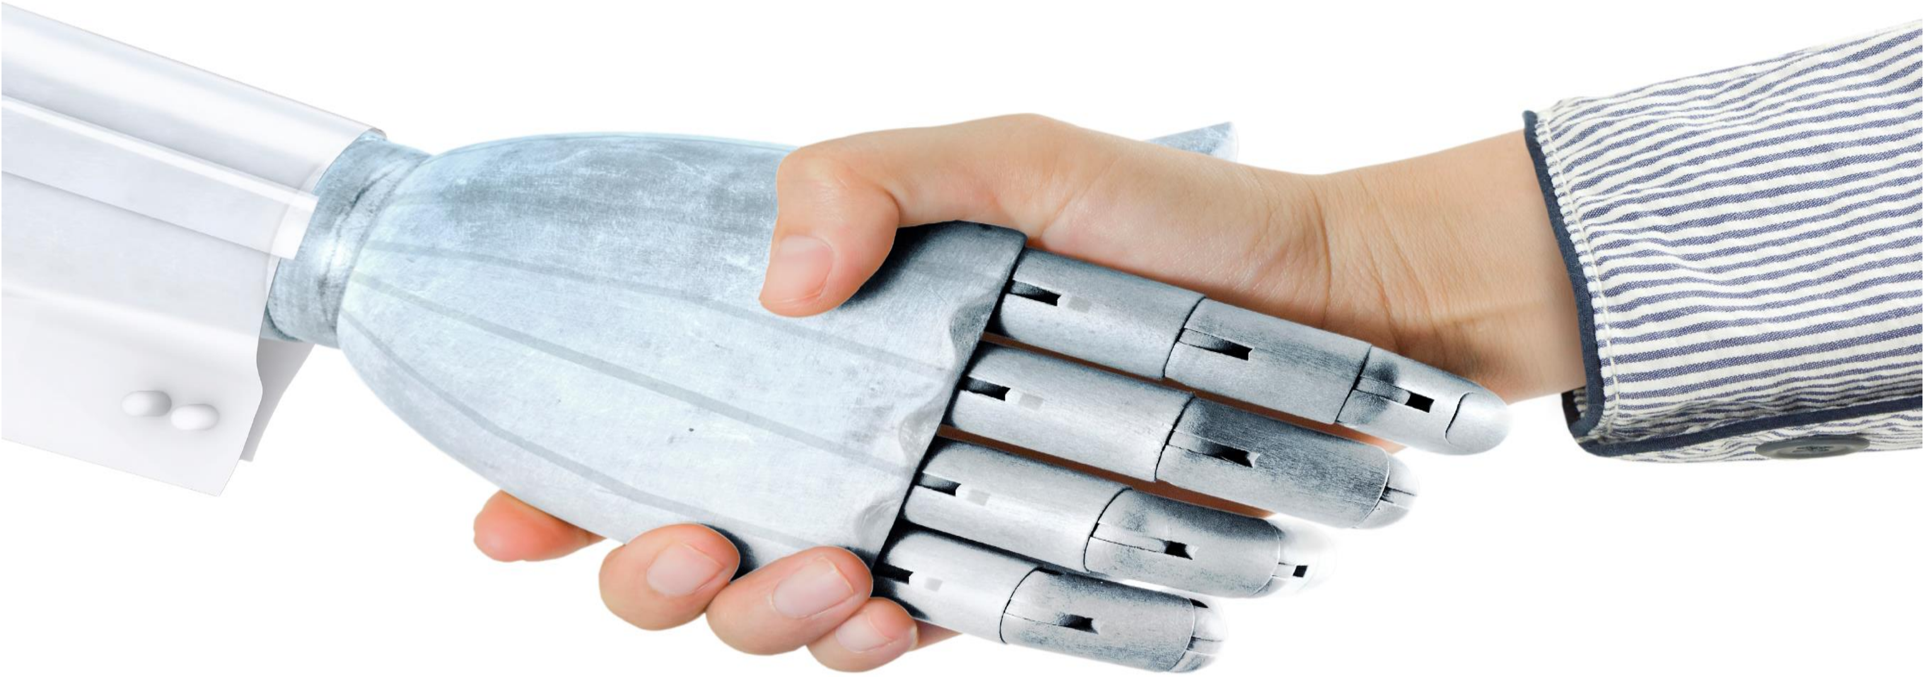

Humans and Machines of Like Mind: Augmenting Humans Through Collaboration in Decision-Making Tasks

Julie Shah

Associate Professor

Department of Aeronautics and Astronautics

Computer Science and Artificial Intelligence Lab

Current State in Human-Robot Teaming


Amazon Robotics


BMW Spartanburg, SC

Current State in Human-Robot Teaming

Amazon Robotics

BMW Spartanburg, SC

Penelope Surgical Instrument Server

- Coexistence but *not* Collaboration.

