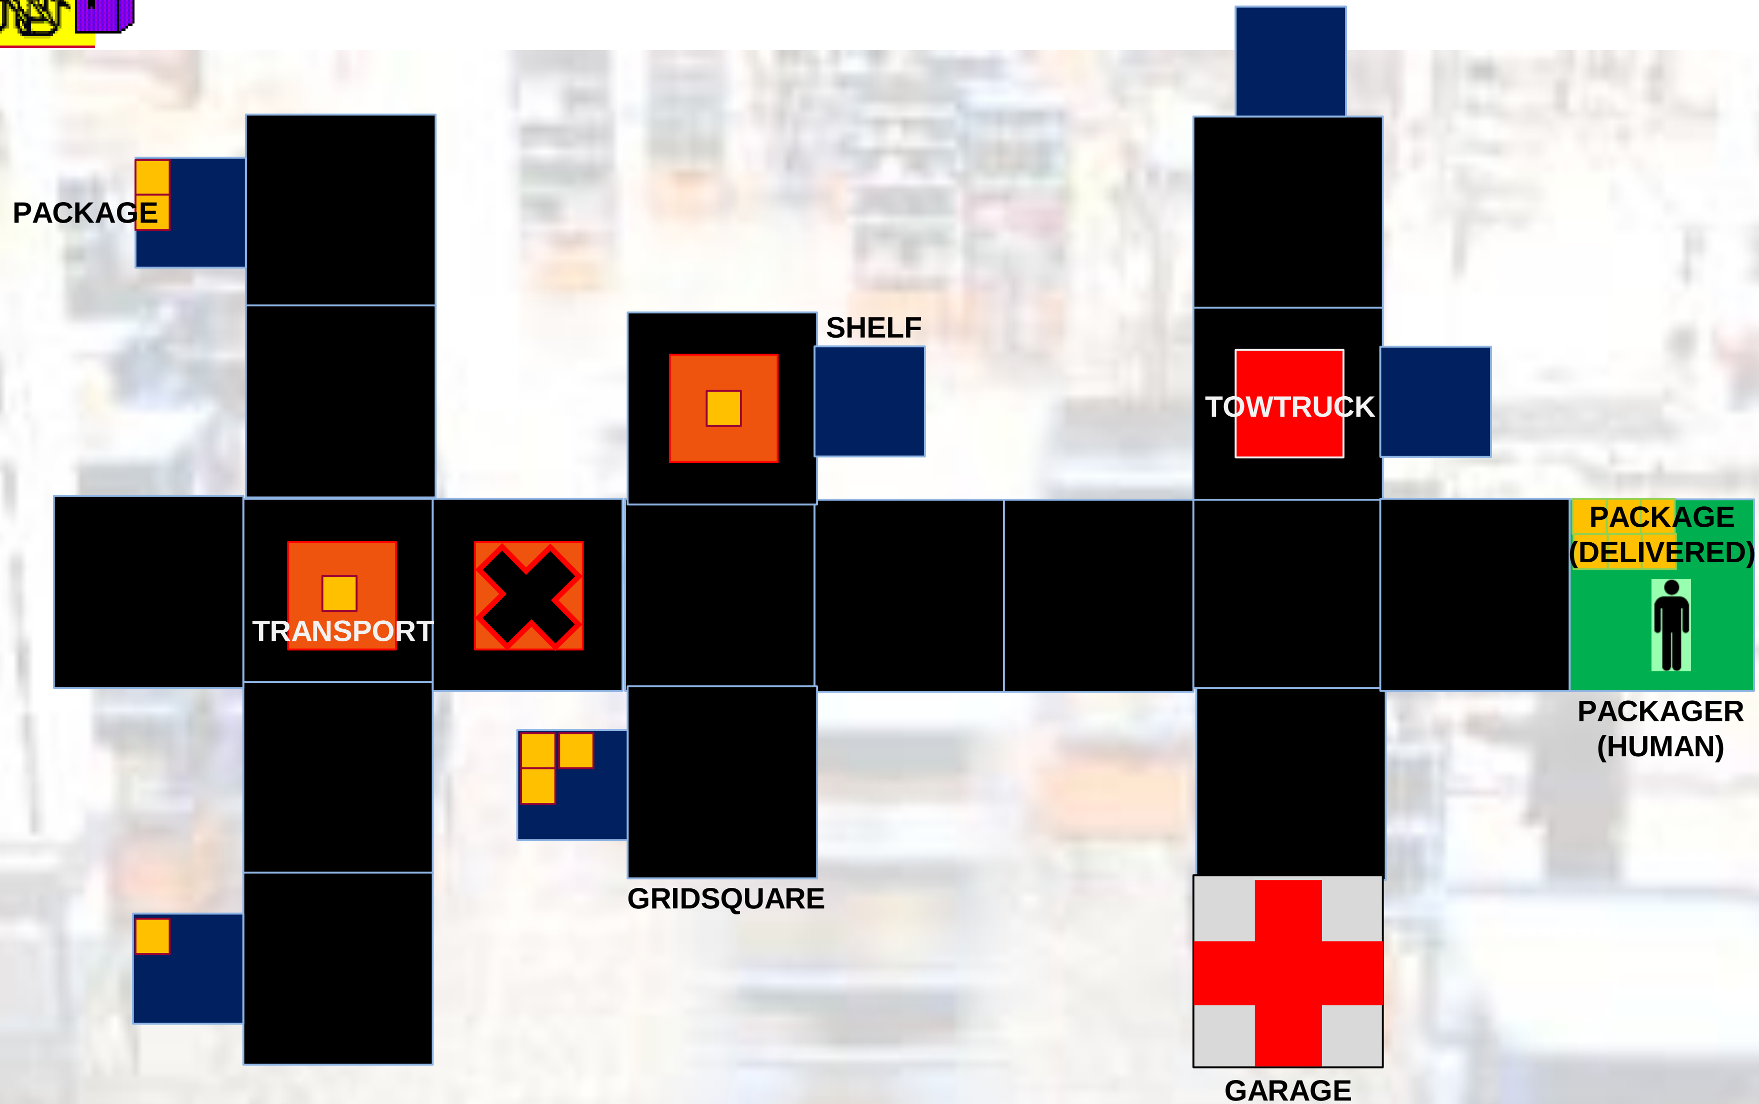

Warehouses: Commitments

1. Transports ***holding*** Packages
2. Towtrucks ***towing*** Transports
3. Packages ***delivered*** to Packager

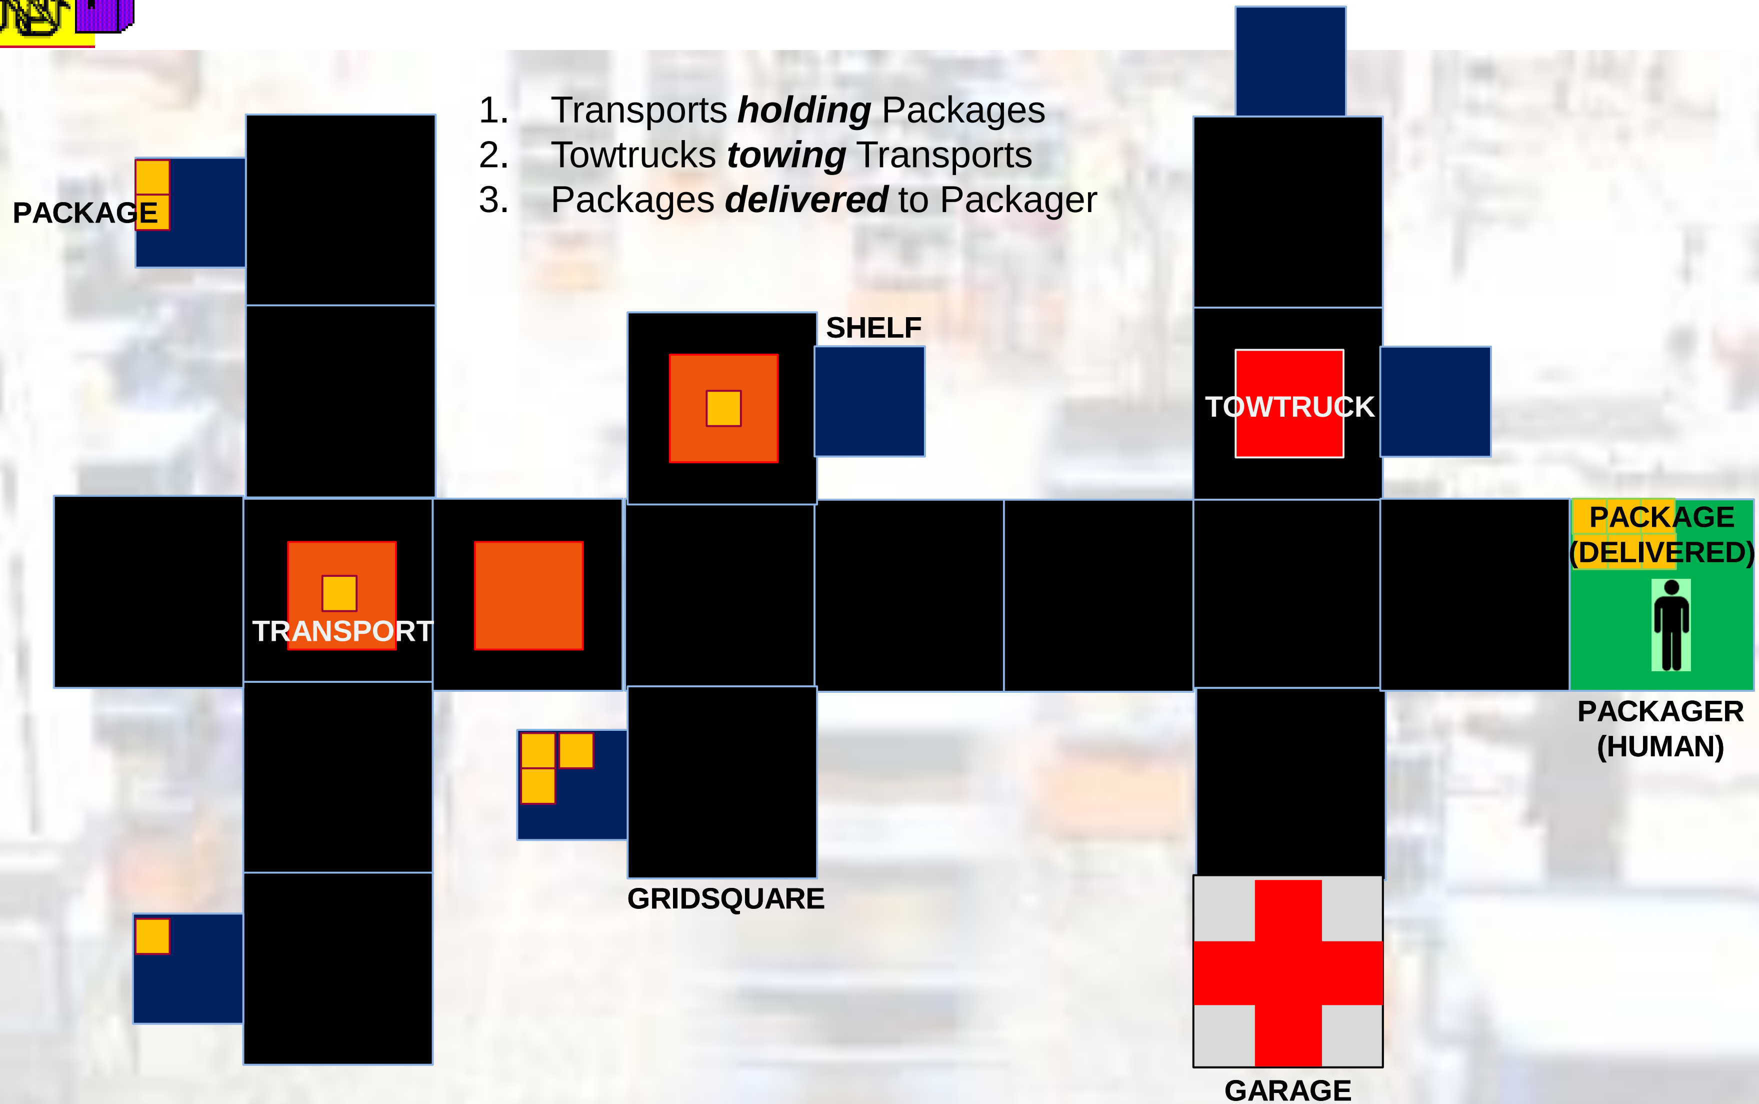

