Planning Challenges in Human-Machine Collaboration

(In Praise of Human-Aware AI)

Subbarao Kambhampati

Arizona State University

Association for Advancement of Artificial Intelligence (AAAI)

Search “rao asu”
WeChat: Subbarao2z

Funding from ONR, ARO and NSF gratefully acknowledged

CCAI as a Rosetta Stone

Can Rao, who knows AI, learn Chinese by listening to CCAI talks?

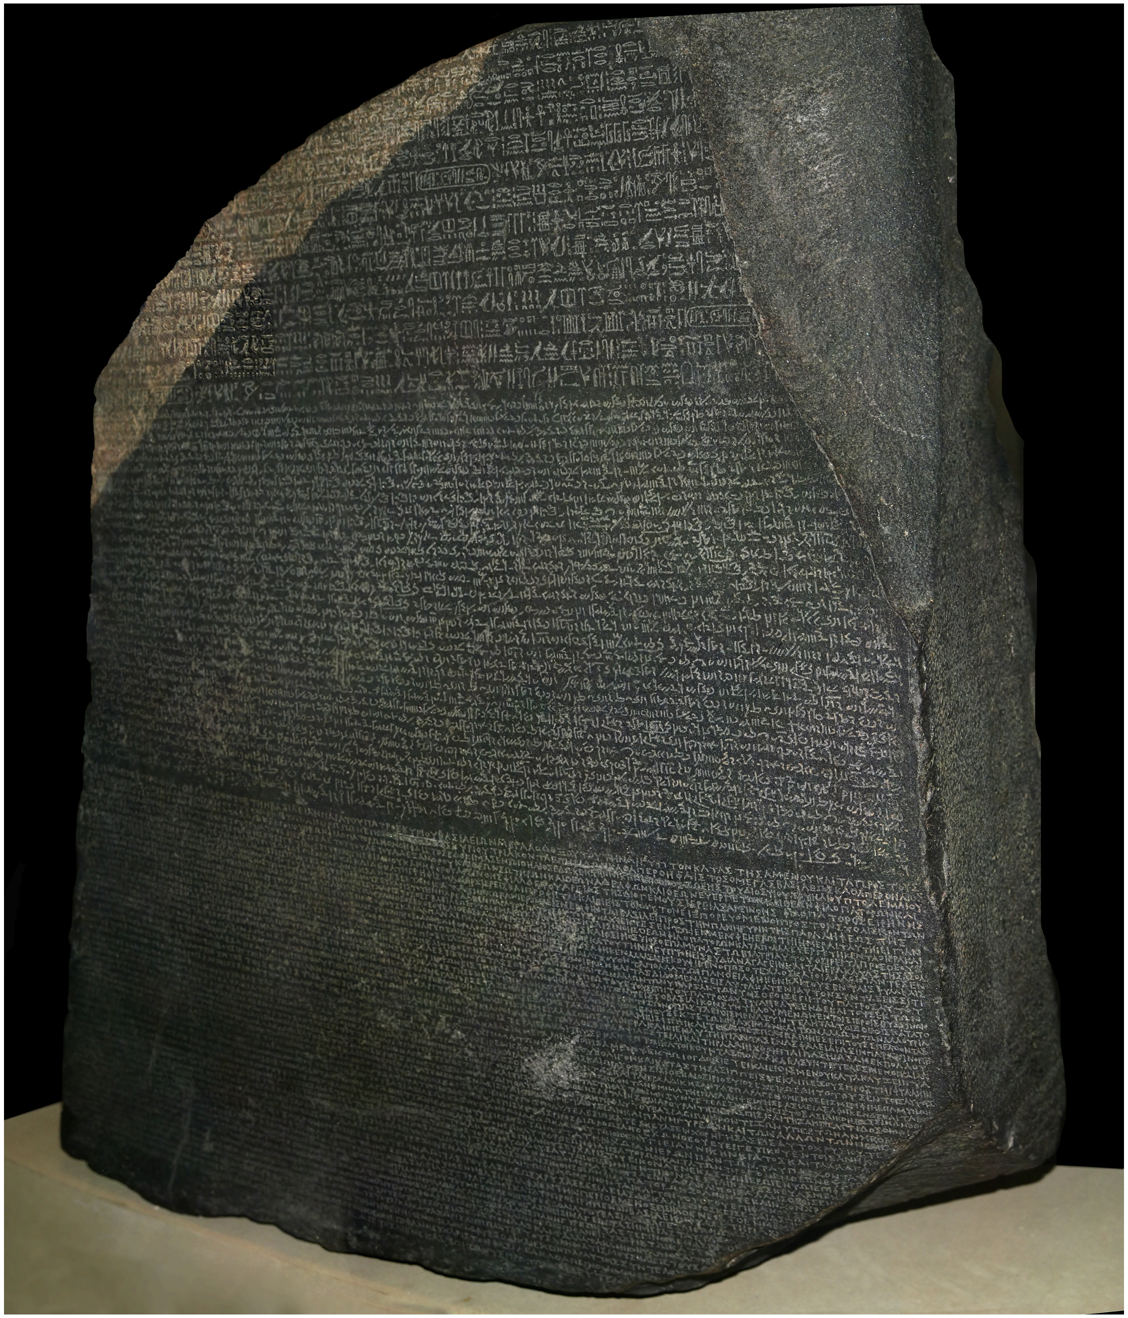

What AI needs to do right away..

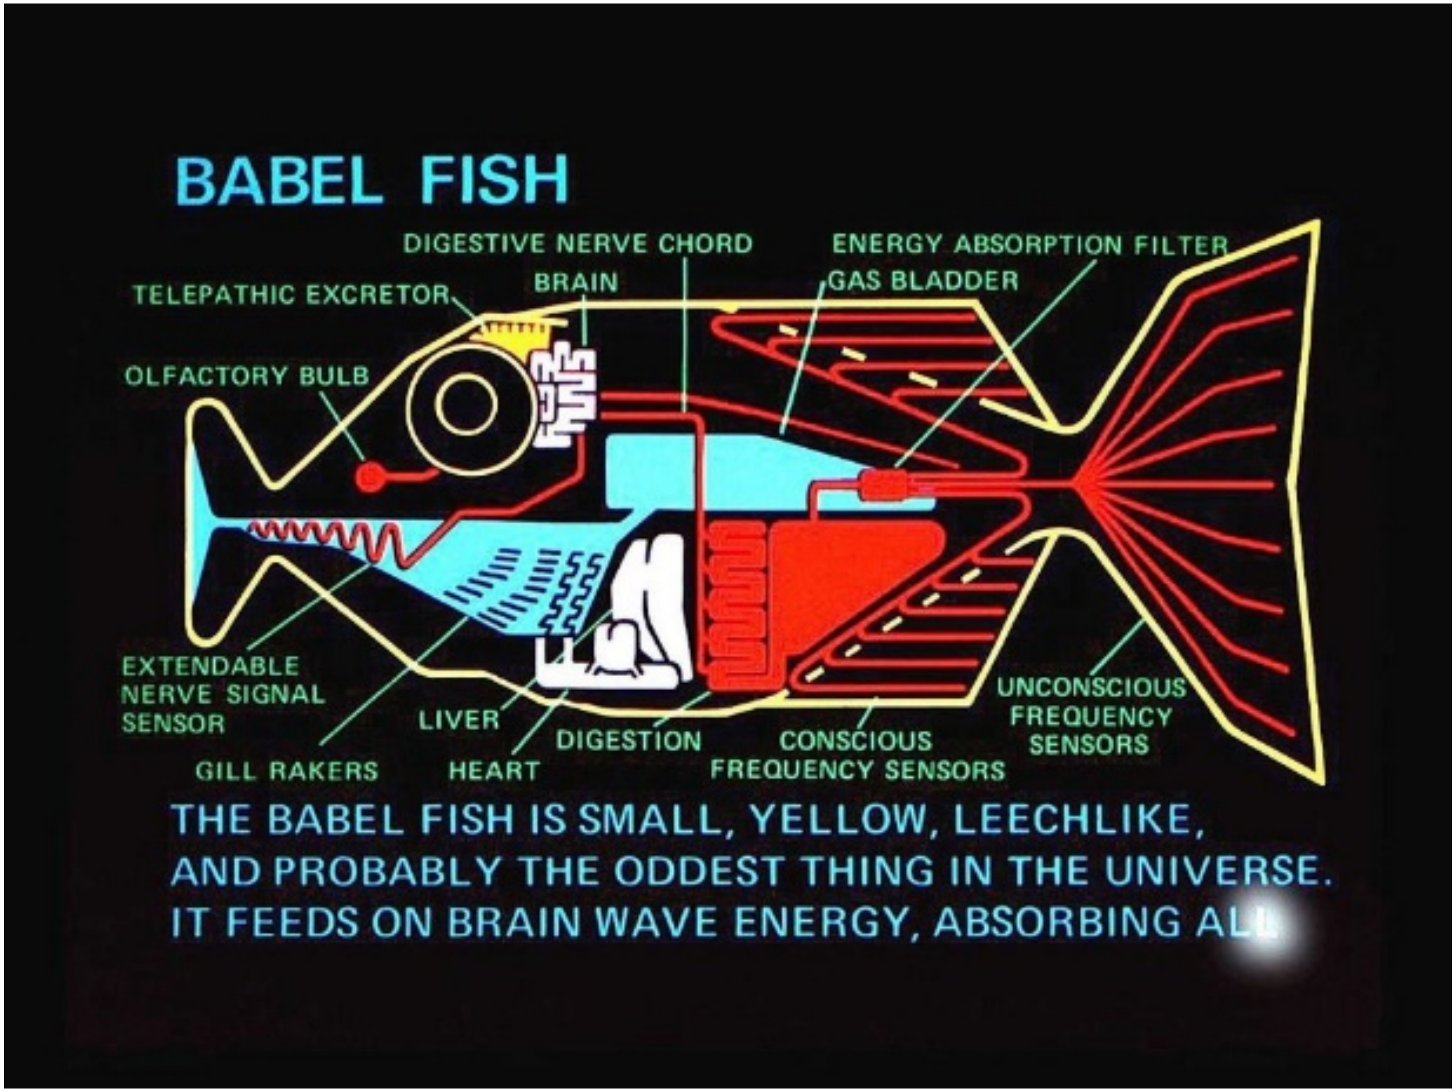

[Hitchhiker's Guide To the Galaxy]

AAAI & China AI Community

- Researchers from China are a formidable force in AI and AAAI
  - One of the top countries in terms of number of submissions as well as acceptances
- AAAI-17 dates shifted to avoid conflict with the start of the Year of Rooster!
- We want more prominent representation of China in AAAI
  - Prof. Qiang Yang is now on the Executive Council

AAAI Conference in February 4-9, 2017 in San Francisco!

Agenda

- Part I: The Path to General AI goes through Human-Machine Collaboration
  - ..and it is a good thing!
    - Expands reach and scope of AI enterprise
    - Reduces some of the off-the-top worries about AI
    - Brings up novel research challenges
- Part II: Planning Challenges in Human-Machine Collaboration
  - Brief review of how the planning problem “expands” in the face of interaction/teaming with humans
  - Specific challenges and some ongoing work in my group

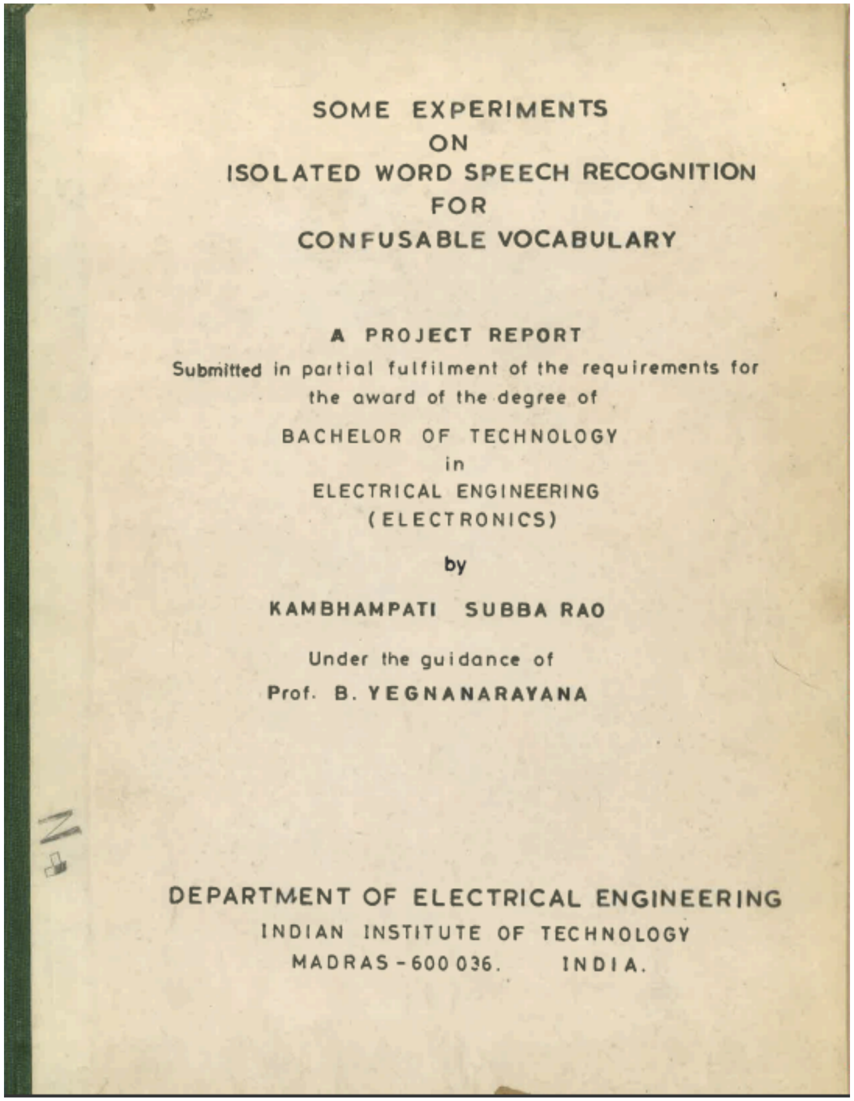

SOME EXPERIMENTS
ON
ISOLATED WORD SPEECH RECOGNITION
FOR
CONFUSABLE VOCABULARY

A PROJECT REPORT
Submitted in partial fulfilment of the requirements for
the award of the degree of
BACHELOR OF TECHNOLOGY
in
ELECTRICAL ENGINEERING
(ELECTRONICS)

by

KAMBHAMPATI SUBBA RAO

Under the guidance of
Prof. B. YEGNANARAYANA

DEPARTMENT OF ELECTRICAL ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY
MADRAS - 600 036. INDIA.

