Planning Challenges in Human-Machine Collaboration

Subbarao Kambhampati
Arizona State University

Funding from ONR, ARO and NSF gratefully acknowledged

Planning Challenges in Human-Machine Collaboration

Subbarao Kambhampati
Arizona State University

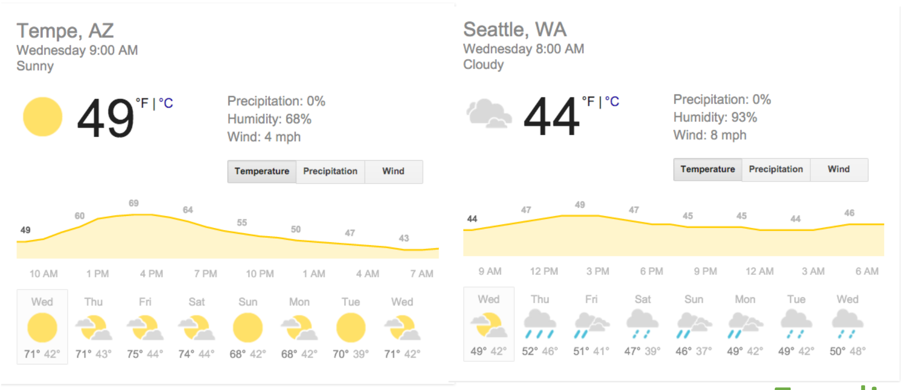

Funding from ONR, ARO and NSF gratefully acknowledged


FLASHBACK

590 A - Research Seminar in Artificial Intelligence **Autumn Quarter 2007** Faculty organizer: Mausam **Day / Time: Fridays 3:30-4:20** **Location: EEB 003**	

Theme: The Future of AI

The theme for CSE590A will be 'The Future of AI'. We will present talks from several seasoned researchers regarding their vision for AI. The hope is that the seminars will excite all of us regarding the role of AI in our future, intrigue and puzzle us by the burning open research questions and hopefully also provoke us in thinking about a broader, long-term vision for our own research. The format will be informal and interactive and we expect to have fun discussions after the talks.

Mailing List

We will not use the cse590a mailing list. Instead, announcements about the seminar will go to uw-ai. If you do not already subscribe to uw-ai, then join by sending mail to uw-ai-request@cs.washington.edu, with the line "subscribe listname" in the body of the message. You are **encouraged** to discuss the presentations on the mailing list.

Calendar

	Speaker	Title
September 28	No class, Welcome TGIF!	
October 5	Eric Horvitz, Microsoft Research	The Future of AI
October 12	Wolfram Burgard, University of Freiburg	The Future of AI: a Robotics Perspective
October 19	Benjamin Grosof, Semantic Technologies, Vulcan Inc.	The Future of AI, with a Semantic and Business Focus
October 26	Pedro Domingos, University of Washington	How We're Going to Solve the AI Problem
November 2	Subbarao Kambhampati, Arizona State University	Future of AI: Darned Humans--can't live with them and can't live without them (Audio)
November 9	Thomas Dietterich, Oregon State University	The Future of AI: Learning, Manipulating, Generating, and Recognizing Activities
November 16	Dieter Fox, University of Washington	Future of AI: Interacting with the physical world
November 23	No class. Happy Thanksgiving!	
November 30	Dan Weld, University of Washington	The Future of AI
December 7	Oren Etzioni, University of Washington	Paradigm Shift in AI

Future of AI

Human-Aware AI

(aka Darned Humans: Can’t Live with them. Can’t Live without them)

Subbarao Kambhampati

Arizona State University

Given at U. Washington on 11/2/2007

AI's Curious Ambivalence to humans..

- Our systems seem happiest
  - either far away from humans
  - or in an adversarial stance with humans

You want to help humanity, it is the people that you just can't stand…

What happened to Co-existence?

- Whither McCarthy's advice taker?
- ..or Janet Kolodner's house wife?
- …or even Dave's HAL?
  - (with hopefully a less sinister voice)

HAAI

Human-aware AI

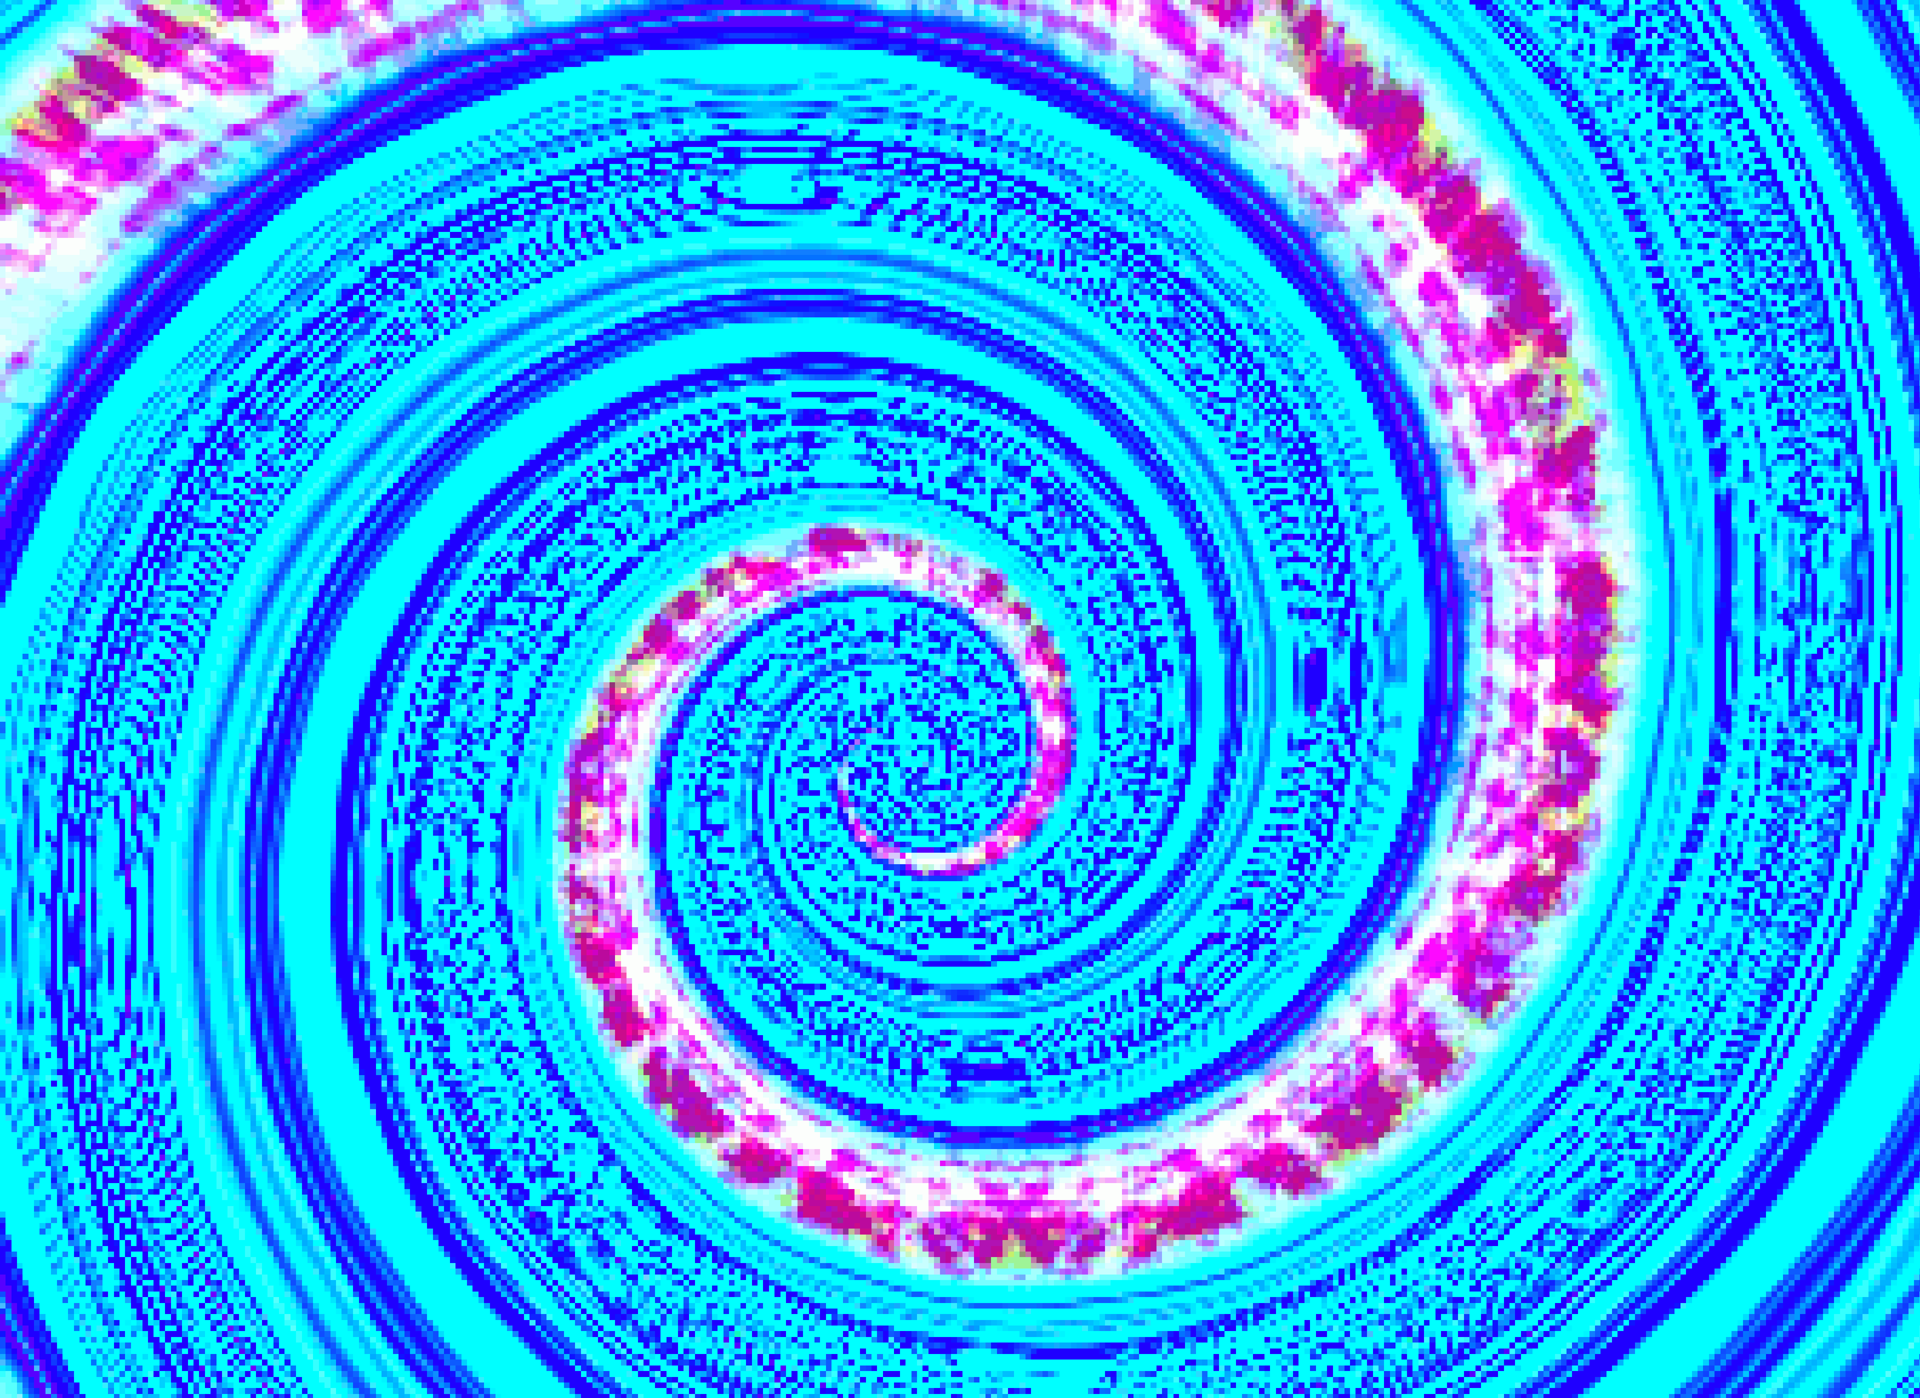


Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons
More than 1,000 experts and
warning of military artificial i
VB
Netflix's Hastings: B
machines and genet
CHRIS O'BRIEN JANUARY 18, 2016 3:41 AM
TAGS: AI, GENETICS, NETFLIX, REED HASTINGS
Image Credit: Flickr/epSos .de
Before Reed Hastings cofounded a little compa
changing the way we watch TV, he was an artifi
AI has come a long way since Hastings got his r
in 1988. But he still follows developments in th
conversation on stage today at the DLD Confer
Hastings said he was far less worried about loo
apocalypse than are many other observers, su
"Some people worry about what happens when machine intelligence is too
strong," Hastings said. "That's like worrying about our Mars colony and people
Zero Gravity Solutions, Inc. Signs Space Act Agreement with NASA Ames Research Center
HAAI
there
ence
loyment
to

September 05, 2007

AI: It's OK Again!

Is AI on the rise again?

· Email · Print
· Discuss · Reprint
add to:
· Del.icio.us · Slashdot
· Digg · Y! MyWeb
· Google · Blink
· Spurl · Furl

(Page 1 of 2)

TABLE OF CONTENTS

1. **Hype and History**
2. Analysis and Synthesis

Michael Swaine

Over the last half century, AI has had its ups and down. But for now, it's on the rise again.

In 1956, John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon organized the

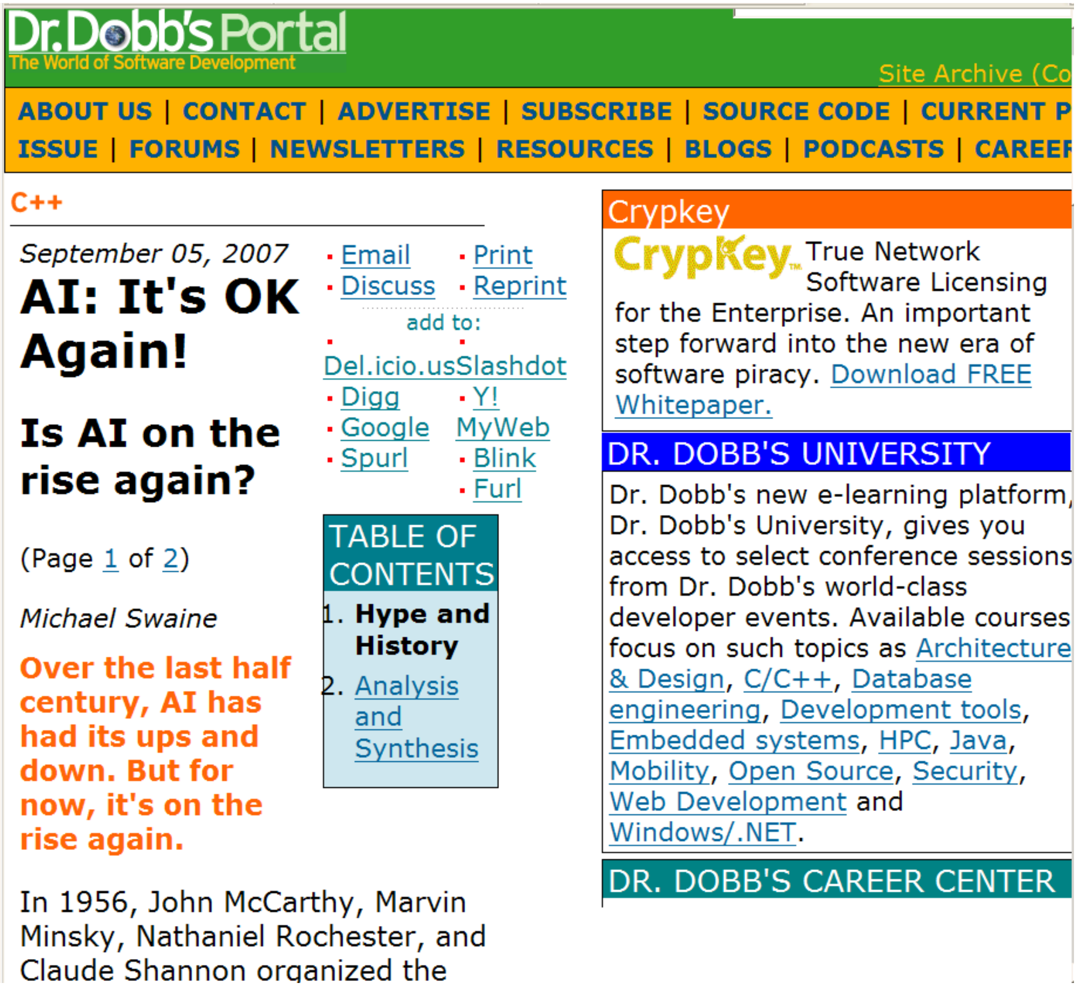
Dr.Dobb's Portal
The World of Software Development
Site Archive (Co
ABOUT US | CONTACT | ADVERTISE | SUBSCRIBE | SOURCE CODE | CURRENT P
ISSUE | FORUMS | NEWSLETTERS | RESOURCES | BLOGS | PODCASTS | CAREEF

C++

September 05, 2007

AI: It's OK Again!

Is AI on the rise again?

(Page 1 of 2)

Michael Swaine

Over the last half century, AI has had its ups and down. But for now, it's on the rise again.

In 1956, John McCarthy, Marvin Minsky, Nathaniel Rochester, and Claude Shannon organized the

- Email
- Print
- Discuss
- Reprint

add to:

- Del.icio.us
- Slashdot
- Digg
- Y!
- Google
- MyWeb
- Spurl
- Blink
- Furl

TABLE OF CONTENTS

1. **Hype and History**
2. Analysis and Synthesis

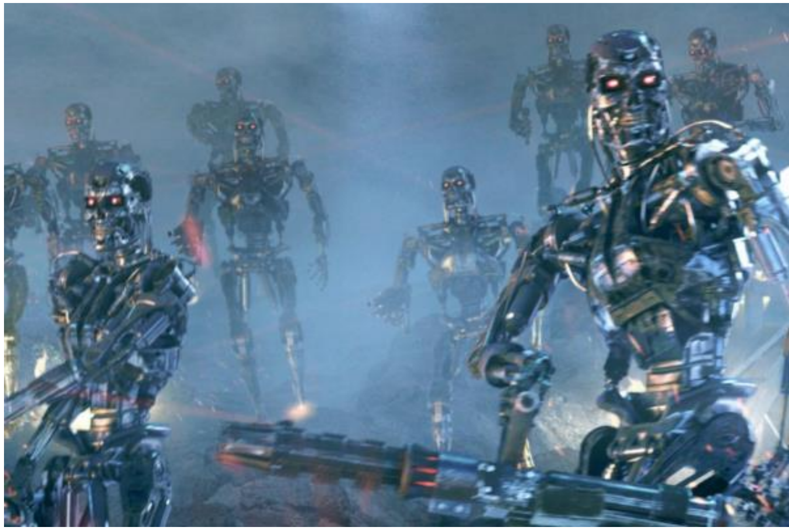

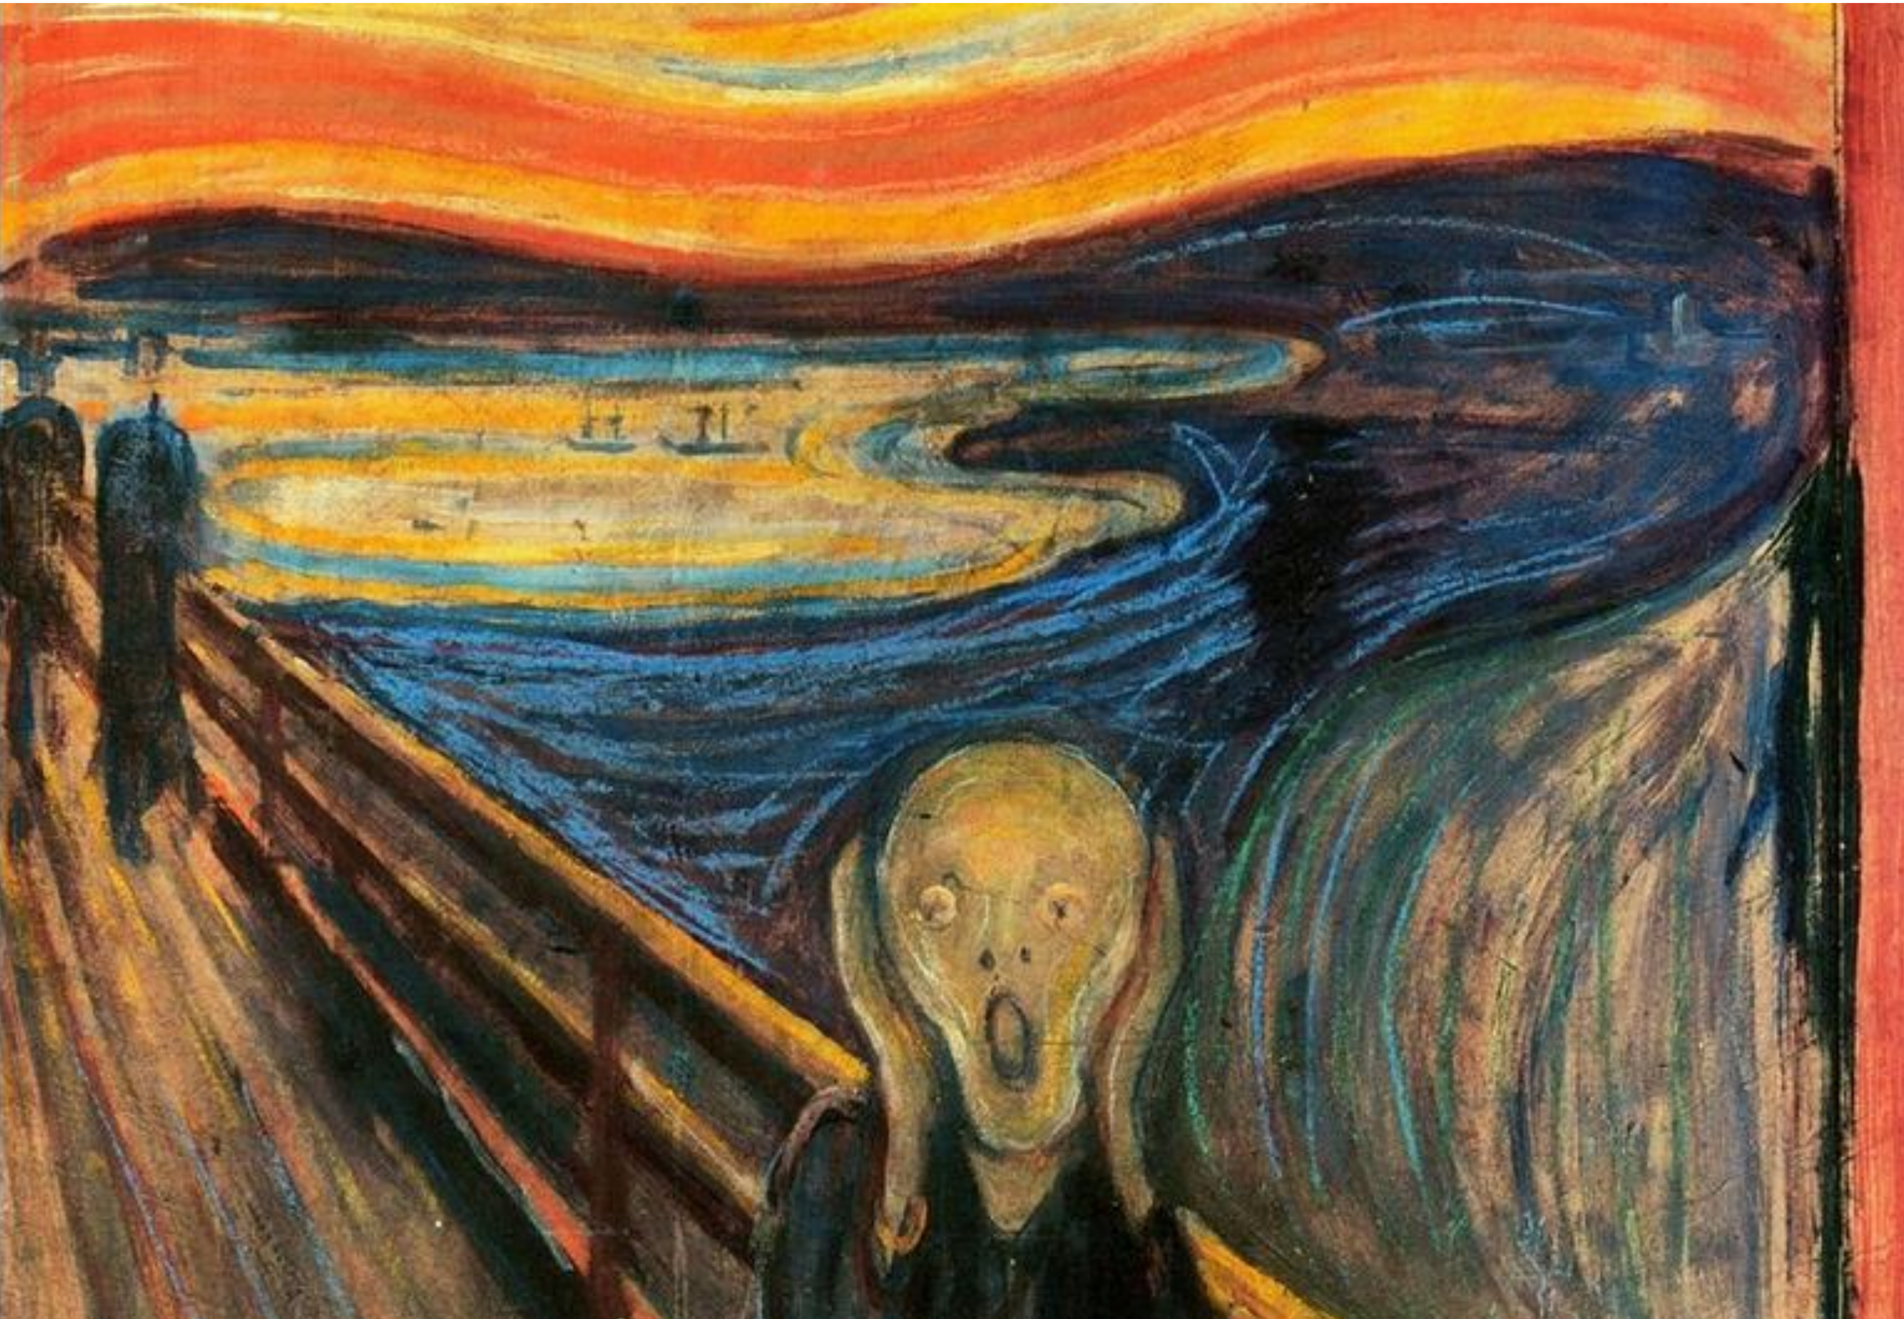

Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons

More than 1,000 experts and leading robotics researchers sign open letter warning of military artificial intelligence arms race

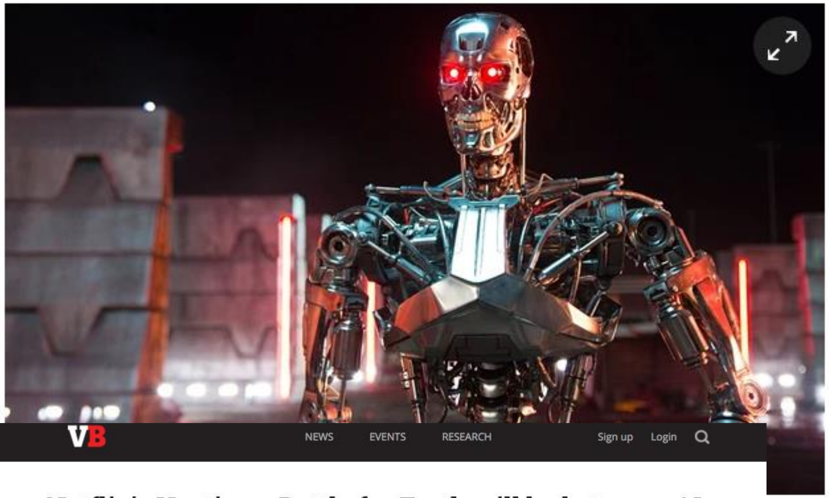

VB NEWS EVENTS RESEARCH Sign up Login

Netflix's Hastings: Battle for Earth will be between AI machines and genetically modified humans

CHRIS O'BRIEN JANUARY 18, 2016 3:41 AM
TAGS: AI, GENETICS, NETFLIX, REED HASTINGS

Image Credit: Flickr/epSos .de

Before Reed Hastings cofounded a little company called Netflix, which is now changing the way we watch TV, he was an artificial intelligence engineer.

AI has come a long way since Hastings got his masters from Stanford University in 1988. But he still follows developments in the field closely. And during a conversation on stage today at the DLD Conference in Munich, Germany, Hastings said he was far less worried about looming threats of an AI-triggered apocalypse than are many other observers, such as Tesla's Elon Musk.

"Some people worry about what happens when machine intelligence is too strong," Hastings said. "That's like worrying about our Mars colony and people

ry AI

January 19, 2016

Stephen Hawking says he believes **the key to saving humanity will be colonizing other planets.** But the renowned physicist, whose recent lecture will be broadcast next week, does not think that will happen soon.

BBC News

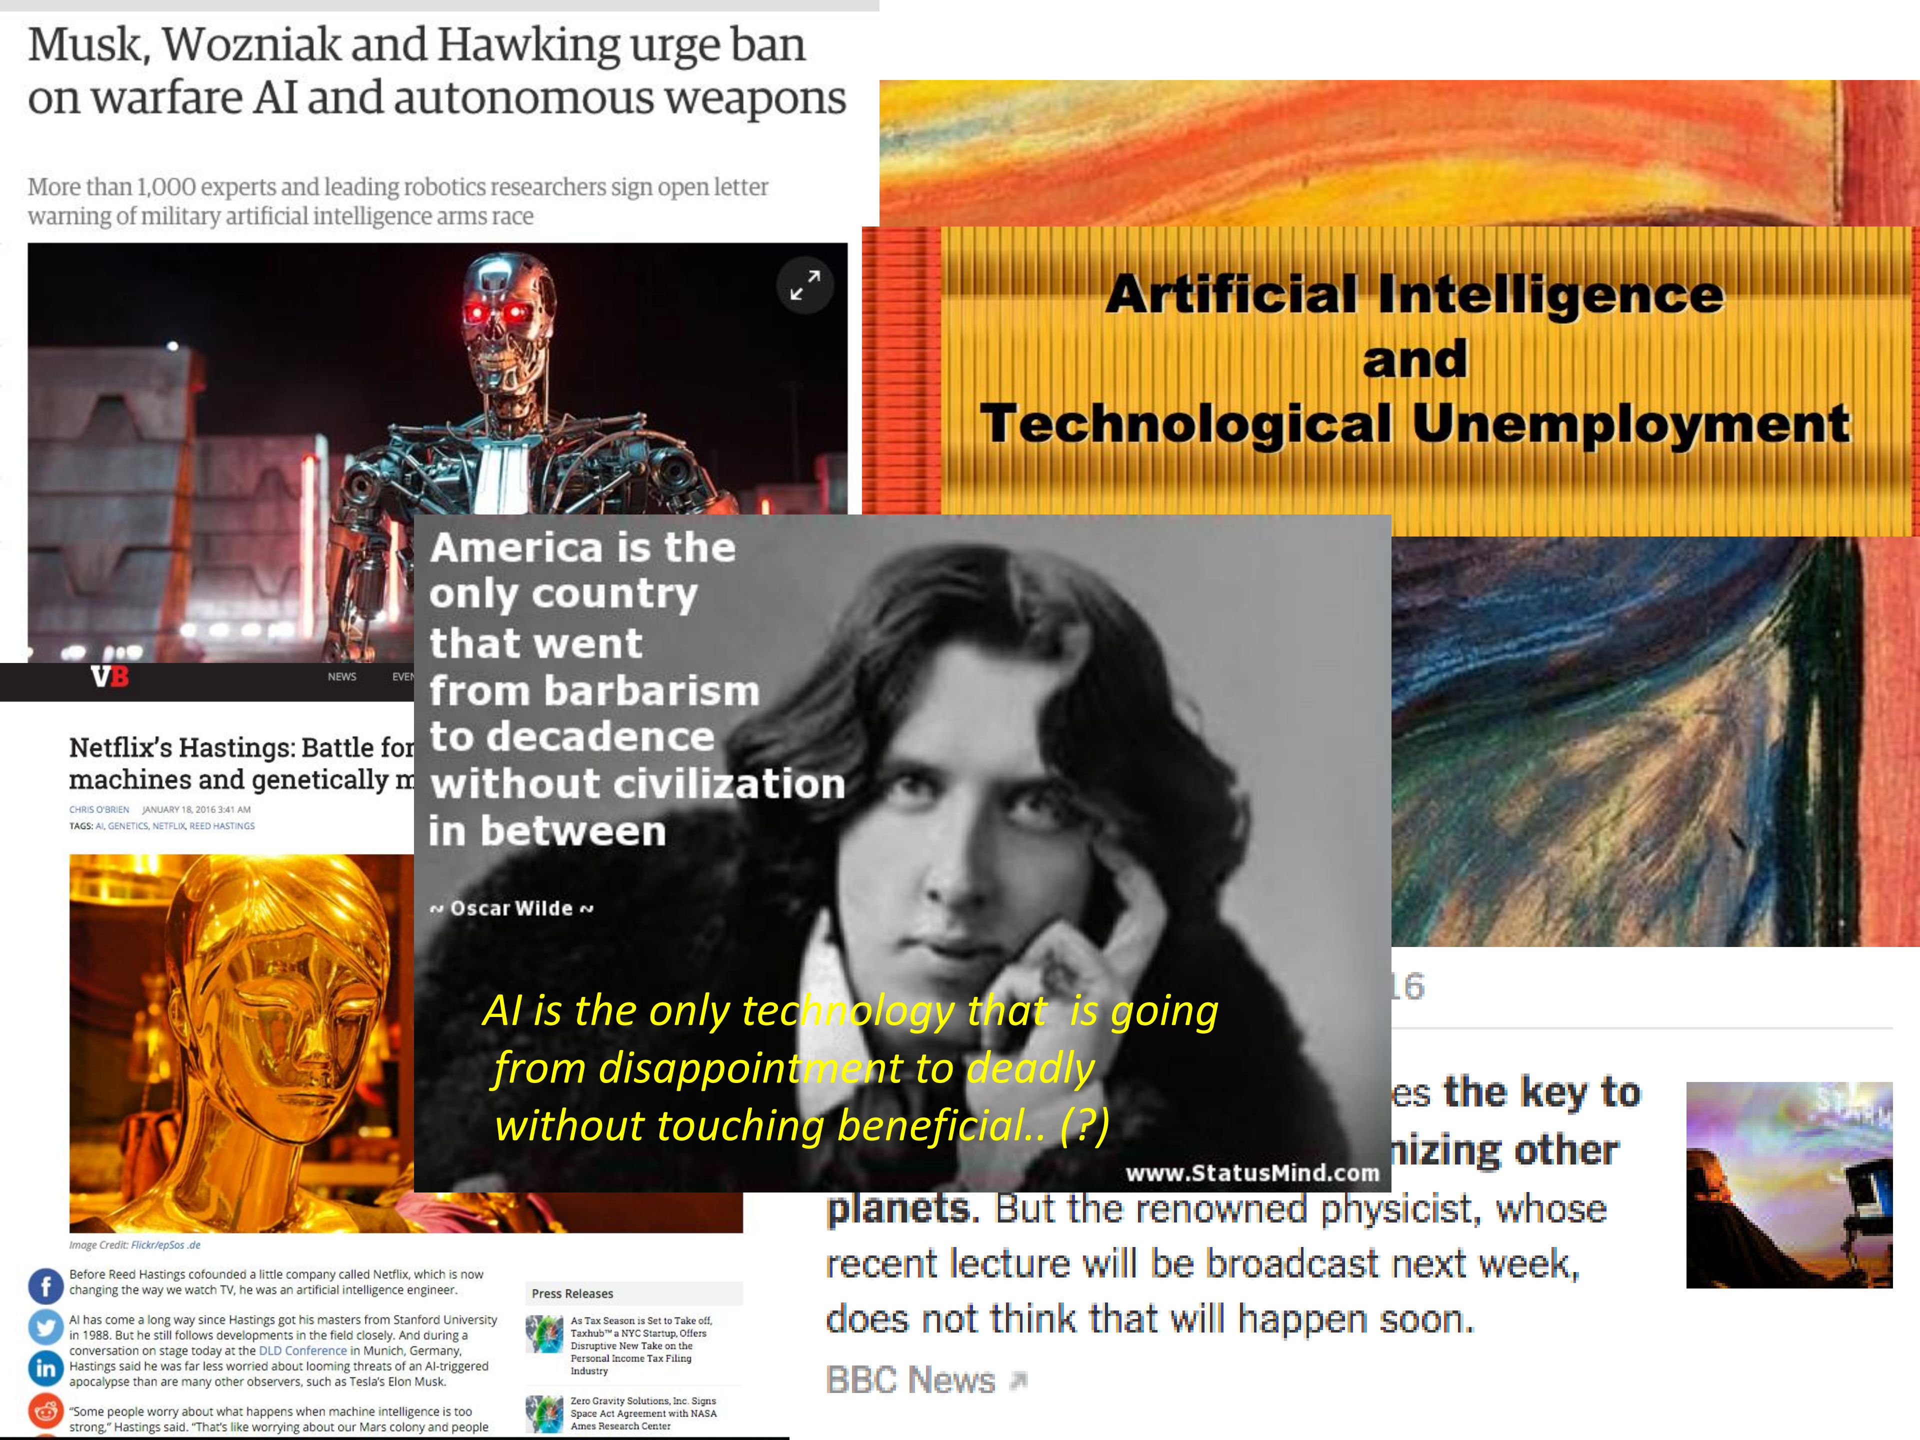

Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons
More than 1,000 experts and leading robotics researchers sign open letter warning of military artificial intelligence arms race
Artificial Intelligence
and
Technological Unemployment
America is the only country that went from barbarism to decadence without civilization in between
~ Oscar Wilde ~
AI is the only technology that is going from disappointment to deadly without touching beneficial.. (?)
www.StatusMind.com
VB
NEWS
EVEN
Netflix's Hastings: Battle for
machines and genetically m
CHRIS O'BRIEN JANUARY 18, 2016 3:41 AM
TAGS: AI, GENETICS, NETFLIX, REED HASTINGS
Image Credit: Flickr/epSos .de
Before Reed Hastings cofounded a little company called Netflix, which is now changing the way we watch TV, he was an artificial intelligence engineer.
AI has come a long way since Hastings got his masters from Stanford University in 1988. But he still follows developments in the field closely. And during a conversation on stage today at the DLD Conference in Munich, Germany, Hastings said he was far less worried about looming threats of an AI-triggered apocalypse than are many other observers, such as Tesla's Elon Musk.
"Some people worry about what happens when machine intelligence is too strong," Hastings said. "That's like worrying about our Mars colony and people
Press Releases
As Tax Season is Set to Take off, Taxhub™ a NYC Startup, Offers Disruptive New Take on the Personal Income Tax Filing Industry
Zero Gravity Solutions, Inc. Signs Space Act Agreement with NASA Ames Research Center
16
es the key to
nizing other
planets. But the renowned physicist, whose recent lecture will be broadcast next week, does not think that will happen soon.
BBC News

Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons

More than 1,000 experts and
warning of military artificial i

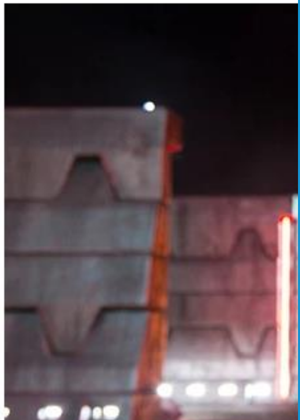

VB

**Netflix's Hastings: B
machines and genet**

CHRIS O'BRIEN JANUARY 18, 2016 3:41 AM
TAGS: AI, GENETICS, NETFLIX, REED HASTINGS

Image Credit: Flickr/epSos .de

Before Reed Hastings cofounded a little compa
changing the way we watch TV, he was an artifi

AI has come a long way since Hastings got his r
in 1988. But he still follows developments in th
conversation on stage today at the DLD Confer
Hastings said he was far less worried about loo
apocalypse than are many other observers, su

"Some people worry about what happens when machine intelligence is too
strong," Hastings said. "That's like worrying about our Mars colony and people

Zero Gravity Solutions, Inc. Signs
Space Act Agreement with NASA
Ames Research Center

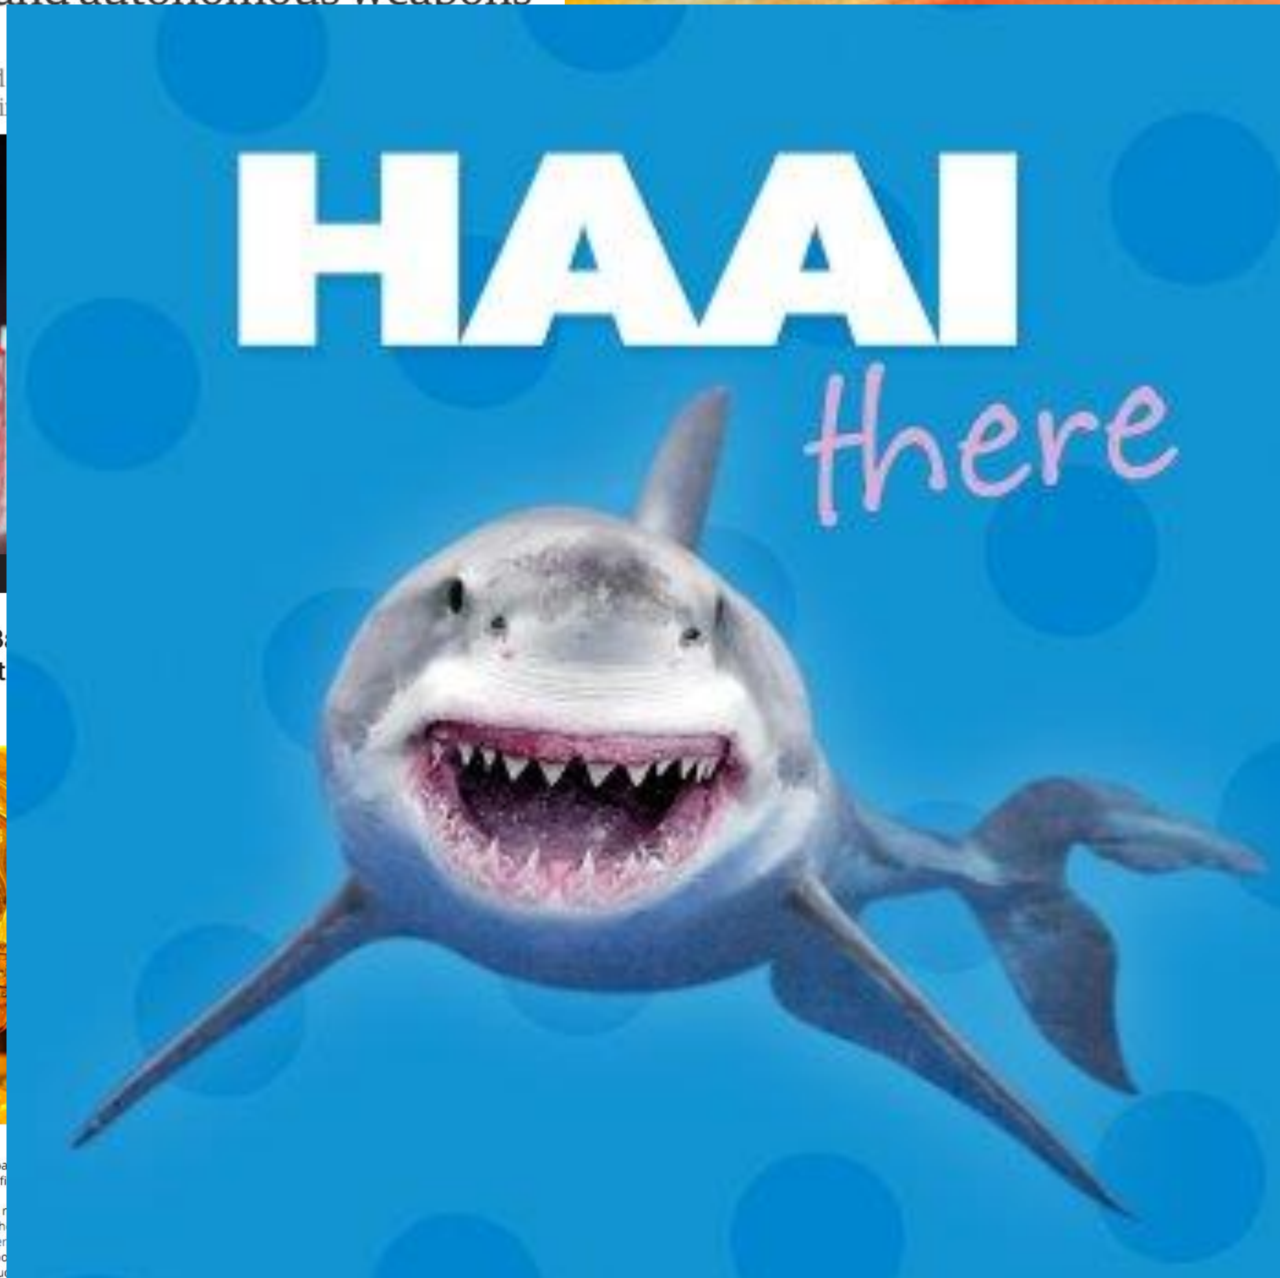

to
r

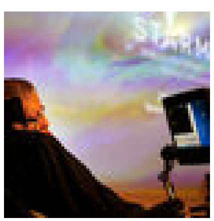

My Plan today: Talk to you about what we have been doing about HAAI

Our specific interest: Understand how the planning & decision-making aspects of AI agents change in Human-Machine cohabitation Scenarios

