Challenges of Human-in-the-Loop Planning and Decision Support

Subbarao Kambhampati
Arizona State University
(google: rao asu)

Talk given at U. Conn. 10/21/2014

Yochan Research Group

Automated Planning

- Foundations of automated planning
- Planning under a variety of domain models
  - Classical, temporal, stochastic, partially-observable
- Current Focus: Human-in-the-loop planning

Information Fusion/Integration

- Integrating structured and semi-structured data to support effective query processing
- Structured data: Probabilistic methods for imputation, rectification..
- Semi-structured data: Trust and Relevance based source selection
- Unstructured data: Analyzing and aligning social media data with event transcripts, sentiment analysis

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

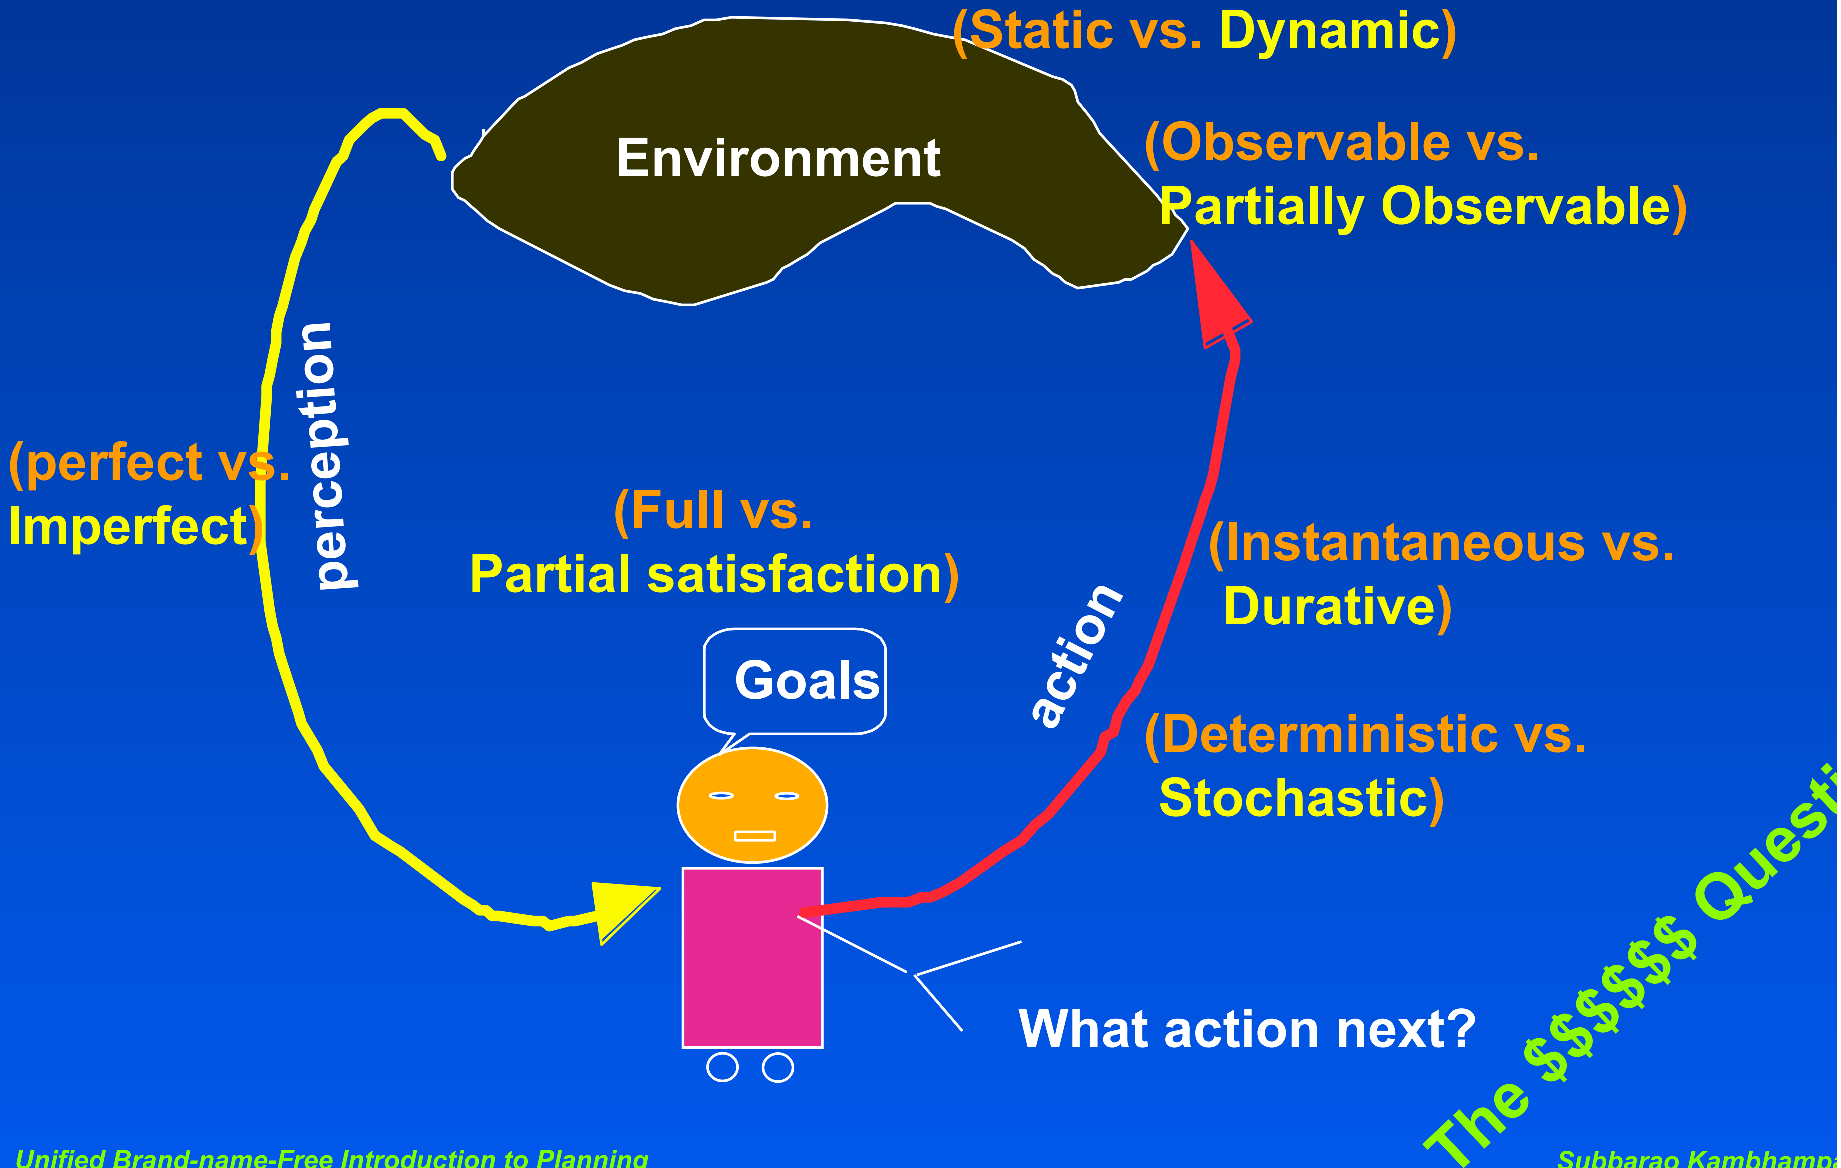

“Classical” Planning

```
(:action pick-up
        :parameters (?ob1)
        :precondition (and (clear ?ob1)
                           (on-table ?ob1)
                           (arm-empty)
                           (block ?ob1))
        :effect
        (and (not (on-table ?ob1))
             (not (clear ?ob1))
             (not (arm-empty))
             (holding ?ob1)))
```

Blocks world

State variables:
Ontable(x) On(x,y) Clear(x) hand-empty holding(x)

Initial state:
Complete specification of T/F values to state variables
--By convention, variables with F values are omitted

Goal state:
A partial specification of the desired state variable/value combinations
--desired values can be both positive and negative

Init:
Ontable(A),Ontable(B),
Clear(A), Clear(B), hand-empty

Goal:
~clear(B), hand-empty

Pickup(x)
Prec: hand-empty,clear(x),ontable(x)
eff: holding(x),~ontable(x),~hand-empty,~Clear(x)

Putdown(x)
Prec: holding(x)
eff: Ontable(x), hand-empty,clear(x),~holding(x)

Stack(x,y)
Prec: holding(x), clear(y)
eff: on(x,y), ~cl(y), ~holding(x), hand-empty

Unstack(x,y)
Prec: on(x,y),hand-empty,cl(x)
eff: holding(x),~clear(x),clear(y),~hand-empty

Ontable(A)
Ontable(B),
Clear(A)
Clear(B)
hand-empty

Pickup(A)

holding(A)
~Clear(A)
~Ontable(A)
Ontable(B),
Clear(B)
~handempty

Pickup(B)

holding(B)
~Clear(B)
~Ontable(B)
Ontable(A),
Clear(A)
~handempty

P-Space Complete

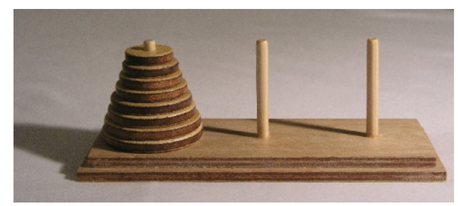

Scalability was the big bottle-neck…
We have figured out how to scale synthesis..

Problem is Search Control!!!

- Before, planning algorithms could synthesize about 6 – 10 action plans in minutes
- Significant scale-up in the last decade
  - Now, we can synthesize 100 action plans in seconds.

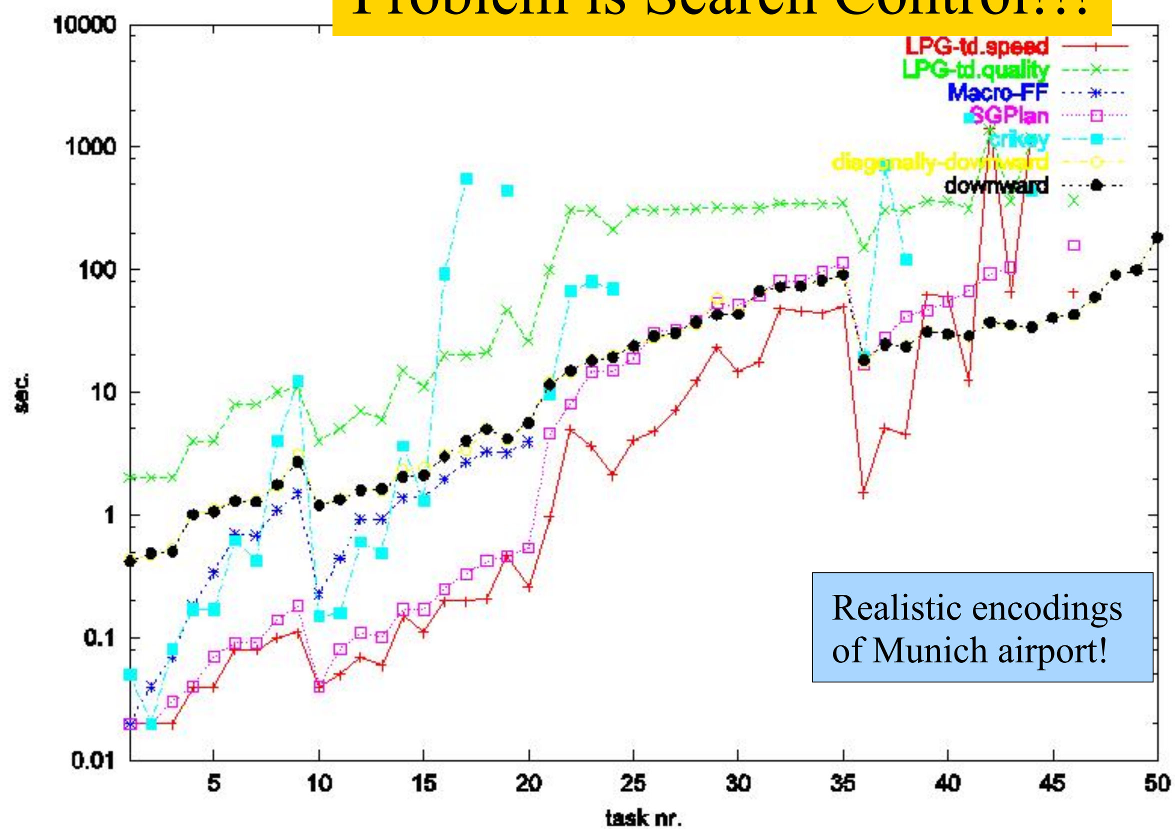

Realistic encodings of Munich airport!