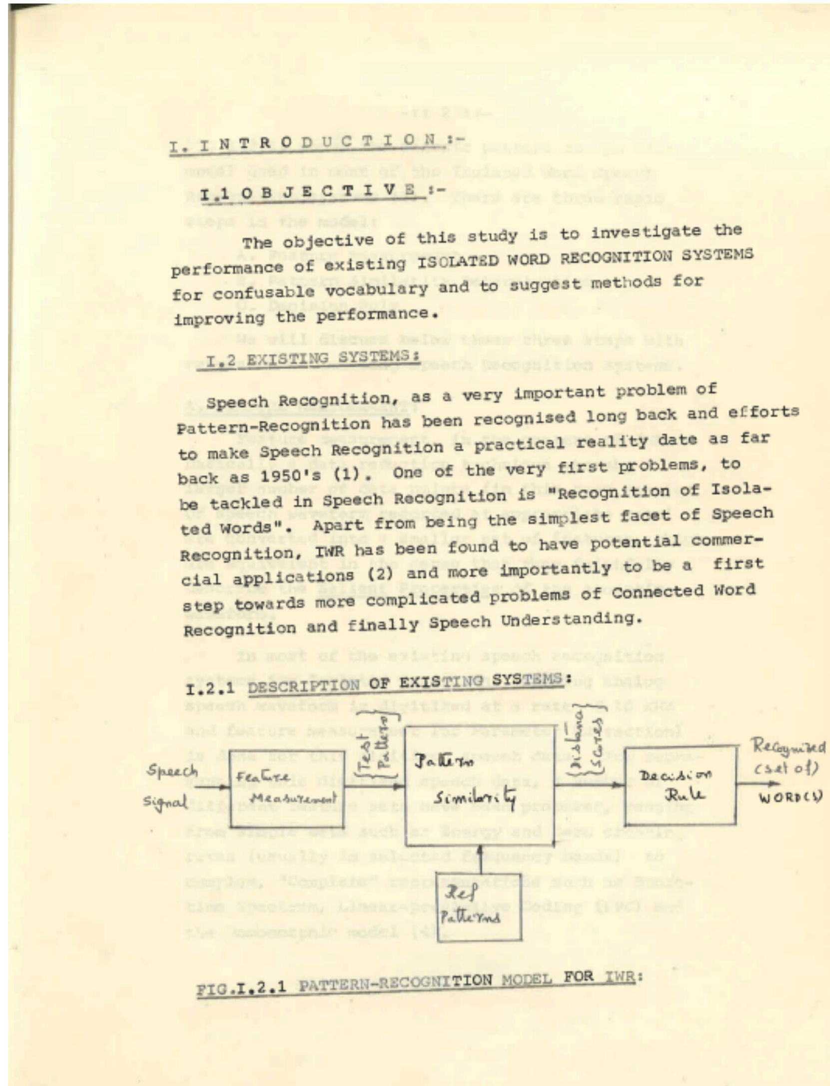

I. INTRODUCTION:-

I.1 OBJECTIVE:-

The objective of this study is to investigate the performance of existing ISOLATED WORD RECOGNITION SYSTEMS for confusable vocabulary and to suggest methods for improving the performance.

I.2 EXISTING SYSTEMS:

Speech Recognition, as a very important problem of Pattern-Recognition has been recognised long back and efforts to make Speech Recognition a practical reality date as far back as 1950's (1). One of the very first problems, to be tackled in Speech Recognition is "Recognition of Isolated Words". Apart from being the simplest facet of Speech Recognition, IWR has been found to have potential commercial applications (2) and more importantly to be a first step towards more complicated problems of Connected Word Recognition and finally Speech Understanding.

I.2.1 DESCRIPTION OF EXISTING SYSTEMS:

FIG.I.2.1 PATTERN-RECOGNITION MODEL FOR IWR:

1983 Bachelors thesis ☺

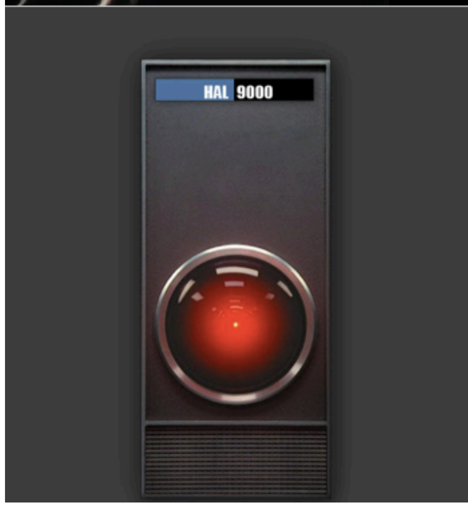

You Can't Do That, Dave!

Collateral Lessons from a Computational Quest to Design HAL

Subbarao Kambhampati
Computer Science & Engineering

The Fundamental Questions Facing Our Age

- Origin of the Universe
- Origin of Life
- Nature of Intelligence

"To know your future you must know your past"

– George Santayana

Predictions are hard,
especially about the future

--Niels Bohr

AI & Life: Thursdays with Rao

- Computational quest to design HAL may not directly teach you pat life lessons of the "Tuesday's with Morrie" type
- But, it gives you a deeper, more nuanced, appreciation of some of life's fundamental tradeoffs…
- ..and I think you will be much the richer for the understanding

Thank You

In this thought experiment, a person in the "Chinese room" is passed questions from outside the room, and consults a library of books to formulate an answer

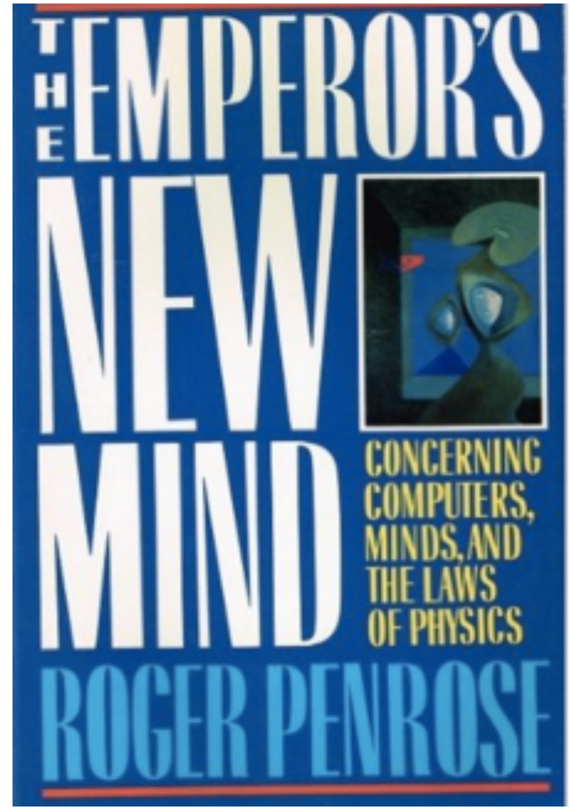

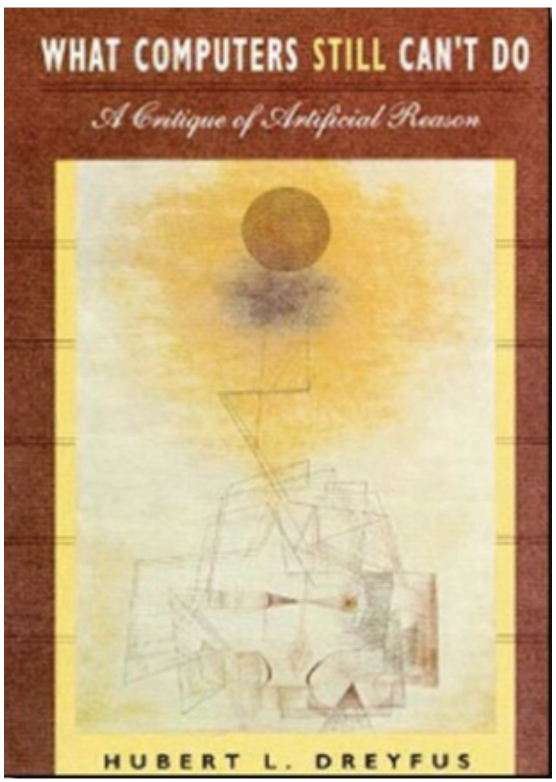

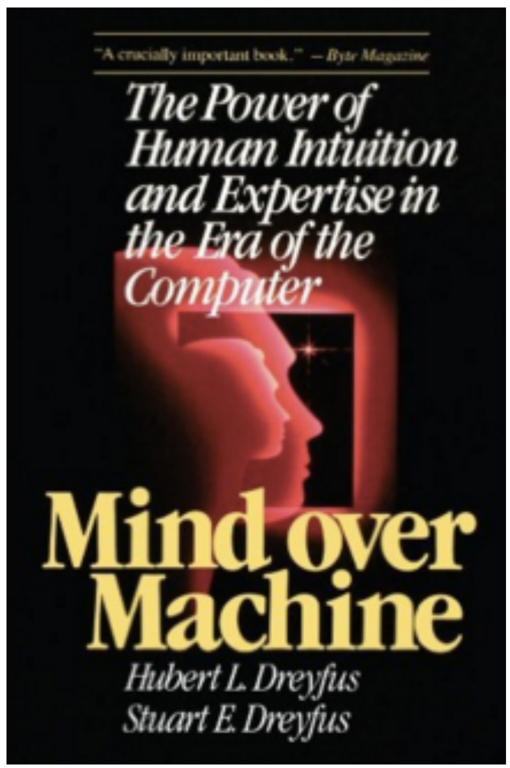

"Physicists and Philosophers united against AI"?

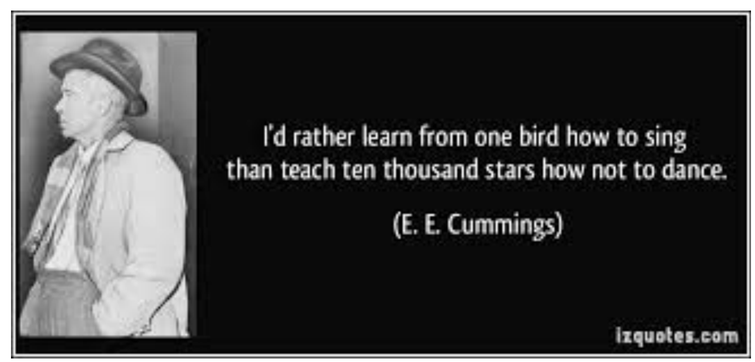

Musk, Wozniak and Hawking urge ban on warfare AI and autonomous weapons

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

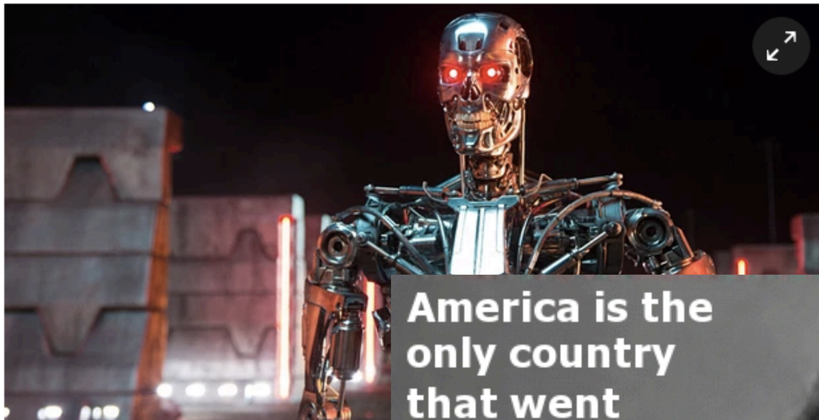

VB NEWS EVEN

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

We can fight ‘em Robots, one by one ☺

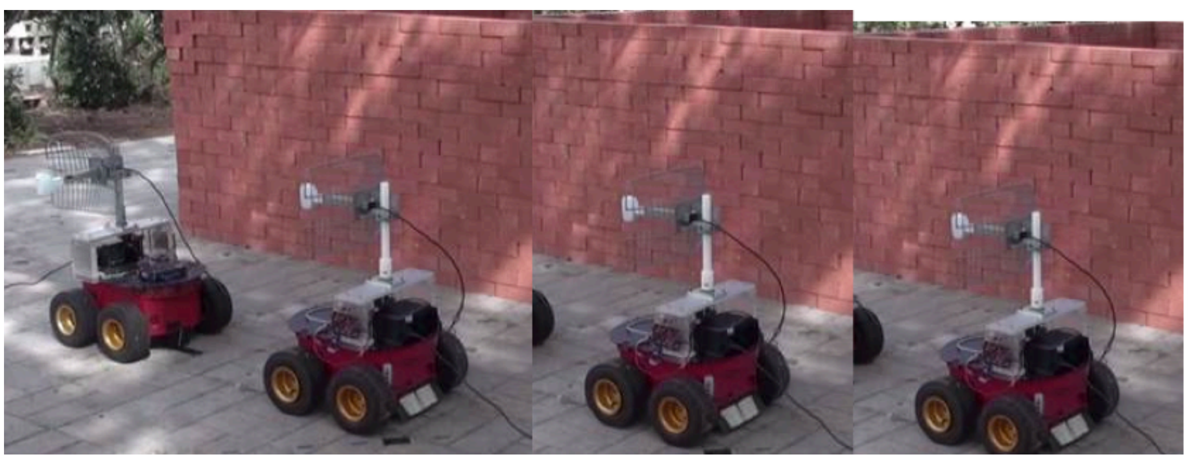

AI's Curious Ambivalence to humans..