A Generalized Model of Replanning

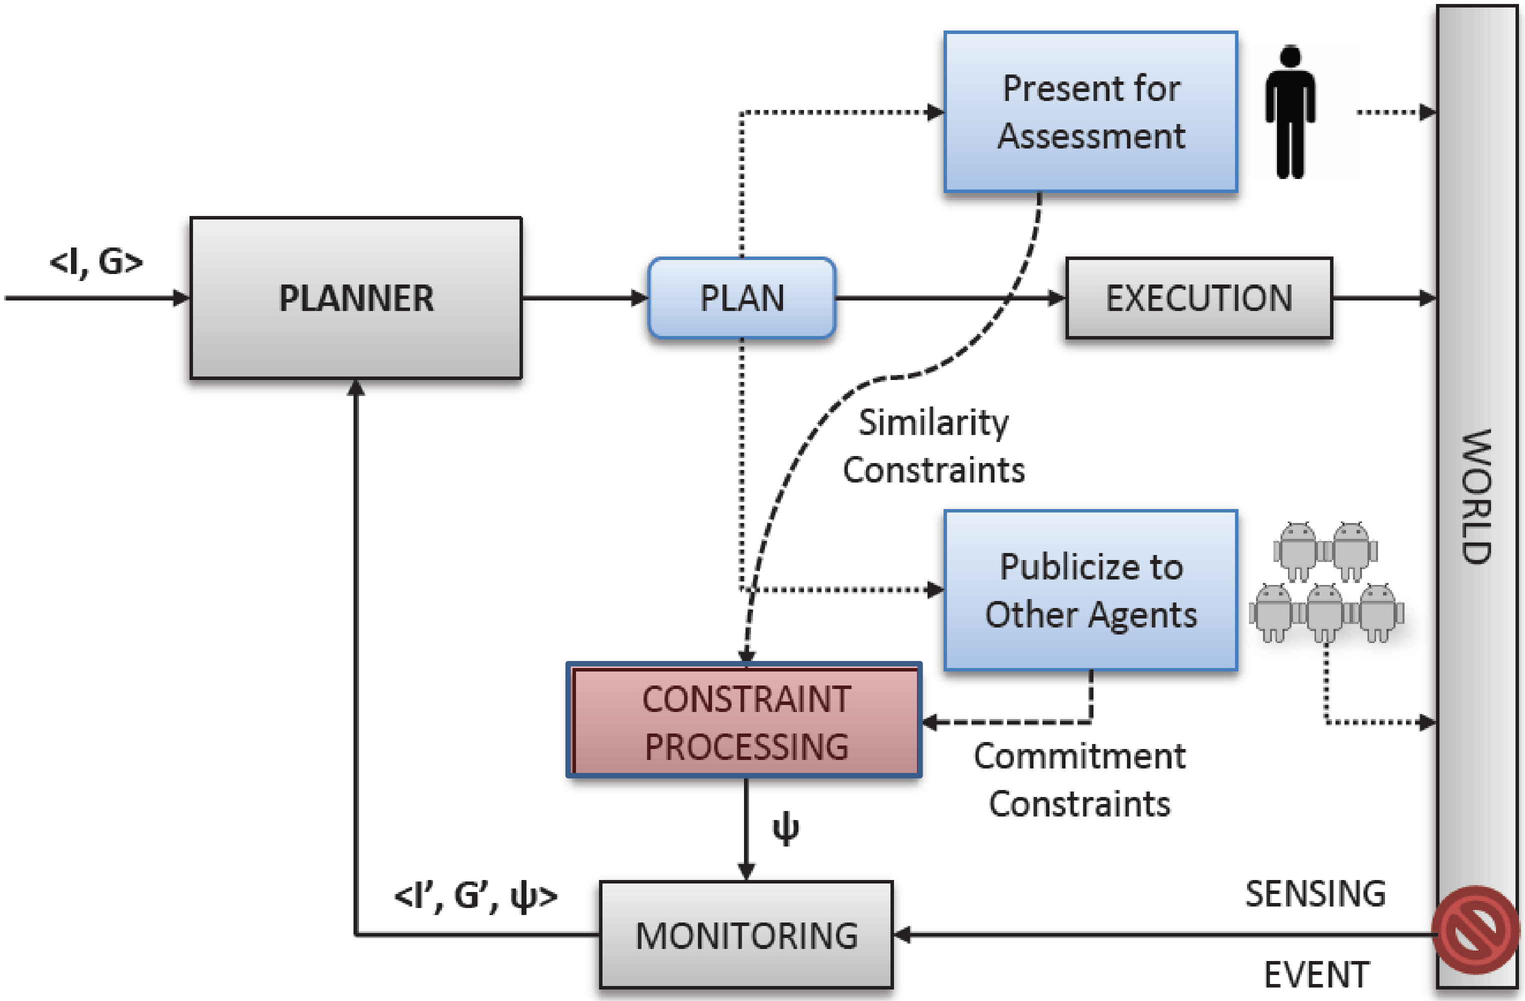

Replanning Constraints

M1 REPLANNING AS RESTART (From scratch)	› **No Constraints**
M2 REPLANNING AS REUSE (Similarity)	› Depends on the **similarity metric** between plans › **ACTION** SIMILARITY $\min \mid \pi \, \Delta \, \pi` \mid$ › **CAUSAL** SIMILARITY $\min \mid CL(\pi) \, \Delta \, CL(\pi`) \mid$
M3 REPLANNING TO KEEP COMMITMENTS	› **Dependencies** between π and other plans › Project down into **commitments** that π\` must fulfill › Exact nature of commitments depends on π › E.g.: **Multi-agent** commitments (between rovers)