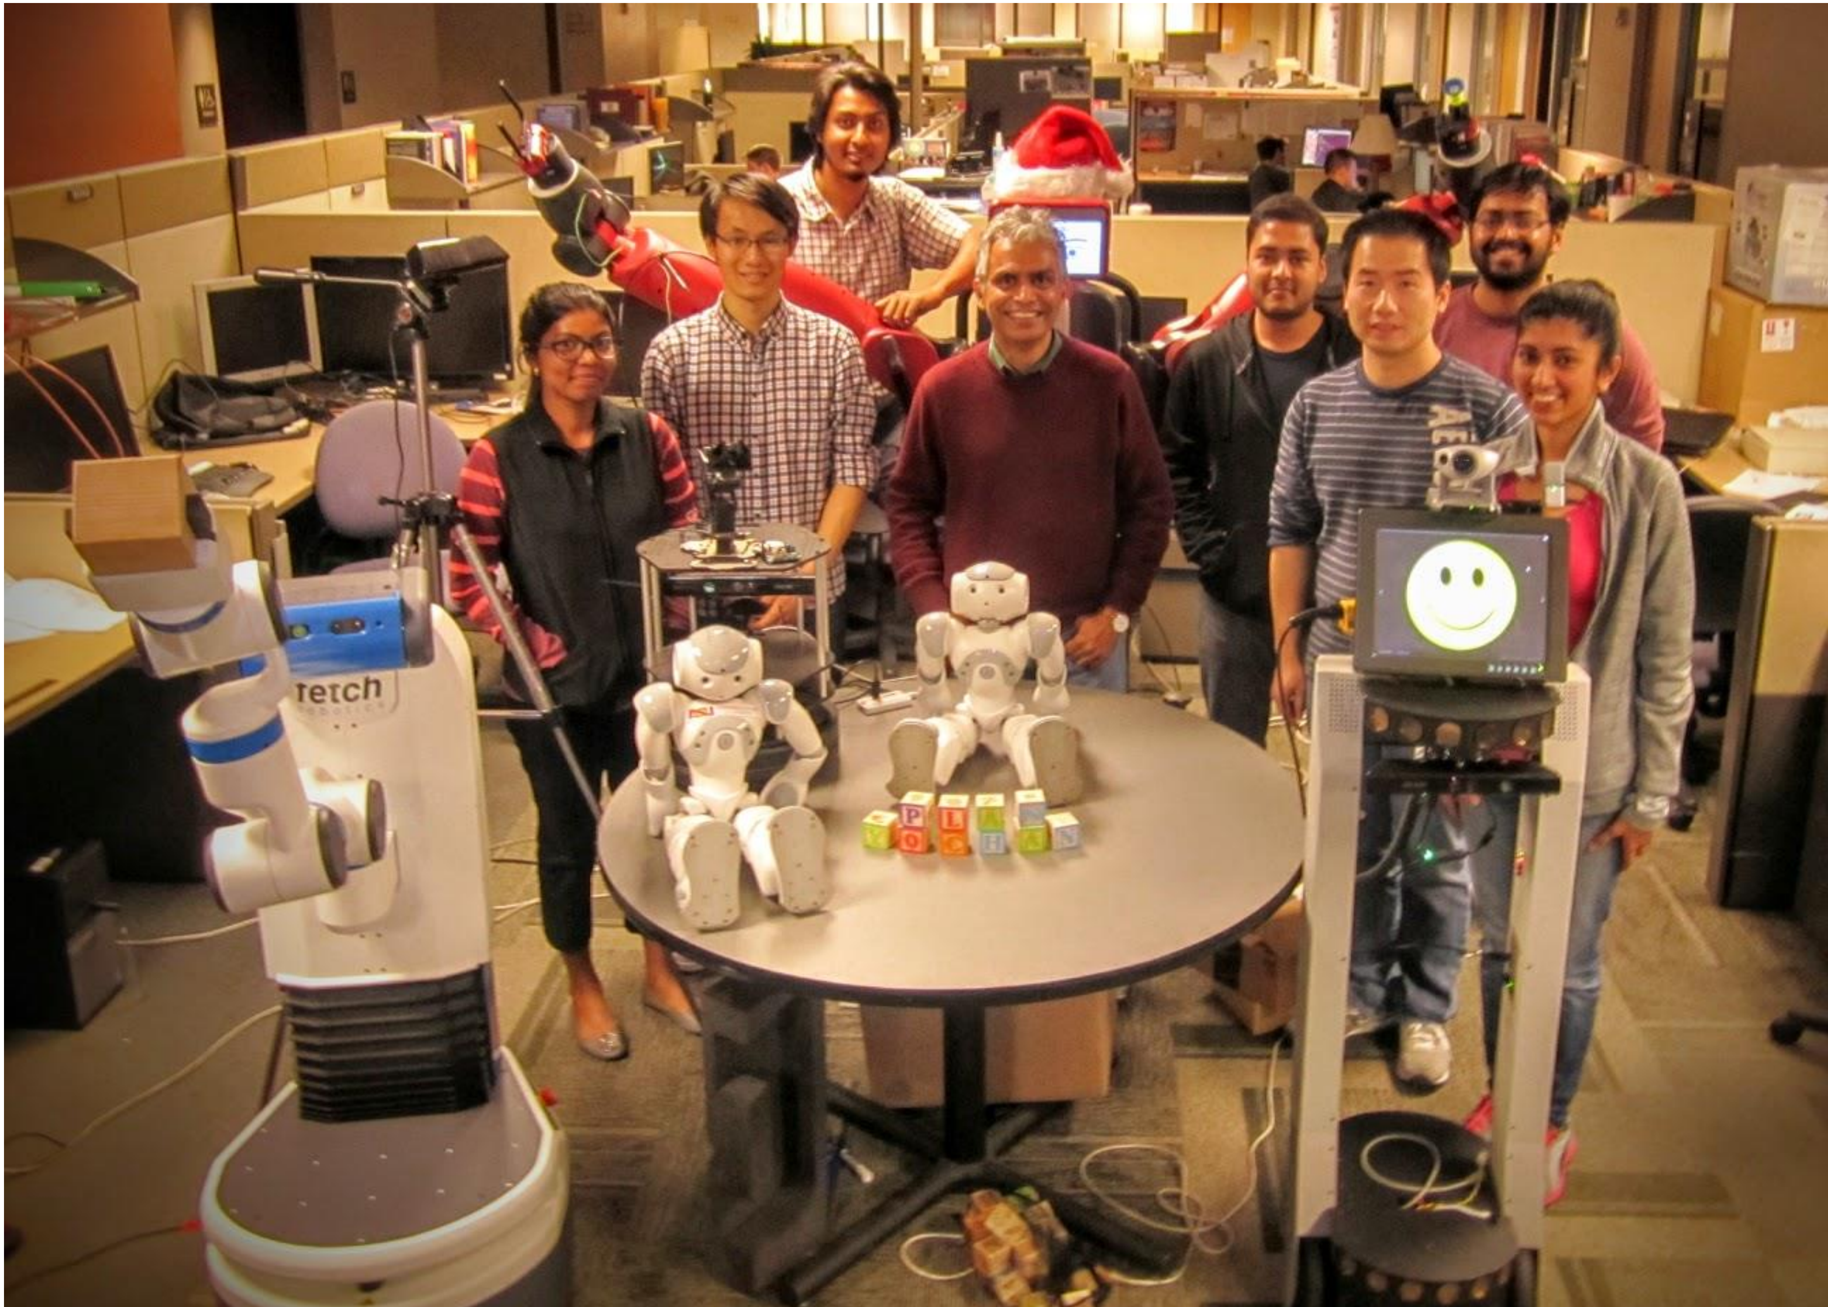
fetch

AI’s Curious Ambivalence to humans..

- Our systems seem happiest
  - either far away from humans
  - or in an adversarial stance with humans

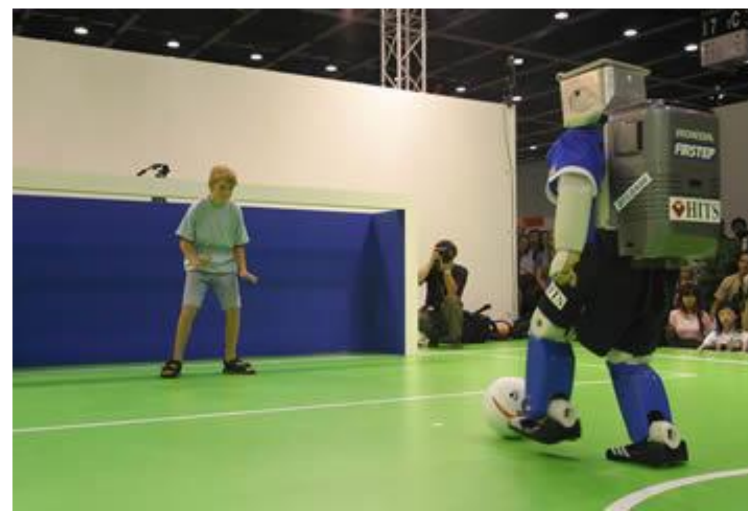

You want to help humanity, it is the people that you just can’t stand…

Planning: The Canonical View

Full
Problem
Specification

PLANNER

Fully Specified
Action Model

Fully Specified
Goals

Completely Known
(Initial) World State

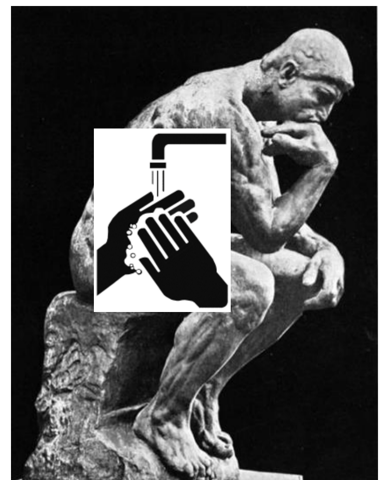

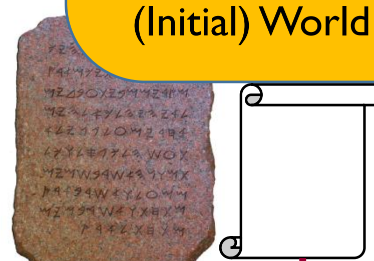

Plan (Handed off
for Execution)

Assumption:
→Complete Action Descriptions
→Fully Specified Preferences
→All objects in the world known up front
→One-shot planning
Allows planning to be a pure inference problem

☹ But humans in the loop can ruin a really a perfect day ☹

Human-in-the-Loop Planning

- In many scenarios, humans are part of the planning loop, because the planner:
  - Needs to plan to avoid them
    - Human-Aware Planning
  - Needs to provide decision support to humans
    - Because “planning” in some scenarios is too important to be left to automated planners
    - “Mixed-initiative Planning”; “Human-Centered Planning”; “Crowd-Sourced Planning”
  - (May need) help <u>*from*</u> humans
    - Mixed-initiative planning; “Symbiotic autonomy”
  - Needs to team with them
    - Human-robot teaming; Collaborative planning

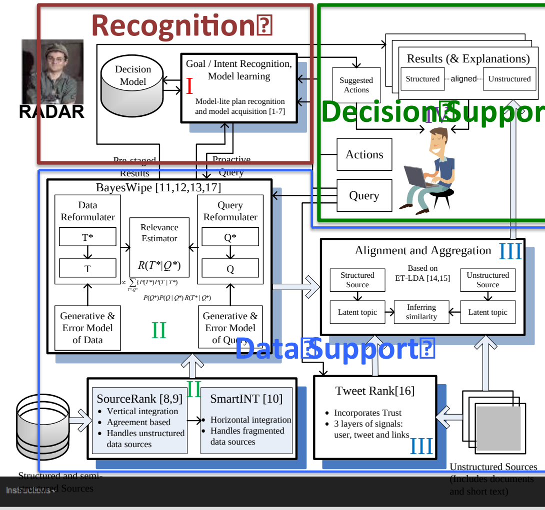

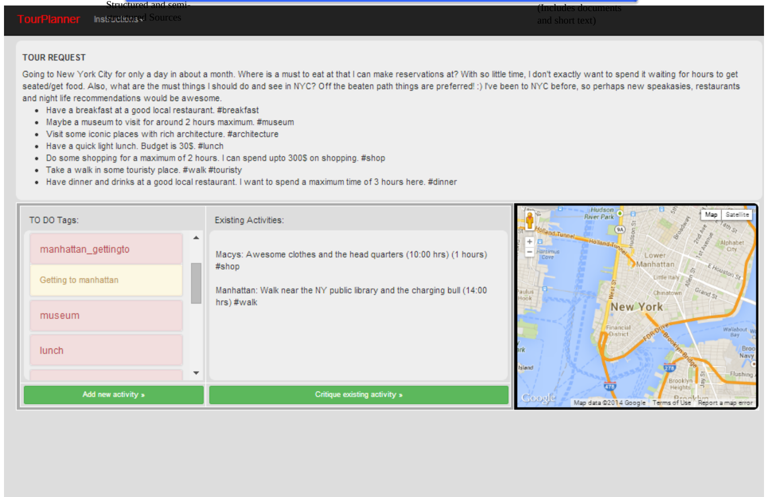

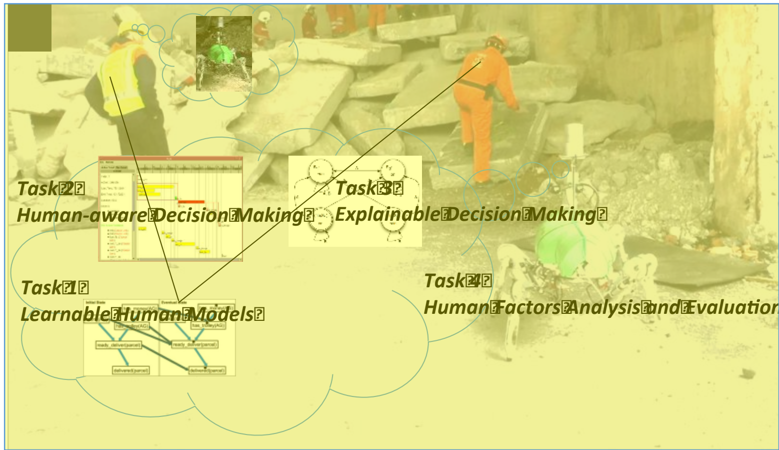

Planning: The Classical View

Full Problem Specification

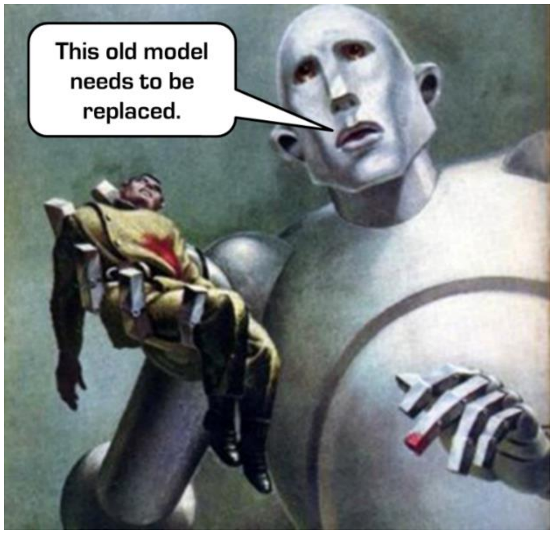

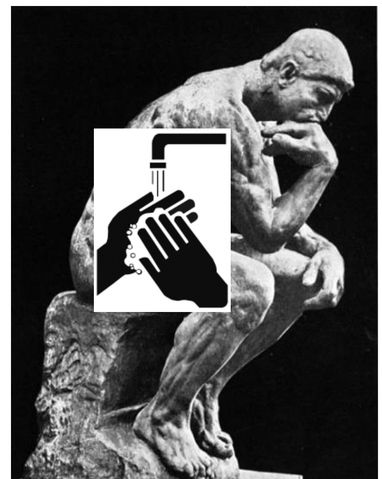

PLANNER

Fully Specified Action Model

Fully Specified Goals

Completely Known (Initial) World State

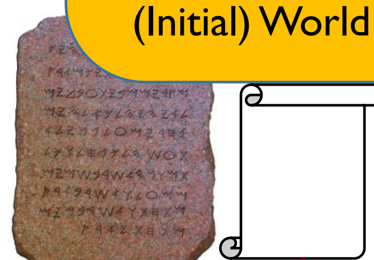

Plan (Handed off for Execution)

Violated Assumptions:
- →Complete Action Descriptions (Split knowledge)
- →Fully Specified Preferences (uncertain users)
- →Packaged planning problem (Plan Recognition)
- →One-shot planning (continual revision)

Planning is no longer a pure inference problem ☹

☹ But humans in the loop can ruin a really a perfect day ☹

Human-in-the-Loop Planning & Decision Support

AAAI 2015 Tutorial

rakaposhi.eas.asu.edu/hilp-tutorial

Subbarao Kambhampati
Arizona State University

Kartik Talamadupula
IBM T.J. Watson Research Center

AAAI-15 Austin, Texas USA
The First *Winter* AI Conference!