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

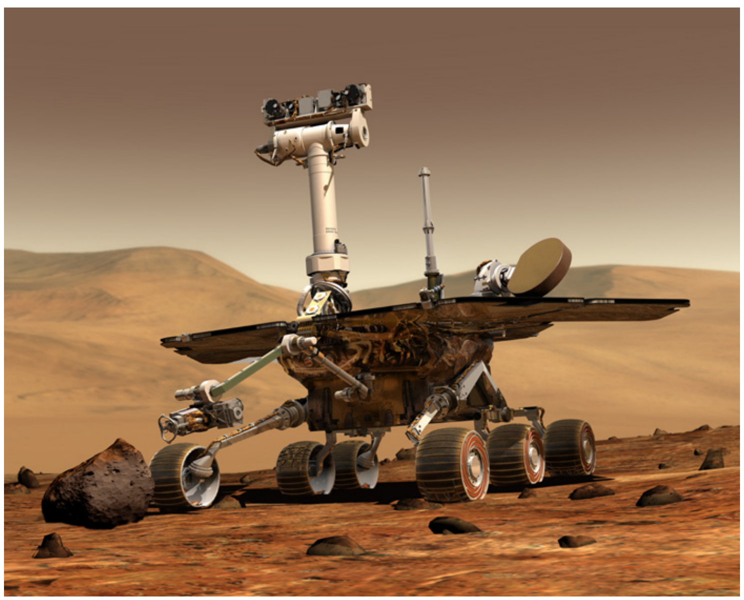

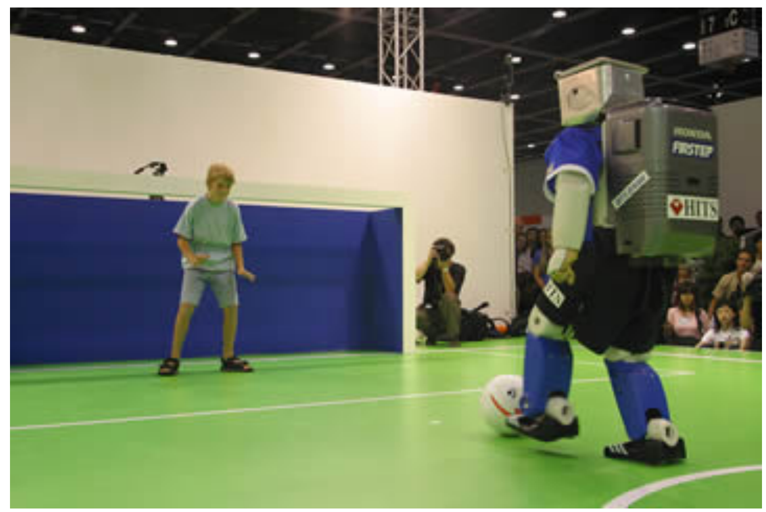

You want to help humanity, it is the people that you just can't stand…

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

But isn’t this cheating?

- Doesn’t putting human in the loop dilute the AI problem?
- Won’t it be cheating?
  - Like the original Mechanical Turk…

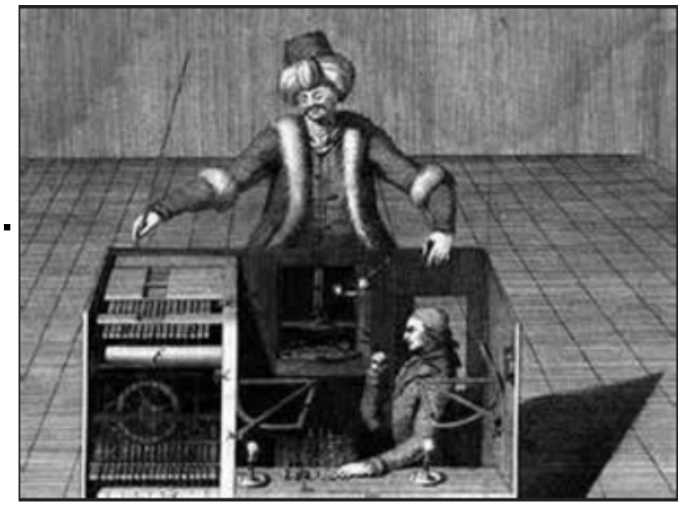

Many Intelligences..

- Perceptual & Manipulation intelligence that seem to come naturally to us
  - Form the basis for the Captchas..
    - But rarely form the basis for our own judgements about each other's intelligence
- Emotional Intelligence
- Social Intelligence

- Cognitive/reasoning tasks
  - That seem to be what we get tested in in SAT etc.

What is AI?

Spock or Kirk?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

- By dubbing "acting rational" as the definition of AI, we carefully separated the AI enterprise from "psychology", "cognitive science" etc.
- But pursuit of HAAI pushes us right back into these disciplines (and more)
  - Making an interface that improves interaction with humans requires understanding of human psychology..
    - E.g. studies showing how programs that have even a rudimentary understanding of human emotions fare much better in interactions with humans

Architecture of an Intelligent Agent

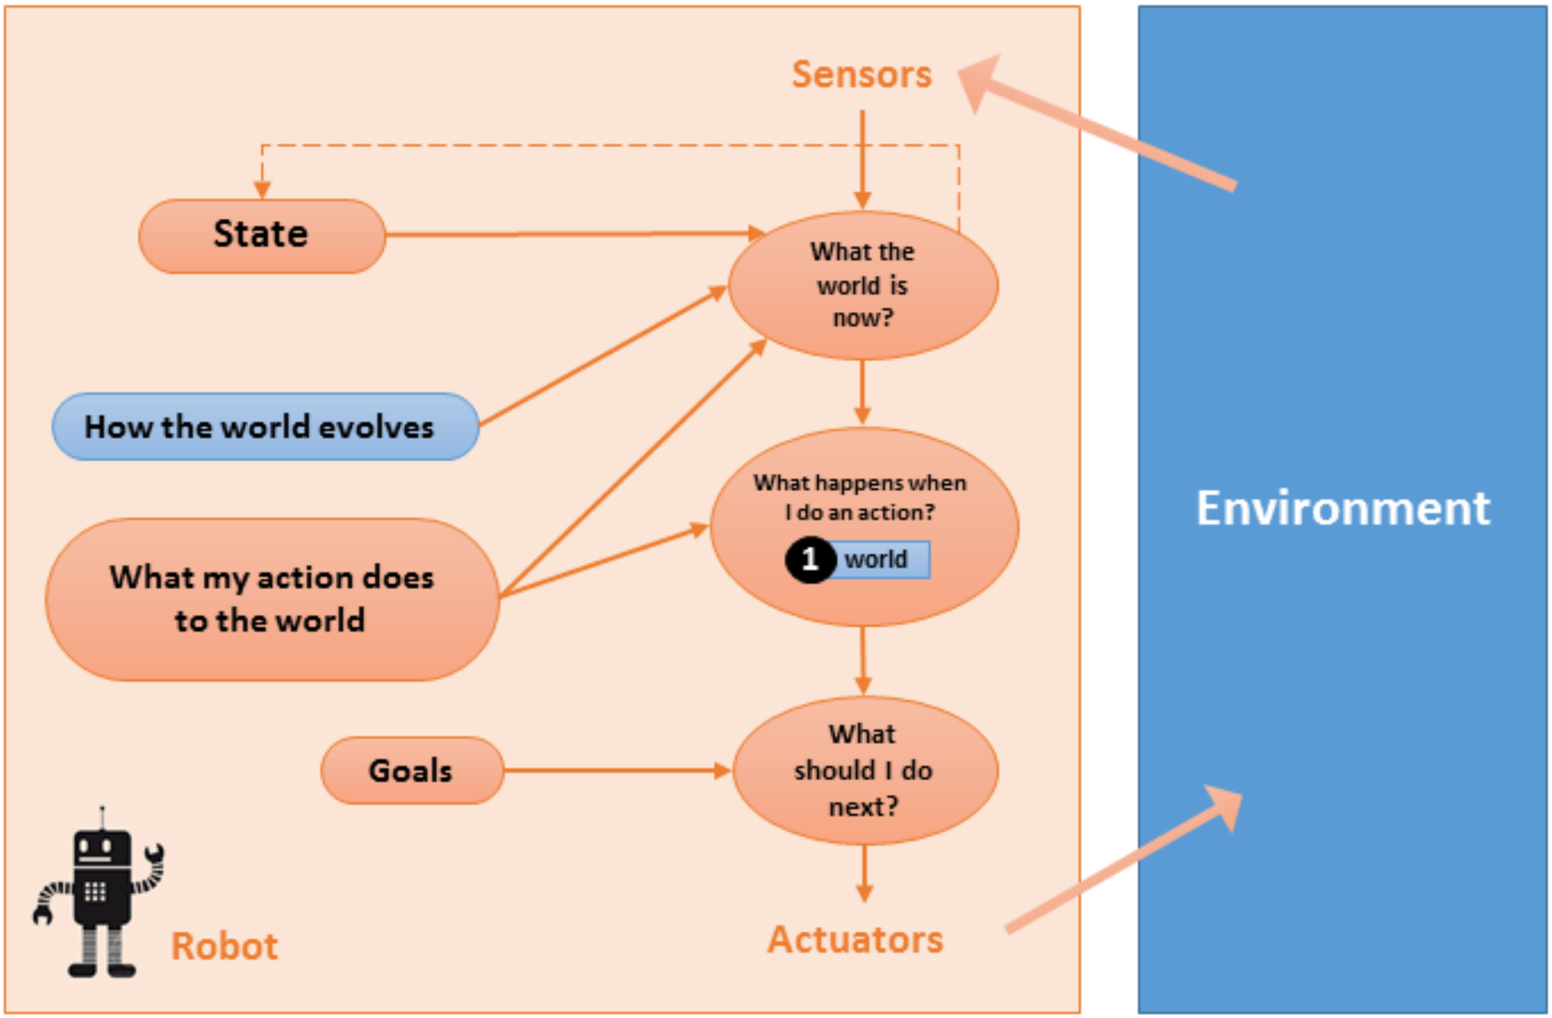

Architecture of an Intelligent Agent teaming with a human

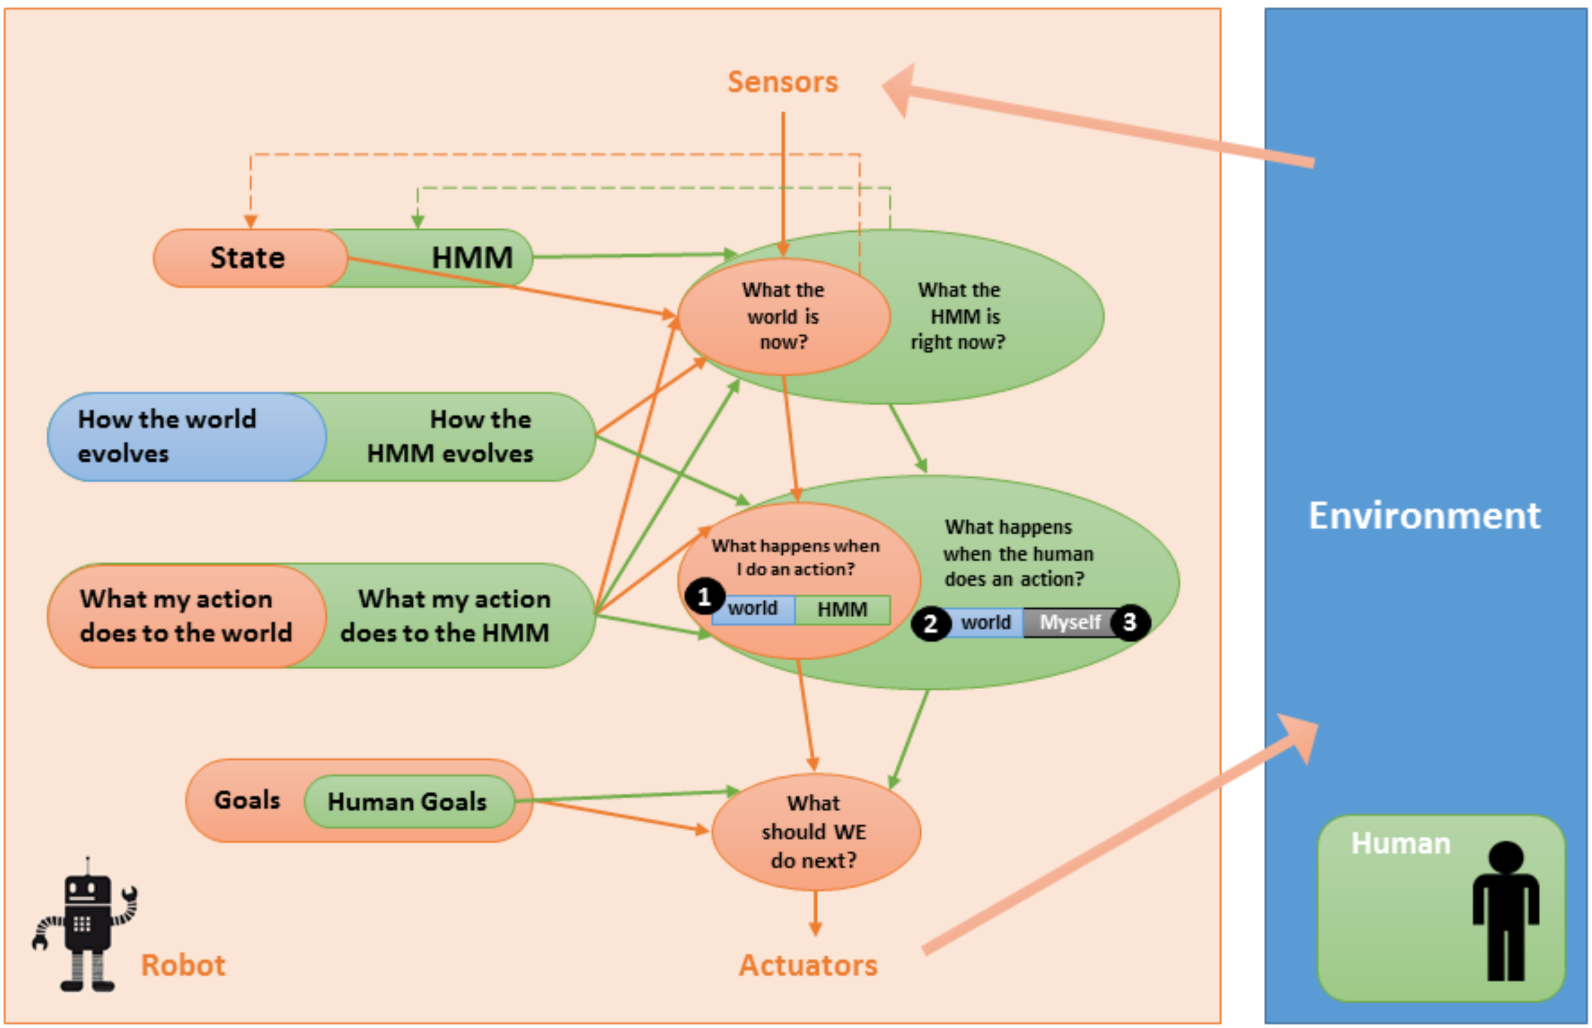

HMM= Human Mental Model

Use Case Scenarios: One Robot & One Human

Prediction: H is about to break a door open when R notices H's intention and predicts that breaking the door open will cause a board will fall on H. R thus moves to catch the board preventively.