Replanning: Solution Techniques

M1 REPLANNING AS RESTART (From scratch)	**CLASSICAL PLANNING**	› Solve new instance [I`,G`] for π` using classical planner
M2 REPLANNING AS REUSE (Similarity)	**ITERATIVE PLAN REPAIR** (Local Search)	› Start from π › Minimize differences while finding a candidate π` › Stop when [I`,G`] satisfied
M3 REPLANNING TO KEEP COMMITMENTS	**COMPILATION** (Partial Satisfaction Planning)	› Commitments are *constraints* on plan generation process › Commitments = Soft Goals G_s › Add G_s to G` → G`` › Run PSP planner with [I`,G``]

Research Question

There exist multiple replanning solution techniques, founded in addressing different constraints during the replanning process.

1. *To what extent do the constraints imposed by one type of replanning formulation act as a surrogate in tracking the constraints of another?*

2. *Are the different replanning metrics good surrogates of each other?*

[Talamadupula, Smith et al., Submitted 2014]

Experimental Setup

1. Generate randomized problem instances of increasing complexity

2. Set up replanning constraints for each replanning metric
    a. Speed: No constraints
    b. Similarity: Number of differences with previous plan
    c. Commitment Satisfaction: Enumerate commitment violations

3. Perturb the initial problem instance; create a perturbed instance for each case (2a, 2b, 2c)

4. Run problem instances with a PSP or preference based planner

Experimental Results

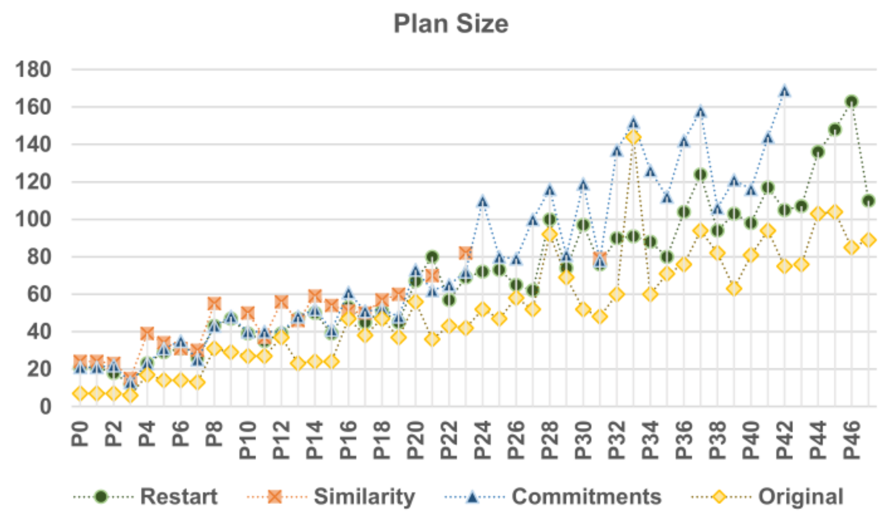

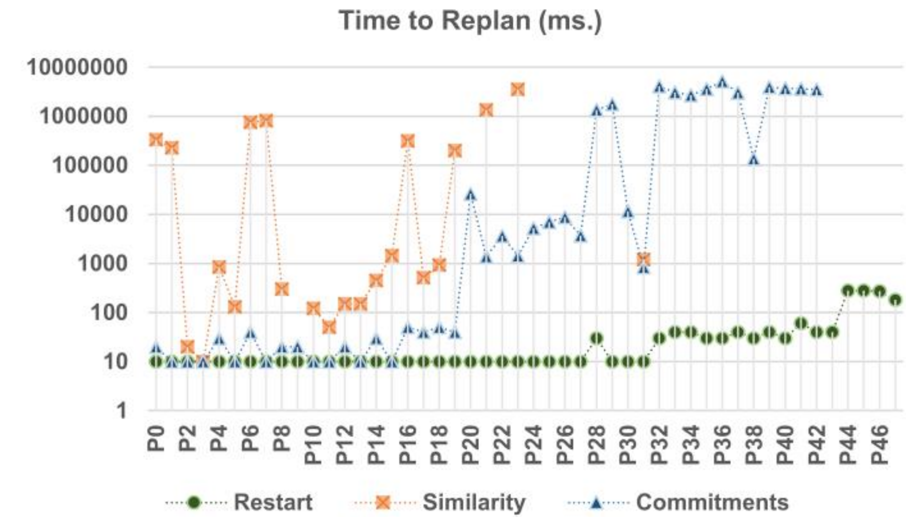

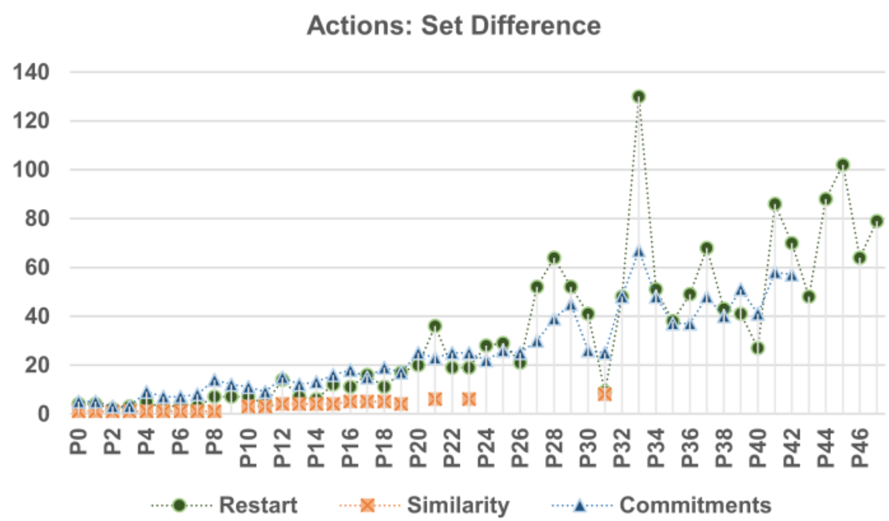

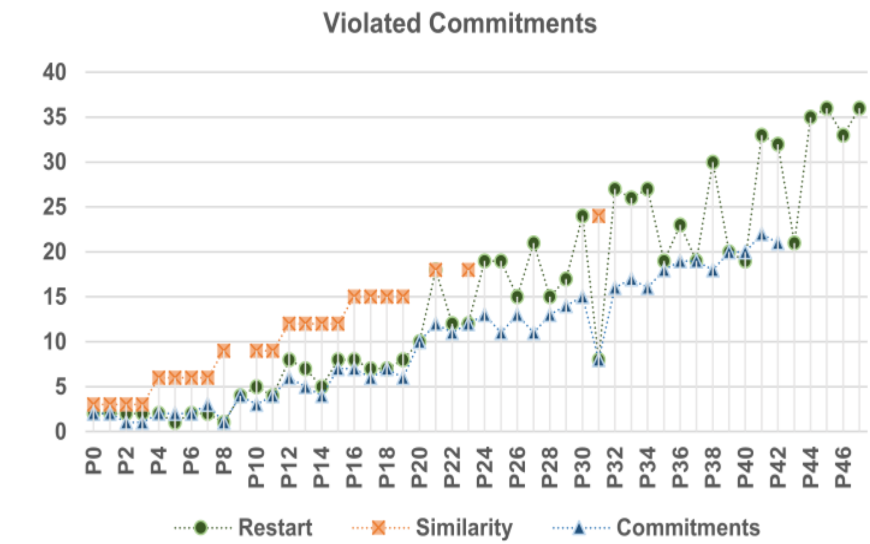

Limitations & Extensions

- **Coverage**: IPC Benchmark Domains
  - Additional experimental conditions
- Modeling **Execution Failures**
  - Currently initial state is perturbed
    - Approximation of execution failure
    - Solution: Perturb state where execution stopped
- **Compilation** to **Classical Planning**
- Replanning **Metrics**
  - Realistic cost and penalty estimates

Broader Impact: HIL Planning

VOTED ICAPS 2014 BEST DEMO BY … THE CROWD!