The primary revolution in planning in the recent years has been **methods to** scale up plan synthesis

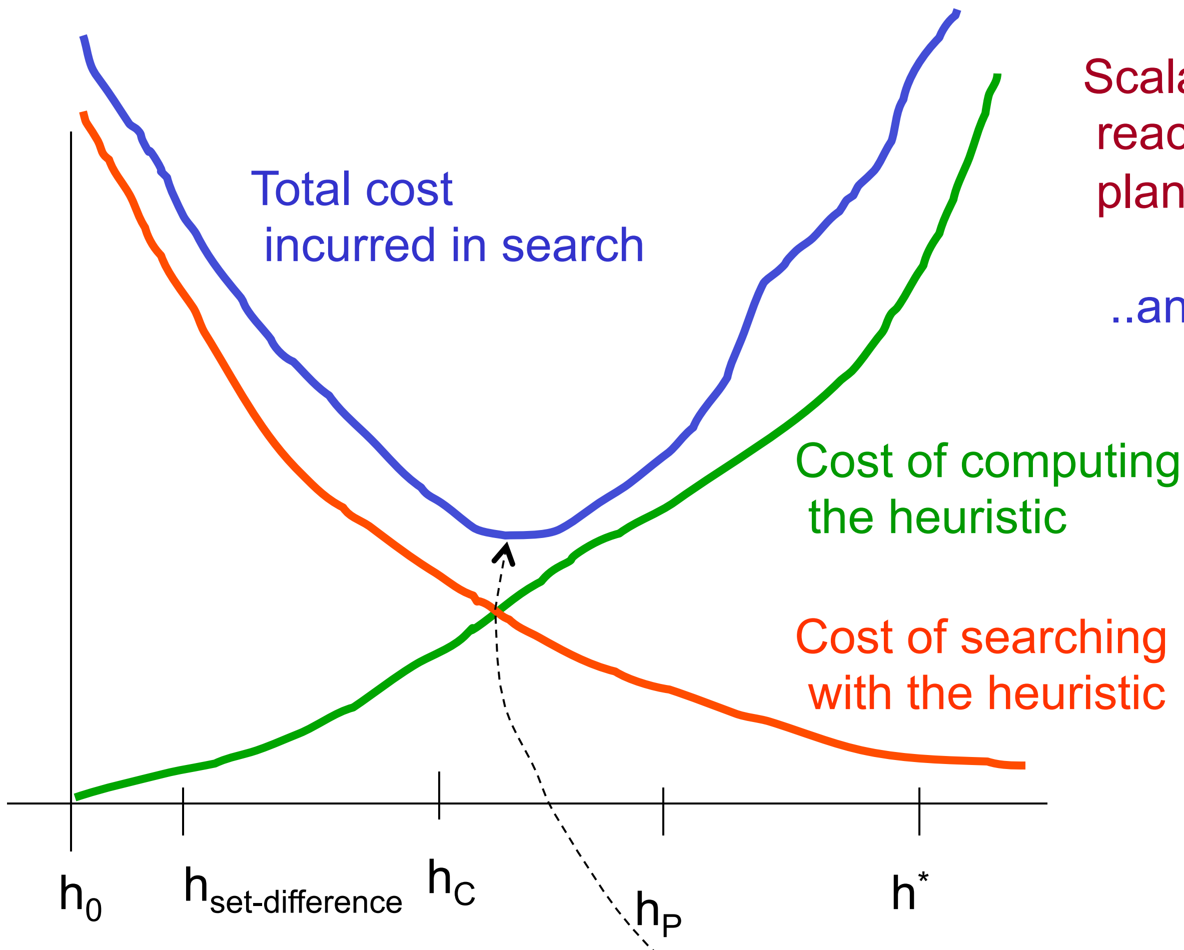

Scalability came from sophisticated reachability heuristics based on planning graphs..

..and not from any hand-coded domain-specific control knoweldge

"Optimistic projection of achievability"

Not always clear where the total minimum occurs

- Old wisdom was that the global min was closer to cheaper heuristics
- Current insights are that it may well be far from the cheaper heuristics for many problems
    - E.g. Pattern databases for 8-puzzle
    - Plan graph heuristics for planning

Planning Graph and Projection

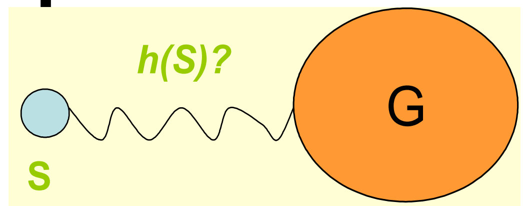

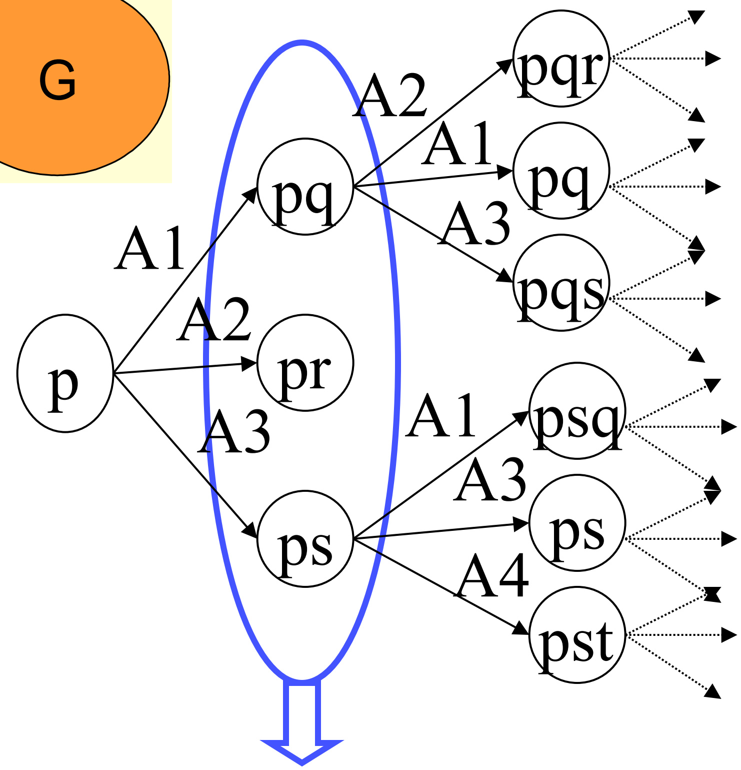

- Envelope of Progression Tree (Relaxed Progression)
  - Proposition lists: Union of states at k^{th} level
  - Mutex: Subsets of literals that cannot be part of any legal state
- Lowerbound r [illegible] information

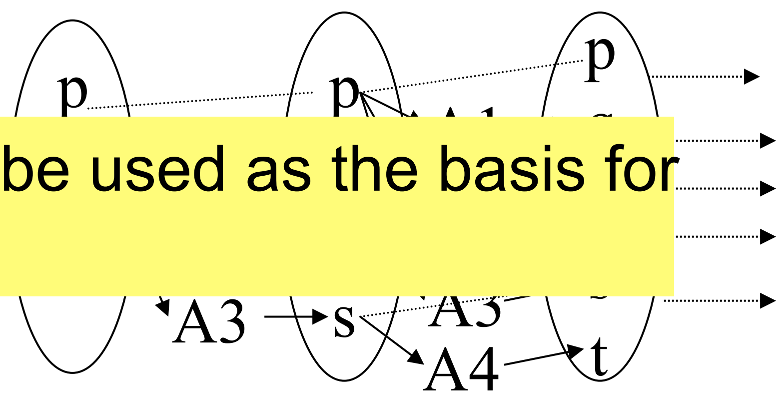

Planning Graphs can be used as the basis for heuristics!

[Blum&Furst, 1995] [ECP, 1997][AI Mag, 2007]

Heuristics for Classical Planning

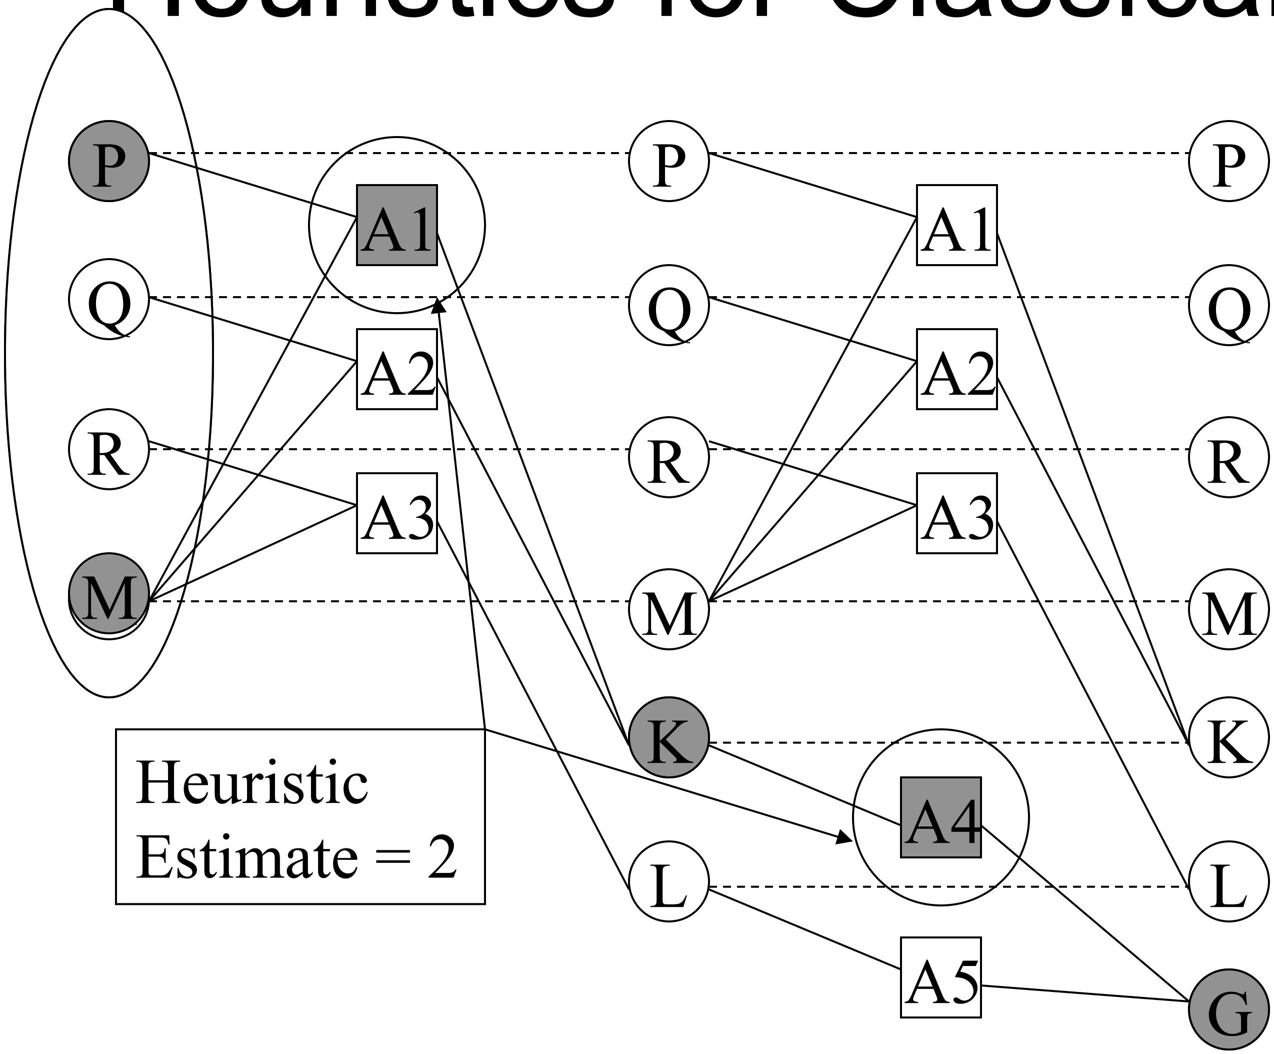

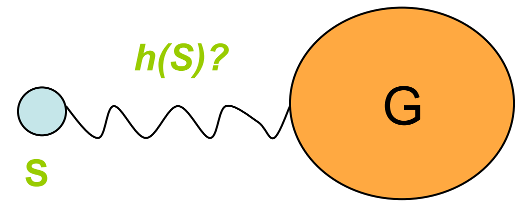

Relaxed plans are solutions for a relaxed problem

How far have we got?

How planners from past IPCs would have performed in IPC-8?

- In Sequential Satisficing track, LAMA-11 (winner of Sequential Satisficing track of IPC-7) would have been **12th** out of 21.
- In Agile track, LPG and FF would have been, respectively, **13th** and **17th** out of 17.

..and we have done our fair bit…

So why the continued fascination with classical planning?

- ..of course, the myriad applications for classical STRIPS planning ☺
- But more seriously, because classical planners have become a customized substrate for "compiling down" other more expressive planning problems
  - Effective approaches exist for leveraging classical planners to do partial satisfaction planning, conformant planning, conditional planning, stochastic planning etc.

So, what next?