Capability models: R notices that a heavy object blocks the entrance to a hallway that H wants to explore. Based on its capability model of H (i.e., what H can and cannot lift) and H's goal, R decides to interrupt its current activity and move the block out of the way.

Mitigate Risk: H needs to search a building for wounded people but is uncertain about the structural integrity and worried that parts might collapse. After communicating this to the robot, the robot propose a plan that has the robot go in first to assess the risk better and then to split the search in ways that minimize human risk

Anticipation: R is tasked to wait outside a building and watch out for enemies while H is performing a search inside. After more time has elapsed than it would take to perform the search, the R decides to go inside and find H to help H in case H has encountered any problems.

Normative behavior: R has an order to deliver medical supplies to H, but notices on the way a wounded victim that needs medical attention. R decides that caring for the victim is most important that delivering the supplies and notifies H of the delay.

Coordination based on Mental Model Inference: R learns that H needs to get a medical kit to be able to triage a recently discovered victim. R knows that H is aware that a medical kit is located in a particular room, but infers that H is unaware that R has a medical kit that H could use. While R cannot directly deliver its medical kit or wait for H due to other commitments, it can place its medical kit along the hallway where it expects H to go in order to get the first medical kit, thus relying on H's ability to notice the available medical kit and pick it up instead of the other remotely located kit.

[With Matthias Scheutz]

Use Case Scenarios: Multiple Robots & Humans

Belief revision: H1 and H2 are working in two different areas each assuming that the other will take care of a third area. R detects that discrepancies in its mental models of both humans in conjunction with its observations and decides to work on the third area (alternatively, R informs both H1 and H2 about the discrepancy).

Mental state inference: R notices that H1 cannot see H2 is approaching with equipment that H1 needs. R further observes that H1 is about to talk to another person and infers that this might be to order the urgently needed equipment from another person. Hence, R contacts H1 directly and informs H1 of H2's arrival.

Workload: R knows that H1 has currently high workload (e.g., from running simulations of H1's current activities based on R's model of H1's performance obtained from prior training) and thus does not interrupt H1 with a request from H2 that can wait, but communicates to H2 that it will take care of the request later.

Social regulation: R notices an escalation in the interaction between H1 and H2 about how to best proceed, where H1 and H2 each propose different plans. To mitigate, R proposes a compromise plan that contains elements of both H1's and H2's proposals (a "social" solution)

Activity coordination through shared mental models: Two humans H1 and H2 each work with a robotic teammate R1 and R2 in a first responder scenario in the "hot zone" of a natural disaster. H1 and R1 work on the fair side of the designated area, while DHT2 and R2 begin working in an area closer to the boundary. When R1 arrives with H1 at the designed area, R1 notices that not all the necessary equipment is available and communicates with R2 about the availability of the missing items. R2 quickly predicts equipment needs and anticipates that those items are not needed for a while. After quickly getting the OK from H2 to lend the equipment to R1, R2 drives off to meet R1 half-way, exchanges the equipment, and R1 returns to H1 in time to be able to continue triaging the victims with the missing equipment (which H1 did not even notice). Once the equipment is no longer needed, R1 meets up with R2 again, returning the equipment in time for H2 to have it available.

[With Matthias Scheutz]

Symbols or Neurons?

- "A physical symbol system has the necessary and sufficient means for general intelligent action.

  --Allen Newell &

  Herbert Simon

- Symbols are Luminiferous Aether of AI

  —Geoff Hinton

Interpretable AI…
It is not just for Safety!
(Symbols/Neurons Redux)

- We humans may be made of neurons, but we seem to care a “lot” about comprehensibility and “explanations”
- If we want AI systems to work with us, they better handle this
  - This is an important challenge for the neural architectures
    - What do those middle layers represent?
      - DARPA Initiative on **XAI**..
- Not just explanations, but explicable behavior!

Why intentionally design a dystopian future and spend time being paranoid about it?

Factoid: The country with the largest number of paper submissions: China!

China was also a close second in number of papers accepted!

Agenda

- Part I: The Path to General AI goes through Human-Machine Collaboration
  - ..and it is a good thing!
    - Expands reach and scope of AI enterprise
    - Reduces some of the off-the-top worries about AI
    - Brings up novel research challenges
- Part II: Planning Challenges in Human-Machine Collaboration
  - Brief review of how the planning problem “expands” in the face of interaction/teaming with humans
  - Specific challenges and some ongoing work in my group

Planning Involves Deciding a Course of Action to achieve a desired state of affairs

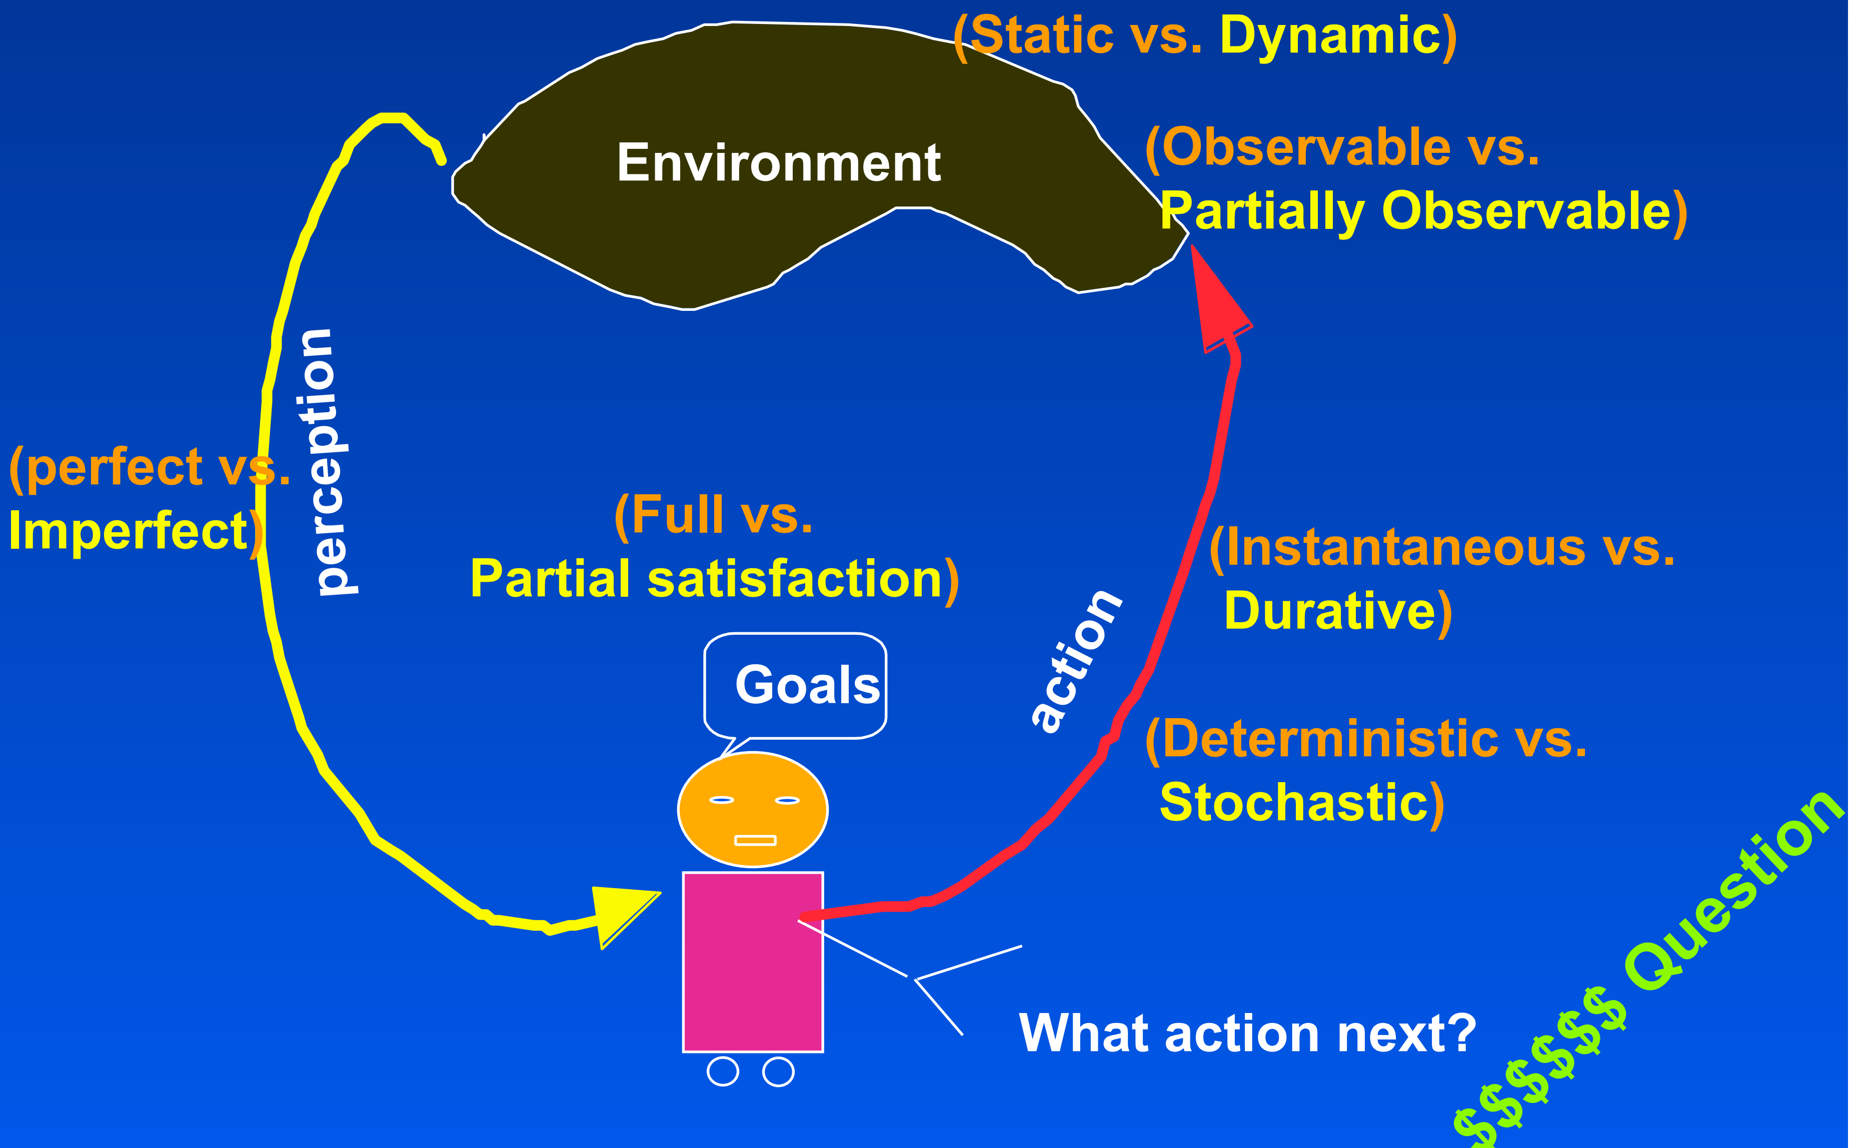

Planning: The Canonical View

Full Problem Specification

PLANNER

Fully Specified Action Model

Fully Specified Goals

Completely Known (Initial) World State

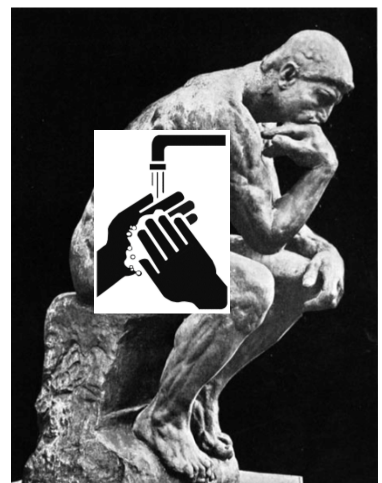

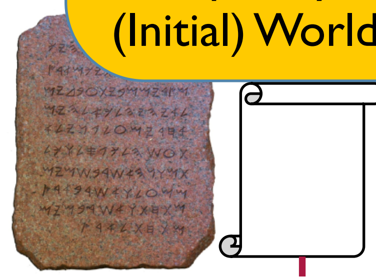

Plan (Handed off for Execution)

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 *from* humans
    - Mixed-initiative planning; “Symbiotic autonomy”
  - Needs to team with them
    - Human-robot teaming; Collaborative planning