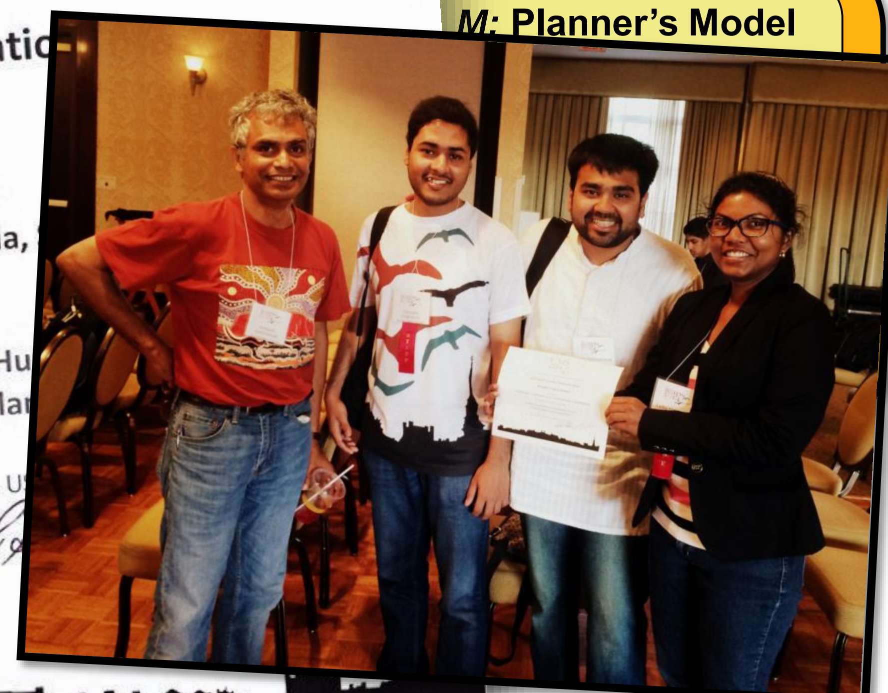

ATION

M: Planner's Model

[HCC

Other Work

Planning for Network Security

- Apply automated planners to the Strategic Planning problem [arXiv:1305.2561]
(Work done as part of an IBM internship)

Foundations of Automated Planning

- Required Concurrency (in Temporal Planning domains) [ICAPS07]
- Search Space Plateaus [ICAPS10]
- Compilation of Replanning Techniques [DMAP13, arXiv14]

Analyzing Tweet Content

- Analyzing language content to detect formalness [ICWSM13]
- Predicting user engagement with real-world events [Submitted]

Information Retrieval on Twitter

- Improving Twitter Search using source & content trustworthiness [CIKM13, AAAI-LBP13, Submitted]
- Hashtag rectification problem

Collaborators

- **Arizona State University**
  - Subbarao Kambhampati
  - J. Benton (SIFT)
  - William Cushing (UC Berkeley)
  - Yuheng Hu (IBM Almaden)
  - Srijith Ravikumar (Amazon)
  - Raju Balakrishnan (Groupon)
  - Lydia Manikonda
  - Tathagata Chakraborti
  - Sumbhav Sethia
  - Sushovan De (Google)
  - Paul Reesman
  - Hankz Hankui Zhuo (Sun-Yat Sen U.)
  - Yu Zhang
  - Nancy Cooke (ASU Poly)
  - Cade Bartlett (ASU Poly)
- **Tufts University**
  - Matthias Scheutz
  - Gordon Briggs
- **NASA Ames Research Center**
  - David E. Smith
- **Indiana University**
  - Paul Schermerhorn (SOARTech)
  - Rehj Cantrell (Nuance Communications)
- **IBM Research (TJ Watson)**
  - Anton V. Riabov
  - Octavian Udrea
  - Anand Ranganathan
- **IBM Research (India)**
  - Shalini Kapoor (IBM Global Services)
  - Shachi Sharma (IRL Delhi)
  - Biplav Srivastava (IRL Delhi)
- **University of Washington**
  - Daniel S. Weld
  - Mausam (IIT Delhi)
- **University of Freiburg**
  - Patrick Eyerich
  - Robert Mattmueller