Funding from ONR, ARO and NSF gratefully acknowledged

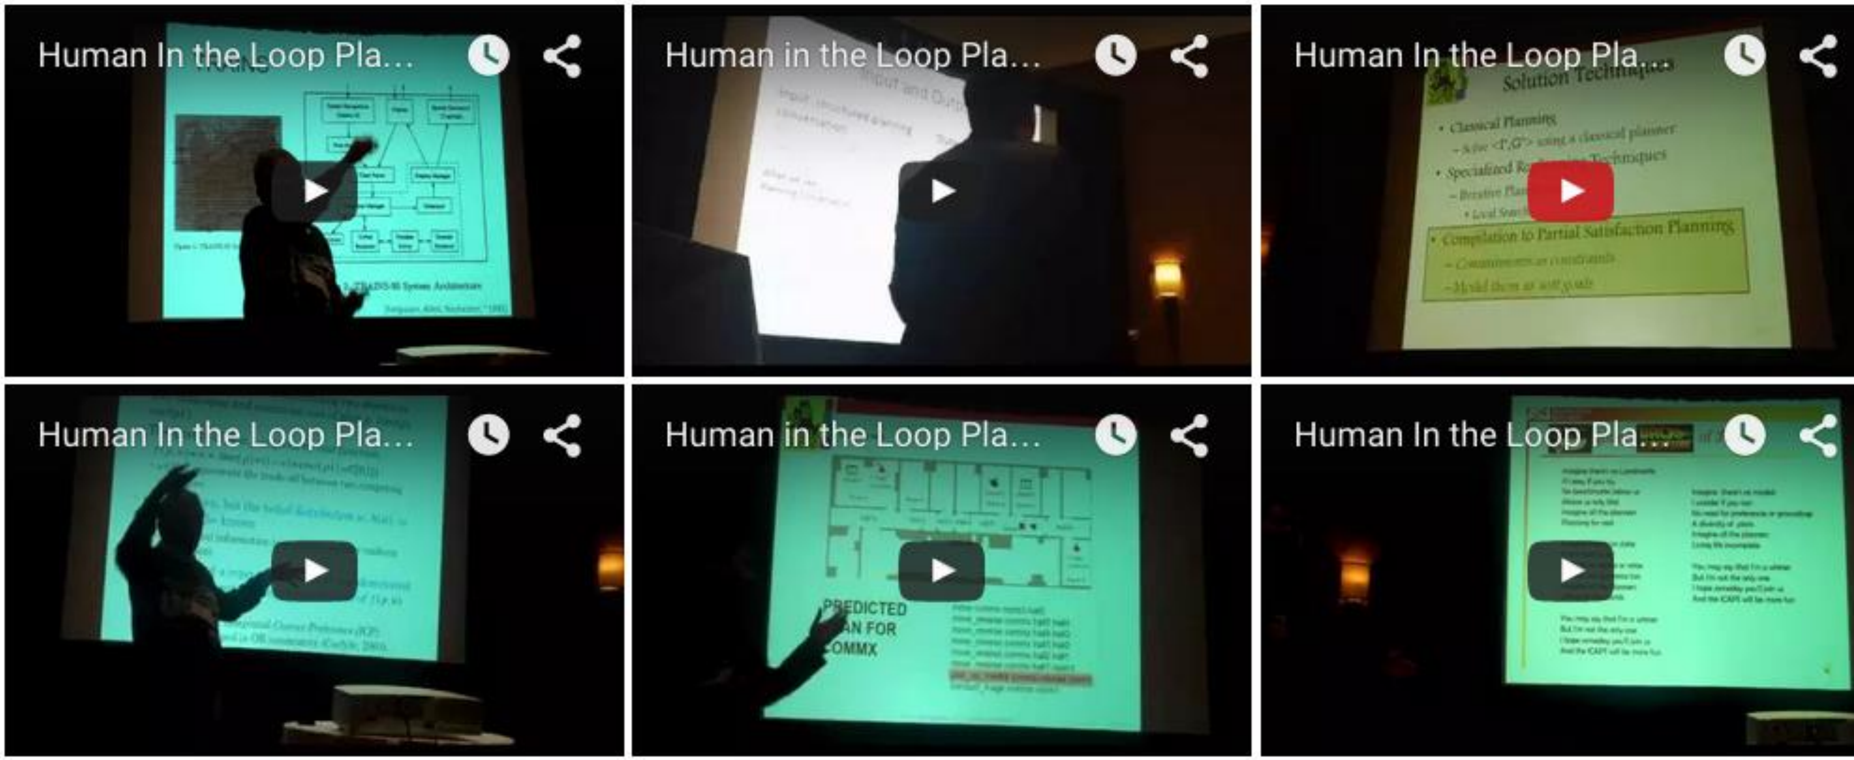

Materials

[Tutorial Slides (Final version, as given) [PDF]]

gratefully acknowledged

Dimensions of Variation in Human in the Loop Planning

- Cooperation Modality
  - Awareness, Interaction, Teaming
- Communication Modality
  - Stigmergic, Custom Interfaces, Speech/NLP
- What is Communicated
  - Goals, preferences, plan constraints, new goals
- Knowledge Level (Who knows what)
  - Incomplete knowledge about human's goals as well as capabilities

Challenges in Human-in-the-loop Planning

- Interpret what humans are doing based on incomplete human and domain models (Modeling)
  - Plan/goal/intent recognition
- Plan with incomplete domain models (Decision Making)
  - Robust planning/execution support with “lite” models
  - Proactive teaming support
- Explanations/Excuses (Interaction/Communication)
  - How should the human and robot coordinate
- Understand effective interactions between humans and machines (Evaluation)
  - Human factor study

Human-in-the-Loop Planning

- In many scenarios, humans are part of the planning loop, because the planner:
  - Needs to plan to avoid them
    - Human-Aware Planning
  - Needs to provide decision support to humans
    - Because “planning” in some scenarios is too important to be left to automated planners
    - “Mixed-initiative Planning”; “Human-Centered Planning”; “Crowd-Sourced Planning”
  - (May need) help *from* humans
    - Mixed-initiative planning; “Symbiotic autonomy”
  - Needs to team with them
    - Human-robot teaming; Collaborative planning

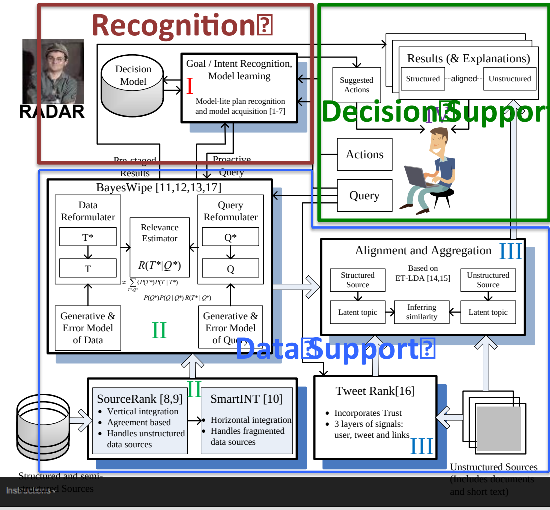

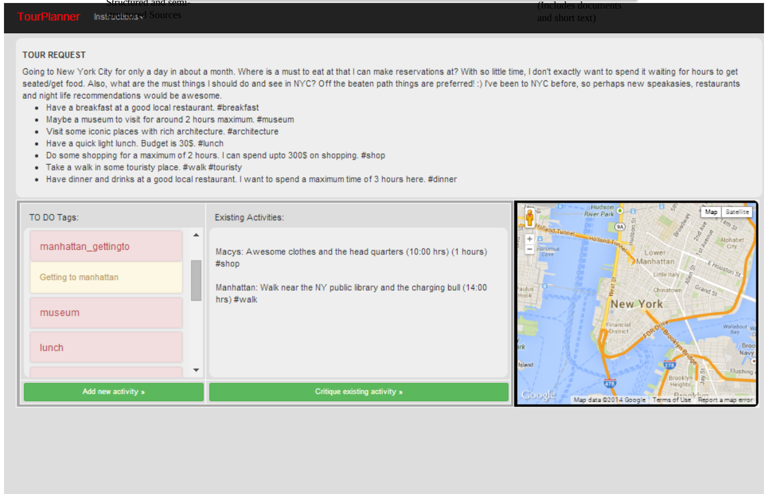

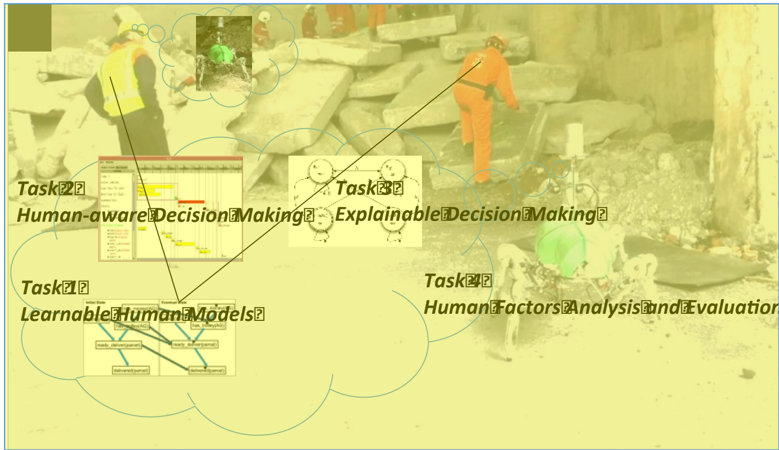

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

Manipulative (proximal) vs. Cognitive (remote) Teaming

- Much of the work in human-robot teaming has been focused on manipulation tasks where the human and the robot are in close proximity
  - Here the plans are mostly path planning/manipulator planning.
- Our focus has been on tasks that require cognitive (in addition to manipulative) decisions—as is typically the case with remote human-robot collaboration in urban search and rescue scenarios.

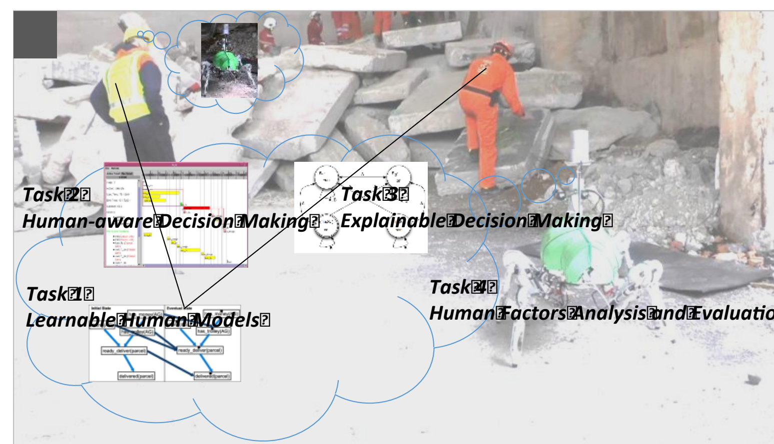

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

How do we get the Planning Models? (e.g. of the human in the loop)

- Typically multi-agent planning methods assume all agents use similar models
  - E.g. All agents with STRIPS action models
- Unreasonable to expect similar sorts of action models for the robot and the human..
  - Human models (from the Robot's point of view) are likely to be highly incomplete (as, of course, Robot's model from the human point of view)
- So how do we represent (and handle) incomplete models of human capabilities?

Challenges in acquiring Human Models

- The temptation is to go with existing action models & introduce incompleteness
  - Atomic: MDP/POMDP
  - Factored: STRIPS, RDDL, HTN etc
    - Example work by Garland&Lesh(2002)
- While they are fine if someone hand-specifies them, they are much harder to learn, given the kinds of information that is likely to be available.
  - Significant incompleteness in observations
    - Sensor occlusion, noisy observations,
      - [Zhuo & Kambhampati, IJCAI 2013]
  - There may be significant gaps between observations

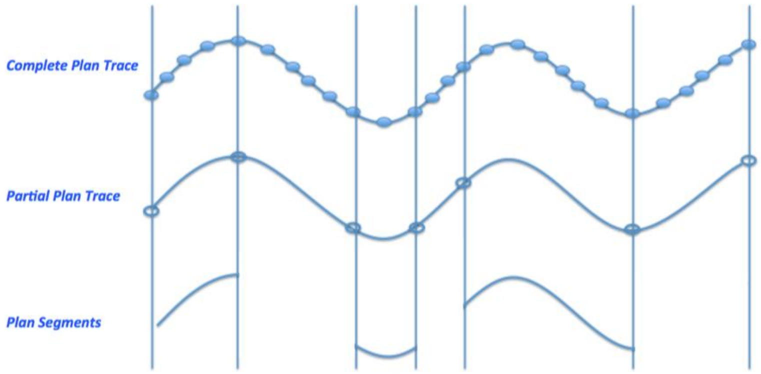

Spectrum of Domain Models

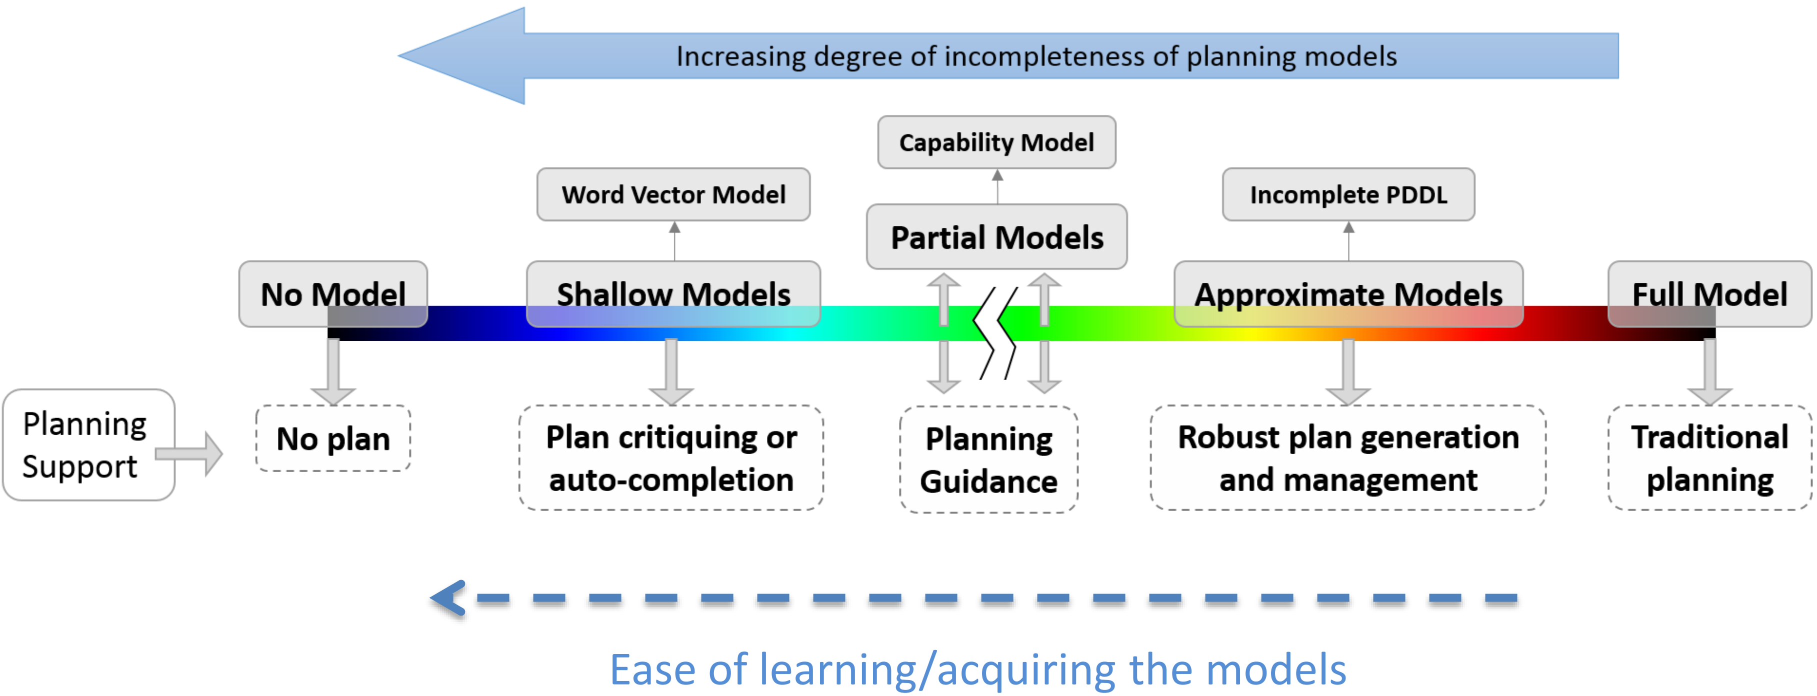

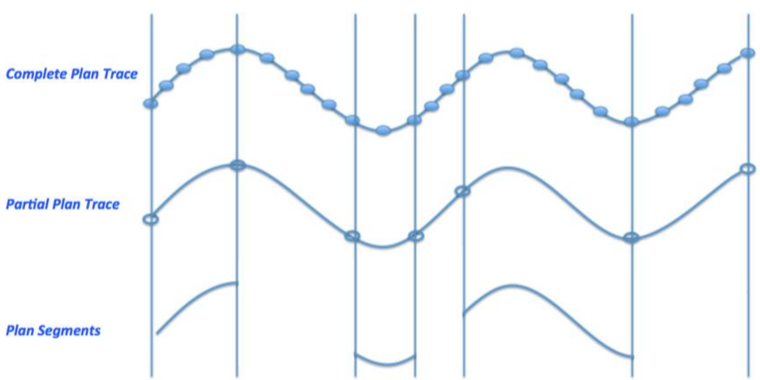

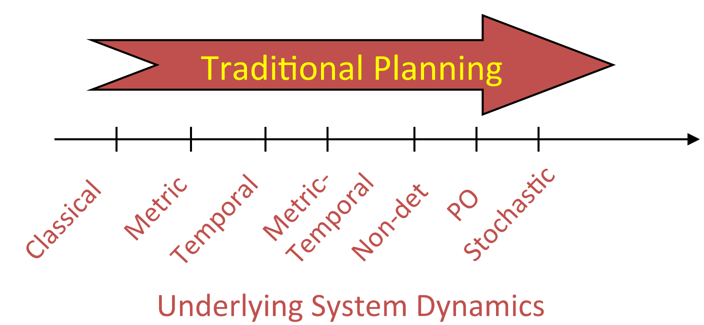

Partial PDDL Domain Models

Each action **a** is associated with *possible* precond and effects (in addition to the normal precond/eff):

- **PreP(a)** *[p]:* set of propositions that **a** *might* depend on during execution
- **AddP(a)** *[p]:* : set of propositions that **a** *might* add after execution
- **DelP(a)** *[p]:* : set of propositions that **a** *might* delete after execution

Example: An action **a** that is known to depend on **p1**, add p4 and delete **p3**. In addition, it might have **p3** as its precondition, might add **p2** and might delete **p1** after execution.

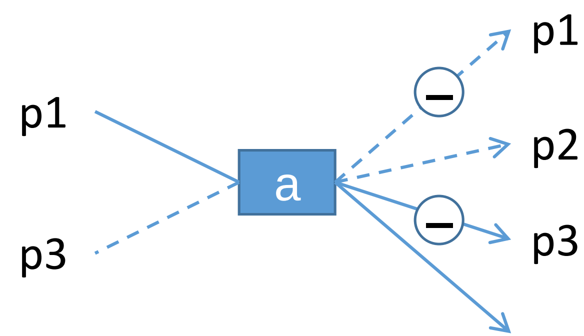

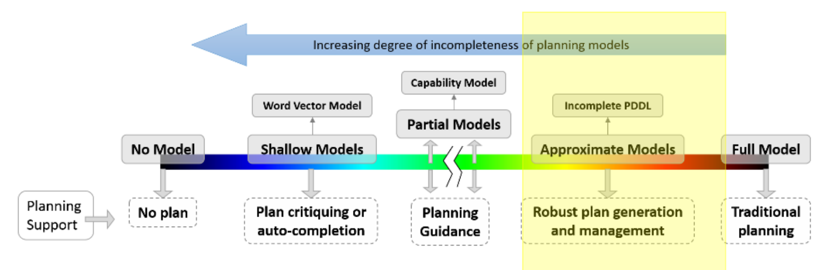

There are known knowns; there are things we know that we know. There are known unknowns; that is to say, there are things that we now know we don't know. But there are also unknown unknowns; there are things we do not know we don't know.

Solution Concept: Robust Plans

- Solution concept:
  - Robust plan
  - Plan is highly robust if executable in large number of most-likely candidate models
- Robustness measure
  - Set of candidate domain models **S** (consistent with the given deterministic partial domain model **D**)
  - A complete but unknown domain model **D***
    - Can be any model in **S**

$$R(\pi) = \frac{|\Pi|}{2^K}$$

$|\Pi|$ Number of candidate models with which the plan succeeds

$$K = \sum_a \mathrm{PreP}(a) + \mathrm{AddP}(a) + \mathrm{DelP}(a)$$

Easily generalized to consider model likelihood

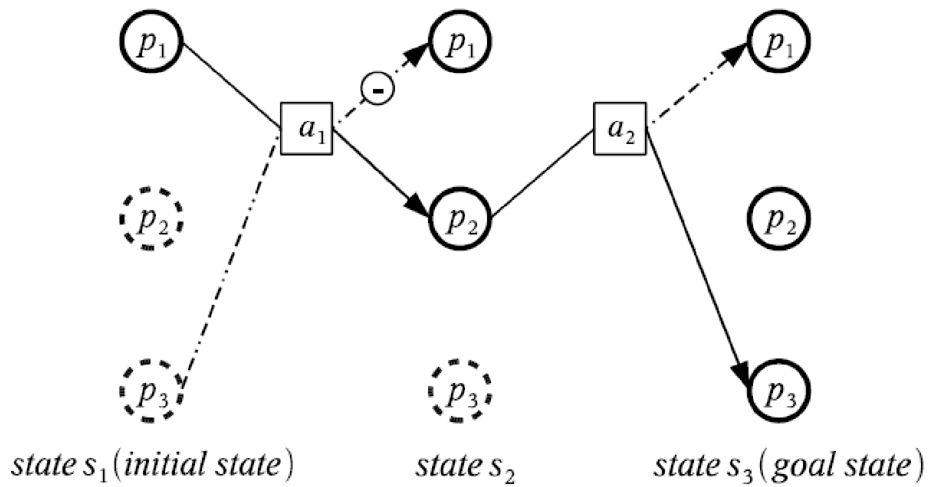

Candidate models of plan	1	2	3	4	5	6	7	8
a_1 relies on p_3	yes	yes	yes	yes	no	no	no	no
a_1 deletes p_1	yes	yes	no	no	yes	yes	no	no
a_2 adds p_2	yes	no	yes	no	yes	no	yes	no
Plan status	fail	fail	fail	fail	succeed	fail	succeed	succeed

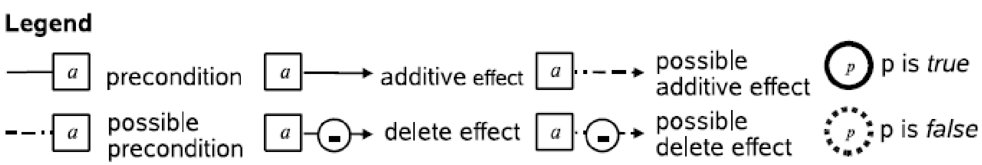

Robustness value: 3/8

Generating Robust Plans

- Compilation approach: Compile into a *(Probabilistic) Conformant Planning* problem
  - One “unobservable” variable per each possible effect/precondition
    - Significant initial state uncertainty
  - Can adapt a probabilistic conformant planner such as POND [JAIR, 2006; AIJ 2008]
- Direct approach: Bias a planner’s search towards more robust plans
  - Heuristically assess the robustness of partial plans
    - Need to use the (approximate) robustness assessment procedures
    - A novel extension to relaxed planning heuristics to take robustness into account

D. Bryce et al. / Artificial Intelligence 172 (2008) 685–715

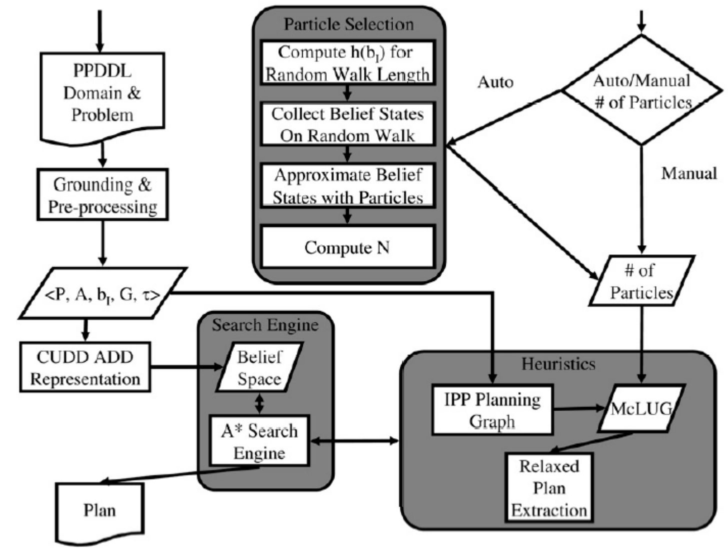

Fig. 6. *POND* architecture.