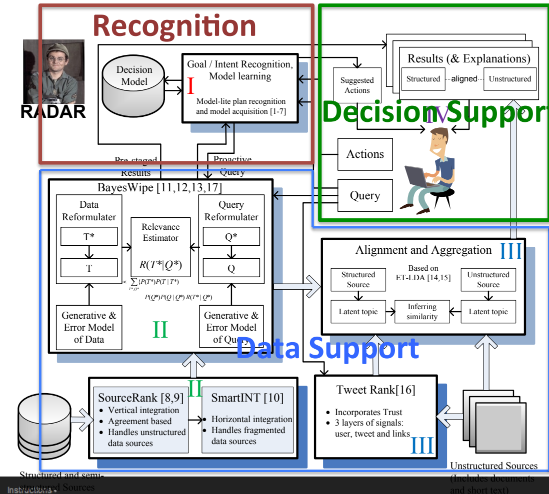

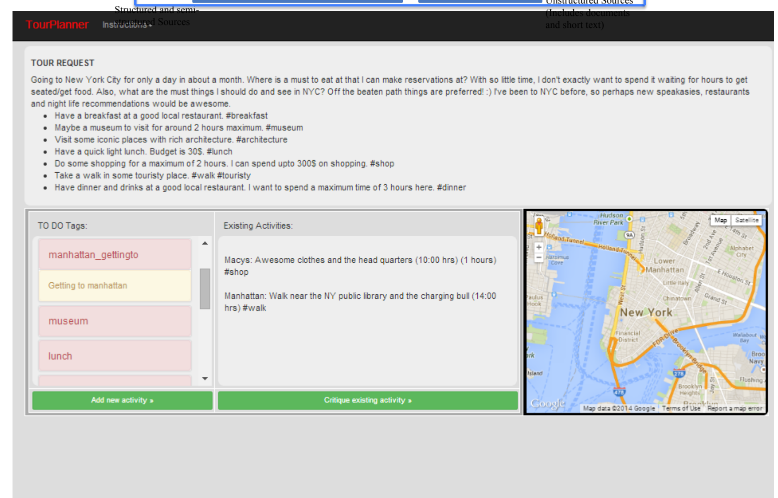

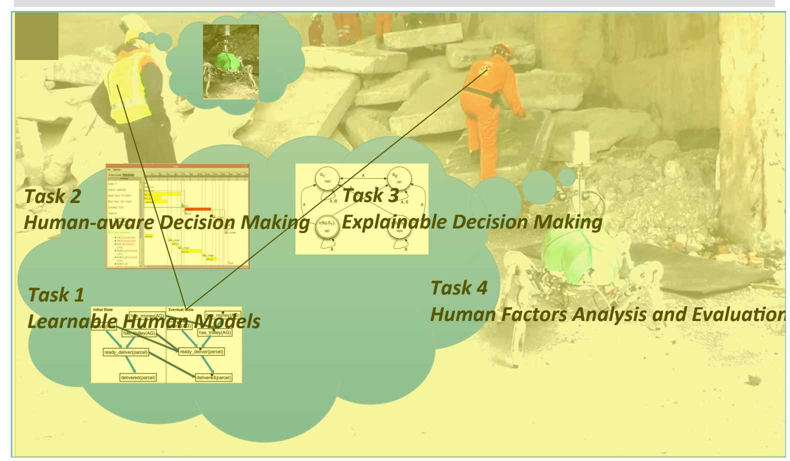

~~Planning: The Classical View~~

Fully Problem Specification

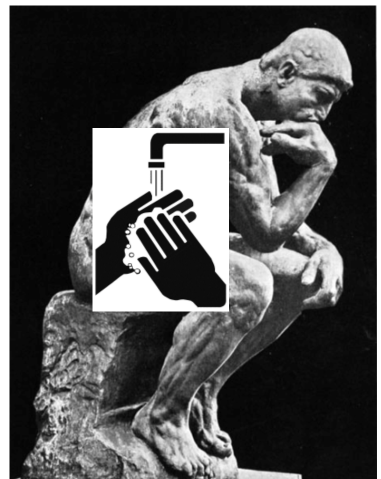

PLANNER

Fully Specified Action Model

Fully Specified Goals

Completely Known (Initial) World State

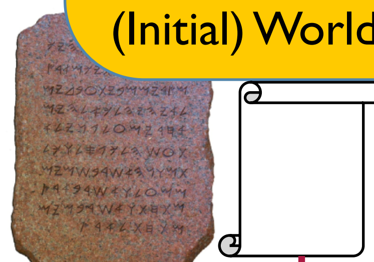

Plan (Handed off for Execution)

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

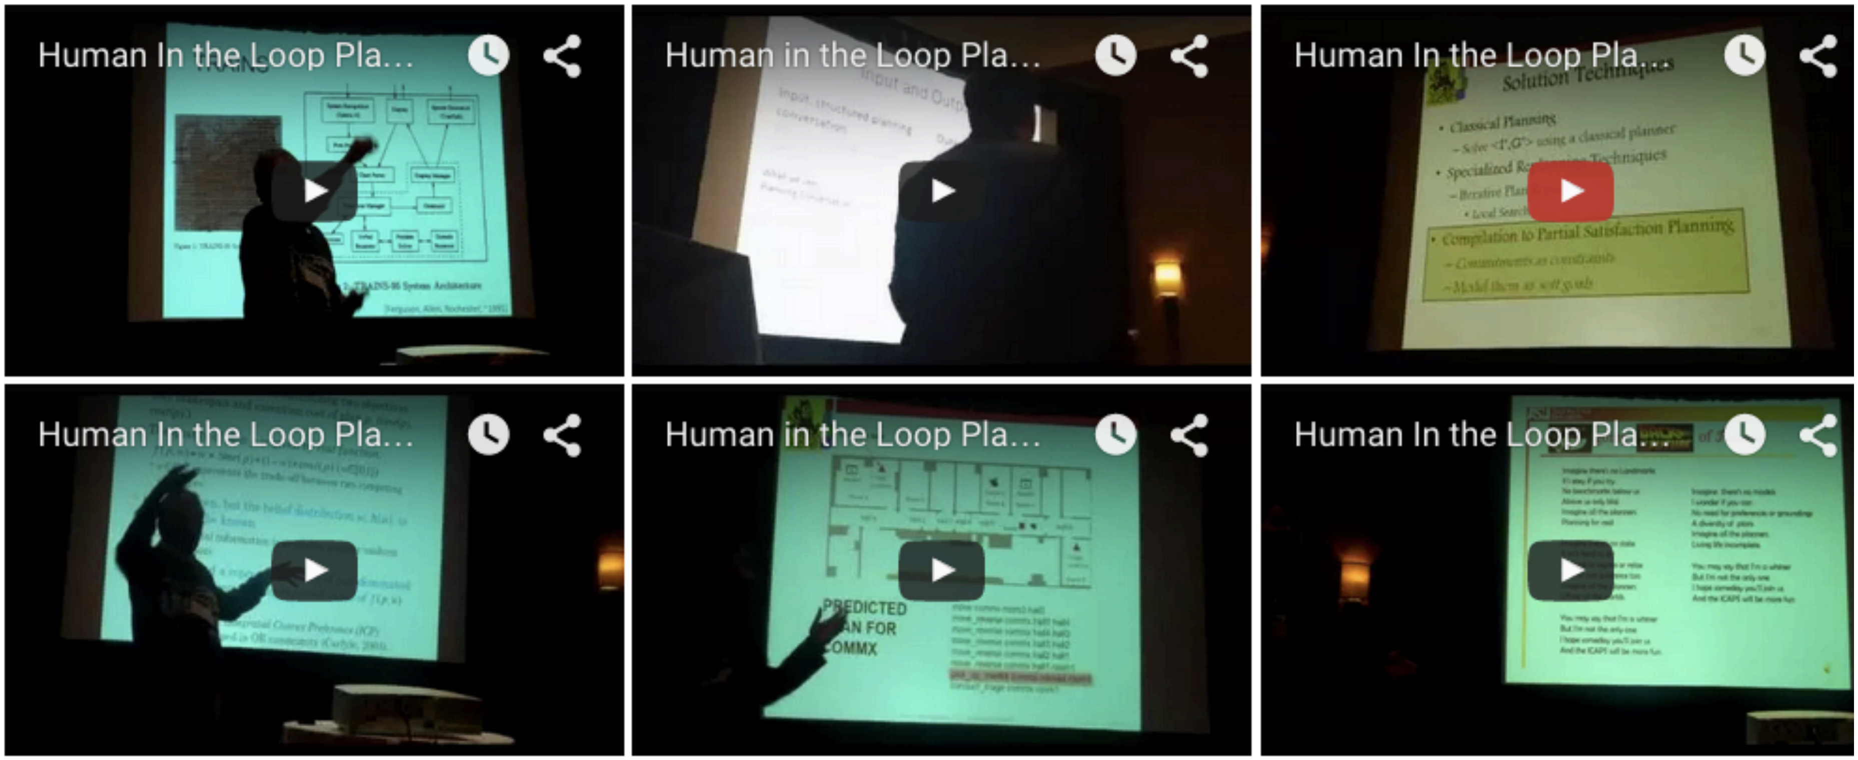

Materials

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

gratefully acknowledged

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
- Explicable Behavior, Explanations/Excuses (Interaction/Communication)
  – How should the human and robot coordinate
- Understand effective interactions between humans and machines (Evaluation)
  – Human factor study

Overview of our ongoing work

- 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 explicable 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.

Overview of our ongoing work

- 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 explicable 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” ☹

Effective ways to handle the more expressive planning problems by exploiting the deterministic planning technology

--Classical planners have become the *de facto* substrates for P-Space Complete problems..