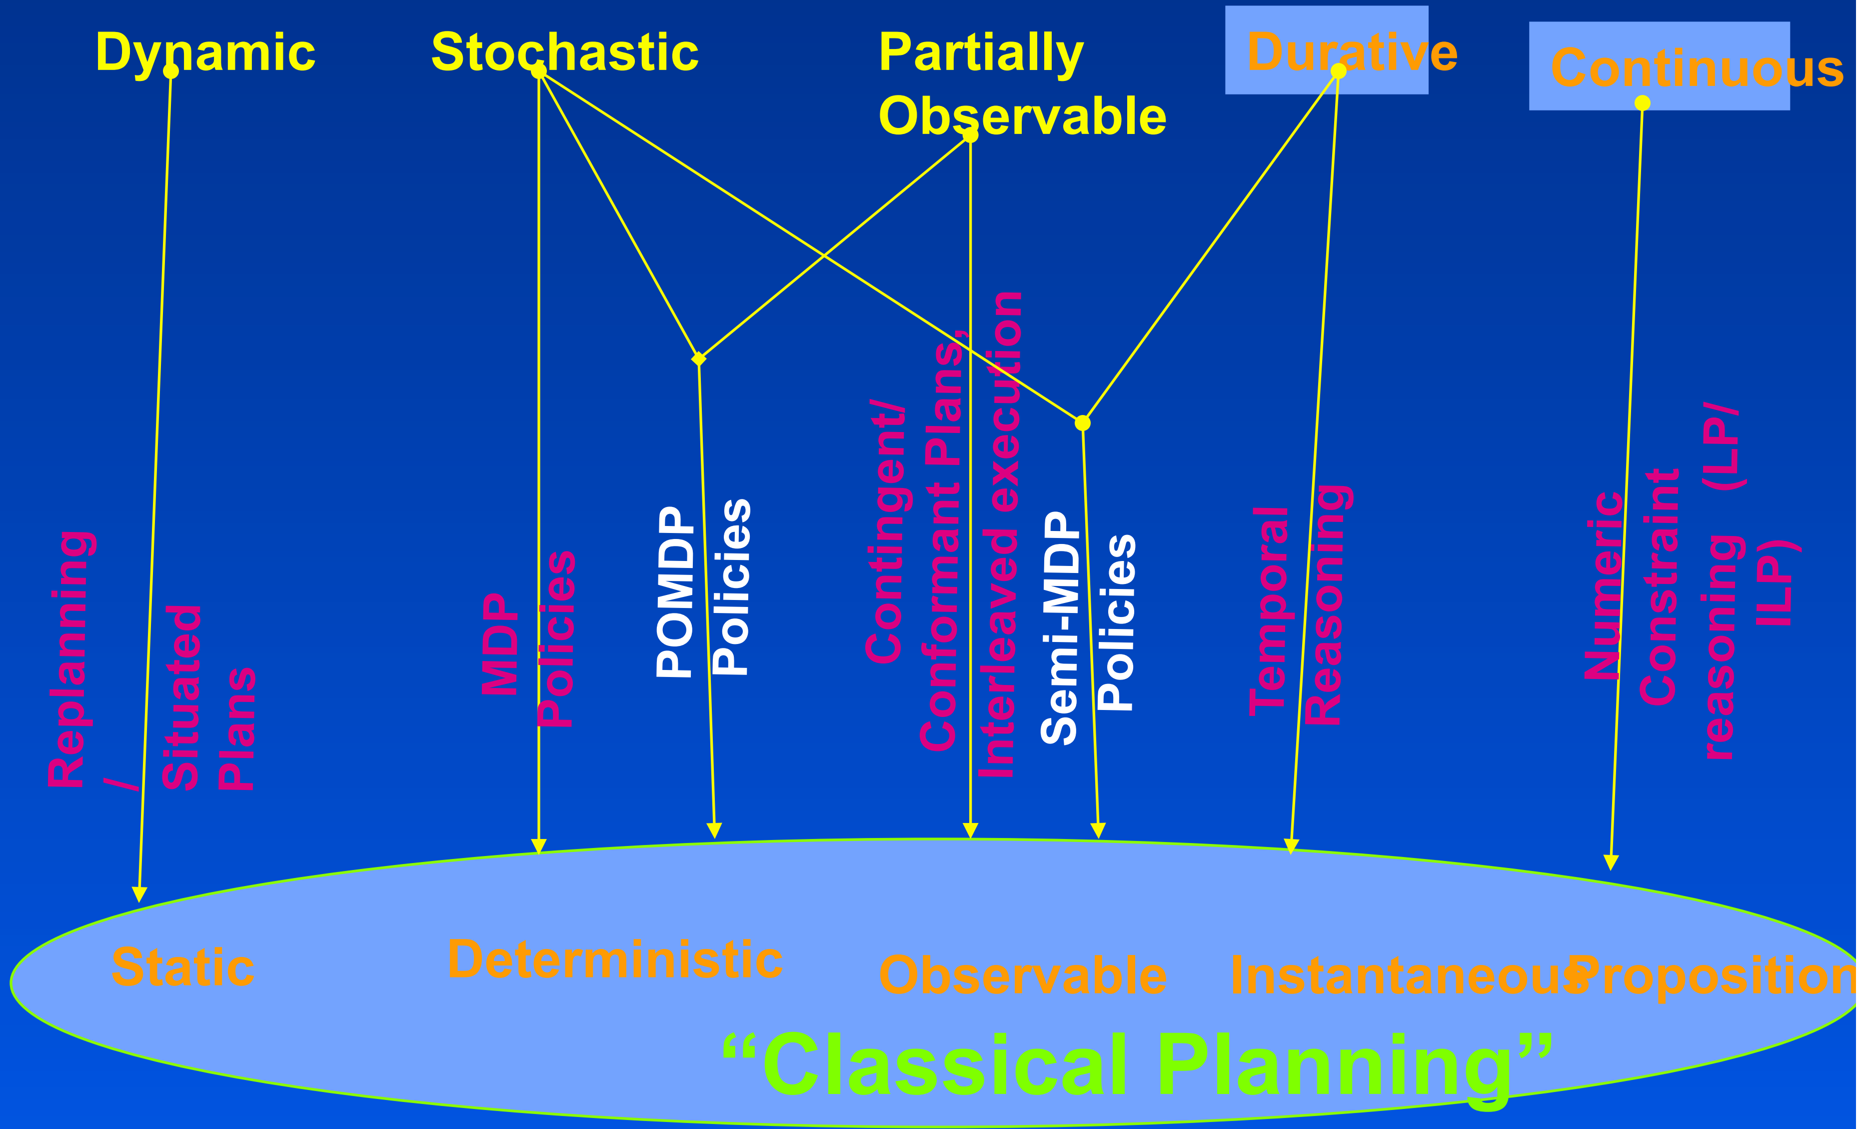

Dynamic
Stochastic
Partially Observable
Durative
Continuous
Replanning/ Situated Plans
MDP Policies
POMDP Policies
Contingent/ Conformant Plans, Interleaved execution
Semi-MDP Policies
Temporal Reasoning
Numeric Constraint reasoning (LP/ ILP)
Static
Deterministic
Observable
Instantaneous
Propositional
“Classical Planning”

Compilation Substrates for Planning

SAT

- First of the substrates
  - Kautz&Selman got the classic paper award honorable mention
- Continued work on fast SAT solvers
- Limited to bounded length planning
- (Not great for metric constraints)

IP/LP

- Followed closely on the heels of SAT
- Can go beyond bounded length planning
  - Allows LP Relaxation
  - (Has become the basis for powerful admissible heuristics)
- IP solvers evolve much slower..

(Classical) Planning

- Tremendous progress on heuristic search approaches to classical planning
- A currently popular approach is to compile expressive planning problems to classical planning
  - Conformant planning, conditional planning
  - (even plan recognition)

Applications—sublime and mundane

- Mission planning (for rovers, telescopes)
- Military planning/scheduling
- Web-service/Work-flow composition
- Paper-routing in copiers
- Gene regulatory network intervention

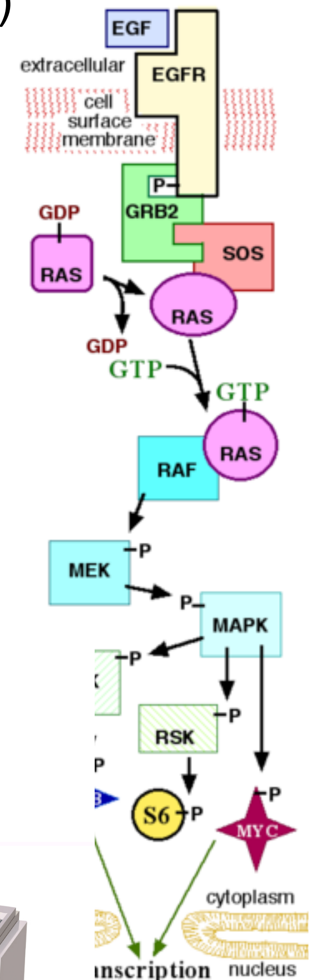

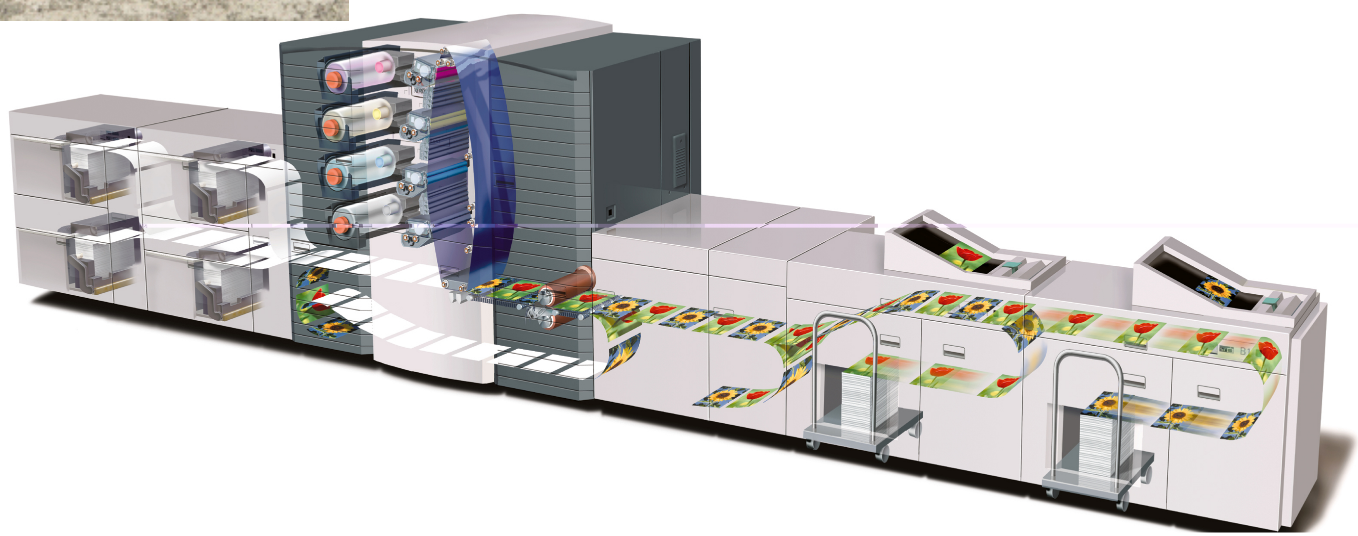

Planning: The Canonical View

A fully specified
 problem
 --Initial state
 --Goals
 (each non-negotiable)
 --Complete Action Model

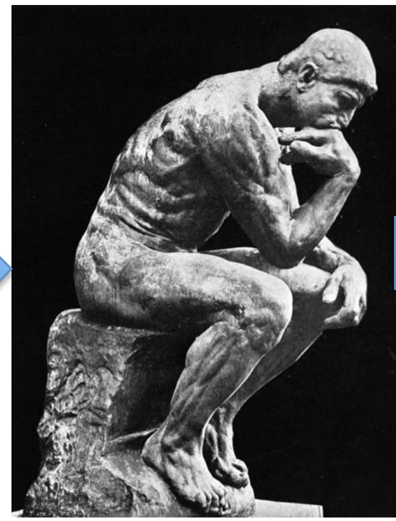

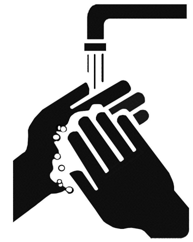

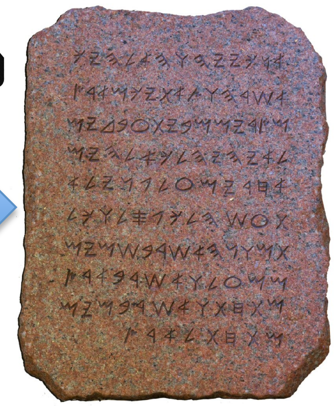

The Plan

Need for Human-in-the-Loop/Human-Aware Planning &Decision Support

- Planners are increasingly embedded in systems that include both humans and machines
  - Human Robot Teaming
    - Petrick et al, Veloso et al, Williams et al, Shah et al, Kambhampati et al
  - Decision support systems; Crowd-planning systems; Tutorial planning systems
    - Allen et al, Kambhampati et al; L
- Necessitates Human-in-the-Loop Planning
  - But, isn't it just "Mixed-Initiative Planning"?
    - ..a lot of old MIP systems had the "Humans in the land of Planners" paradigm (the humans help planners)
    - In effective human-aware planning, planners realize they inhabit the land of humans..

Human-Robot Teamin

- ➢Search and report (rescue)
- ➢Goals incoming on the go
- ➢World is evolving
- ➢Model is changing

- ➢Infer instructions from Natural Language
- ➢Determine goal formulation through clarifications and questions