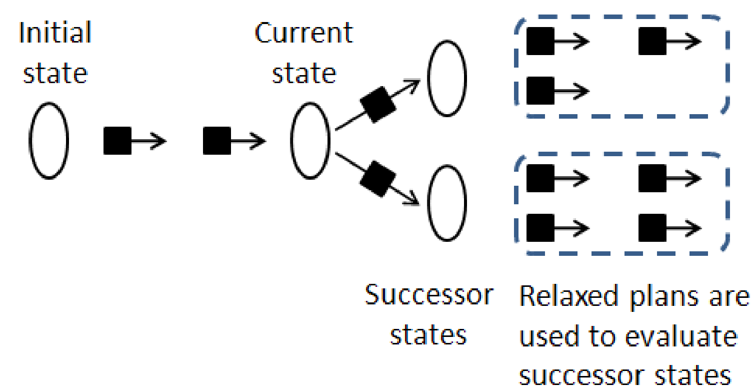

[Nguyen et al; NIPS 2013; Nguyen & Kambhampati, ICAPS 2014]

Synthesizing Robust Plans: A Compilation

Incomplete model
Complete world state

$\rightarrow$

Complete model
Belief state

(Conformant Probabilistic Planning)

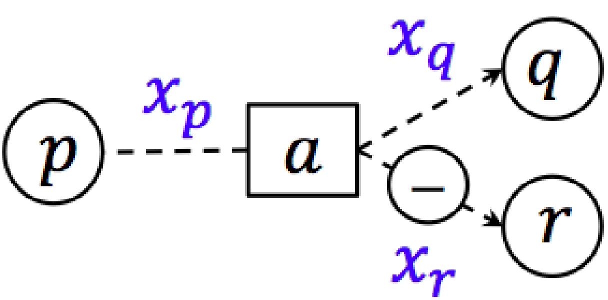

x_p (0.5) x_q (0.7) x_r (0.2)

Resulting action $\boldsymbol{a}'$ with eight conditional effects.

Cond: $x_p \wedge p \wedge x_q \wedge x_r$ **Eff:** $q \wedge \neg r$

[NIPS 2013]

Synthesizing Robust Plans: A Heuristic Search

❖ **Anytime approach**

1. Initialize: $\delta = 0$
2. **Repeat**
   - ❖ Find plan π s.t. $R(\pi) > \delta$
   - ❖ If plan found: $\delta = R(\pi)$

   Until time bound reaches
3. Return π and $R(\pi)$ if plan found

$h(s, \delta)$: how far it is approximately from s to a goal state so that the resulting plan has approximate robustness > δ.

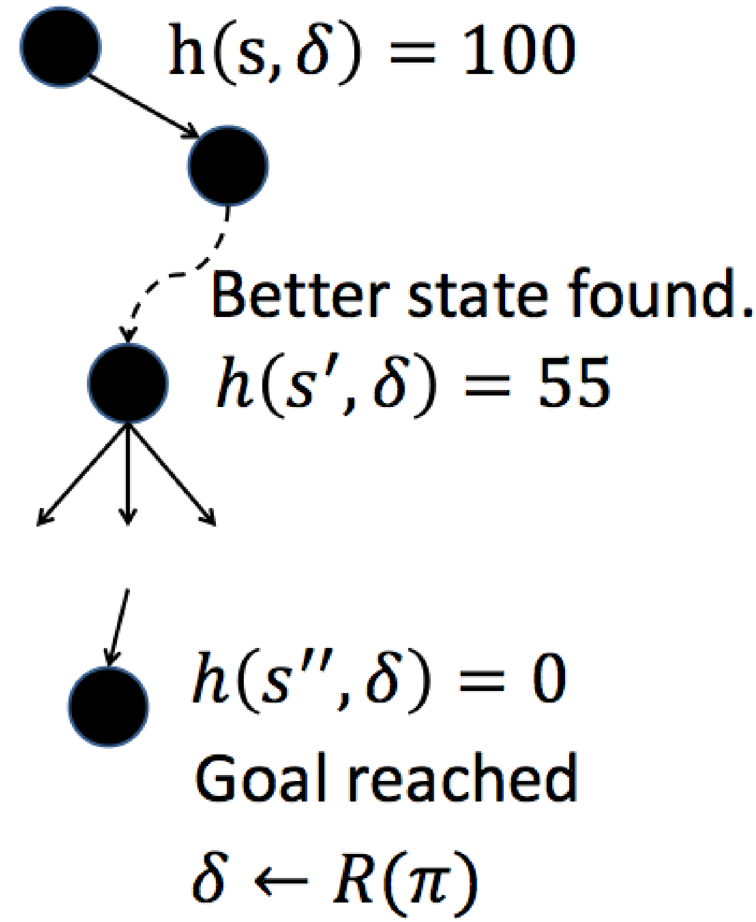

[ICAPS 2014]

- **Approximate plan robustness**
  - **Lower bound**

$$l(\Sigma) = \prod_{c \in \Sigma} \Pr(c) \; \le \; WMC(\Sigma)$$

$$I \;\; \bullet \xrightarrow[\pi_k]{\Sigma_{\pi_k}} \bullet \xrightarrow[\text{Relaxed plan } \tilde{\pi}]{\Sigma_{\tilde{\pi}}} \bullet \qquad l\left(\Sigma_{\pi_k} \wedge \Sigma_{\tilde{\pi}}\right) > \delta$$

  - **Upper bound**: divide Σ into independent sets Σ^{i}

$$u(\Sigma) = \prod_{\Sigma^i} \min_{c \in \Sigma^i} \Pr(c) \ge \; WMC(\Sigma)$$

$$I \;\; \bullet \xrightarrow[\pi_k]{\Sigma_{\pi_k}} \bullet$$

If $u\left(\Sigma_{\pi_k}\right) > \delta$

then compute $WMC(\pi_k)$

Capability Model

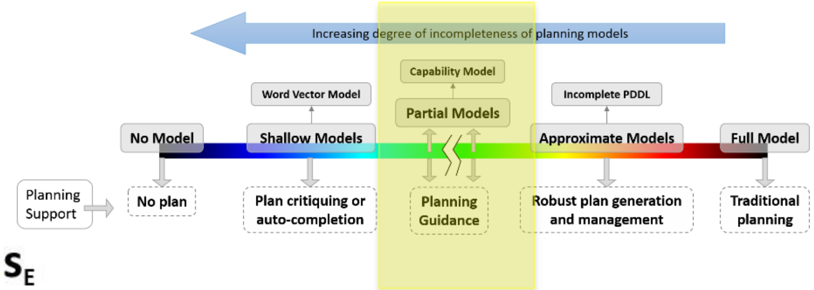

A capability:

$$P(\dot{X}_\phi = s_E \,|\, X_\phi = s_I) \longleftrightarrow S_I \Rightarrow S_E$$

A conditional probability
(specified by a partial initial and eventual state)

T-gap capability model

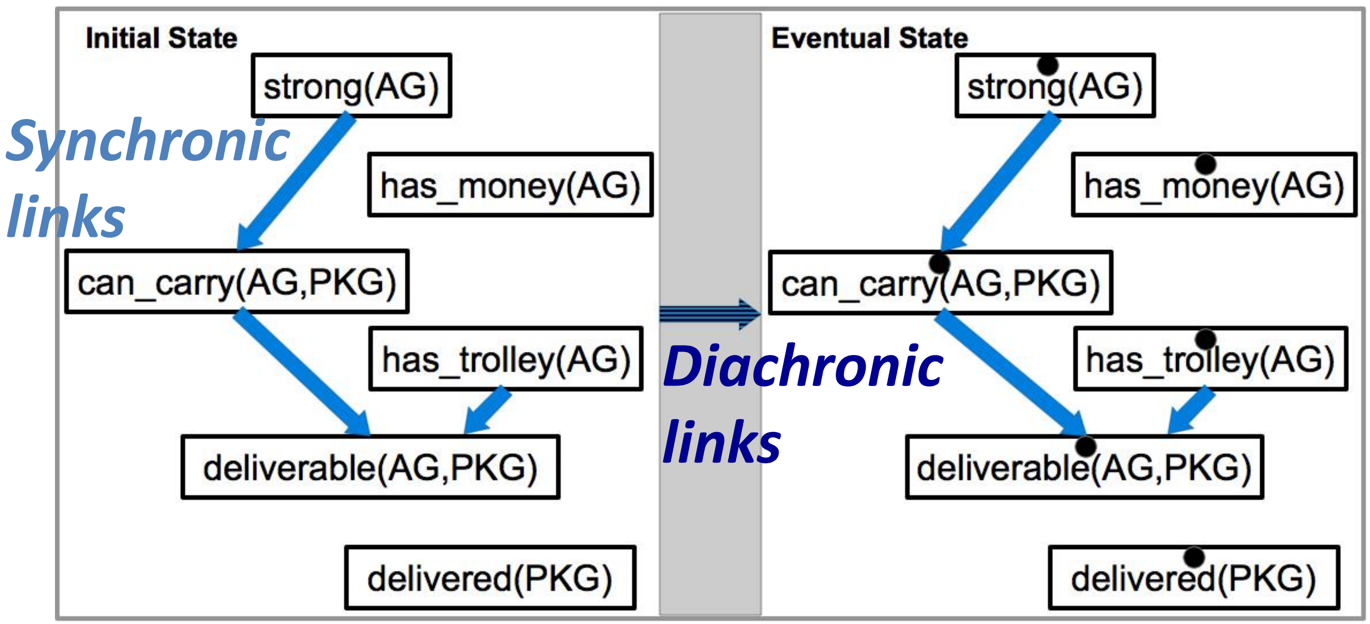

[AAMAS 2015]

(Generalization of 2-TBN model used in RDDL)
(Imperfect analogy to) HTN Models. A capability can be thought of as an abstract task

Capability Models

We start with the "default assumption" that domain models are incomplete

- ***DEFINITION (CAPABILITY)*** **– Given an agent, a capability is a mapping $\boxed{S_\phi} \times \boxed{S_\phi} \rightarrow [0, 1]$, which is an assertion about the probability of the existence of a plan in fewer than or equal to *T* atomic state changes that can connect the two partial states.**

Partial states

Bound on the gaps between observations

->: denote an atomic state change

{has_water(AG), has_coffee_beans(AG)}
-> {has_boilling_water(AG), has_coffee_beans(AG)}
-> {has_boilling_water(AG), has_ground_coffee_beans(AG)}
-> {has_coffee(AG)}

When T = 2 {
has_water(AG) => has_ground_coffee_beans(AG)
has_boiling_water(AG) => has_coffee(AG)…

When T = 3 {
… (including all capabilities when T = 2)
has_water(AG) => has_coffee(AG)

Parameter Learning

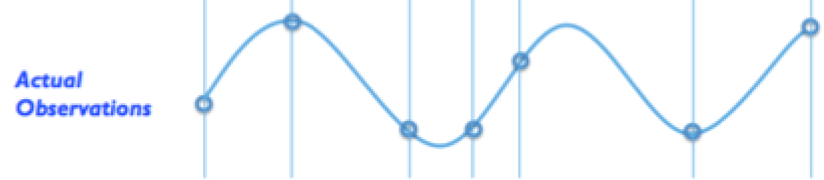

We assume that the maximum number of missing state observations between any two observations in the partial plan trace is upper bounded by T

DEFINITION (T-GAP PARTIAL PLAN TRACE)*. A T-gap partial plan trace is a partial plan trace in which all *$k_{[1, 2\ldots]}$ <= T

$$\mathcal{T} = \langle s_i, s_{i+k_1}, s_{i+k_2}, \cdots \rangle$$

Learning samples

Apply Bayesian learning (assuming beta distributions):

$$\rho(f_{ij}|D) = beta(f_{ij}; a_{ij} + s_{ij}, b_{ij} + t_{ij})$$

Planning with Capability Models

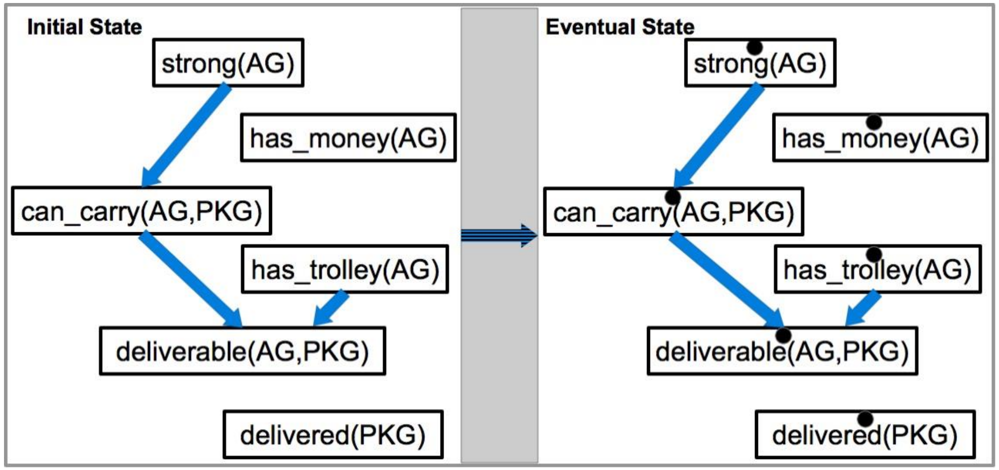

T-gap capability model

- Any planning state is a set of complete states: a belief state

  {(complete state 1), (complete state 2)…}

- Select a capability to apply: **s_I => s_E =** $P(\dot{X}_\phi = s_E \mid X_\phi = s_I)$

- For each s* in the belief state,
  - Applicable $s_I \sqsubseteq s^*$

    Success: compute a set of resulting states s, $s_E \sqsubseteq s,$

    $$P(s) = \frac{P(s^* \Rightarrow s)}{P(s^* \Rightarrow s_E)} = \frac{P(\dot{X}_\phi = s | X_\phi = s^*)}{P(\dot{X}_\phi = s_E | X_\phi = s^*)}$$

    Failure: no change

    $$\sum_{s \in S} P(s) = 1$$

  - Inapplicable – no change to s*

Action Vector Models

- View observed action sequences as “sentences” in a language whose “words” are the actions
- Apply skip-gram models to these sequences and embed the action “words” in a higher dimensional space
  - The proximity of the action words in that space is seen as their “affinity”
- Use the action affinities as a way to drive planning and plan recognition

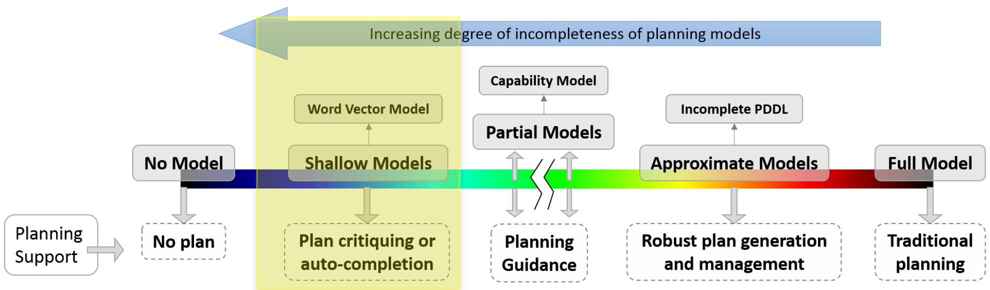

Problem Formulation

- The recognition problem defined by: (L, O, A)
  - L: a plan library, e.g.,

plan 1: *pick-up-B stack-B-A pick-up-D stack-D-C*
plan 2: *unstack-B-A put-down-B unstack-D-C put-down-D*
plan 3: *pick-up-B stack-B-A pick-up-C stack-C-B pick-up-D stack-D-C*

  - O: a sequence of observations, e.g.,

pick-up-B NULL unstack-D-C put-down-D NULL stack-C-B NULL NULL

  - A: A set of actions
- Task: find a plan to best explain O:

pick-up-B stack-B-A unstack-D-C put-down-D pick-up-C stack- C-B pick-up-D stack-D-C

Note that:

- without initial states/goals/intermediate states in L
- |p| = |O|
- p is not necessarily in L

Learn vectors of actions

$$\overbrace{a_{11}\ a_{12}\ \ldots\ a_{1n}}^{\text{Plan 1}}\quad \overbrace{a_{21}\ a_{22}\ \ldots\ a_{2n}}^{\text{Plan 2}}\quad \ldots\quad \overbrace{a_{T1}\ a_{T2}\ \ldots\ a_{Tn}}^{\text{Plan T}}$$

Learn vectors w_i for a_i in A by optimizing

$$\frac{1}{T}\sum_{t=1}^{T}\sum_{-c\leq j\leq c,j\neq 0}\log \underline{p(w_{t+j}|w_t)}$$

- T = |L|
- c is the window size of action context

The basic probability defined by hierarchical softmax, [cf. Mikolov et al. NIPS-13]

Action Vector Models can be used to Recognize Plans

With the learnt vectors w_i, we can predict the target plan (as the most consistent with the affinities). We use an EM procedure to speedup the prediction.

$$\mathcal{F}(\tilde{p}) = \sum_{k=1}^{M} \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{k+j} | w_k)$$

- M = |the target plan|

The target plan to be recognized

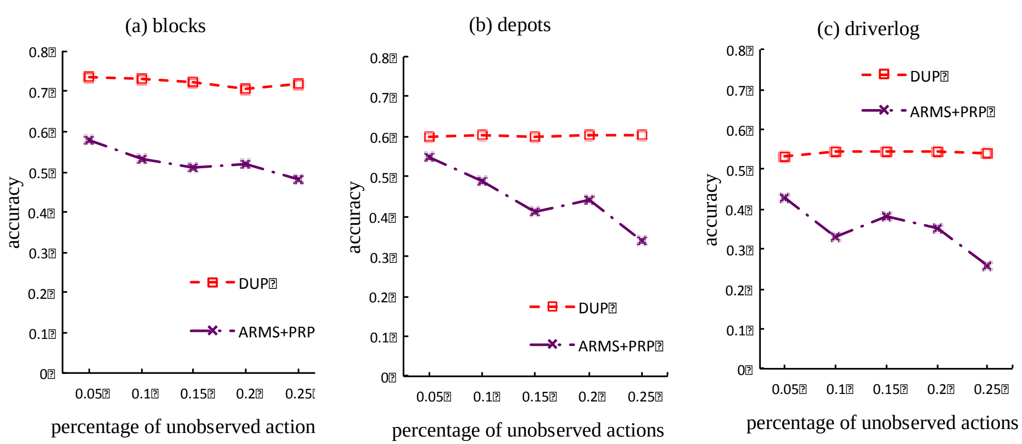

Algorithm 1 Framework of our DUP algorithm

Input: plan library $\mathcal{L}$, observed actions $\mathcal{O}$

Output: plan $\tilde{p}$

1: learn vector representation of actions
2: initialize $\Gamma_{o,k}$ with $1/M$ for all $o \in \bar{A}$, when k is an unobserved action index
3: **while** the maximal number of repetitions is not reached **do**
4: sample unobserved actions in $\mathcal{O}$ based on Γ
5: update Γ based on Equation (6)
6: project Γ to [0,1]
7: **end while**
8: select actions for unobserved actions with the largest weights in Γ
9: **return** $\tilde{p}$

[AAMAS16]

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

How to plan to be useful?

- Depends on the modality of interaction between the humans and the robot
  - Are they in an explicit team vs. cohabiting the same environment?
  - Are they communicating or is it stigmergic collaboration?
- Our early work focused on issues in explicit teaming and full communication

Planning for Human-Robot Teaming

~~Full~~ Problem Specification

Problem Updates
[TIST10]

Assimilate Sensor Information

PLANNER
Sapa Replan

~~Fully Specified~~ Action Model

~~Fully Specified~~ Goals

Goal Manager

Open World Goals

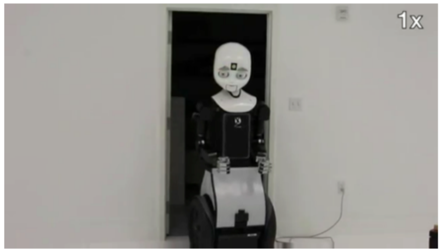

- When to start sensing?
  - Indicator to start sensing
- What to look for?
  - Object type
  - Object properties
- When to stop sensing?
  - When does the planner know the world is closed?
- Why should the robot sense?
  - Does the object fulfill a goal?
  - What is the reward? Is it a bonus?

Talamadupula, Benton et al., ACM TIST 2010]

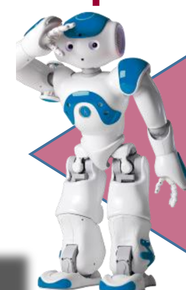

Planning for

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]

[IROS14]

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)

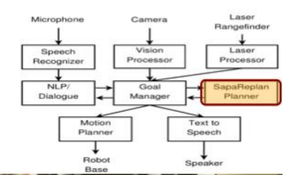

[Cantrell, Talamadupula et al., HRI 2012]

[In collaboration with hrilab, Tufts University]

Human-Robot Cohabitation

Behavior Modeling – Human Aware Planning

- Humans and robots sharing workspace (not necessarily as a team).
- Need for human-aware planning for modeling a robot's interactions with its human colleagues.