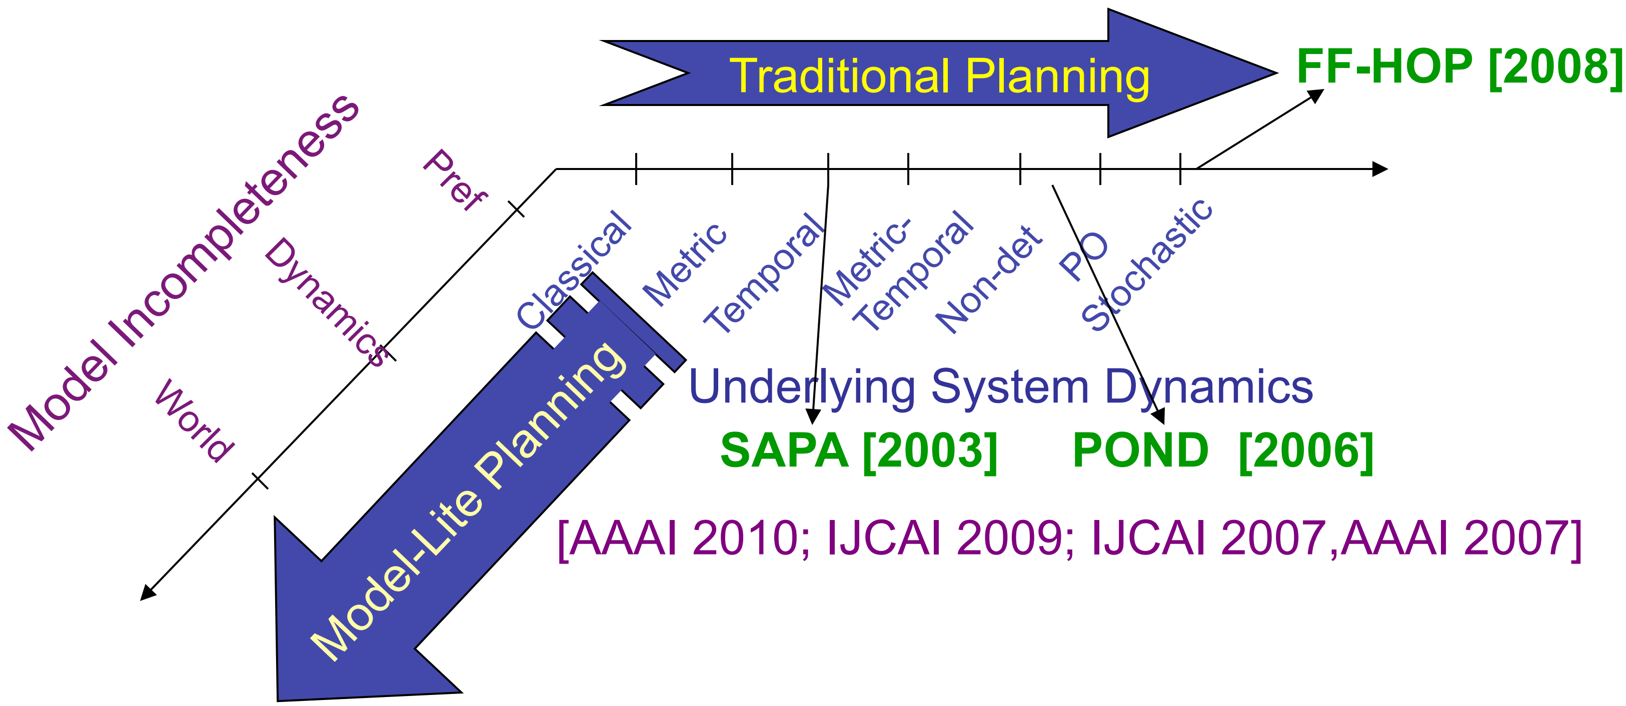

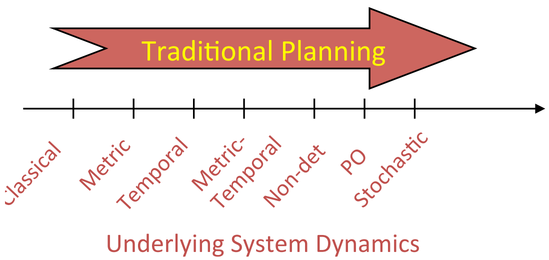

Spectrum of Domain Models

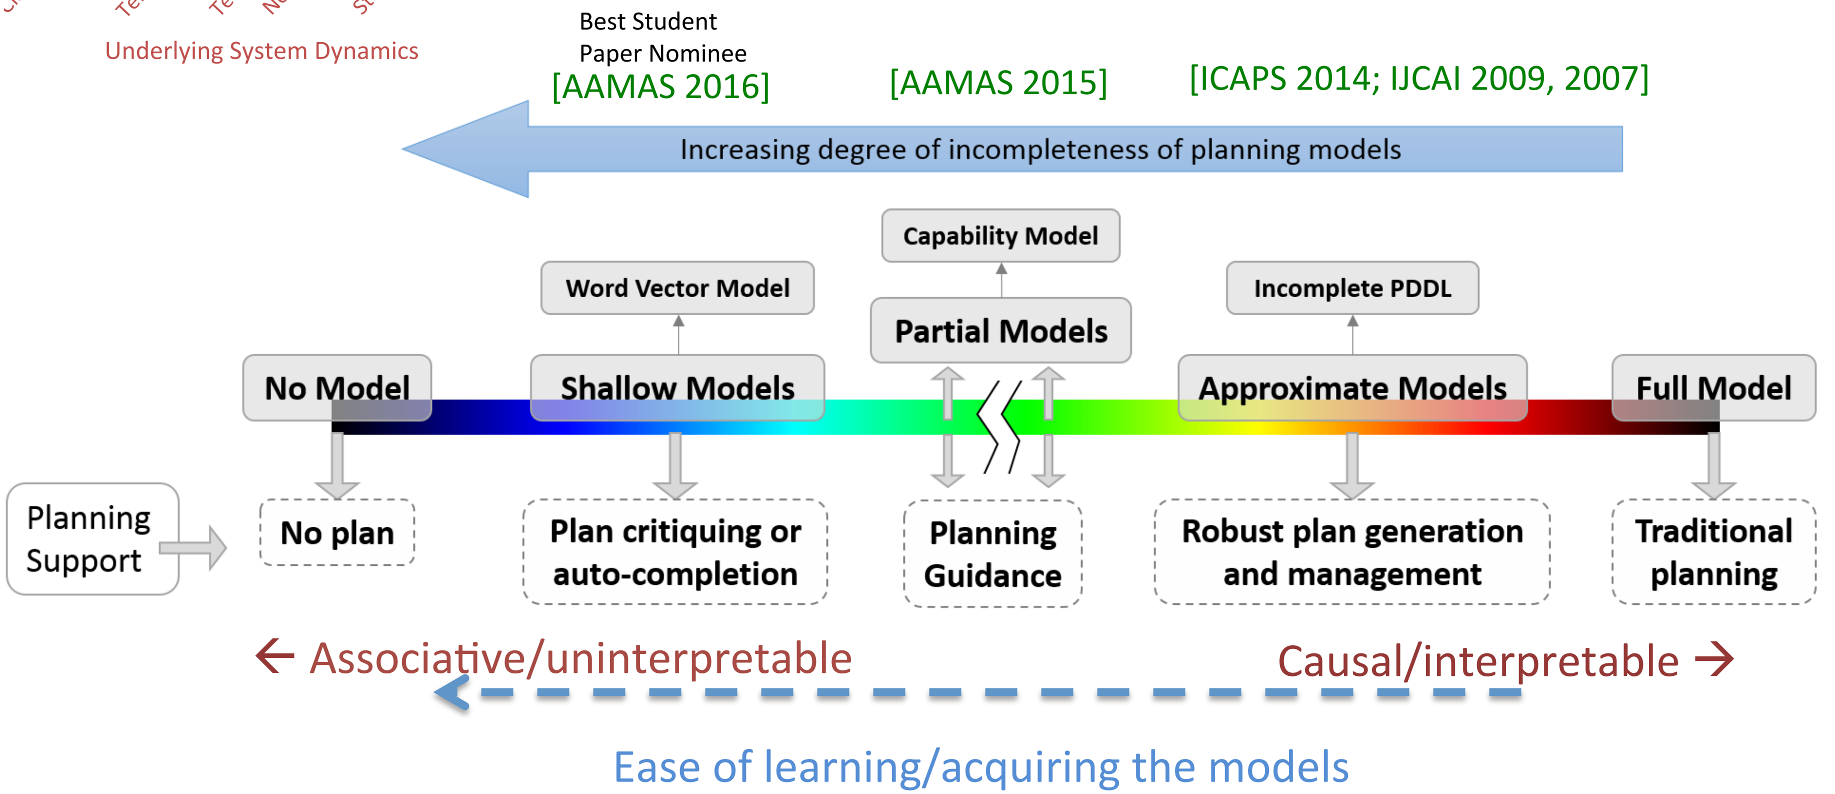

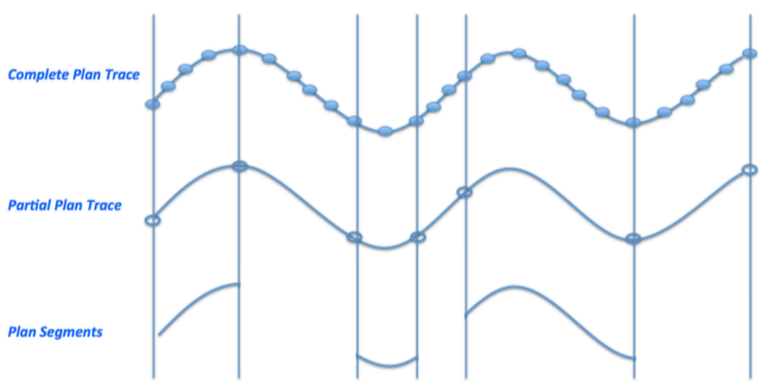

Note the contrast to ML research where the progress is going from uninterpretable/non-causal models *towards* interpretable and causal models.
So we might meet in the middle!

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

<ins>Example</ins>: 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.

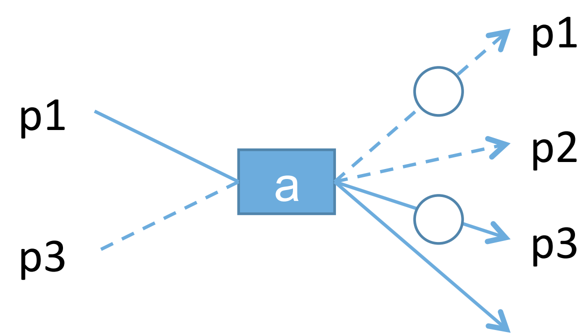

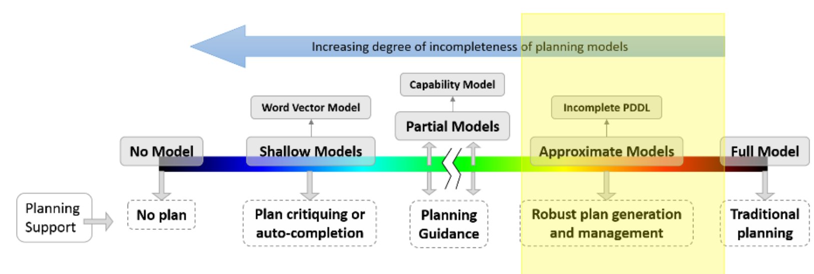

There are known knowns; there are things we know that we know. There are known unknowns; that is to say, there are things that we now know we don't know. But there are also unknown unknowns; there are things we do not know we don't know.

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

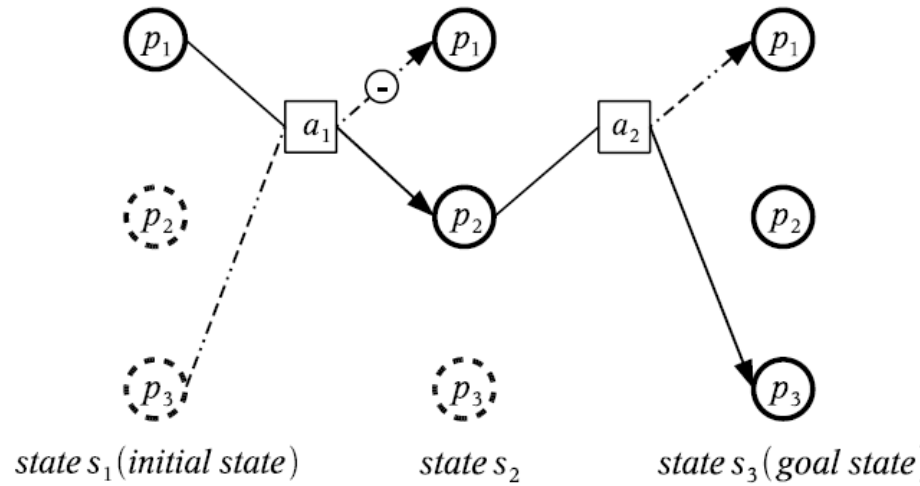

Candidate models of plan	1	2	3	4	5	6	7	8
a_1 relies on p_3	yes	yes	yes	yes	no	no	no	no
a_1 deletes p_1	yes	yes	no	no	yes	yes	no	no
a_2 adds p_2	yes	no	yes	no	yes	no	yes	no
Plan status	fail	fail	fail	fail	succeed	fail	succeed	succeed

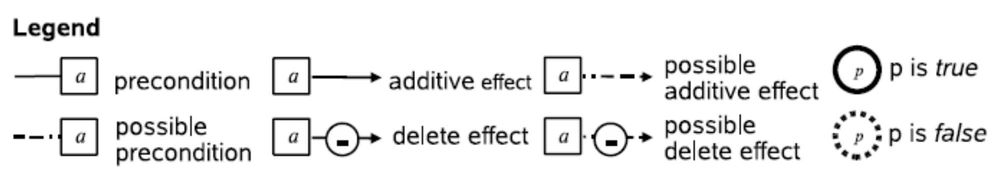

Robustness value: 3/8

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

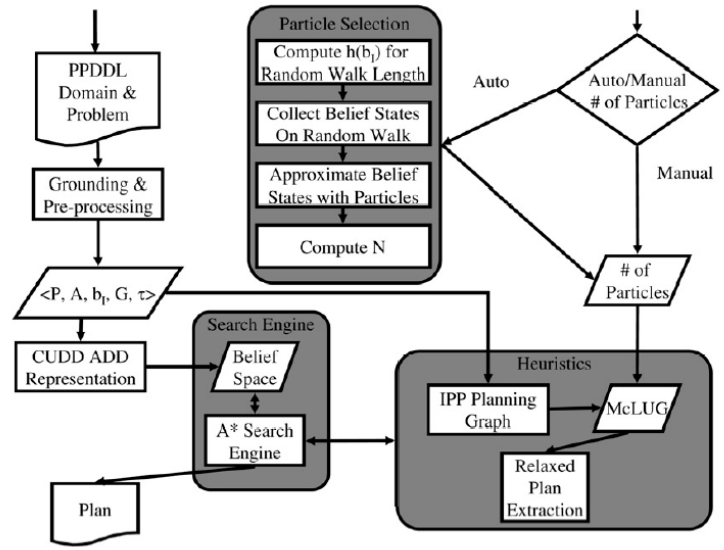

Fig. 6. *POND* architecture.

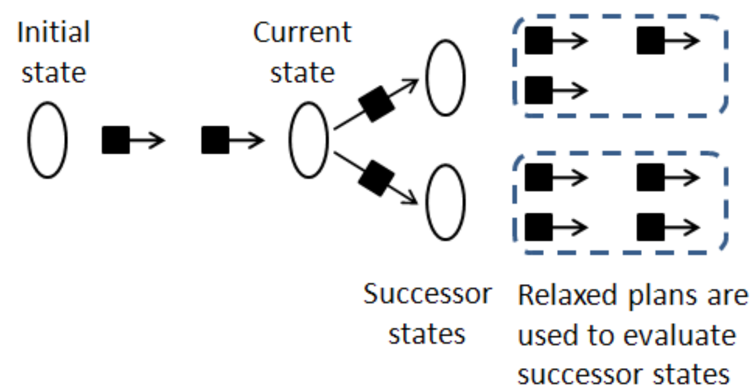

[Nguyen et al; NIPS 2013; Nguyen & Kambhampati, ICAPS 2014]

Synthesizing Robust Plans: A Compilation

Incomplete model
Complete world state

→

Complete model
Belief state

(Conformant Probabilistic Planning)

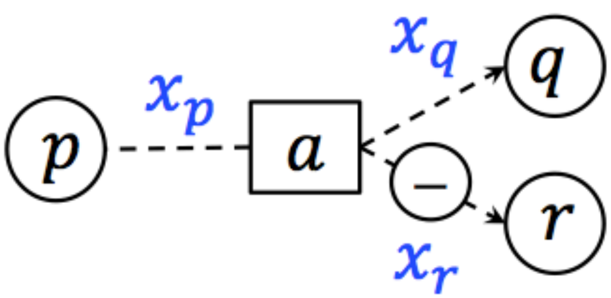

x_p (0.5) x_q (0.7) x_r (0.2)

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.