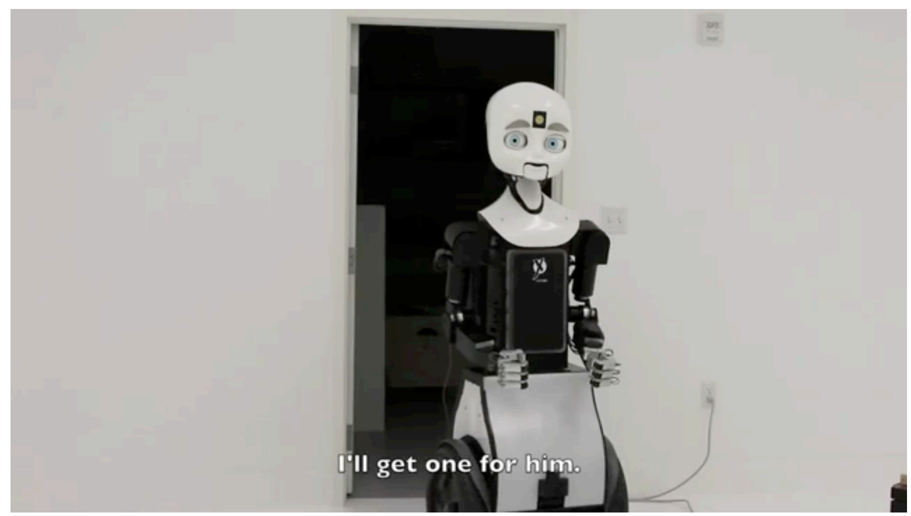

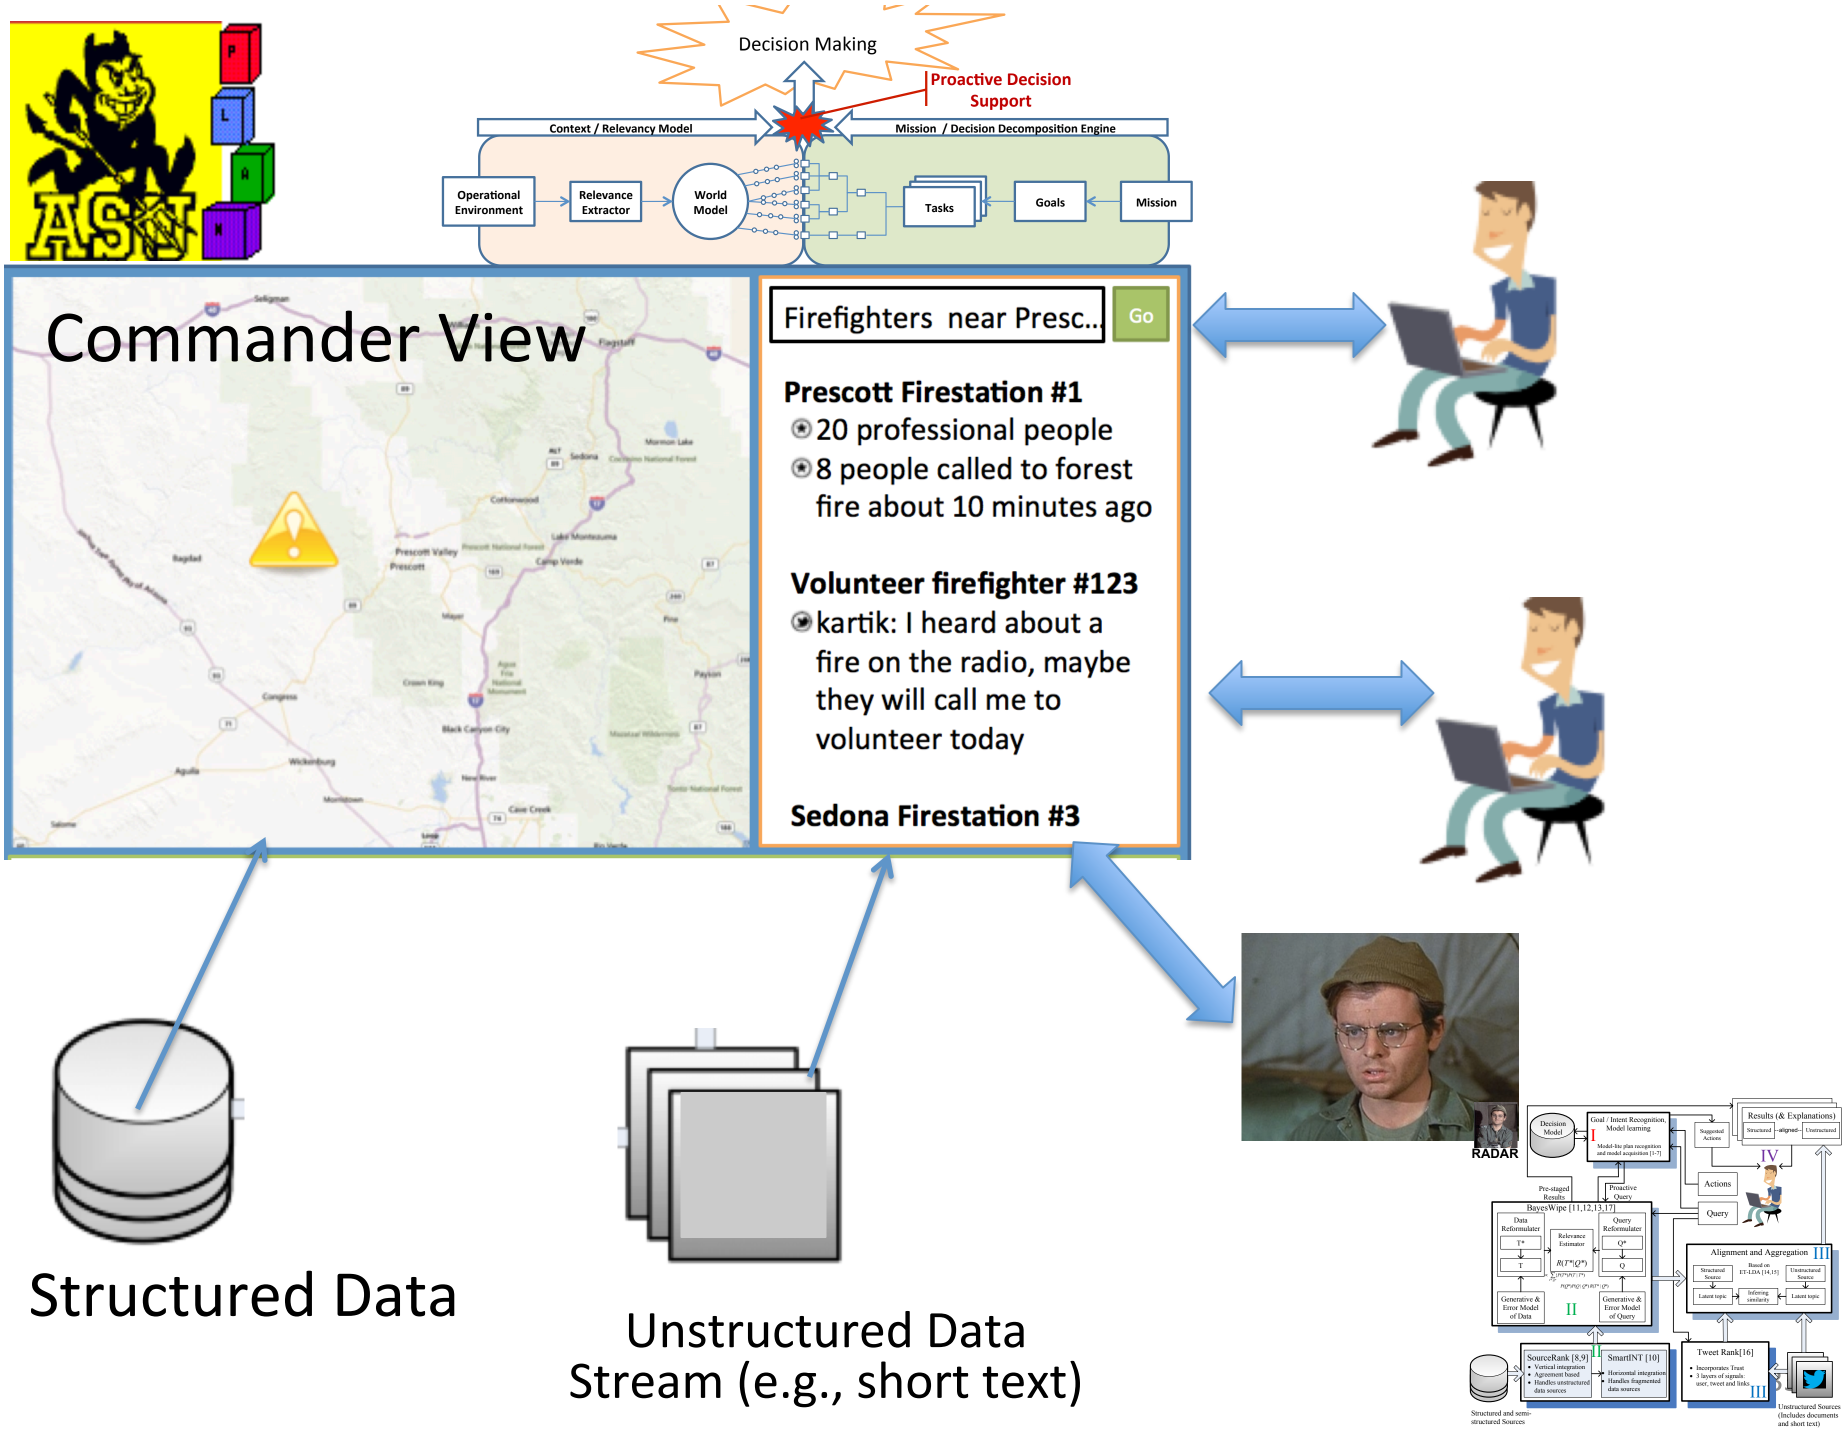

Decision Making
Proactive Decision Support
Context / Relevancy Model
Mission / Decision Decomposition Engine
Operational Environment
Relevance Extractor
World Model
Tasks
Goals
Mission
Commander View
Firefighters near Presc...
Go
Prescott Firestation #1
20 professional people
8 people called to forest fire about 10 minutes ago
Volunteer firefighter #123
kartik: I heard about a fire on the radio, maybe they will call me to volunteer today
Sedona Firestation #3
RADAR
Structured Data
Unstructured Data Stream (e.g., short text)

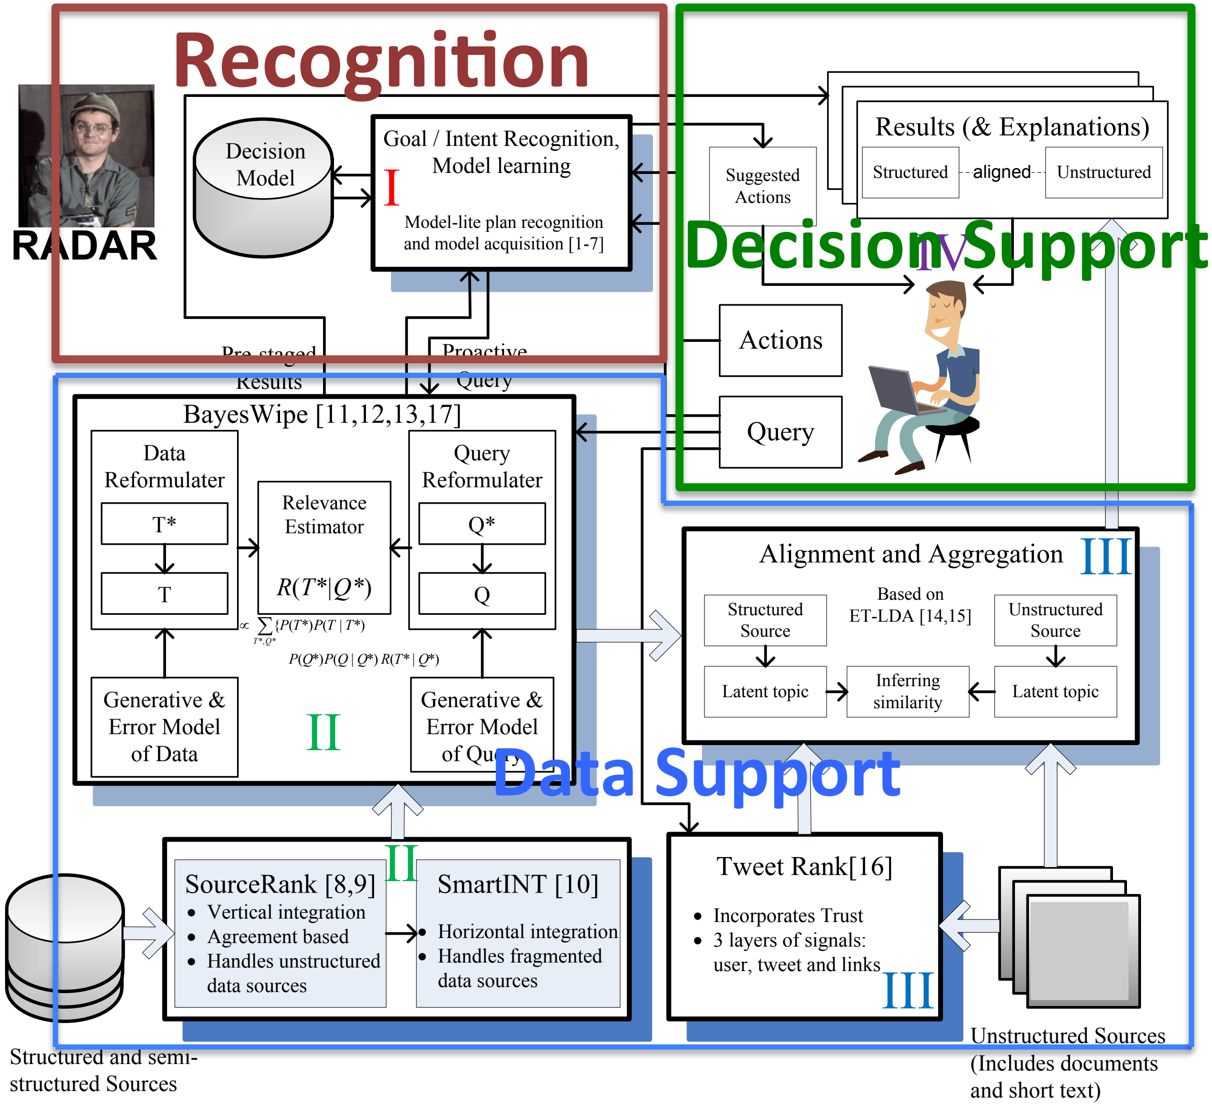

RADAR
Recognition
Decision Model
Goal / Intent Recognition, Model learning
I
Model-lite plan recognition and model acquisition [1-7]
Pre-staged Results
Proactive Query
Decision Support
Suggested Actions
Results (& Explanations)
Structured
aligned
Unstructured
IV
Actions
Query
BayesWipe [11,12,13,17]
Data Reformulater
T*
T
Relevance Estimator
R(T*|Q*)
Query Reformulater
Q*
Q
Generative & Error Model of Data
II
Generative & Error Model of Query
Data Support
Alignment and Aggregation
III
Structured Source
Based on ET-LDA [14,15]
Unstructured Source
Latent topic
Inferring similarity
Latent topic
SourceRank [8,9]
Vertical integration
Agreement based
Handles unstructured data sources
II
SmartINT [10]
Horizontal integration
Handles fragmented data sources
Tweet Rank[16]
Incorporates Trust
3 layers of signals: user, tweet and links
III
Structured and semi-structured Sources
Unstructured Sources (Includes documents and short text)