Stigmergic Collaboration

in human robot cohabitation

- The robot directly interacts with the human’s plans to assist/coordinate by making positive interventions
  - e.g. planning for serendipity
- The robot coordinates it’s own behavior to suit the human’s predicted plans to minimize conflicts
  - e.g. planning with conflicts on shared resources

Much of the planning challenge is about defining the interaction constraints that affect the robot’s planning process.

Current Use Case

Urban Search and Rescue (USAR) scenario

- Commander can perform triage (needs to get a medkit to do so)
- The Robot can also conduct triage or deliver medkits if requested
    - The medkits are the shared resources here – the robot must de-conflict its plans to use the medkit with that of the human's.

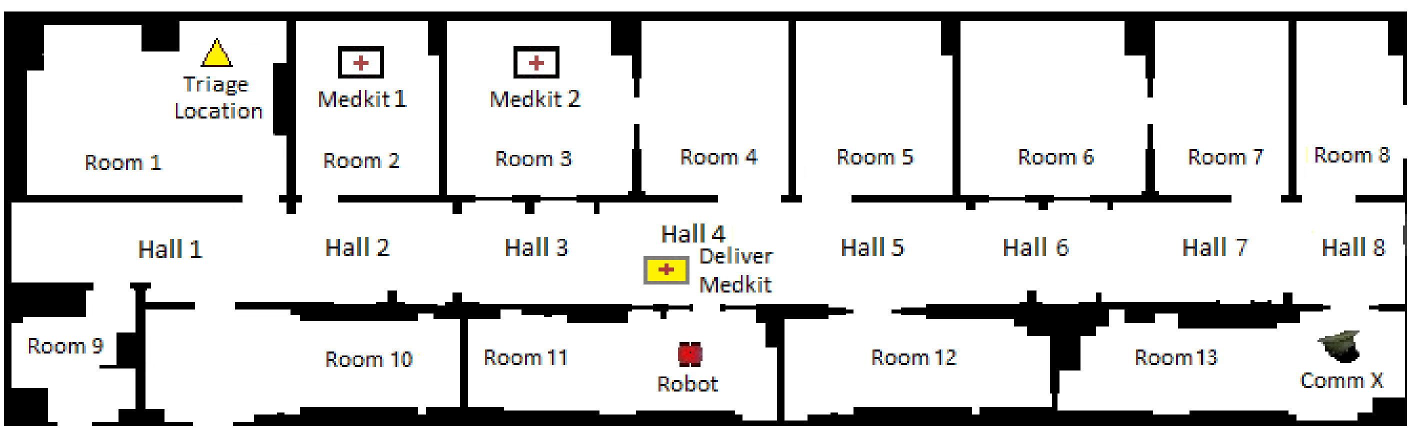

Planning for Serendipity

A running example

`CommX` has to conduct triage in `room1`.

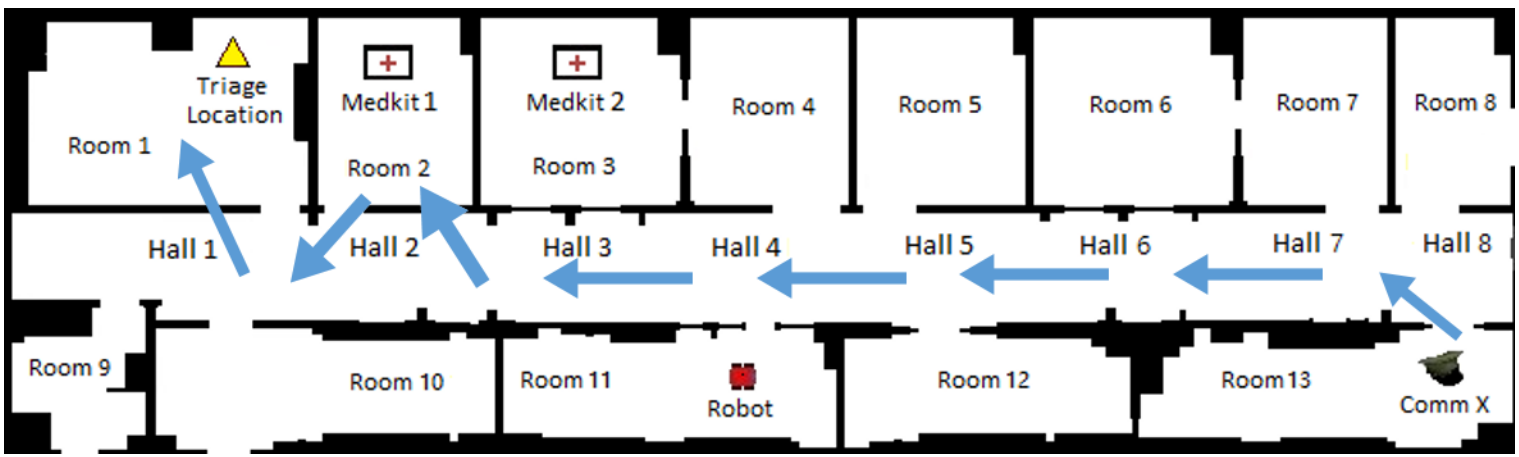

Optimal plan for `CommX` involves picking up `medkit1` in `room2`.

Planning for Serendipity. Tathagata Chakraborti, Gordon Briggs, Kartik Talamadupula, Yu Zhang, Matthias Scheutz, David Smith, Subbarao Kambhampati. IROS 2015, Hamburg.

Planning for Serendipity

A running example

CommX has to conduct triage in room1.

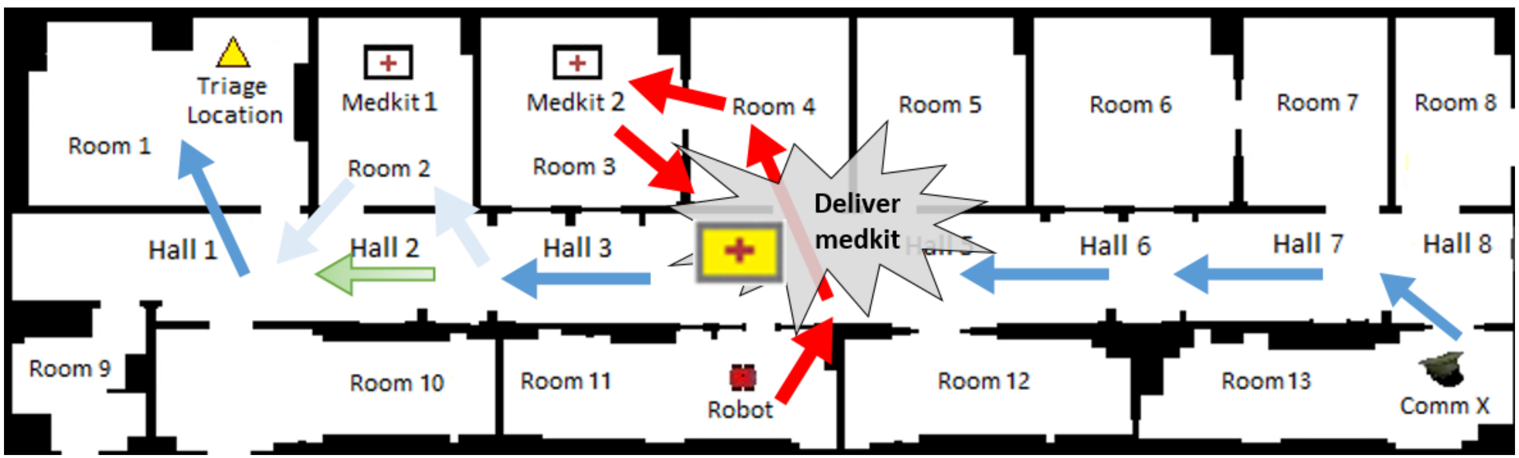

The robot fetches medkit2 from room3 and drops it off in hall3 before CommX passes by, thus saving him the effort to get a medkit himself.

Planning for Serendipity

A running example

`CommX` has to conduct triage in `room1`.

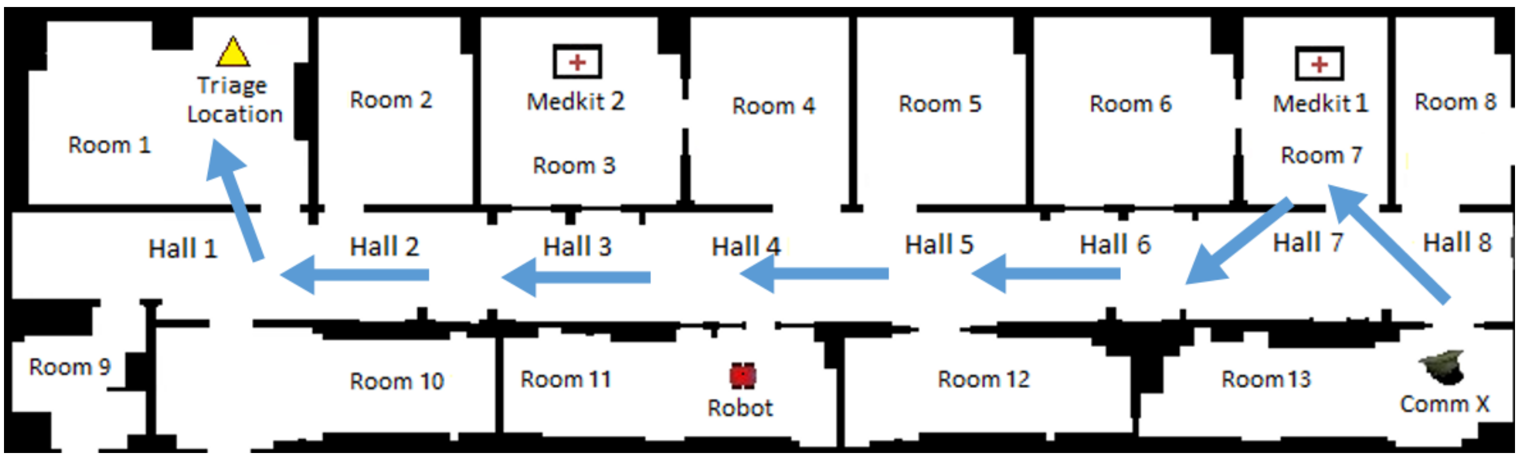

For the current configuration, the optimal plan for `CommX` involves picking up `medkit1` in `room7`.

Planning for Serendipity. Tathagata Chakraborti, Gordon Briggs, Kartik Talamadupula, Yu Zhang, Matthias Scheutz, David Smith, Subbarao Kambhampati. IROS 2015, Hamburg.

Planning for Serendipity

A running example

CommX has to conduct triage in room1.

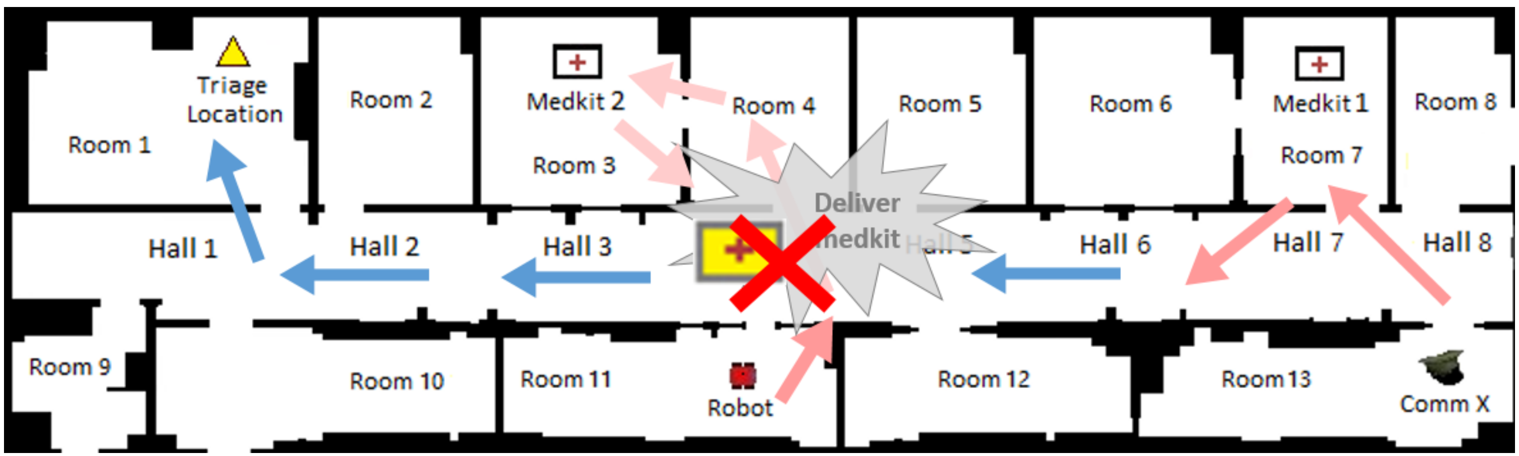

The previous serendipitous intervention becomes redundant here because CommX has already acquired a medkit by the time the robot can intervene.

Planning for Serendipity

A running example

CommX has to conduct triage in room1.

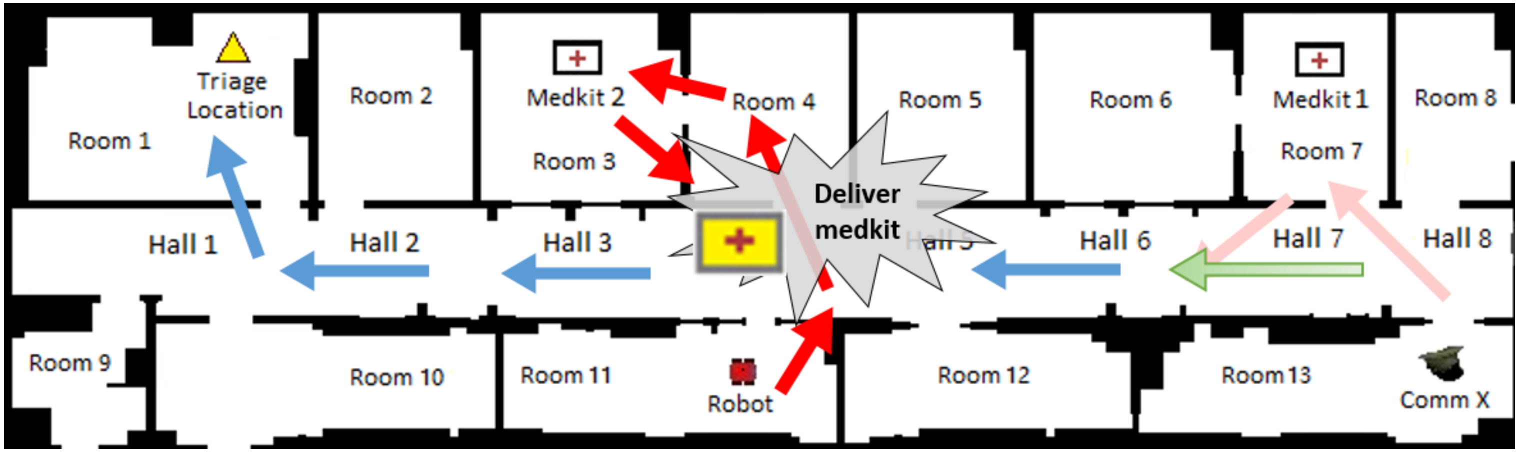

However, if the robot were able to communicate it's intention to intervene, the previous plan for a serendipitous interception still holds.

Plan Interruptibility

Positively removable subplan

Definition 2.0 : If plan $\pi_H = \langle a_1, a_2, \ldots, a_T \rangle$ of the human H with $\delta(\mathbb{I}_H, \pi_H) \models \mathbb{G}_H$, then any subplan $\pi_H^{ij} = \langle a_i, \ldots, a_j \rangle, 1 \leq i < j \leq |\pi_H|$ is *positively removable* iff $\exists \pi_{\mathbb{A}}$ for the set of agents $\mathbb{A} = \{R, H\}$ (R being the robot) such that $\delta'(\bigcup_{\alpha \in \mathbb{A}} \mathbb{I}_\alpha, \pi_{\mathbb{A}}) \models \mathbb{G}_H$ where, for some $i' > i$, $\pi_{\mathbb{A}}(H) = (\subseteq \pi_H[1 : i-1]) \bullet \pi_{\mathbb{A}}(H)[i : i'] \bullet (\subseteq \pi_H[j+1 : |\pi_H|])$ and $C(\pi_{\mathbb{A}}(H)) < C(\pi_H)$ (here $\bullet$ means concatenation).

A plan is *interruptible* iff it has at least one positively removable subplan.

time steps $i \leq t \leq i'$ is when the (serendipitous) interactions can occur

we specify the rest of the plan to be subsequences of the original plan which ensures that the human does not need to go outside his original plan sans the part where the actual interaction occurs.

Serendipitous exceptions here

π_H	a_1	a_2	...		$\pi_H^{ij} = \langle a_i, \ldots, a_j \rangle$			a_T

i ... j

Human's individual plan (as predicted)

Human's component in the composite plan

$$\pi_{\mathbb{A}}(H) = (\subseteq \pi_H[1 : i-1]) \bullet \pi_{\mathbb{A}}(H)[i : i'] \bullet (\subseteq \pi_H[j+1 : |\pi_H|])$$

Plan Preservation

Removable subplans $\neq$ Serendipitous exceptions

- Don't disturb the plan prefix before the serendipitous intervention
  - not necessary if the robot is able to communicate intentions
- The resulting world state after the serendipitous intervention models the original intended state of the human at that point
  - doesn't plan for further assistance

Human's individual plan (as predicted)

i Serendipitous exceptions here j

π_H	a_1	a_2	...		$\pi_H^{ij} = \langle a_i, \ldots, a_j \rangle$			a_T

Human's component in the composite plan

$$\pi_{\mathbb{A}}(H) = (\subseteq \pi_H[1:i-1]) \bullet \pi_{\mathbb{A}}(H)[i:i'] \bullet (\subseteq \pi_H[j+1:|\pi_H|])$$

without communication

$$\pi_{\mathbb{A}}(H)[1:i-1] = \pi_H[1:i-1]$$

where $i = \operatorname{argmin}_i[a = \pi_{\mathbb{A}}(H)[i] \wedge a \notin \pi_H],\ \forall\, a \in A_H$

$$\delta'(\textstyle\bigcup_{\alpha\in\mathbb{A}} \mathbb{I}_\alpha, \pi_{\mathbb{A}}[1:j]) \models \delta(\mathbb{I}_H, \pi_H[1:j])$$

such that $\delta'(\bigcup_{\alpha\in\mathbb{A}} \mathbb{I}_\alpha, \pi_{\mathbb{A}}) \models \mathbb{G}_H$

and $C(\pi_{\mathbb{A}}(H)) < C(\pi_H)$

Planning with Resource Conflicts

Overview & System Components

- Information from the predicted plans concisely represented as resource profiles and fed to the planning stage.

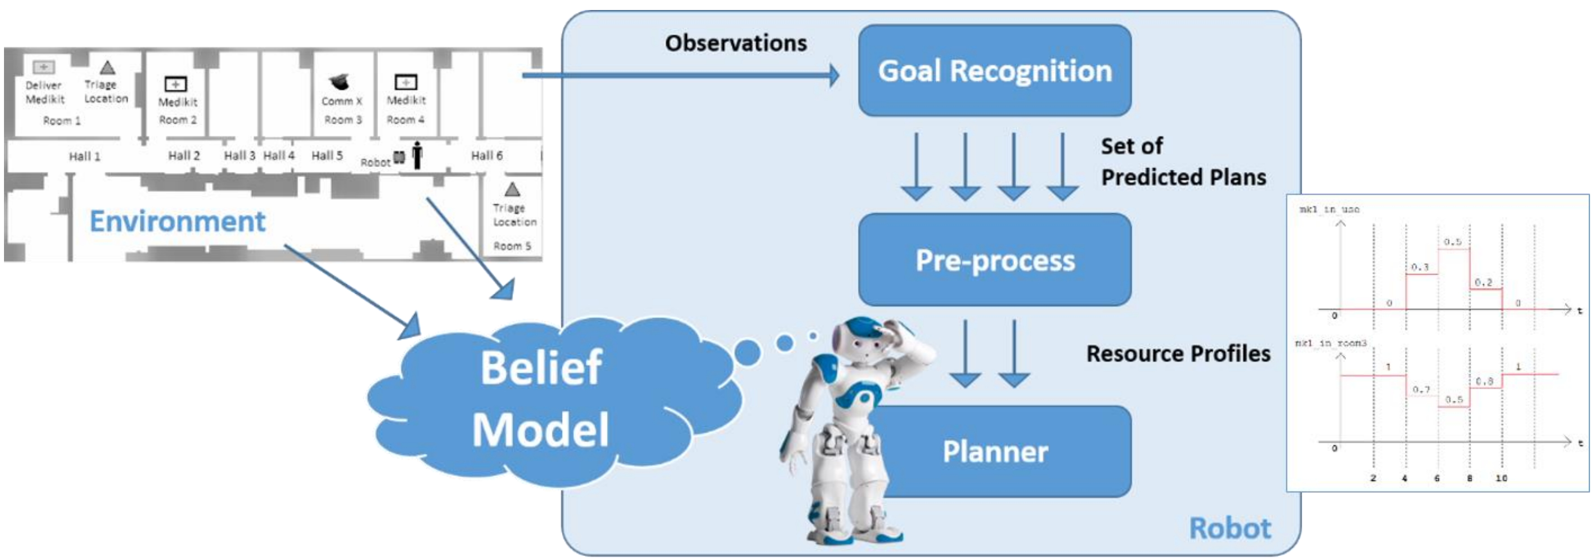

Planning with Stochastic Resource Profiles: An Application to Human-Robot Co-habitation. Tathagata Chakraborti, Yu Zhang, David Smith, Subbarao Kambhampati. ICAPS PlanRob 2015, Jerusalem.