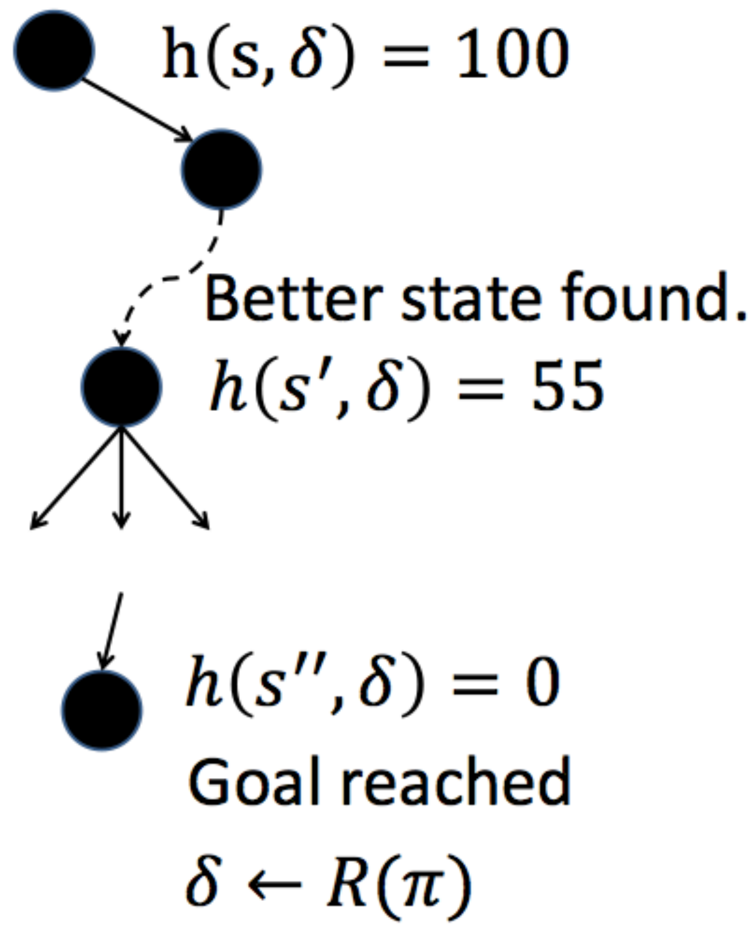

[ICAPS 2014]

Capability Model

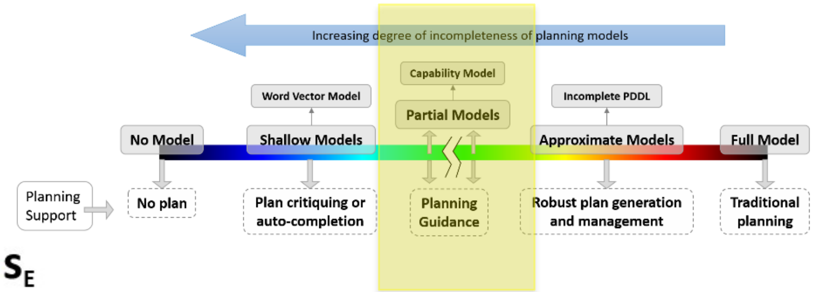

A capability:

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

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

T-gap capability model

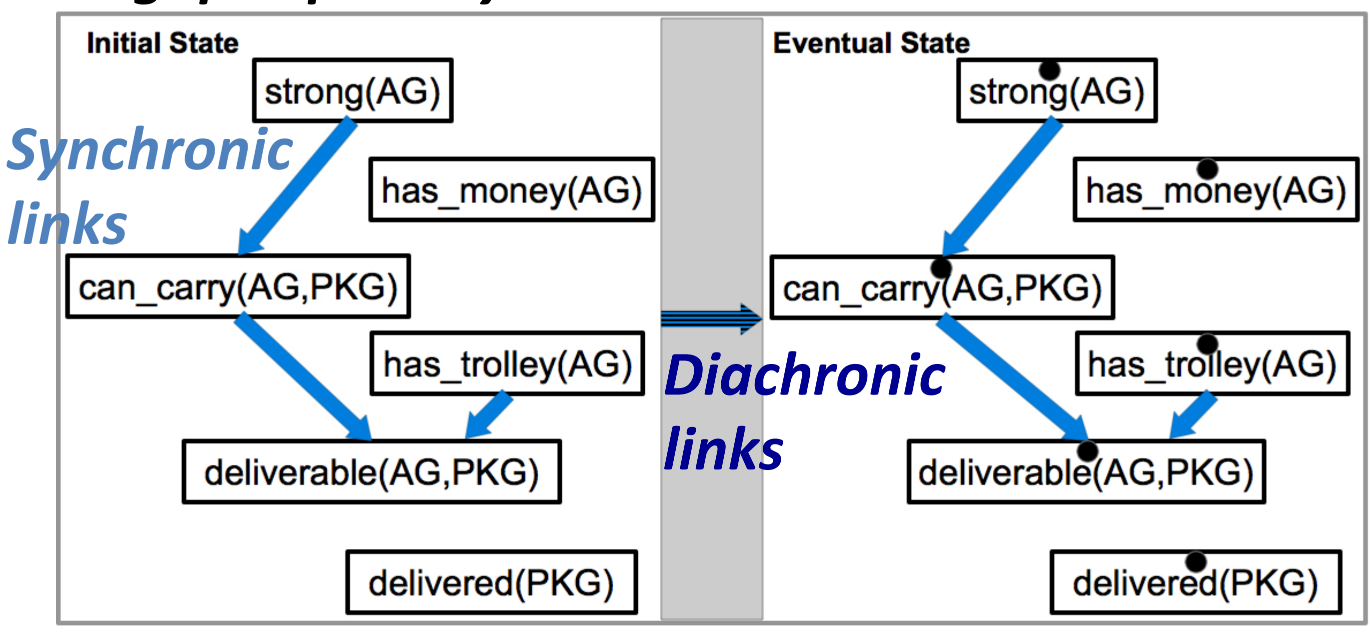

[AAMAS 2015]

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

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

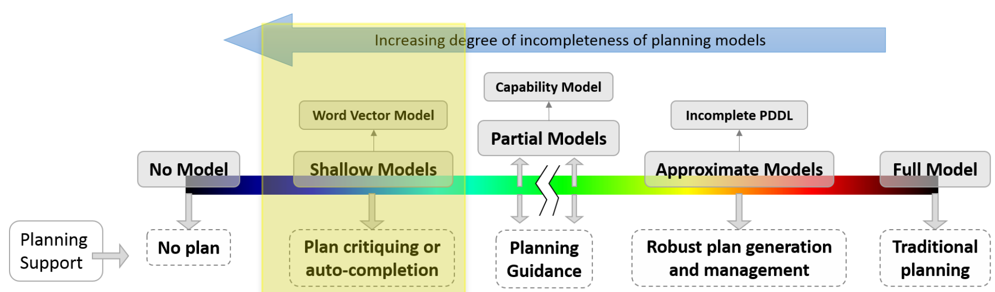

Learn vectors of actions

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

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

Learn vectors w_i for a_i in A by optimizing

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

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

Learning shallow models can avoid overfitting!!

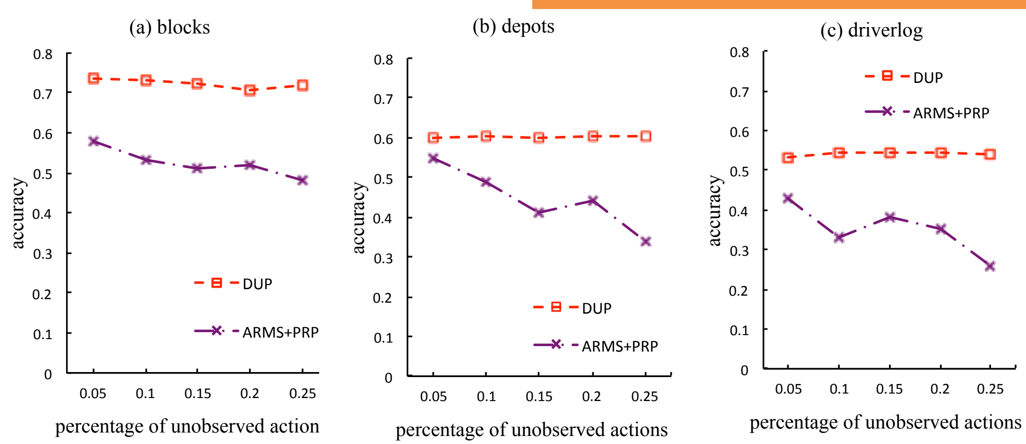

Algorithm 1 Framework of our DUP algorithm

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 $o \in \bar{\mathcal{A}}$, 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** $\tilde{p}$

Nominated for Best Student Paper Award at [AAMAS16]

Overview of our ongoing work

- 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” ☹

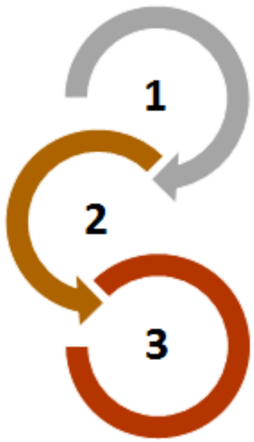

Model and predict human's intentions
- belief modeling and plan recognition

Obtain and represent relevant information efficiently
- resource profiles

Inform the robot's planning process with this information
- Interaction constraints in IP-based planner

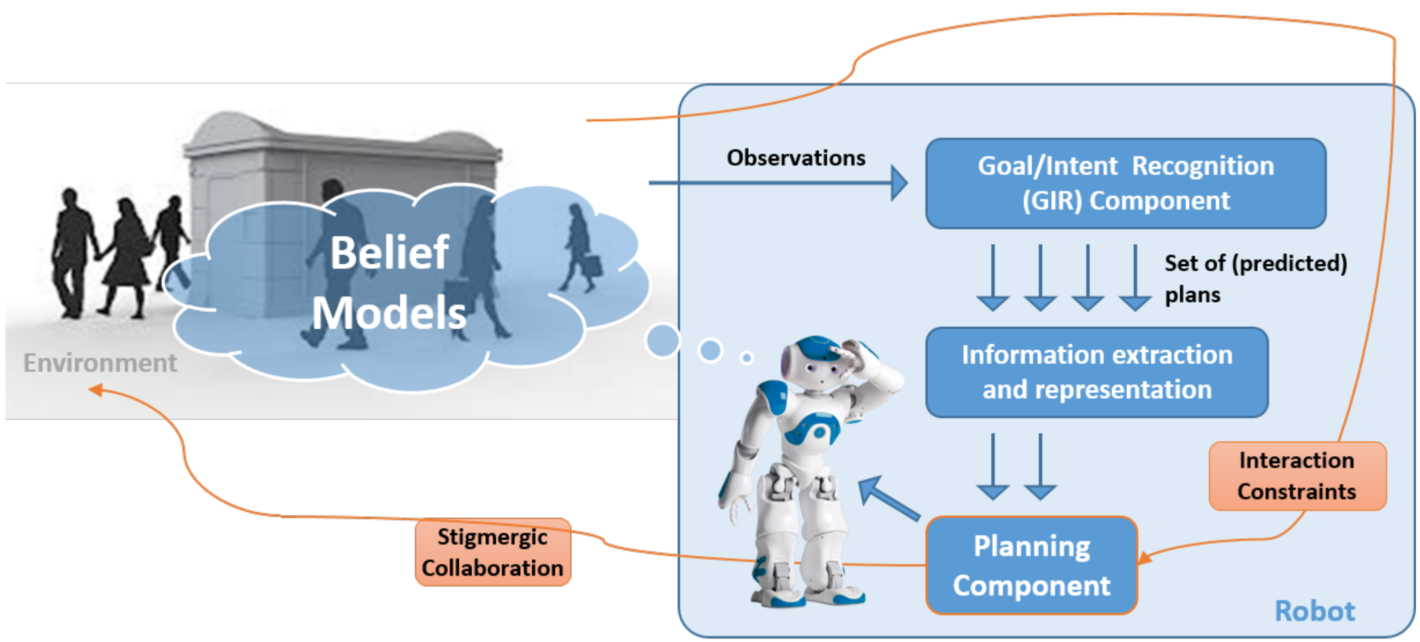

Planning for Human-Robot Teaming

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

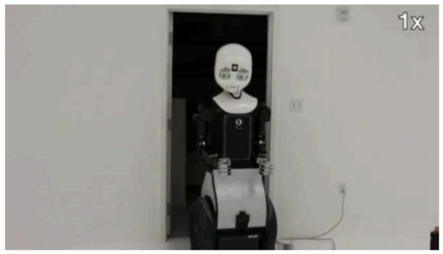

- When to stop sensing?
  - When does the planner know the world is closed?
- Why should the robot sense?
  - Does the object fulfill a goal?
  - What is the reward? Is it a bonus?

Talamadupula, Benton et al., ACM TIST 2010]

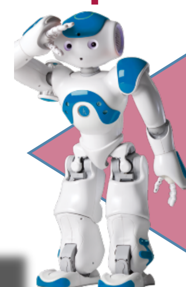

Planning for

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]

Coordinate with Humans
[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"

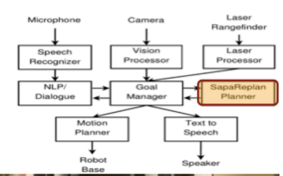

- NLP Module
  - i. Reference resolution
  - ii. Parsing
  - iii. Background knowledge
  - iv. Action submission (to planner)

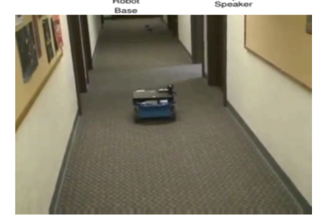

[Cantrell, Talamadupula et al., HRI 2012]

[In collaboration with hrilab, Tufts University]

Overview of our ongoing work

- 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 explicable 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" ☹

Explicability

Explicable behavior increases team fluency by reducing the need for post-facto explanations

Explicability → **Human's expectation of agent behavior** → **Hidden in human's mind**

Goal: Clear a cluttered table

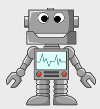

Move to the table

Push the items off the table

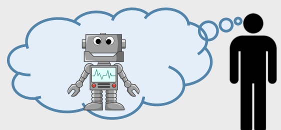

Move to the table

Keep books on shelf

Keep the coffee mug

Keep pen in pen holder

Table is clear

Explicable Plan Generation for Human-Robot Interaction

Anagha Kulkarni, Sarath Sreedharan, Tathagata Chakraborti, Yu Zhang and Subbarao Kambhampati

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!