Crowd-Sourced Planning

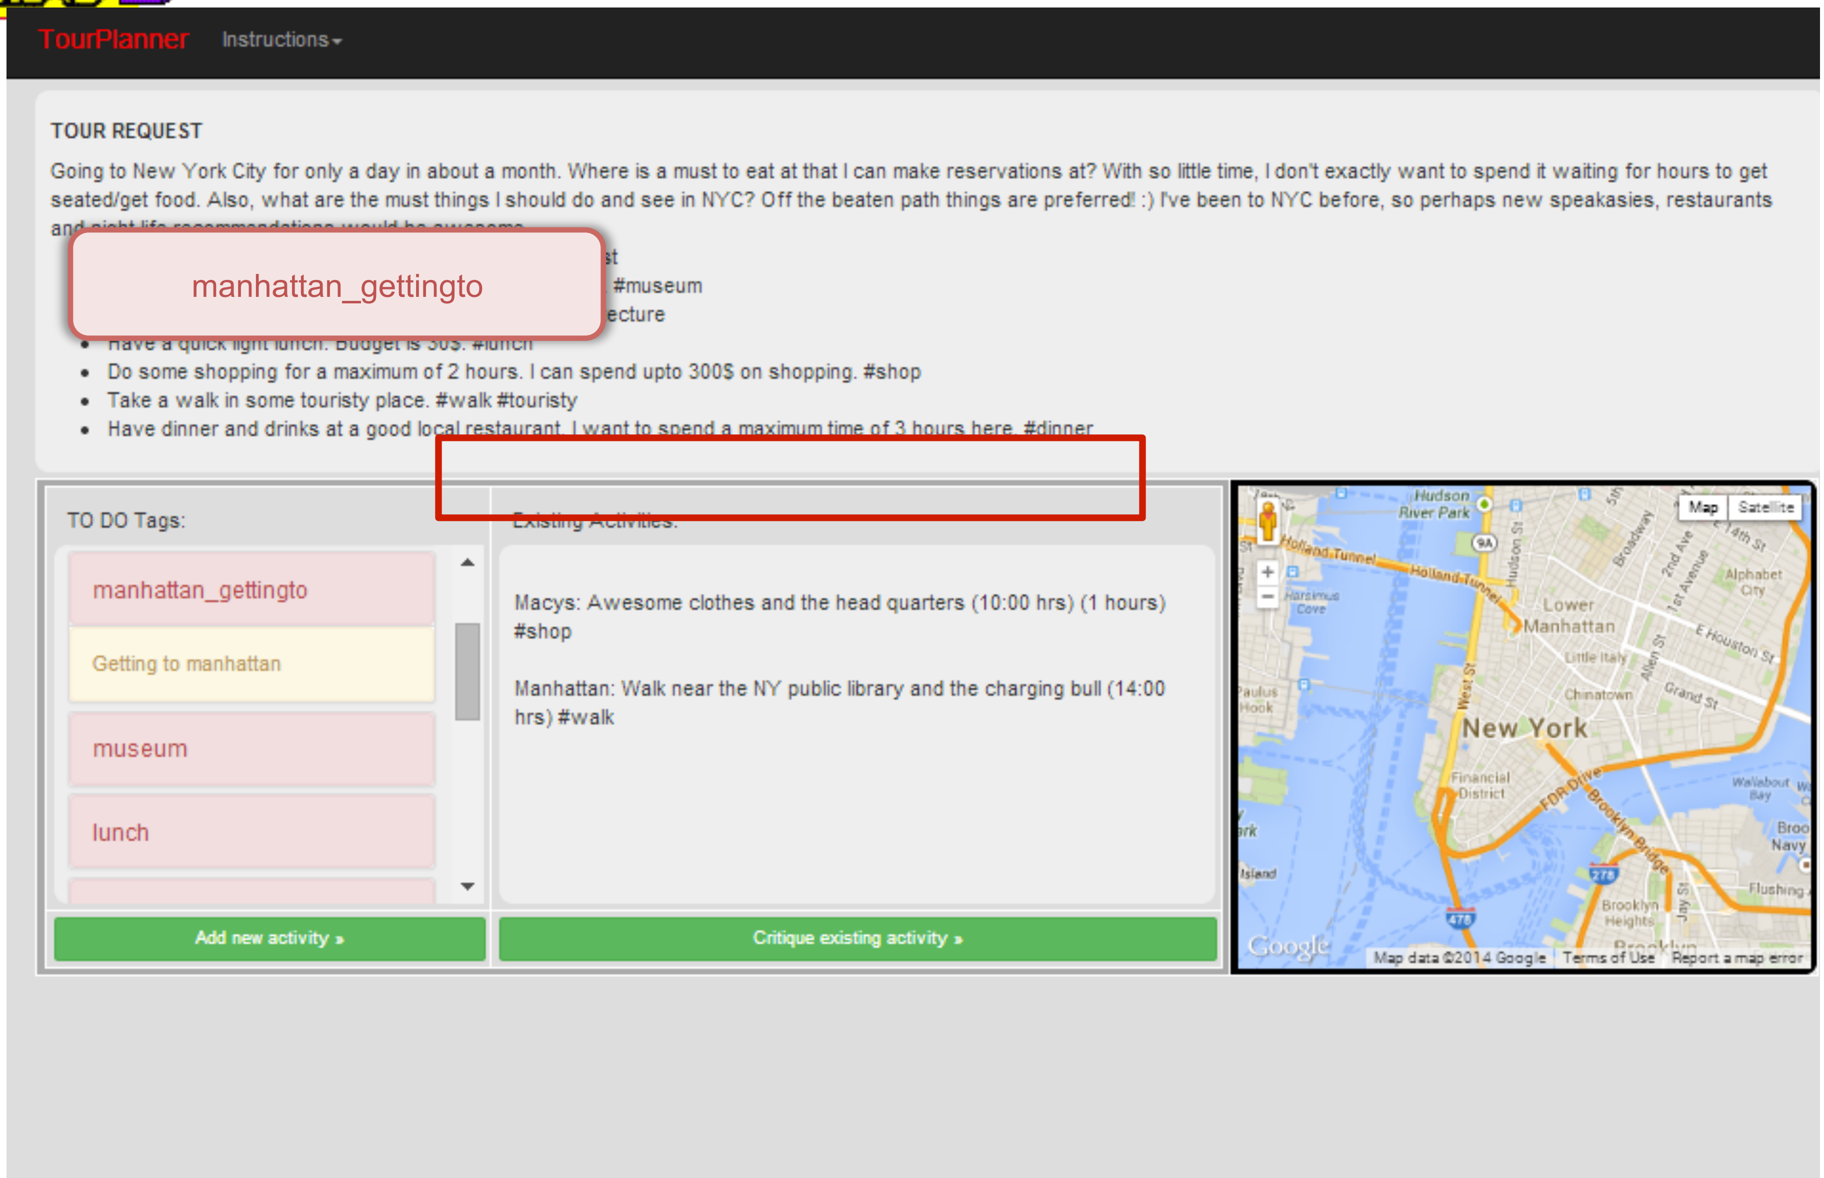

Results: Role of Planner Module

VOTED ICAPS 2014 BEST DEMO BY … THE CROWD!

-GOAL

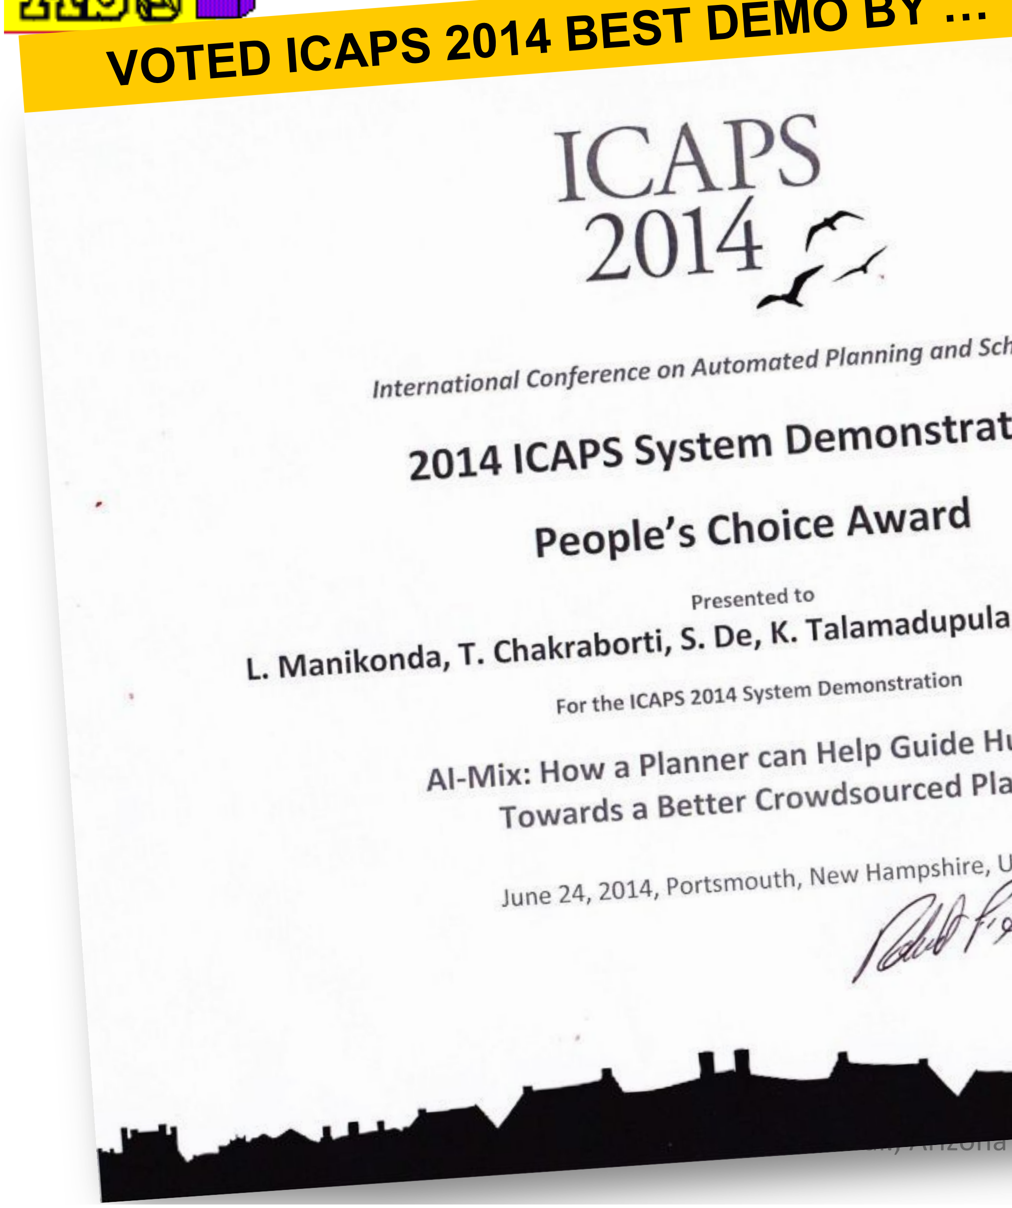
ICAPS 2014

International Conference on Automated Planning and Scheduling

2014 ICAPS System Demonstratio

People's Choice Award

Presented to

L. Manikonda, T. Chakraborti, S. De, K. Talamadupula, S

For the ICAPS 2014 System Demonstration

AI-Mix: How a Planner can Help Guide Hu

Towards a Better Crowdsourced Pla

June 24, 2014, Portsmouth, New Hampshire, US

Planning: The Canonical View

A fully specified
problem
--Initial state
--Goals
(each non-nego
--Complete Action

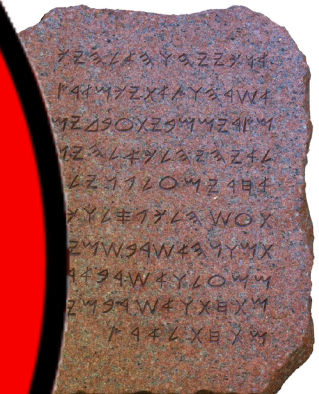

he Plan

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

Violated Assumptions:

→~~Complete~~ Action Descriptions (fallible domain writers)
→~~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 ☹

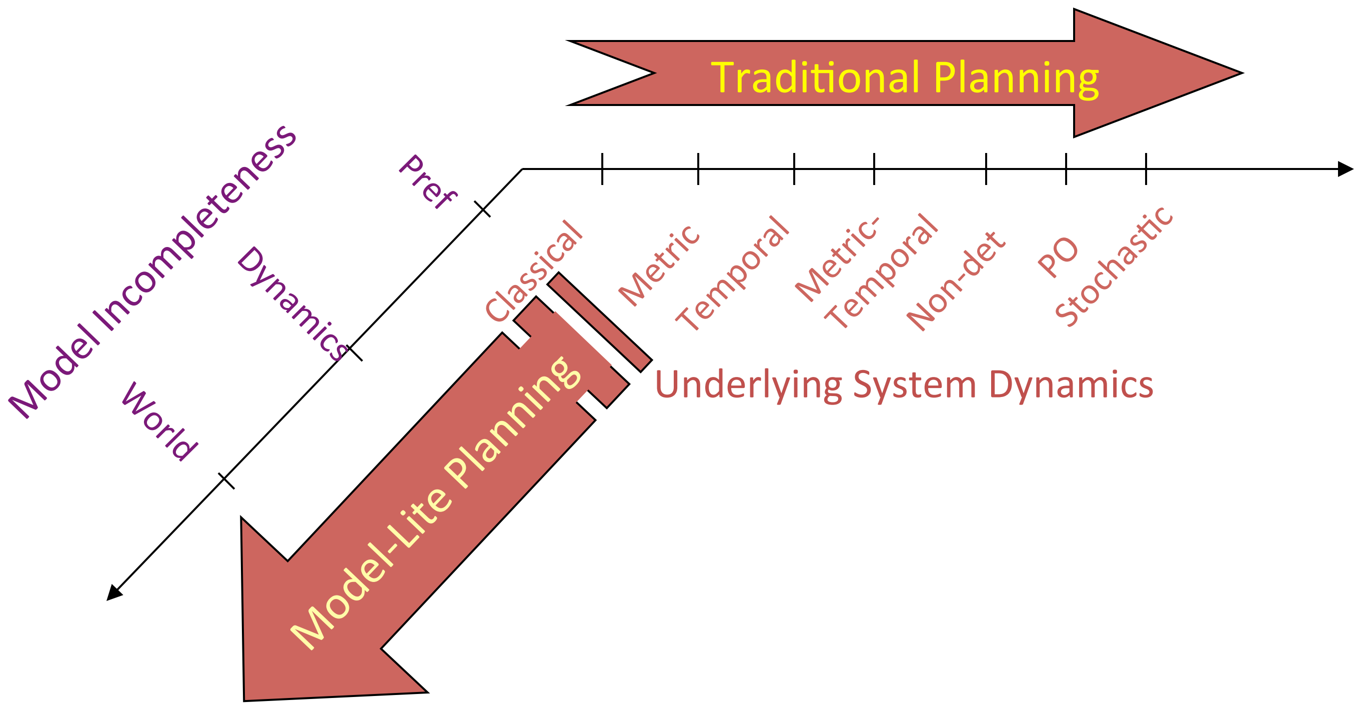

Effective ways to handle the more expressive planning problems by exploiting the deterministic planning technology

Learning is useful but not the (sole) answer..

- A tempting way to handle incompleteness is to say that we should wait until the full model is obtained
  - Either through learning
    - (We do have work in Model Learning: E.g. Zhuo et al. IJCAI 2013)
  - Or by the generosity of the domain writer..
- Problem: Waiting for complete model is often times not a feasible alternative
    - The model may never become complete…
    - We need to figure out a way of maintaining incomplete models, and planning with them (pending learning..)