[AAMAS16]

Resource Profiles

different levels of abstraction

- We can have profiles at different levels of abstraction to reason about different aspects of the plan
  - Yes/no of resource usage
  - Profiles over actual groundings of the resource variables

I = {at(commX, room1), at(mk1, room3), connected(room1, room2), connected(room2, room3), connected(room1,hall1), connected(hall1, room2)}

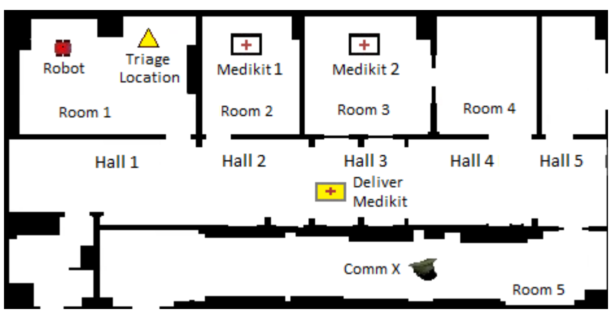

Plan 1 - p(π_1) = 0.6
0 sec: move(room1, room2)
2 sec: move(room2, room3)
4 sec: pick-up(mk1, room3)
5 sec: triage(room3)
8 sec: ~end~

Plan 2 - p(π_2) = 0.4
0 sec: move(room1, hall1)
2 sec: move(hall1, room2)
4 sec: move(room2, room3)
6 sec: pick-up(mk1, room3)
7 sec: triage(room3)
10 sec: ~end~

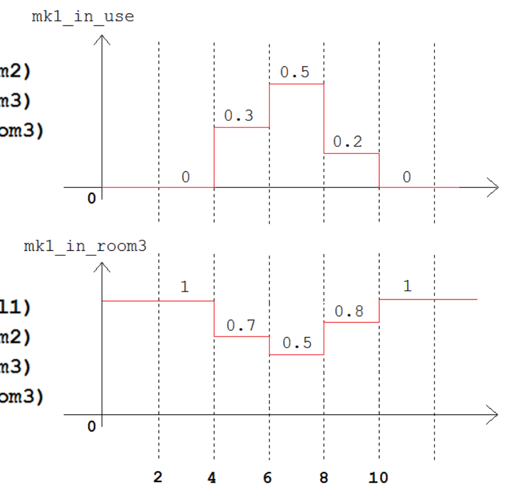

Planning with Stochastic Resource Profiles: An Application to Human-Robot Co-habitation. Tathagata Chakraborti, Yu Zhang, David Smith, Subbarao Kambhampati. ICAPS PlanRob 2015, Jerusalem.

Modeling Behavior – Compromise

Robot settles for a suboptimal plan

`CommX` has to do triage in `room1`, `Robot` is tasked to conduct triage in `hall3` – optimal plans require `medkit1` from `room2` for both agents.

```
01 - MOVE_ROBOT_ROOM1_HALL1
02 - MOVE_ROBOT_HALL1_HALL2
03 - MOVE_ROBOT_HALL2_HALL3
04 - MOVE_ROBOT_HALL3_HALL4
05 - MOVE_REVERSE_ROBOT_HALL4_ROOM4
06 - MOVE_REVERSE_ROBOT_ROOM4_ROOM3
07 - PICK_UP_MEDKIT_ROBOT_MK2_ROOM3
08 - MOVE_ROBOT_ROOM3_ROOM4
09 - MOVE_ROBOT_ROOM4_HALL4
10 - MOVE_REVERSE_ROBOT_HALL4_HALL3
11 - CONDUCT_TRIAGE_ROBOT_HALL3
12 - DROP_OFF_ROBOT_MK2_HALL3
```

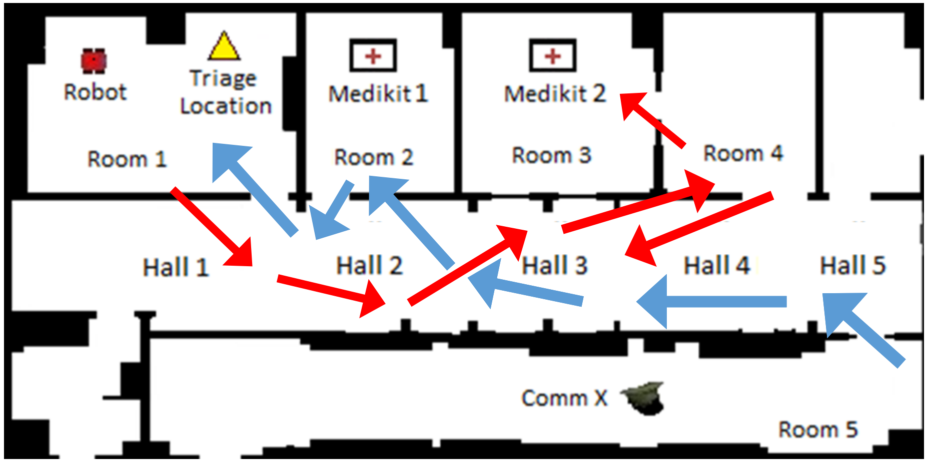

Modeling Behavior – Opportunism

Robot senses favourable turn of events

`CommX` has to do triage in `room1`, `Robot` is tasked to conduct triage in `hall3` – optimal plans require `medkit1` from `room2` for both agents.

```
01 – NOOP
02 – NOOP
03 – NOOP
   ...
12 – NOOP
13 – NOOP
14 – PICK_UP_MEDKIT_ROBOT_MK1_ROOM1
15 – MOVE_ROBOT_ROOM1_HALL1
16 – MOVE_ROBOT_HALL1_HALL2
17 – MOVE_ROBOT_HALL2_HALL3
18 – CONDUCT_TRIAGE_ROBOT_HALL3
19 – DROP_OFF_ROBOT_MK1_HALL3
```

When planning horizon is increased...

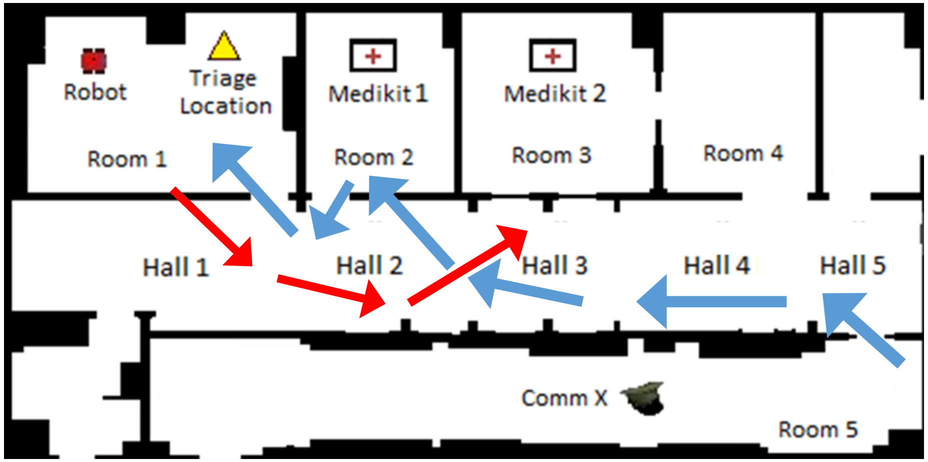

Plan Generation

Integer Programs to model interaction constraints

- We use IP-based planners to model the interaction constraints discussed so far

  - Planning with Resource Conflicts

Minimize overlap between profiles produced by the robot's plans with those predicted from the human's

$$
\begin{aligned}
&min\, k_1 \textstyle\sum_{a \in A_R} \sum_{t \in \{1,2,\ldots,T\}} \mathbb{C}_a \times x_{a,t} \\
&+k_2 \textstyle\sum_{\lambda \in \Lambda} \sum_{f \in \Gamma(\lambda)} \sum_{t \in \{1,2,\ldots,T\}} g_{f,t} \times G^{\lambda}(t) \\
&-k_3 \textstyle\sum_{\lambda \in \Lambda} \sum_{f \in \Gamma(\lambda)} \sum_{t \in \{0,1,\ldots,T-1\}} h_{f,t} \times G^{f^{\lambda}}(t)
\end{aligned}
$$

$$h_{f,t-1} = x_{f,t}\, \forall a\ s.t.\ f \in \mathbb{P}_a, \forall f \in \xi, t \in \{1,\ldots,T\} \tag{5}$$

$$\textstyle\sum_{a \in \Omega_f^+} \sum_t x_{a,t} \leq 1\, \forall f \in \xi, t \in \{1,2,\ldots,T\} \tag{9}$$

$$
\begin{aligned}
&g_{f,t} = \textstyle\sum_{a \in \Omega_f^+} x_{a,t} \\
&+(1 - \textstyle\sum_{a \in \Omega_f^+} x_{a,t} - \sum_{a \in \Omega_f^-} x_{a,t}) \times g_{f,t-1} \\
&\forall f \in \xi, t \in \{1,\ldots,T\}
\end{aligned} \tag{10}
$$

$$h_{f,t} \times G^{f^{\lambda}}(t) \geq \epsilon\, \forall f \in \xi, t \in \{0,1,\ldots,T-1\} \tag{11}$$

  - Planning for Serendipity

Compute positively removable sub-plans that uphold the two preservation constraints

$$Obj\ :\ min \textstyle\sum_{a \in A_{\mathbb{A}}} \sum_{t \in \{1,2,\ldots,T\}} \mathbb{C}_a \times x_{a,t} + K\, ||\xi_2 - \xi_1||$$

$$
\begin{aligned}
&\xi_1 \leq \textstyle\sum_t (t \times x_{a,t})(1 - \sum_t \sum_{a \in \pi_{\alpha_1}} x_{a,t}) \\
&+ T(1 - \textstyle\sum_t x_{a,t}) + T(\sum_t \sum_{a \in \pi_{\alpha_1}} x_{a,t}) \\
&\forall a \in A_H,\ t \in \{1,2,\ldots,T\}
\end{aligned} \tag{7a}
$$

$$x_{a,t} \geq \tfrac{1}{T}(\xi_1 - t)\, \forall a \in \pi_H,\ t \in \{1,2,\ldots,T\} \tag{7b}$$

$$x_{a,t} \leq 1 + \tfrac{1}{T}(\xi_2 - t)\, \forall a \in A_R,\ t \in \{1,\ldots,T\} \tag{8}$$

$$x_{a,t} + x_{a_\phi,t} \geq \tfrac{1}{T}(t - \xi_2)\, \forall a \in \pi_H,\ t \in \{1,2,\ldots,T\} \tag{9}$$

$$
\begin{aligned}
&\textstyle\sum_{a \in A_\alpha} x_{a,t} + \sum_{a \in A_{\mathbb{A}} \setminus \bigcup_{\alpha \in \mathbb{A}} A_\alpha} x_{a,t} \leq 1 \\
&\forall \alpha \in \mathbb{A},\ t \in \{1,\ldots,T\}
\end{aligned} \tag{10}
$$

$$\textstyle\sum_{a \in A_{\mathbb{A}}} \sum_{t \in \{1,2,\ldots,T\}} \mathbb{C}_a \times x_{a,t} \leq cost(\pi_H) \tag{11}$$

$$\xi_1, \xi_2 \in \{1,2,\ldots,T\},\ \xi_2 \leq \xi_1 + 1 \tag{12}$$

$$y_{f,t} \in \{0,1\}\, \forall f \in S_{\mathbb{A}},\ t \in \{0,1,\ldots,T\} \tag{13}$$

$$x_{a,t} \in \{0,1\}\, \forall a \in A_{\mathbb{A}},\ t \in \{1,2,\ldots,T\} \tag{14}$$

Evaluations – Planning for Serendipity

- We compare the reduction in cost of (overall) team plans from individual optimal plans to planning for serendipity, with and without communication.
- The robot's actions costs are discounted with respect to those of the human's to demonstrate how more and more situations become conducive to serendipitous interventions as the robot's actions become relatively cheaper.
  - Number of serendipitous plans indicate that there are plenty of opportunities for such serendipitous interventions.

Discount	w/o comm.	w/ comm.
0%	9.82 (1)	9.72 (13)
10%	9.81 (7)	9.65 (23)
30%	9.79 (7)	9.48 (34)
50%	9.76 (12)	9.25 (40)
70%	9.68 (29)	8.93 (62)
90%	9.55 (32)	8.51 (70)

Average individual plan cost = 9.825

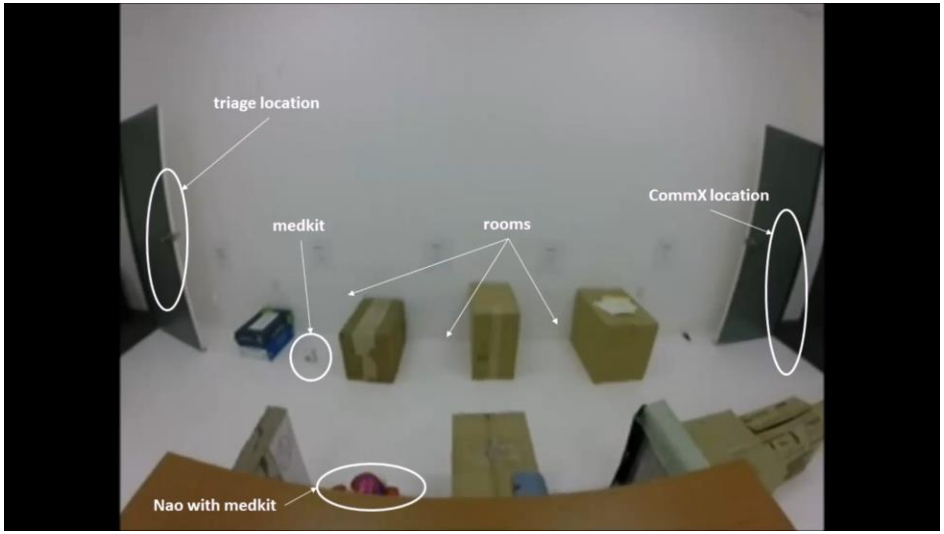

Evaluations – Planning with Resource Conflicts

- We contrast the effect of the parameters of the IP-formulation on the plans produced.
  - Increasing the planning horizon makes the robot more opportunistic
  - Increasing the relative penalty for overlaps in profiles makes the robot more conservative and lowers utility
  - Algorithm is robust to number of observations, but larger hypothesis sets effect the planner negatively as expected
- Complexity of the planner stage only is independent of the number of agents, and size of the hypothesis set – advantage of the modular approach and profile representation of plans.

T	10	13	16	Optimal
C	9.0	5.6	4.53	9.0
U	0.46	0.04	≈ 0	n/a
S	1.0	0.48	0.25	n/a
F	53.3	11.9	6.6	53.3

Table 2: Quality of plans produced w.r.t. T. Opportunities for opportunism explored, conflicts minimized.

k_1/k_3	0.05	0.5	5.0
C	9.47	6.37	6.31
U	0.18	0.17	0.17
S	0.85	0.579	0.578
F	27.5	23.0	21.3

Table 1: Quality of plans produced w.r.t. k_1/k_2. Conservative plans result in lowered utility.

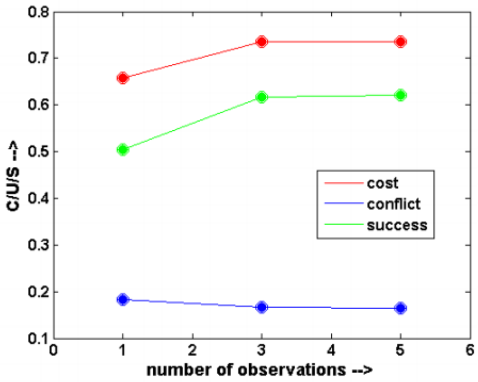

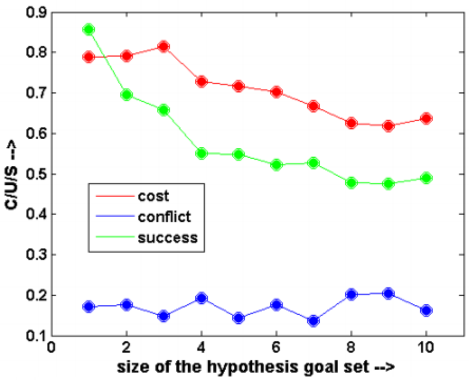

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

When is a plan "Explainable" to the human in the loop?

- The robot generates its plan of action using its model M_R
- The human "interprets" this plan in light of her understanding of the Robot's model M^*_R
- M_R and M^*_R can be quite different..
- Differences can be a result of:
  - Different capabilities (e.g., possible actions)
  - Different knowledge (e.g., level of modeling)
  - Different interpretation of behaviors (e.g., plans) interacting with the world -- more than just trajectory planning!

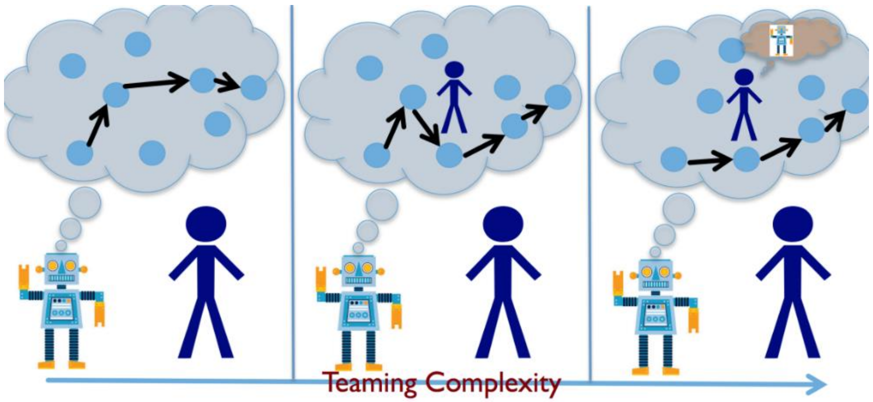

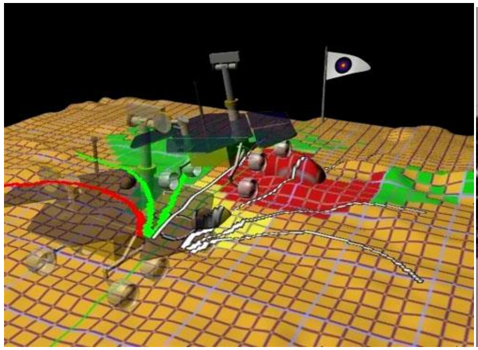

$$\underset{\pi_{M_R}}{\operatorname{argmin}}\, cost(\pi_{M_R}) + \alpha \cdot dist(\pi_{M_R}, \pi_{\mathcal{M}^*_R})$$

But, alas, M^*_R is not known!

Explainability Labeling

Problem: M^*_R is not known

Solution: Learn it, but indirectly as a labeling scheme..

$$\underset{\pi_{M_R}}{\operatorname{argmin}}\, cost(\pi_{M_R}) + \alpha \cdot \boxed{dist(\pi_{M_R}, \pi_{\mathcal{M}^*_R})}$$

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

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

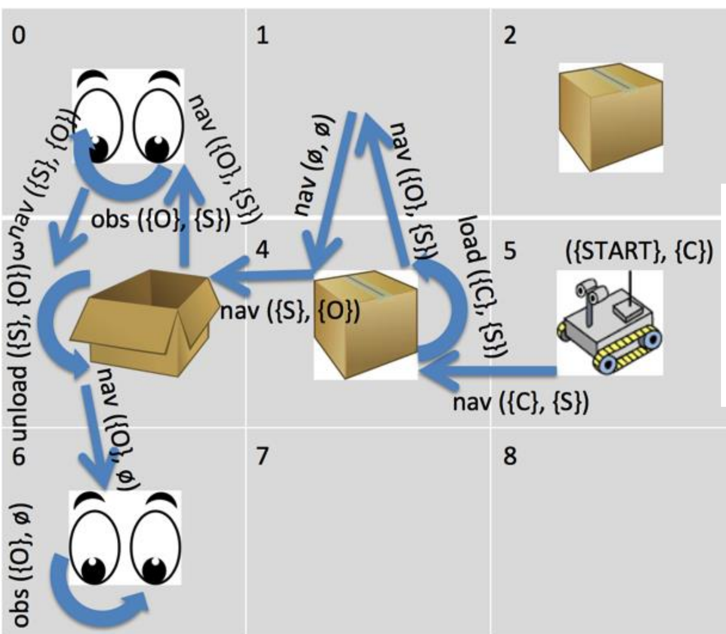

Analogy: Think of learning how to write address labels so the postal carrier can understand..

- Task labels (to associate with actions). For example:
  - Collect
  - Store
  - Observe

More than one label is allowed for actions

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

Learning the Labeling Scheme using CRF

- **Model:**
  - Conditional Random Fields (CRF)

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

- **Features:**
  - Plan features: e.g., at rover L5
  - Action/trajectory Features: e.g., action type
  - Interaction features: e.g., distance to the human

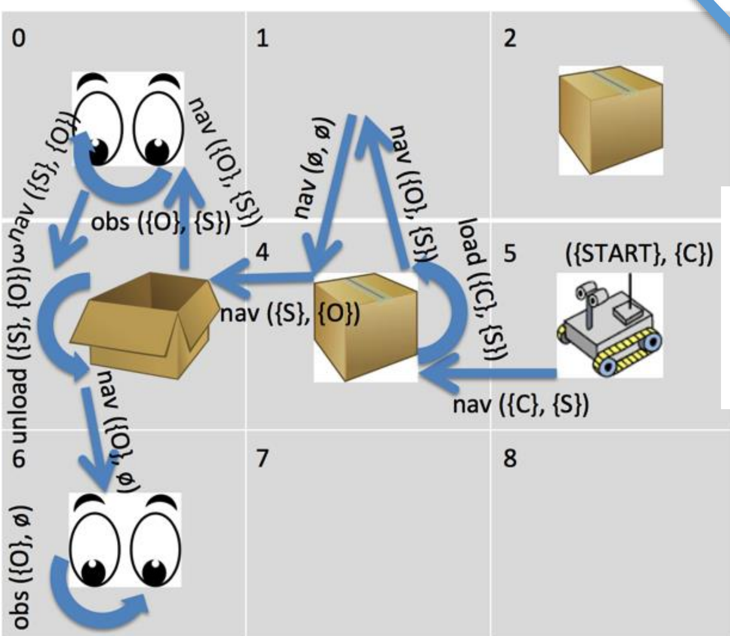

- Task labels (to associate with actions). For example:
  - Collect
  - Store
  - Observe

More than one label is allowed for actions

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

Using Explainability Model

Preliminary results indicate that such a scheme is reasonably effective in picking explainable plans..