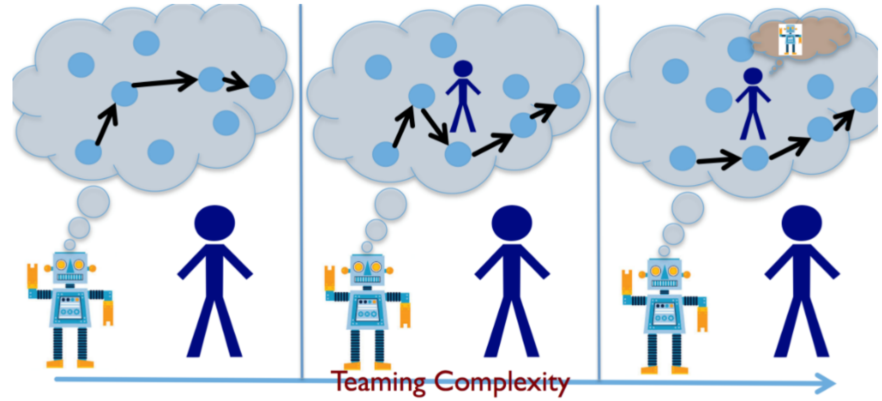

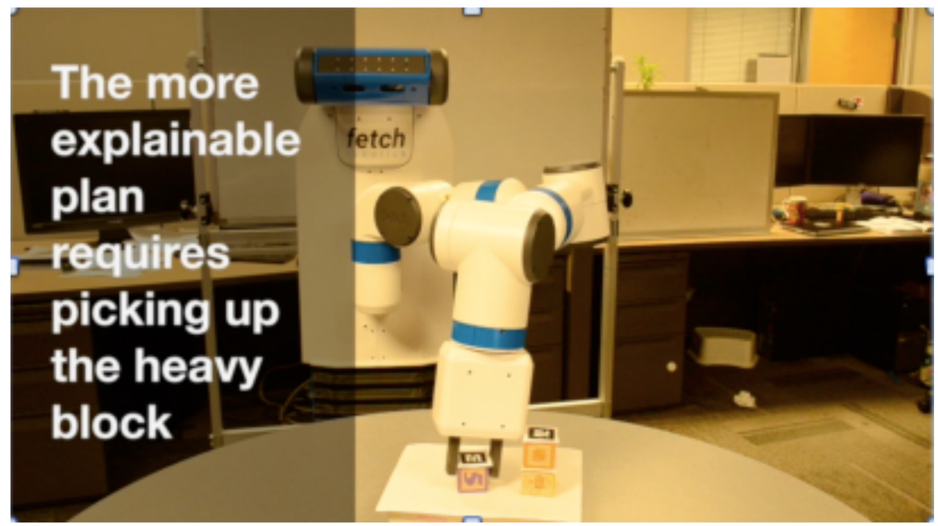

$$\underset{\pi_{M_R}}{\operatorname{argmin}}\, cost(\pi_{M_R}) + \alpha \cdot dist(\pi_{M_R}, \pi_{M^*_R})$$

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

Learning Human Expectation via Explicability Labeling

Problem: M^*_R is not known
Solution: Learn it, but indirectly
as a labeling scheme..

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

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

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

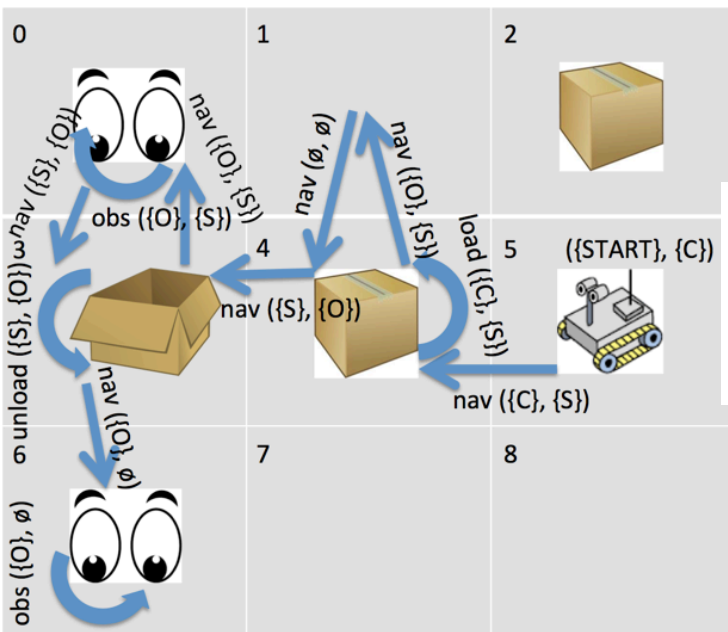

Analogy: Think of learning how to write address labels so the postal carrier can understand..

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

More than one label is allowed for actions

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

Learning the Labeling Scheme using CRF

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

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

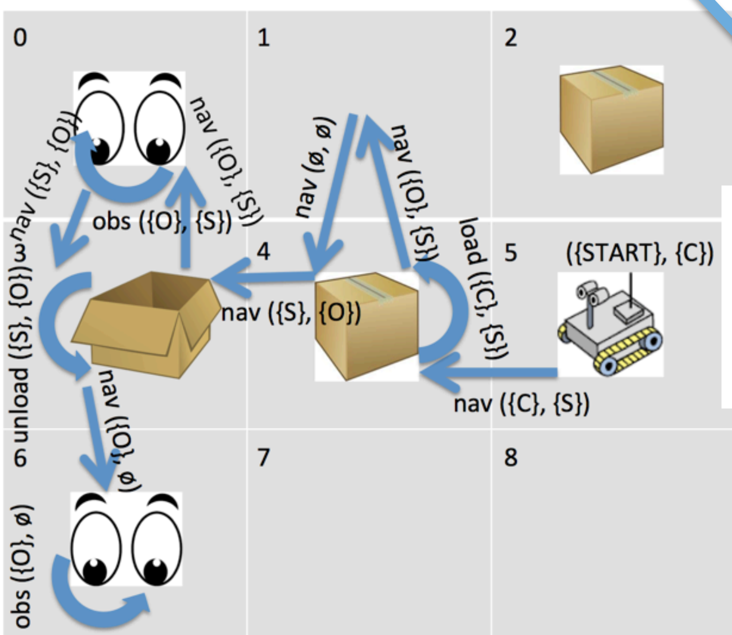

- **Features:**
  - Plan features: e.g., at rover L5
  - Action/trajectory Features: e.g., action type
  - Interaction features: e.g., distance to the human

- 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 Learned Model of Explicability

Preliminary results indicate that such a scheme is 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

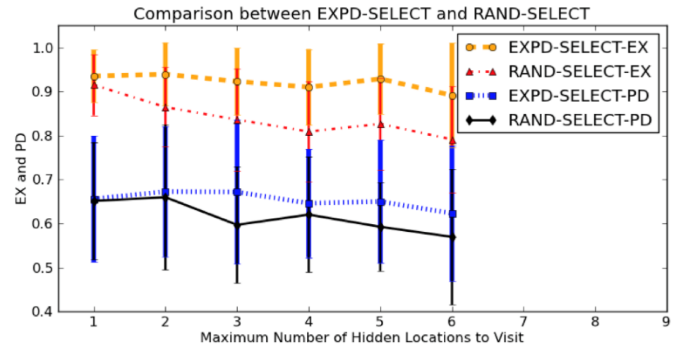

Figure 5: Comparison of EXPD-SELECT and RAND-SELECT

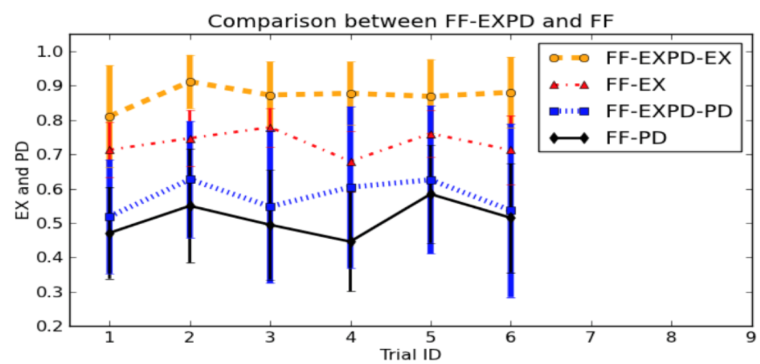

Figure 6: Comparison of FF-EXPD and FF considering $uexp$ in Alg. 1.

Bi-LSTM as Task Predictor for Plan Explicability

Motivation:

1. Consider future inputs.
2. Break Markov Property.

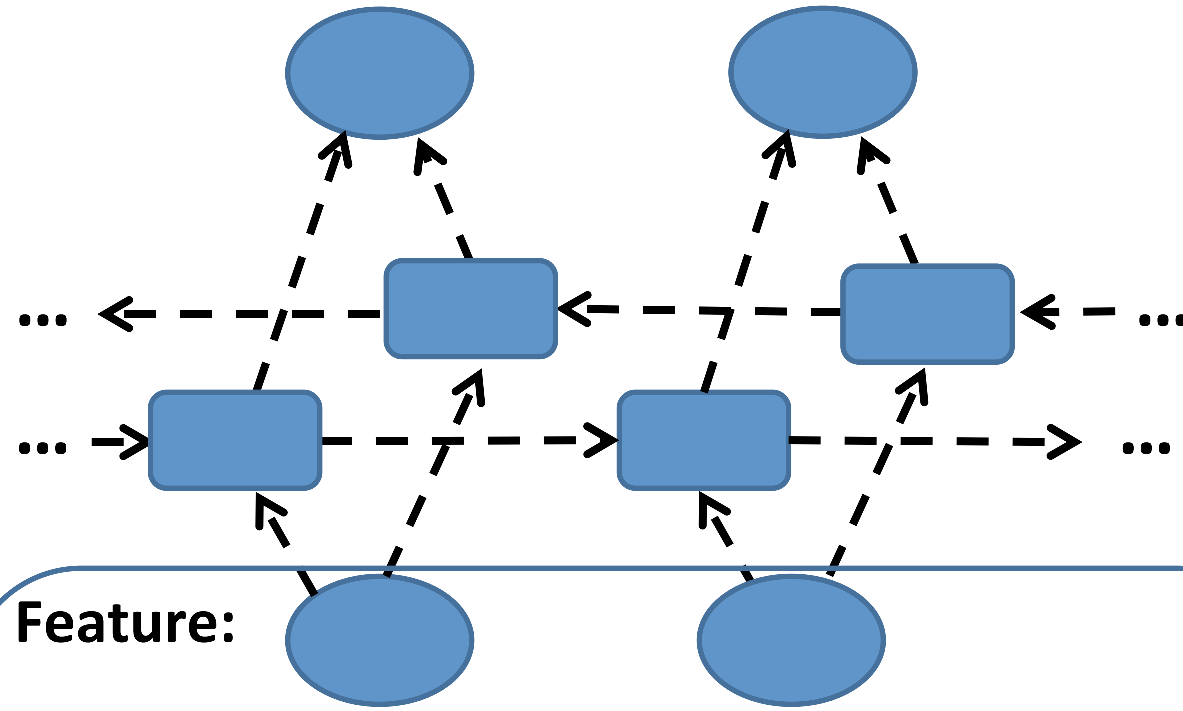

Feature:

61001001010000 1001

Action (0~N) + Executor (0-Human/1-Robot/ 2-Neither) + State (0010...)

noop	near-r1	at-b6-r1	at-b1-r1	...
0	0	1	0	...

Testing Accuracy: 90.76%

Overview of our ongoing work

- 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 explicable 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!

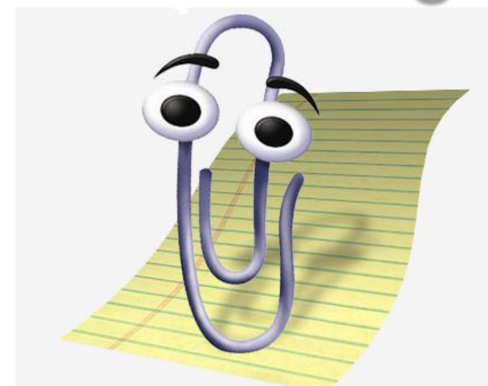

Human Factor Studies

- 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][HRI 2015]

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

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’:

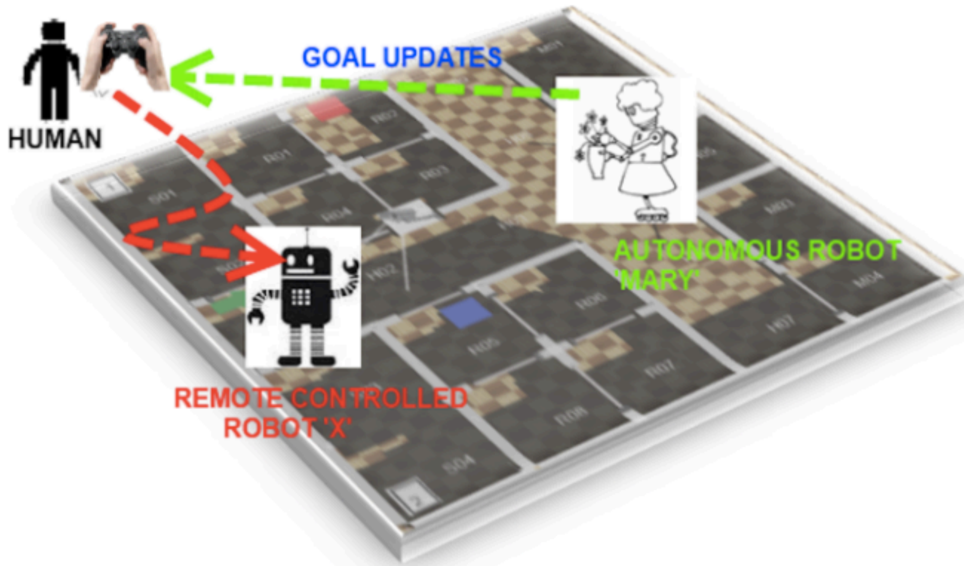

Findings

Robot 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

Summary of our ongoing work

- 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 explicable 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 of the talk

If you have questions,
WeChat:
Subbarao2z

- **Part I: The Path to General AI goes through Human-Machine Collaboration**
  - **..and it is a good thing!**
    - **Expands reach and scope of AI enterprise**
    - **Reduces some of the off-the-top worries about AI**
    - **Brings up novel research challenges**
- **Part II: Planning Challenges in Human-Machine Collaboration**
  - **Brief review of how the planning problem "expands" in the face of interaction/teaming with humans**
  - **Specific challenges and some ongoing work in my group**


If you have questions,
WeChat:
Subbarao2z
fetch

DETAIL SLIDES

(Not covered in the presentation)