MODELS V. PLANNING CAPABILITIES

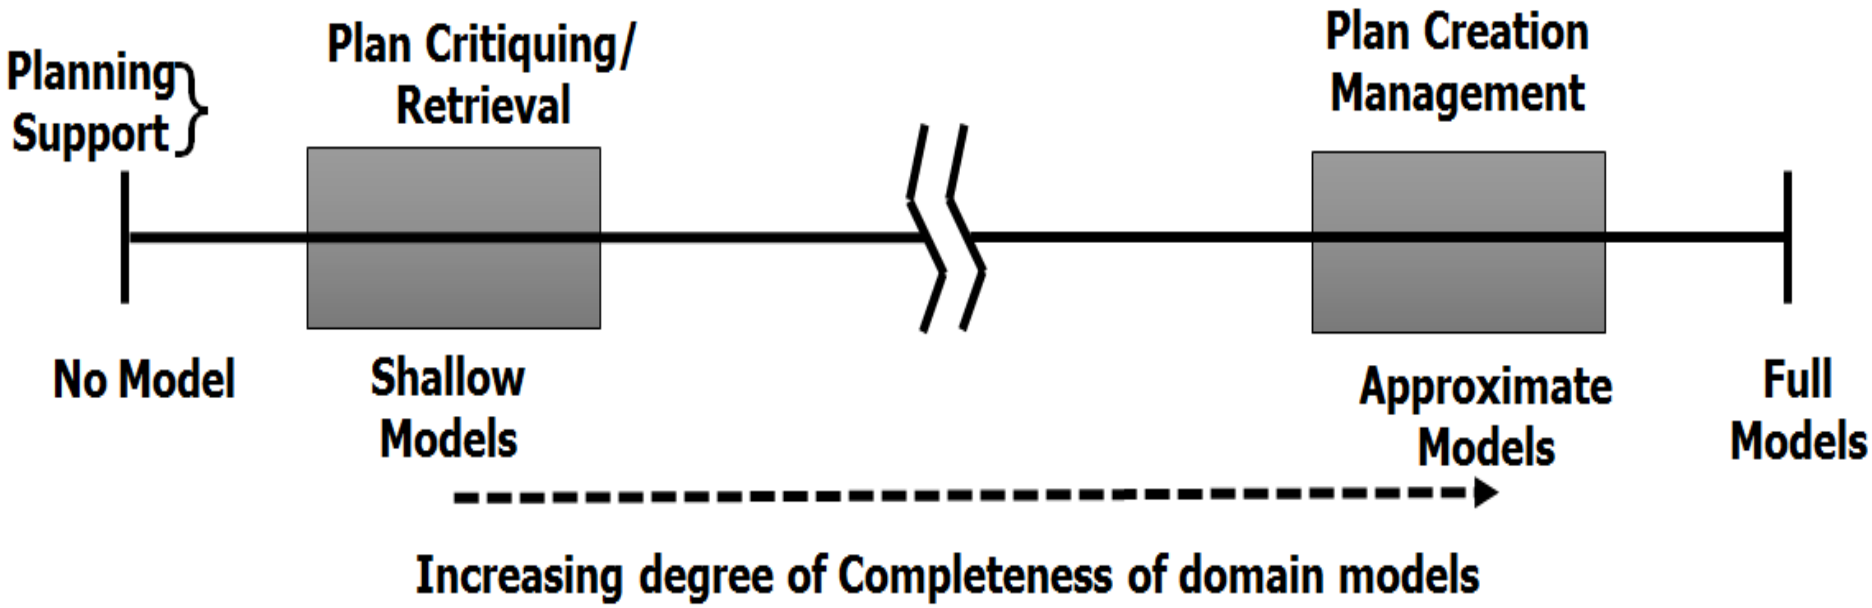

I/O types
Task dependency
(e.g. workflows management, web service composition)

Missing some preconditions/ effects of actions
(e.g. Garland & Lesh, 2002)

Approaches for Planning with Incomplete Models (1)

Incompleteness annotations are available

- One way to make-up for model incompleteness is to expect annotations circumscribing the extent of incompleteness
- In this case, we can explicitly reason with the correctness of candidate plans over all possible models
  - Nguyen et. al. ICAPS 2014; NIPS 2013

Problem Formulation

- ❖ Incomplete domain: $\widetilde{D} = \langle F, A \rangle$
- ❖ Each partially specified action $a \in A$:
  - ➢Known preconditions, effects: $Pre(a), Add(a), Del(a)$
  - ➢Possible preconditions, effects: $\widetilde{Pre}(a), \widetilde{Add}(a), \widetilde{Del}(a)$
  - ➢Optional weights: $w_a^{pre}(p), w_a^{add}(p), w_a^{del}(p) \in (0,1)$
- ❖ Exponential number of candidate complete models: $\langle\!\langle \widetilde{D} \rangle\!\rangle$
  - ➢ One of which is the *true* model.
- ❖ Planning problem: $\widetilde{P} = \langle \widetilde{D}, I, G \rangle$

Incomplete ≠ Stochastic

Robustness measure for plans

$$R(\pi, \widetilde{P}: \langle \widetilde{D}, I, G \rangle) = \sum_{D_i \in \langle\!\langle \widetilde{D} \rangle\!\rangle, \gamma(\pi, I, D_i) \models G} \Pr(D_i)$$

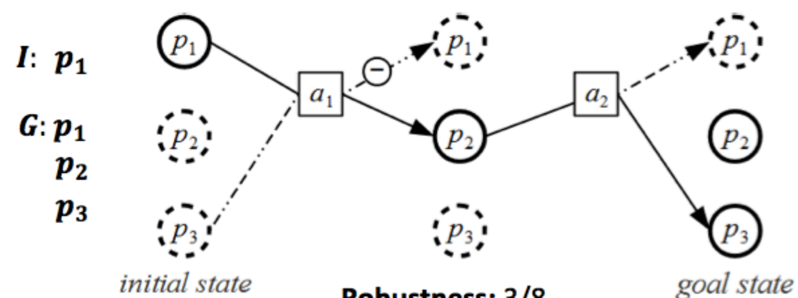

Robustness: 3/8

A Spectrum of robust planning problems

- ❖ Robustness assessment
- ❖ Maximally robust plan generation
- ❖ Generating plan with desired level of robustness
- ❖ Cost-sensitive robust plan generation
- ❖ Incremental robustification

Robustness Assessment as Weighted Model Counting

- **Causal-proof based correctness constraints Σ**

Establishment constraints

- *Unsupported known preconditions must be supported:* $\bigvee_{C_p^i \le k \le i-1, p \in \widetilde{Add}(a_k)} p_{a_k}^{add}$
- *So do possible preconditions, if realized:* $p_{a_i}^{pre} \Rightarrow \bigvee_{C_p^i \le k \le i-1, p \in \widetilde{Add}(a_k)} p_{a_k}^{add}$

Protection constraints

- *Known preconditions deleted by realized effects must be re-established:* $p_{a_m}^{del} \Rightarrow \bigvee_{C_p^i \le k \le i-1, p \in \widetilde{Add}(a_k)} p_{a_k}^{add}$
- *So do possible preconditions, if realized:* $p_{a_i}^{pre} \Rightarrow (p_{a_m}^{del} \Rightarrow \bigvee_{C_p^i \le k \le i-1, p \in \widetilde{Add}(a_k)} p_{a_k}^{add})$

- Plan robustness = Weighted model counting WMC(Σ).
- **Complexity:** Assessing plan robustness is *#P*-complete.

Robust Plan Synthesis: A Compilation Approach

Incomplete model
Complete world state

→

Complete model
Belief state

(Conformant Probabilistic Planning)

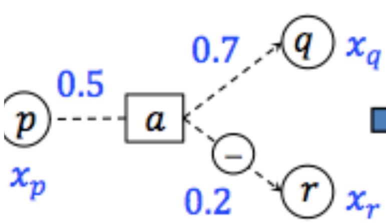

x_p (0.5) x_q (0.7) x_r (0.2)

Initial belief state

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

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

Robust Plan Synthesis: A Heuristic Approach

- **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

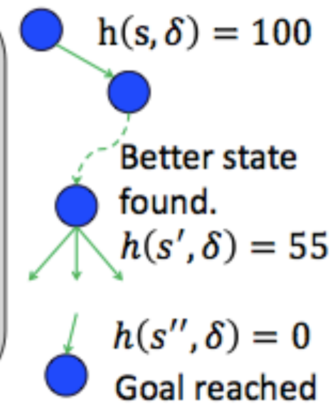

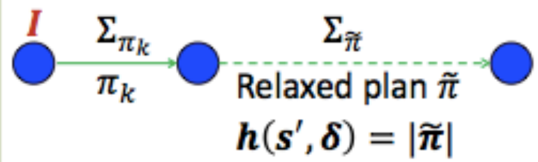

$\boldsymbol{WMC(\Sigma_{\pi_k} \cup \Sigma_{\tilde{\pi}}) > \delta}$

$\boldsymbol{h(s', \delta) = |\tilde{\pi}|}$

Costly!

- **Approximate plan robustness**
  - **Lower bound: Σ as monotone clauses**

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

$l(\Sigma_{\pi_k} \wedge \Sigma_{\tilde{\pi}}) > \delta$

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

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

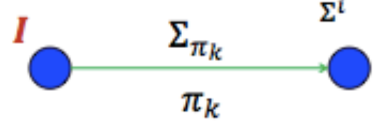

If $u(\Sigma_{\pi_k}) > \delta$
then compute $WMC(\pi_k)$

- **Evaluation**

Approaches for Planning with Incomplete Models (2)

Library of cases is available

- Sometimes, we may have access to "cases"/previous successful plans
- ML-CBP exploits cases directly during planning (by transferring case fragments into a skeletal plan generated w.r.t. M')
  - Zhuo Et al AAAI 2013
- An alternative approach would be to use the cases C to *refine* the model M' into a more accurate model M'' (where M'' is a better approximation of M*)
  - M'' contains both primitive and macro-operators
  - Zhuo et. Al. IJCAI 2013

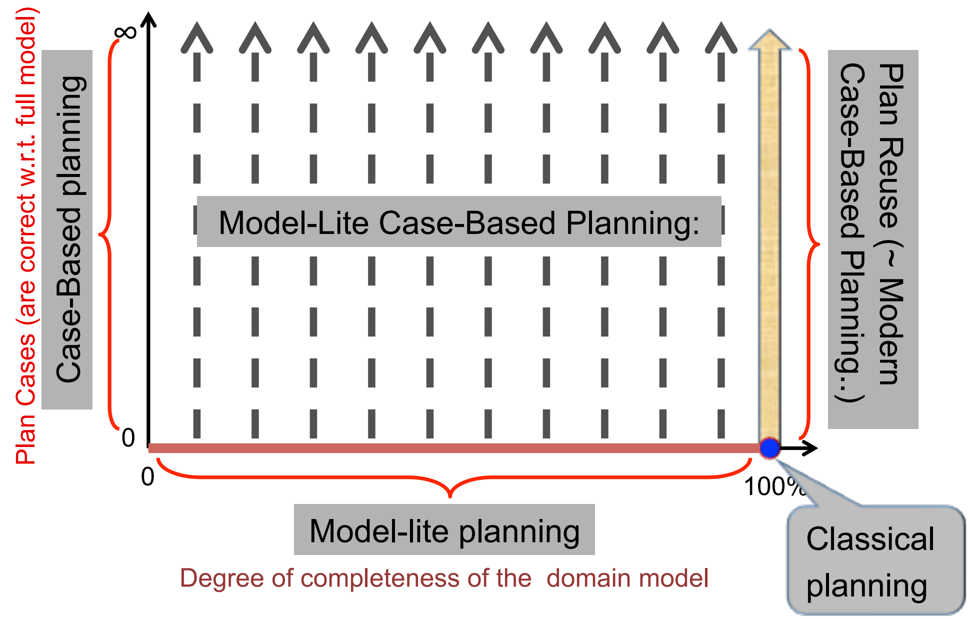

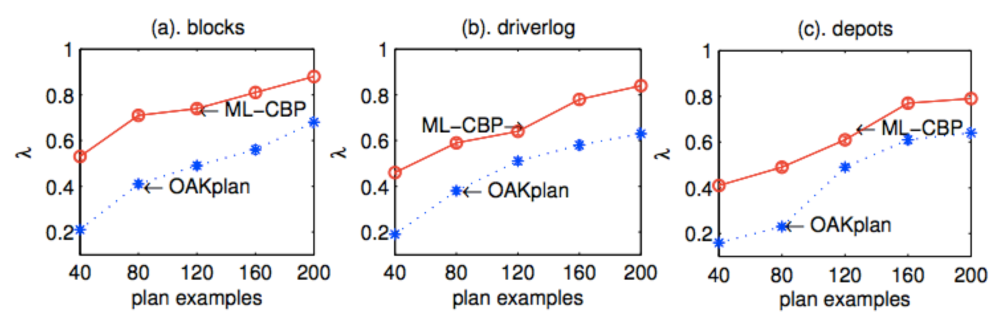

Figure 3: Accuracy w.r.t. number of plan cases.

Approaches for Planning with Incomplete Models (2)

Library of cases is available

- Sometimes, we may have access to "cases"/previous successful plans
- ML-CBP exploits cases directly during planning (by transferring case fragments into a skeletal plan generated w.r.t. M')
  - Zhuo Et al AAAI 2013
- An alternative approach would be to use the cases C to *refine* the model M' into a more accurate model M'' (where M'' is a better approximation of M*)
  - M'' contains both primitive and macro-operators
  - Zhuo et. Al. IJCAI 2013

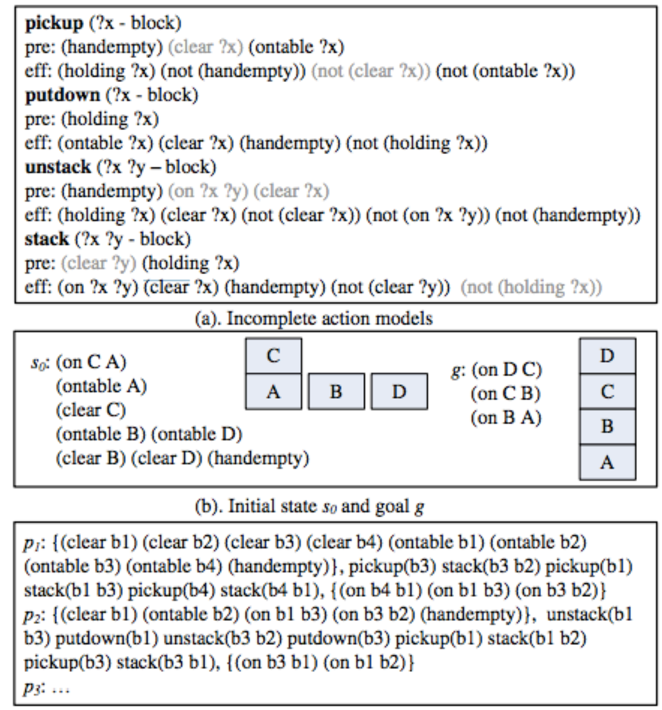

(c). Plan examples

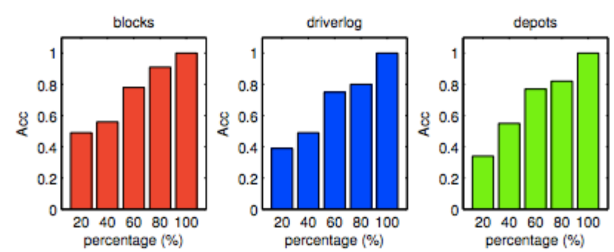

Figure 3: Accuracies w.r.t. completeness of action models.