Plan selection

- Robot can generate a set of plans and select the most explainable/predictable plan

Plan heuristic

- Robots can use it to directly synthesize more explainable/predictable plans

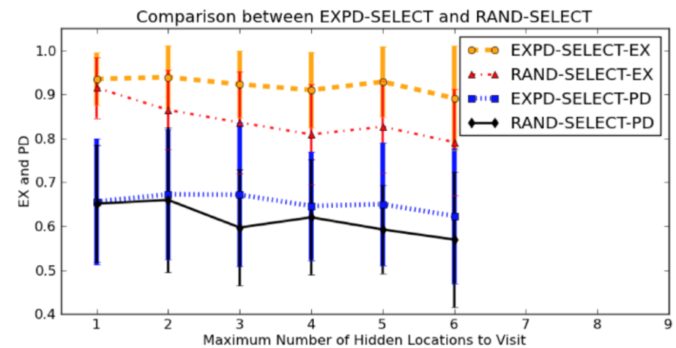

Figure 5: Comparison of EXPD-SELECT and RAND-SELECT

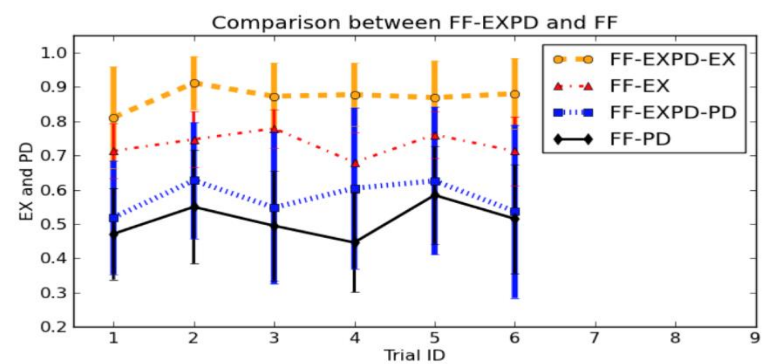

Figure 6: Comparison of FF-EXPD and FF considering $uexp$ **in Alg. 1.**

Agenda for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot's model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the "paper clip" assistant shows, we AI'ers are not great at guessing what humans "like" ☹

Do we really know what (sort of assistance) humans want?

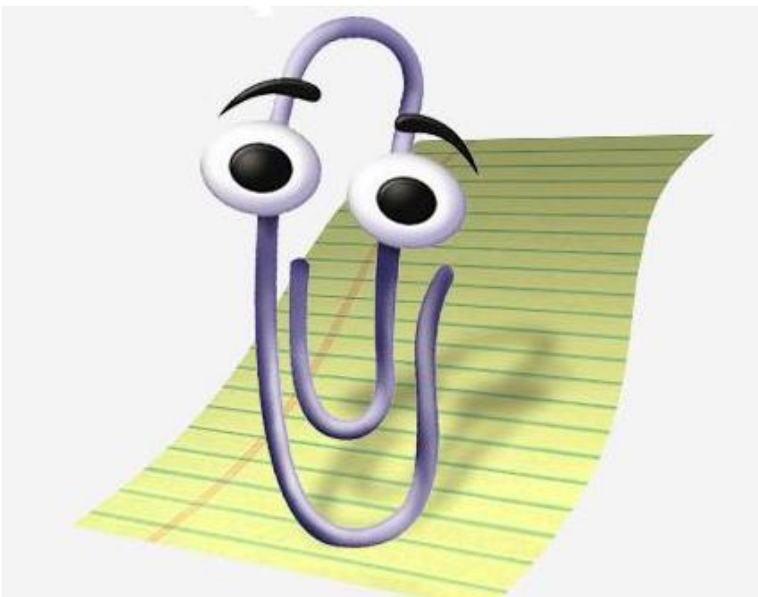

We dance round in a ring and suppose,
But the Secret sits in the middle and knows.

Proactive Help Can be Disconcerting!

Human Factor Studies

- To understand whether human-robot teams perform better with more intelligent/proactive robot teammates or not
- Two studies
  - Wizard-of-Oz Human-Human studies
    - With Cade Bartlett and Nancy Cooke
      - Cade Bartlett's M.S. thesis (in preparation for Journal submission)
  - Human-Planner studies
    - To see if proactive robots that use plan recognition to anticipate human actions help or hinder team performance
      - [IROS 2015][HRI2015]

Human-human Teaming Analysis in Urban Search and Rescue

Simulated search task (Minecraft) with human playing role of USAR robot

- 20 internal/external dyads tested
- Conditions of autonomous/intelligent or remotely controlled robot
- Differences in SA, performance, and communications

Measures

- Performance:
  - $((Rooms\ Marked\ Correctly + Correct\ Presses) - (Repeated\ presses + Incorrect\ Presses))$.
- SA
  - External – Rooms marked correctly
  - Internal – Repeated presses
- Covariates
  - Spatial ability task
  - Demographics
  - Experience
    - Robotics
    - Minecraft
    - Gaming

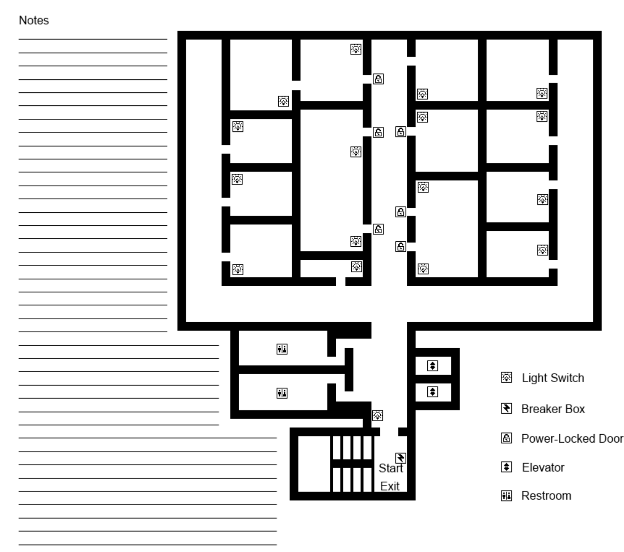

Procedure

1. Random role assignment
2. Consent
3. Spatial ability test
4. Seated with divider between
5. Instructions (according to condition)
6. Search plan
7. Internal training
8. USAR task (15 minutes)
9. Notified of end of time targets (8 minutes in)
10. Demographics/experience/TLX
11. Debrief
12. Compensation

Summary of Key Findings

- Intelligent condition (vs. Remote):
  - Higher dyad performance
  - Lower external workload
  - Less communications (especially from external)
  - Tendency for Higher External SA (non-significant)

    Higher Levels of External SA driven by
    - Greater percent excuses
    - Lower difference in spatial ability between internal and external
    - Higher internal SA (low repeat button presses)
- Communications associated with most effective dyad performance
  - Higher percent excuses (A flag for whether the communication was related to one of the environment's inconsistencies with the provided map)

Human Factor Studies

- To understand whether human-robot teams perform better with more intelligent/proactive robot teammates or not
- Two studies
  - Wizard-of-Oz Human-Human studies
    - With Cade Bartlett and Nancy Cooke
      - Cade Bartlett's M.S. thesis (in preparation for Journal submission)
  - Human-Planner studies
    - To see if proactive robots that use plan recognition to anticipate human actions help or hinder team performance
      - [IROS 2015][HRI 2015]

Analysis of Proactive Support in Human-robot teaming

Simulated search task (Webots) with human remotely controlling a robot while collaborating with an intelligent robot 'Mary':

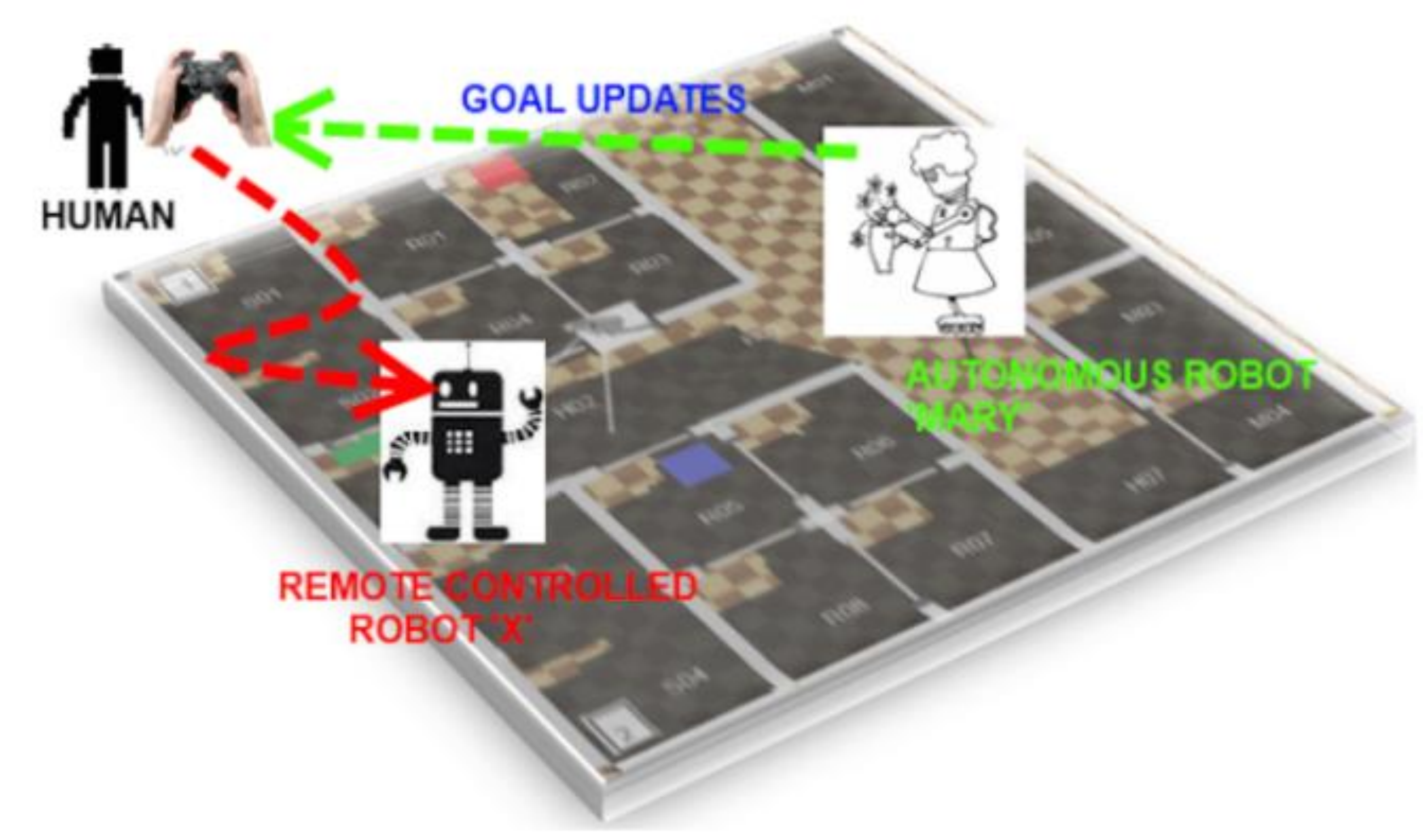

Measures

Performance:

- Time taken to treat the critically injured casualty; time taken to treat both

Task settings:

- Monitoring cameras that can provide observations
- Casualties to be treated using medical kits in medical rooms
- Environment segmented by doors

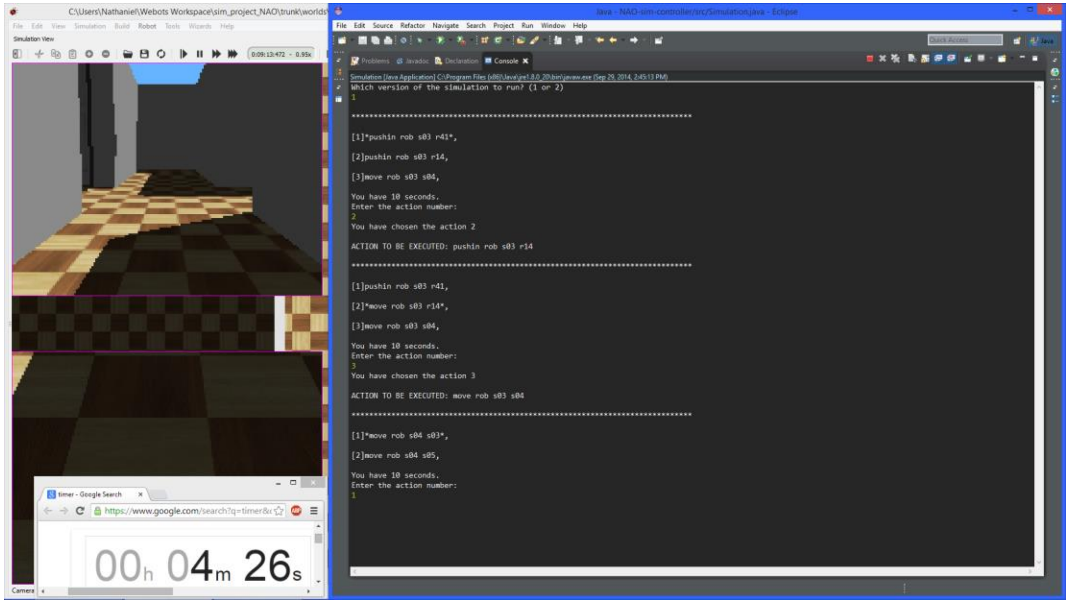

Summary of Key Findings

- Mary with a proactive support capability (vs. without):
    - Higher dyad performance
    - Lower communication
    - Slightly (non-significant) <u>increased mental workload</u>
- Mary with a proactive support capability in our USAR task scenario is generally preferred

[IROS 2015; HRI 2015]

Summary for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- How to make planned behavior explainable to the human in the loop
  - Humans will parse the behavior in terms of their understanding of the Robot’s model
- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

Summary for Today

- How to learn and plan with incomplete domain models
  - Complete--Approximate--Shallow
- How to plan to be useful to the human
  - Avoiding conflicts and offering serendipitous help
- e to
  - eir

Until my next update @2025 AI Seminar

- How to recognize and evaluate what are the desiderata for fluent teaming with humans
  - As the “paper clip” assistant shows, we AI’ers are not great at guessing what humans “like” ☹

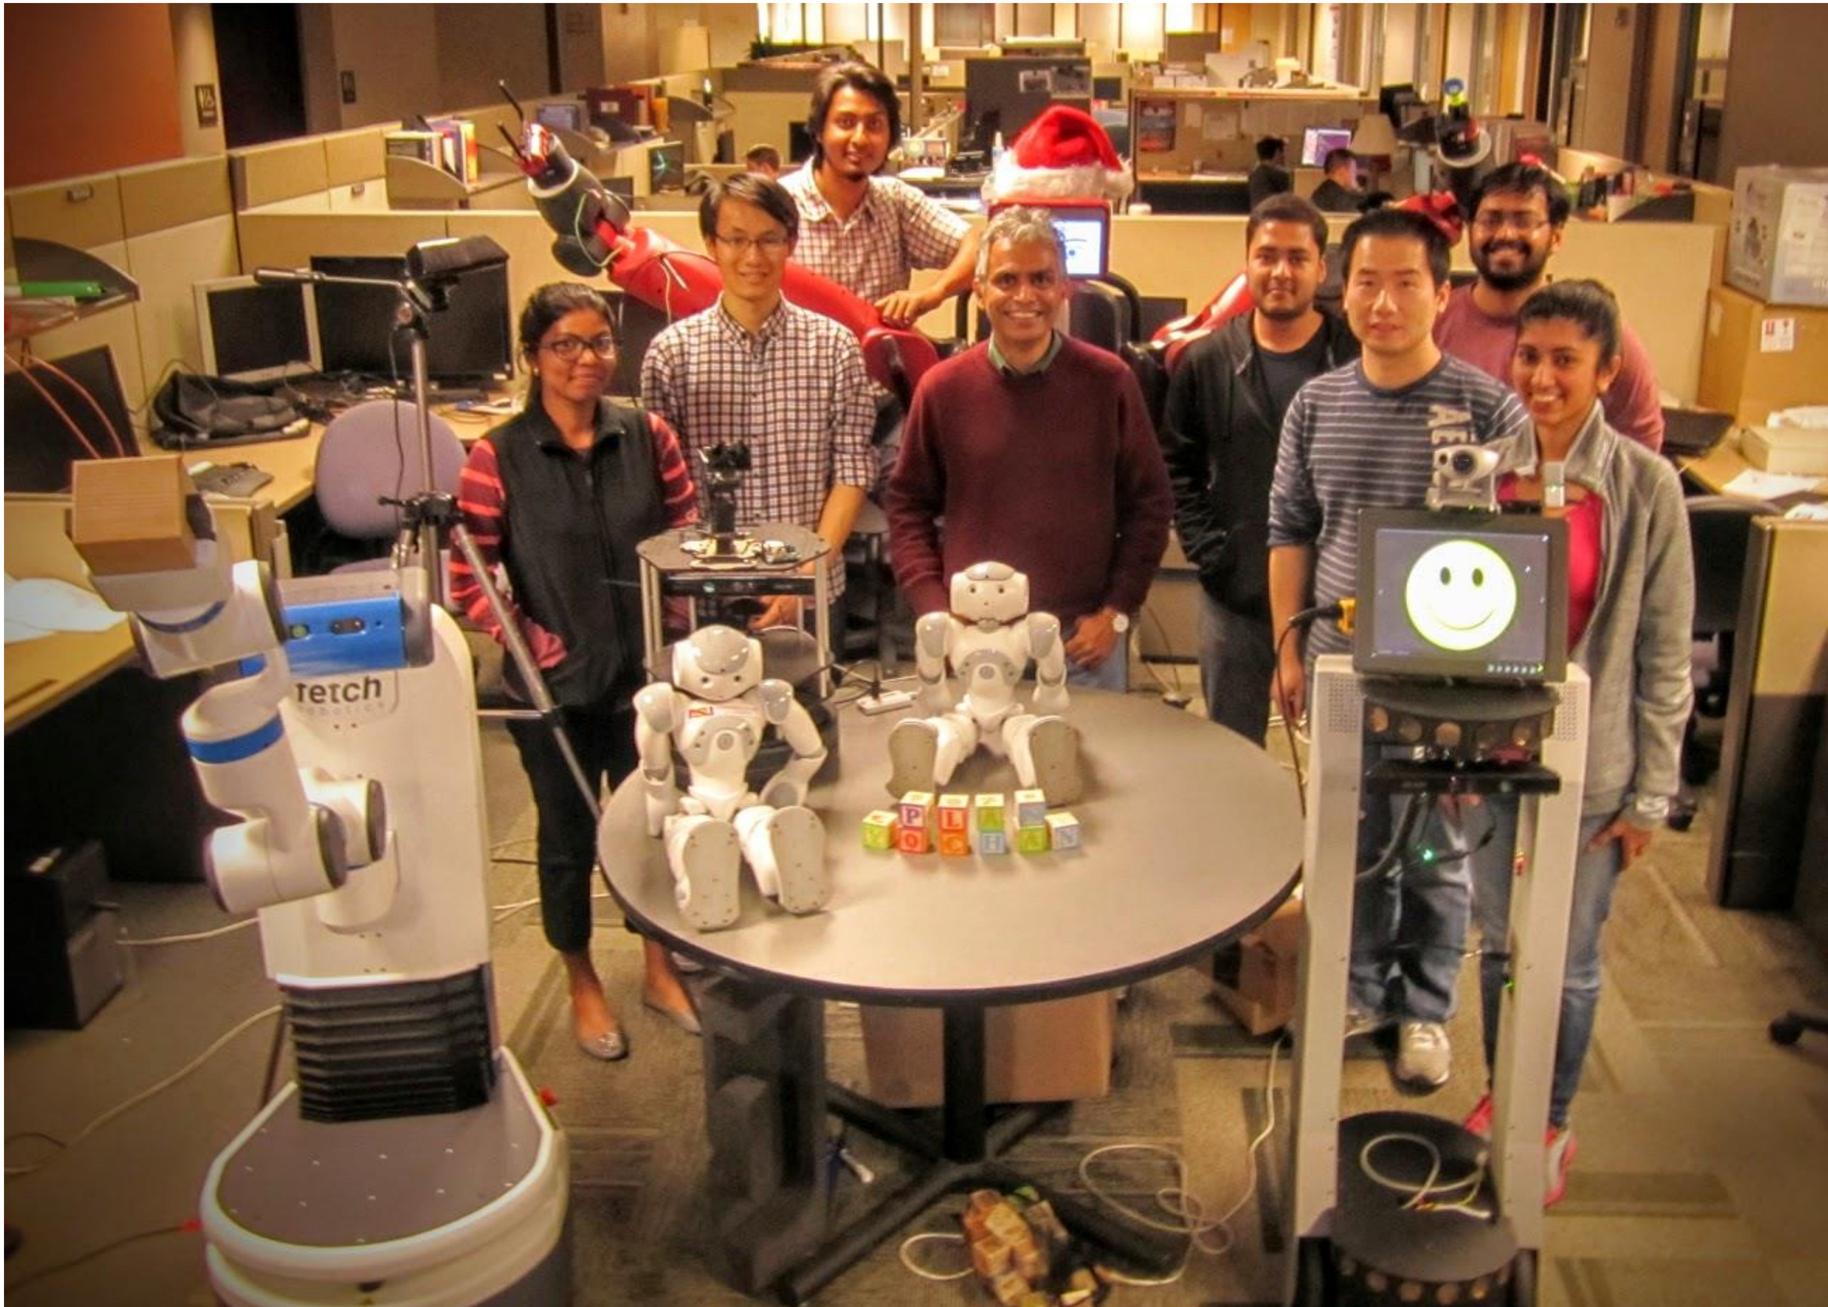
fetch