Conclusion

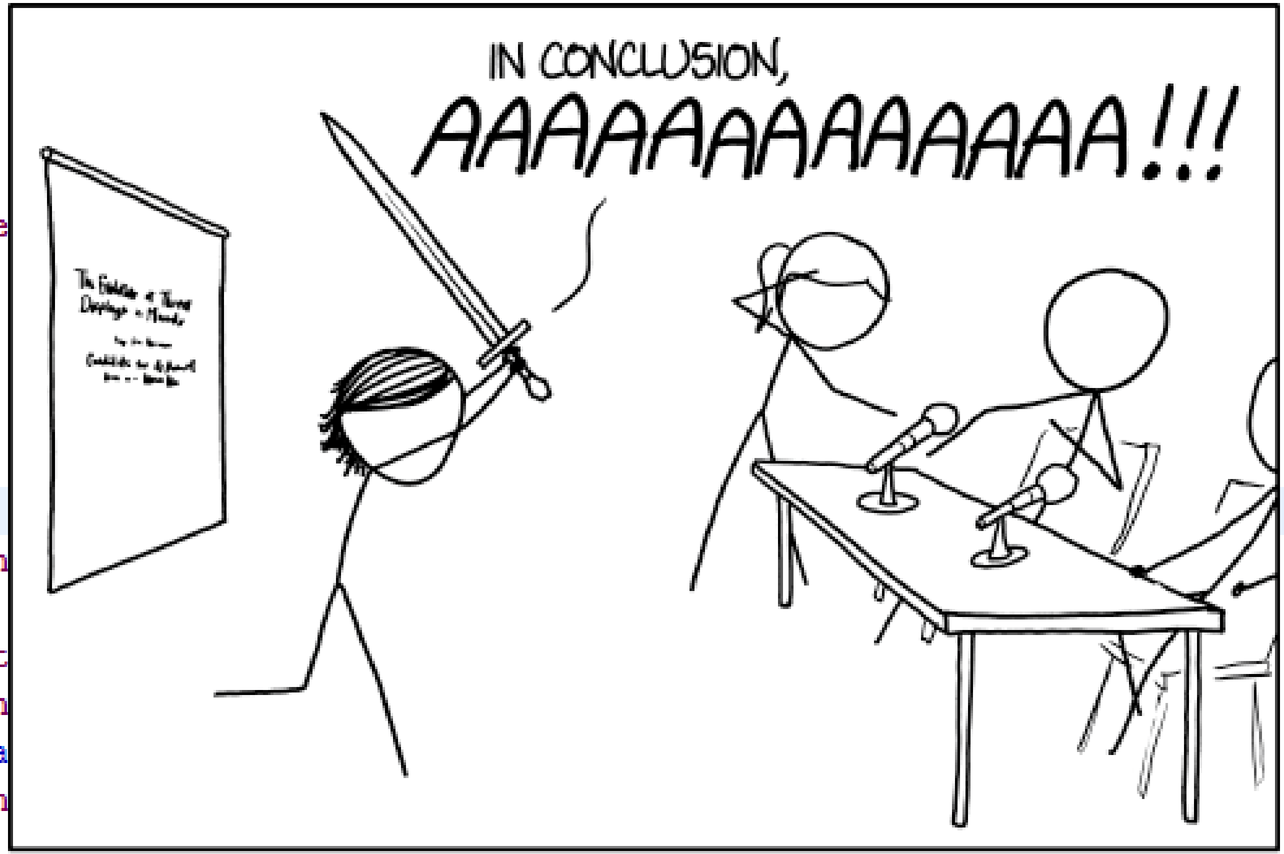

THE BEST THESIS DEFENSE IS A GOOD THESIS OFFENSE.

```
            } e

            }
        } catch
            //
            ret
        } catch
            pla
            syn

            }
            System.err.println("The Legendary Memory Beast of Aaaaarrrrrrggghhh");
            // restart this thread??
        } catch (Exception e) {
        }
    }
}
```

Challenges Addressed

1. **OPEN WORLD GOALS**
   - Provide a way to specify quantified goals on unknown objects
   - Consider a more principled way of handling uncertainty in facts
2. **REPLANNING**
   - Handle state and goal updates from a changing world while executing
   - Present a unified theory of replanning, to analyze tradeoffs
3. **MODEL UPDATES**
   - Accept changes to planner’s domain model via natural language
4. **PLAN RECOGNITION**
   - Use belief models of other agents to enhance planning

Summary

- **Planning for Human-Robot Teaming** (HRT) is an important problem

- Demonstrated the **successful integration** of a planner with an architecture for HRT

- Detailed **techniques** used in that integration, and **novel extensions and analysis** of some of them
  1. Replanning
  2. Plan & Intent Recognition
  3. Open World Quantified Goals
  4. Model Updates

- Broader Implications: **Human-in-the-Loop Planning**