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

COORDINATION IN HUMAN-ROBOT TEAMS USING MENTAL MODELING AND PLAN RECOGNITION

Presented at IROS 2014

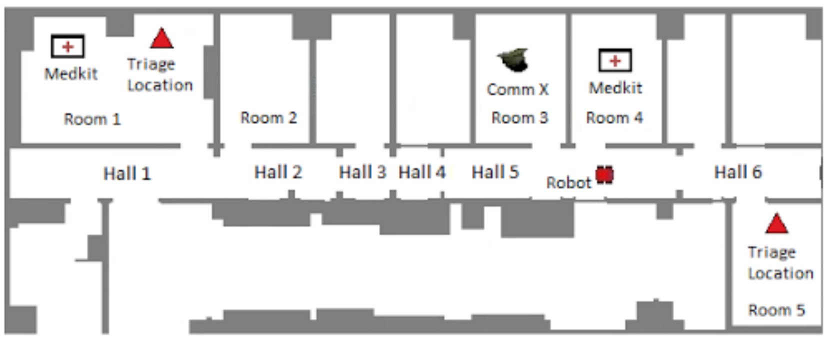

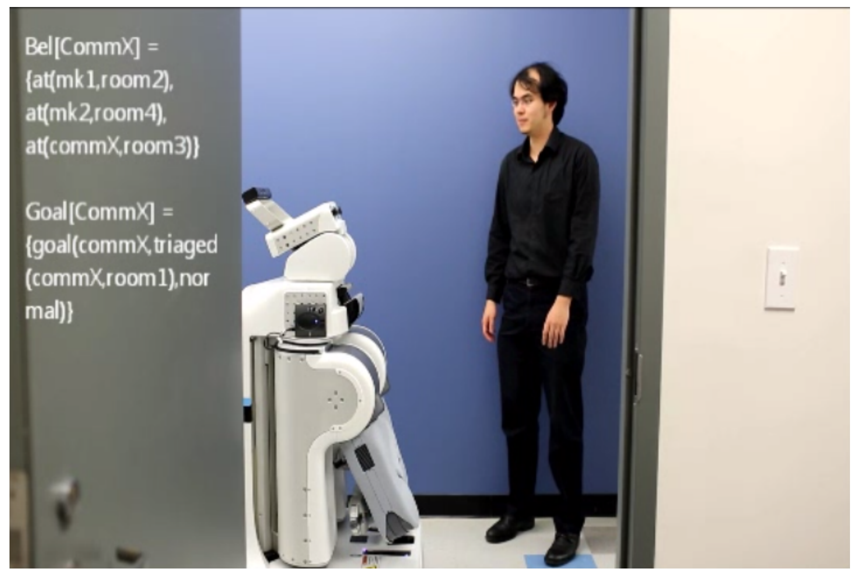

Plan Recognition

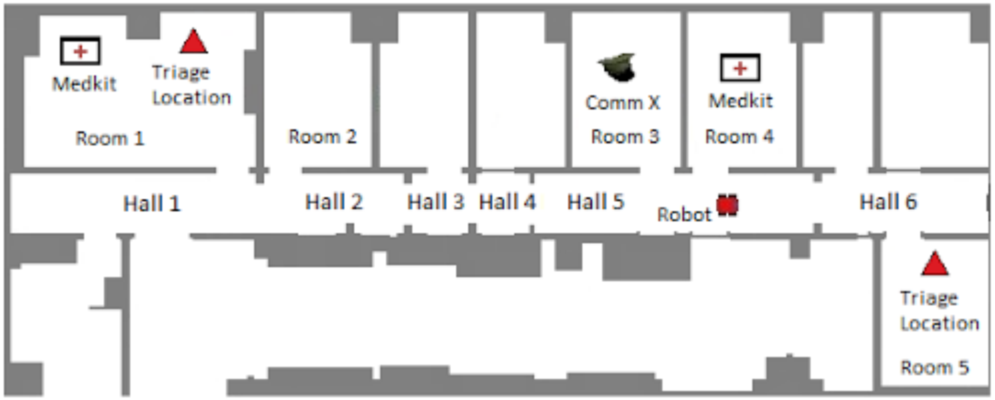

BELIEF IN GOAL

(conducted_triage commX room1)

(conducted_triage commX room5)

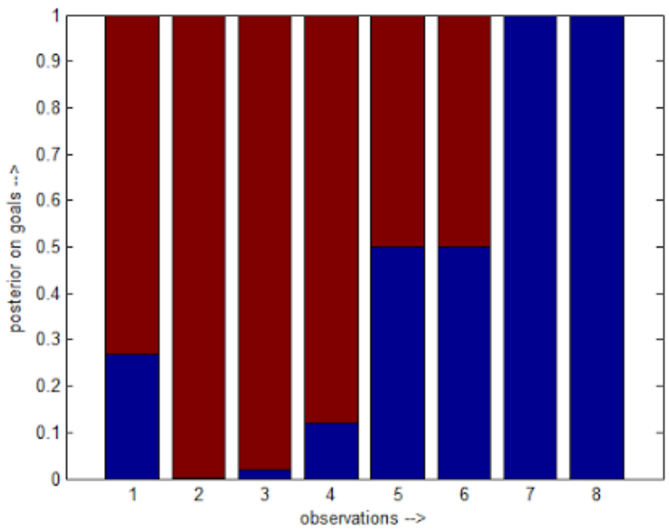

observations -
move commx room3 hall5
move_reverse commx hall5 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

Plan Recognition

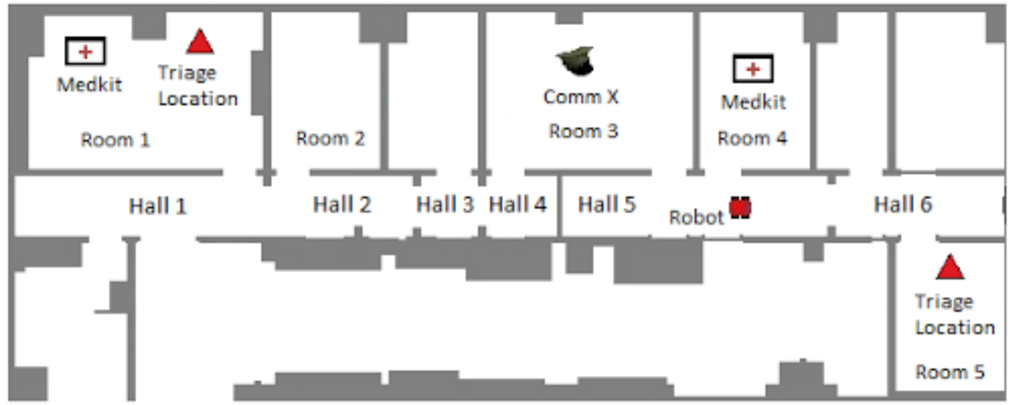

BELIEF IN GOAL

(conducted_triage commX room1)
(conducted_triage commX room5)

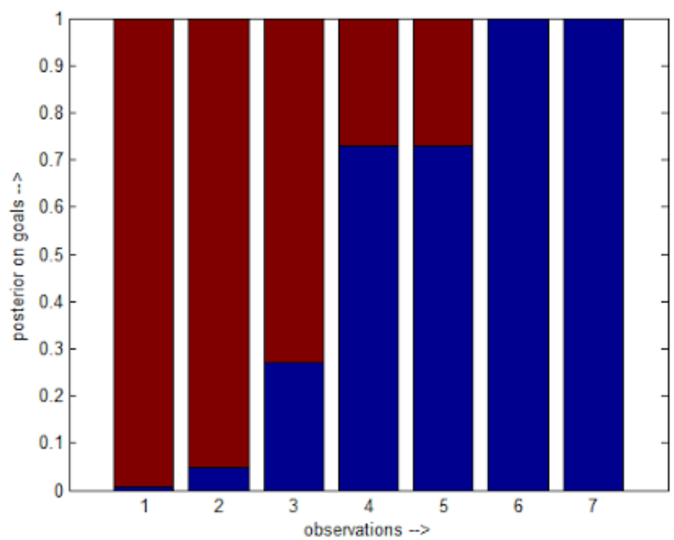

observations -
move commx room3 hall4
move_reverse commx hall4 hall3
move_reverse commx hall3 hall2
move_reverse commx hall2 hall1
move_reverse commx hall1 room1
pick_up_medkit commx mkeast room1
conduct_triage commx room1

AI-MIX: SYSTEM SCHEMATIC

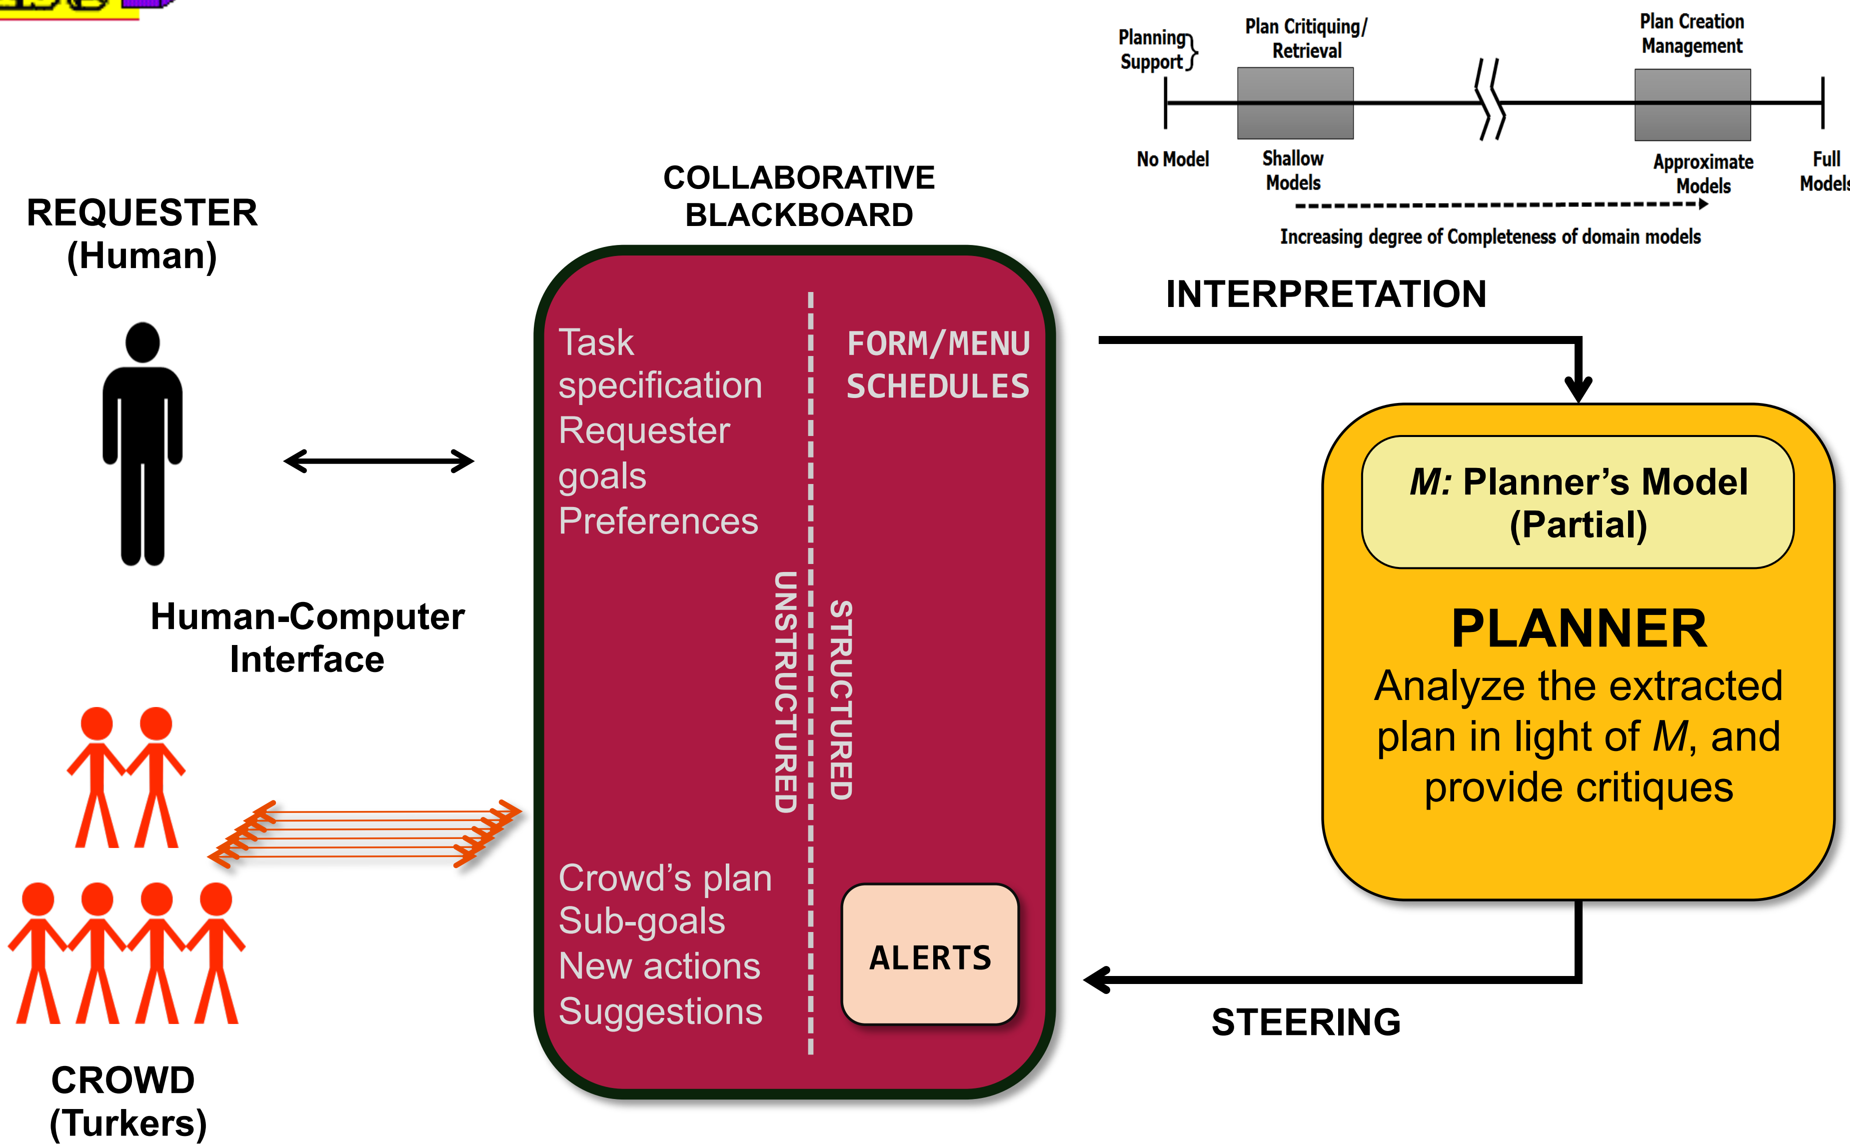

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

REPLANNING EXAMPLE: ROVERS

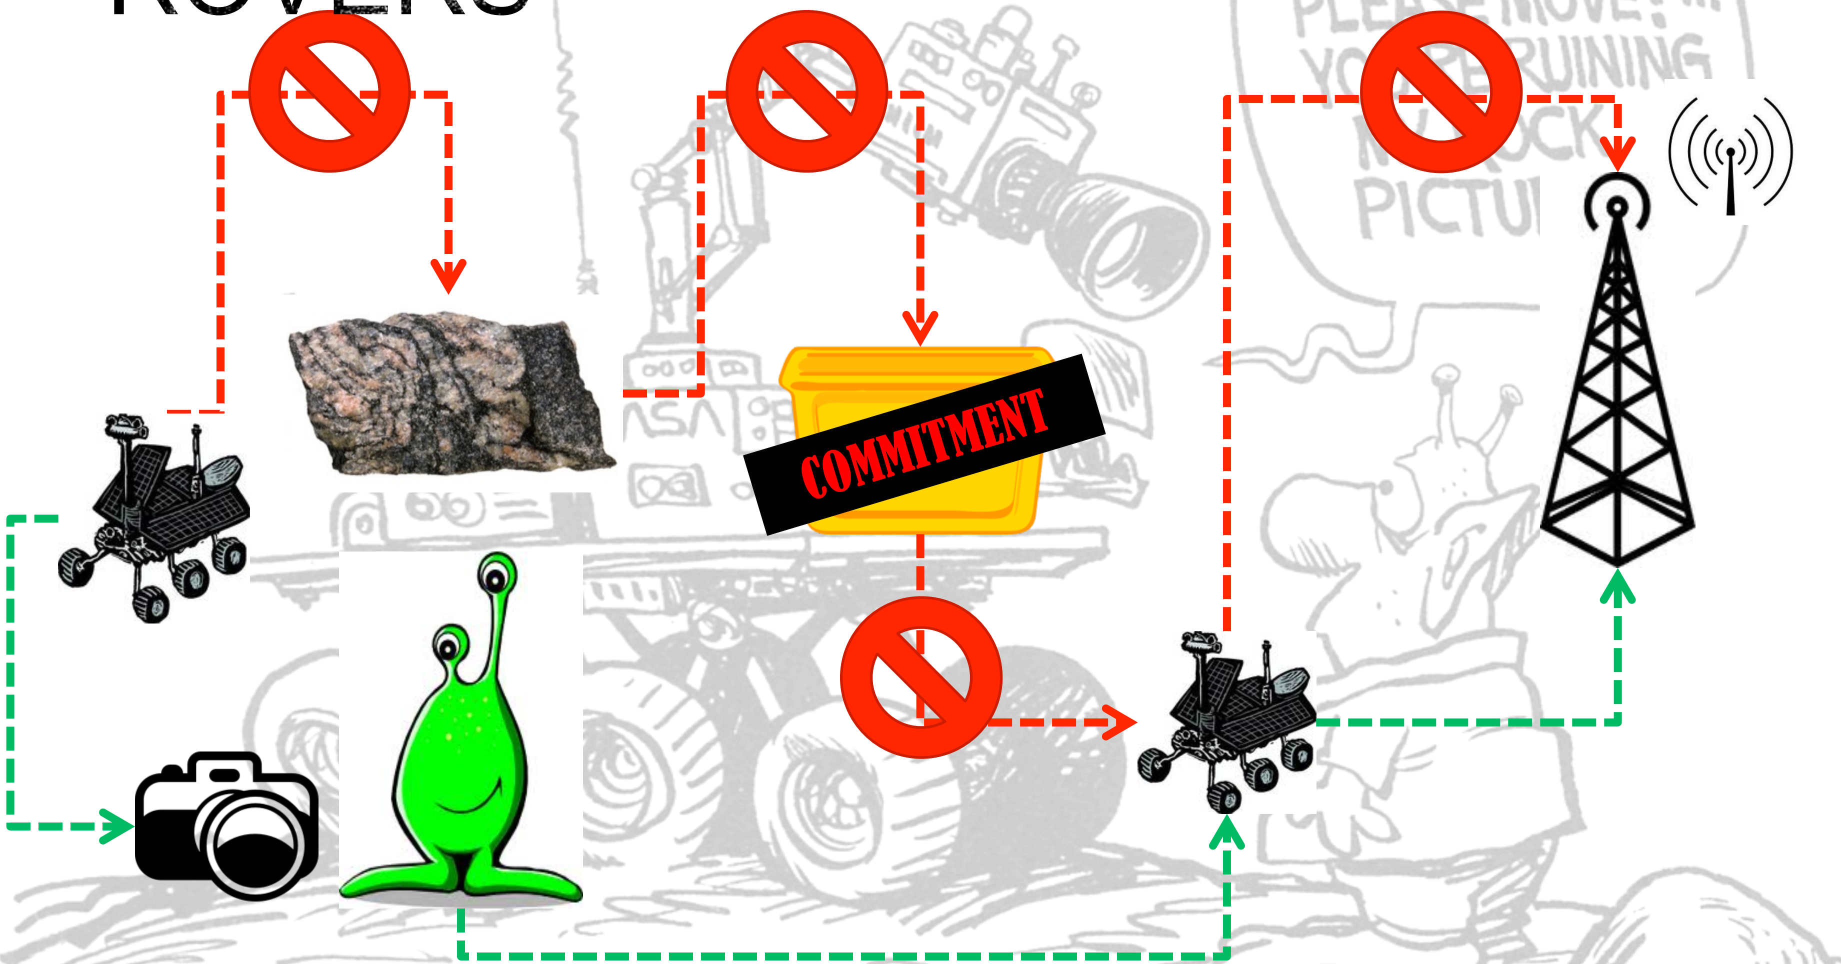

reward for photograph of something interesting far exceeds penalty for violating drop-off commitment

A GENERALIZED MODEL OF REPLANNING

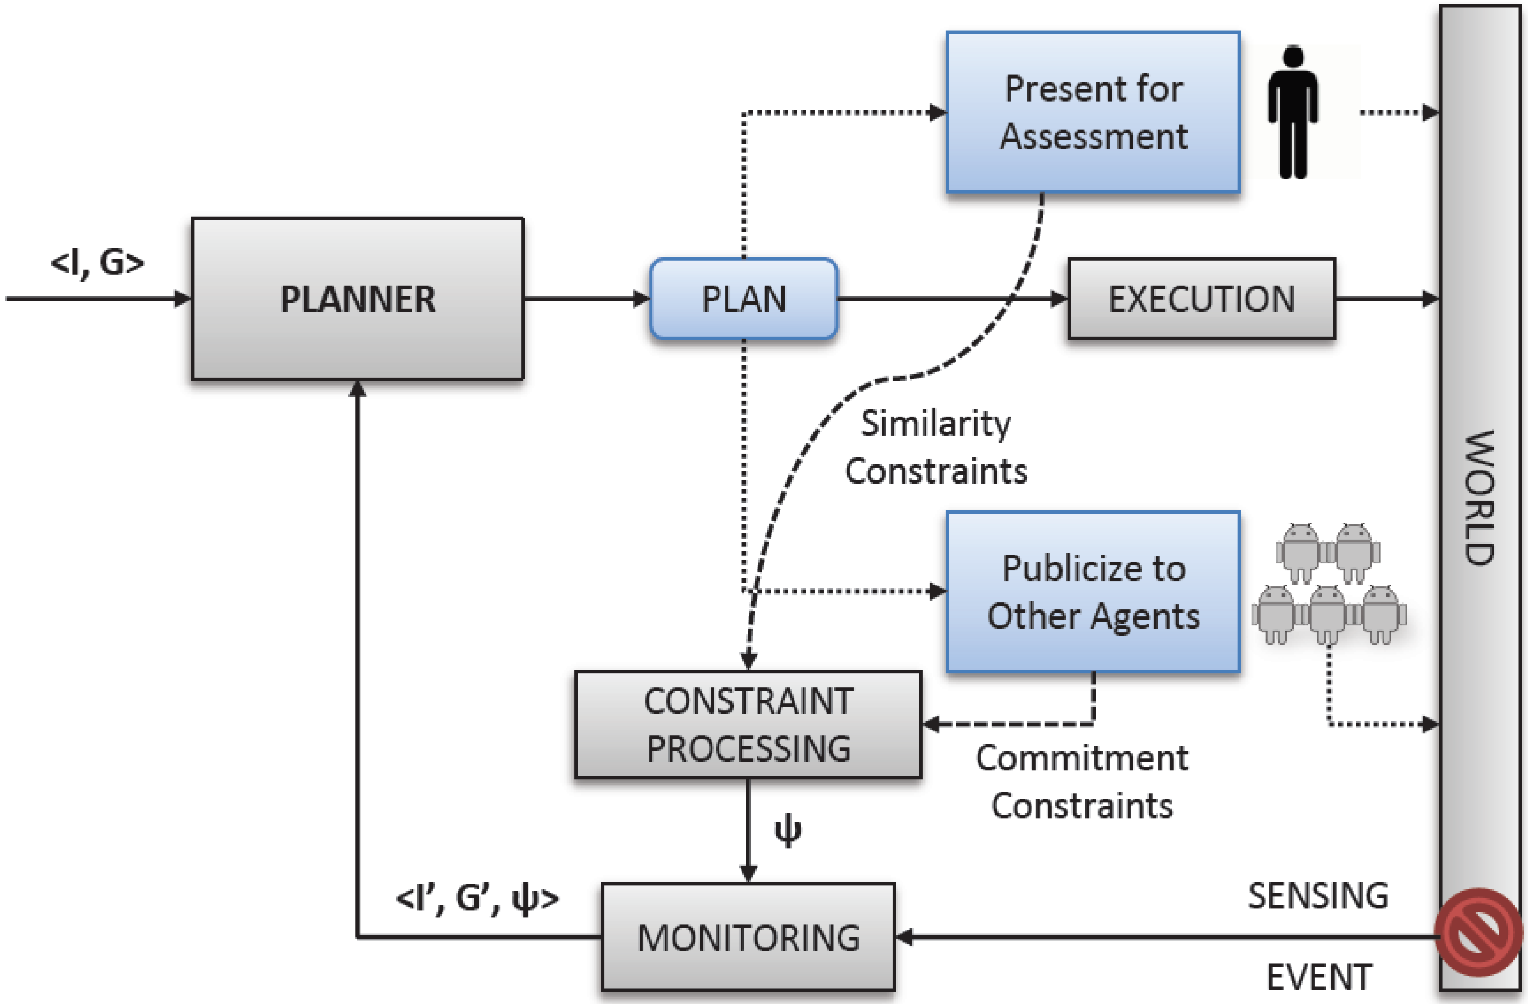

[DMAP13]

REPLANNING CONSTRAINTS

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

[DMAP13]

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

--Human Factors

Tutorial@AAAI 2015

Minecraft USAR Simulation

- Minecraft was chosen for its ease in flexibility and variability in its virtual environment.
- Office environment created in Minecraft for the purpose of simulating a search and rescue task.
- Minecraft allows for complex and changing environments and many internal participants for future iterations.

Example Maps

Team 14 Remote 7-17-2014

Results: Situation Awareness

External Situation Awareness = number of rooms correct (i.e., correct number and color of boxes on the map) slightly lower for Intelligent and more variability

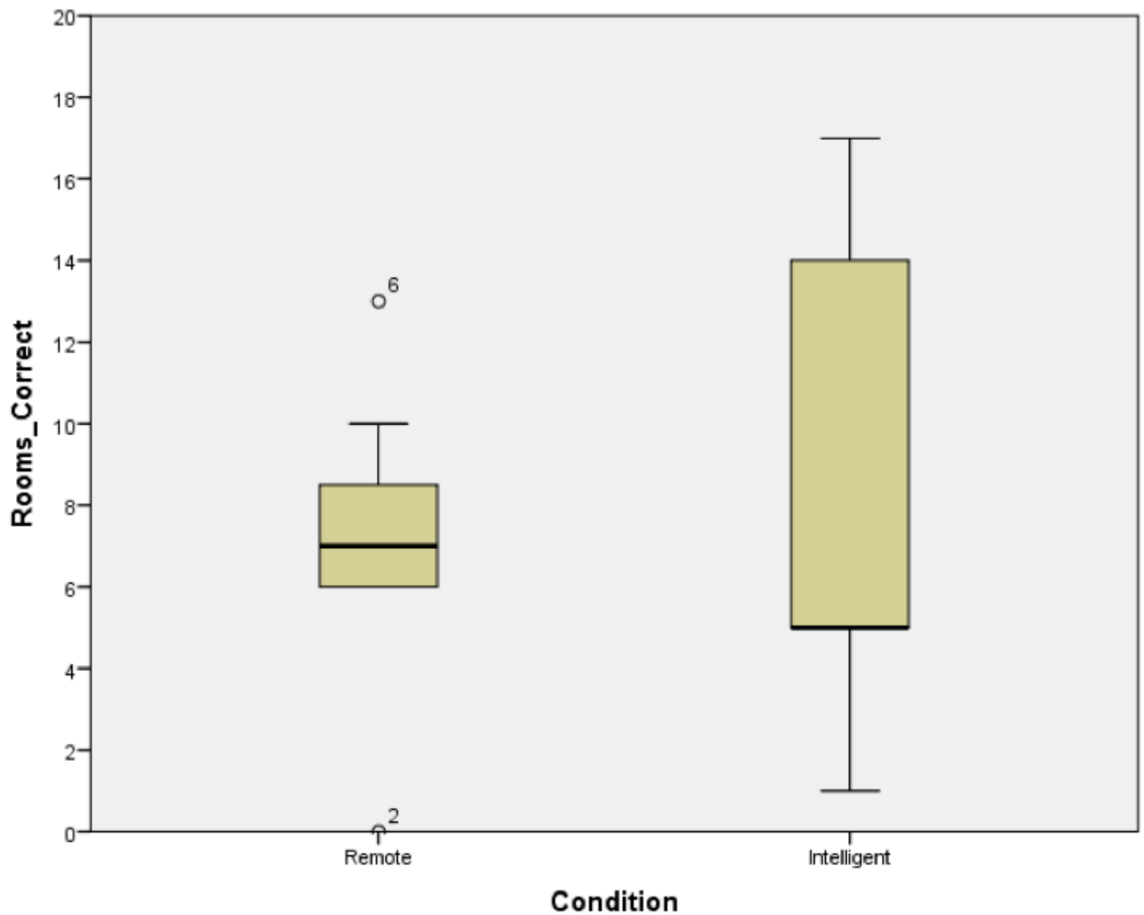

Note: *Data collection is in progress - results are preliminary*

Human-in-the-Loop Planning is making inroads at ICAPS..

- Several papers that handle these challenges of Human-Aware Planning have been presented at the recent ICAPS (and AAAI and IJCAI)
  - Significant help from applications track, robotics track and demonstration track
  - Several planning-related papers in non-ICAPS venues (e.g. AAMAS and even CHI) have more in common with the challenges of Human-aware planning
- ..so consider it for your embedded planning applications

Challenges in Human-in-the-Loop/Human-Aware Planning & Decision Support

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

Eigen Slide

--Human Factors