Inferring hidden mental states enables richer, flexible human-machine teaming

Inferring hidden mental states enables richer, flexible human-machine teaming

- Problem: Current models for HRI and teamwork are based on empirical observation – which works well for highly structured interaction.

Inferring hidden mental states enables richer, flexible human-machine teaming

- Problem: Current models for HRI and teamwork are based on empirical observation – which works well for highly structured interaction.
- Richer collaborations require inferring human mental state.

Inferring hidden mental states enables richer, flexible human-machine teaming

- Problem: Current models for HRI and teamwork are based on empirical observation – which works well for highly structured interaction.
- Richer collaborations require inferring human mental state.

Machine observables $\xrightarrow{?}$ Mental State

Inferring hidden mental states enables richer, flexible human-machine teaming

- Problem: Current models for HRI and teamwork are based on empirical observation – which works well for highly structured interaction.
- Richer collaborations require inferring human mental state.

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision?

- [Hidden State] How does each teammate's prior experience inform their decision-making?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – Do we have consensus?
- [Hidden State] How does each teammate's prior experience inform their decision-making?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – Do we have consensus?

- [Hidden State] How does each teammate's prior experience inform their decision-making? – How do we aggregate data from prior experience?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

What is the current state of our commitment to each decision?

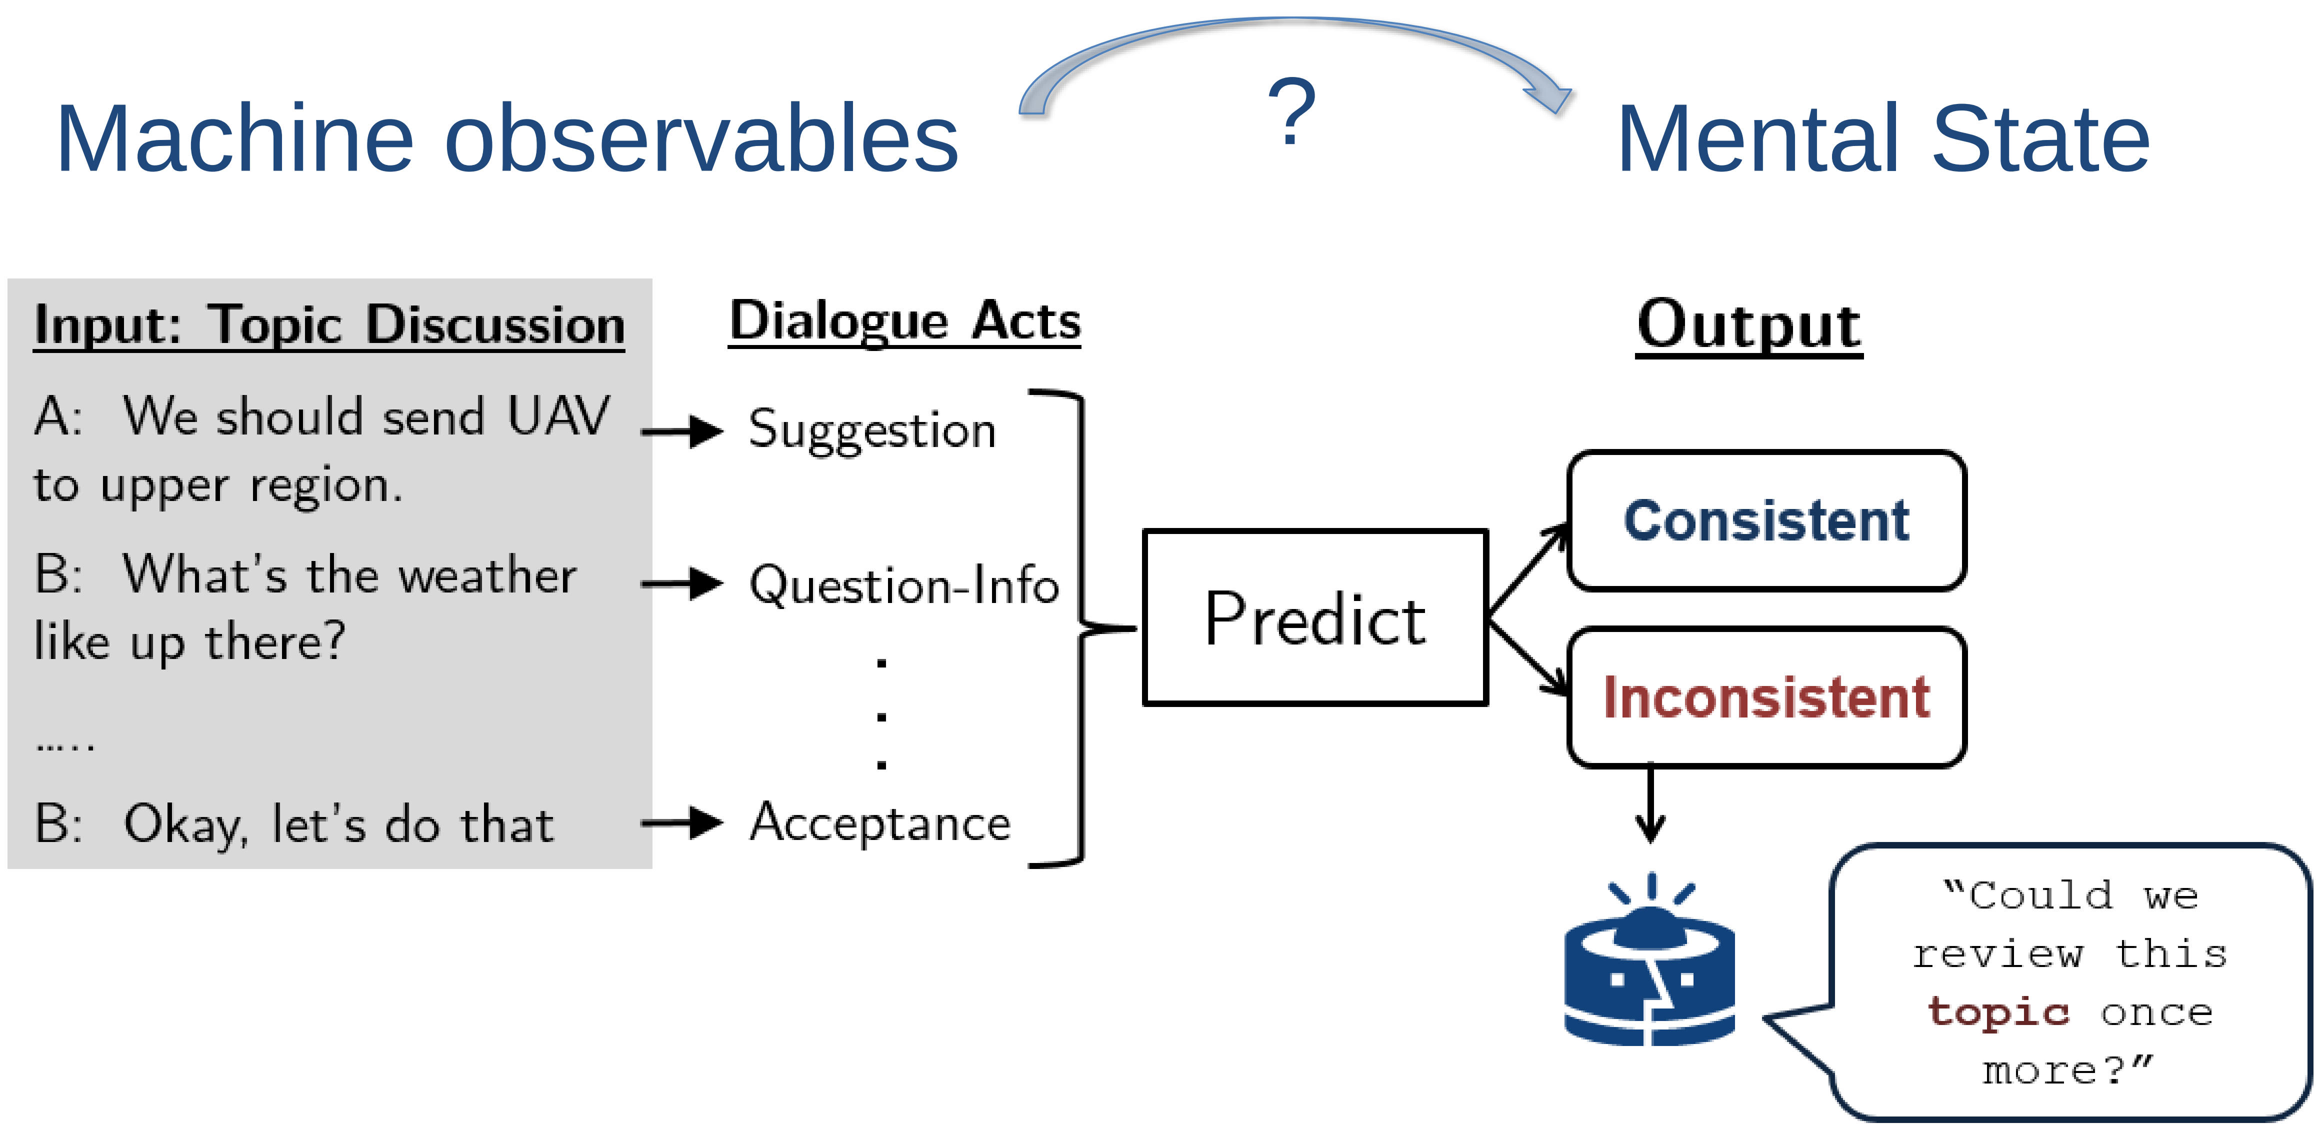

What is the current state of our commitment to each decision?

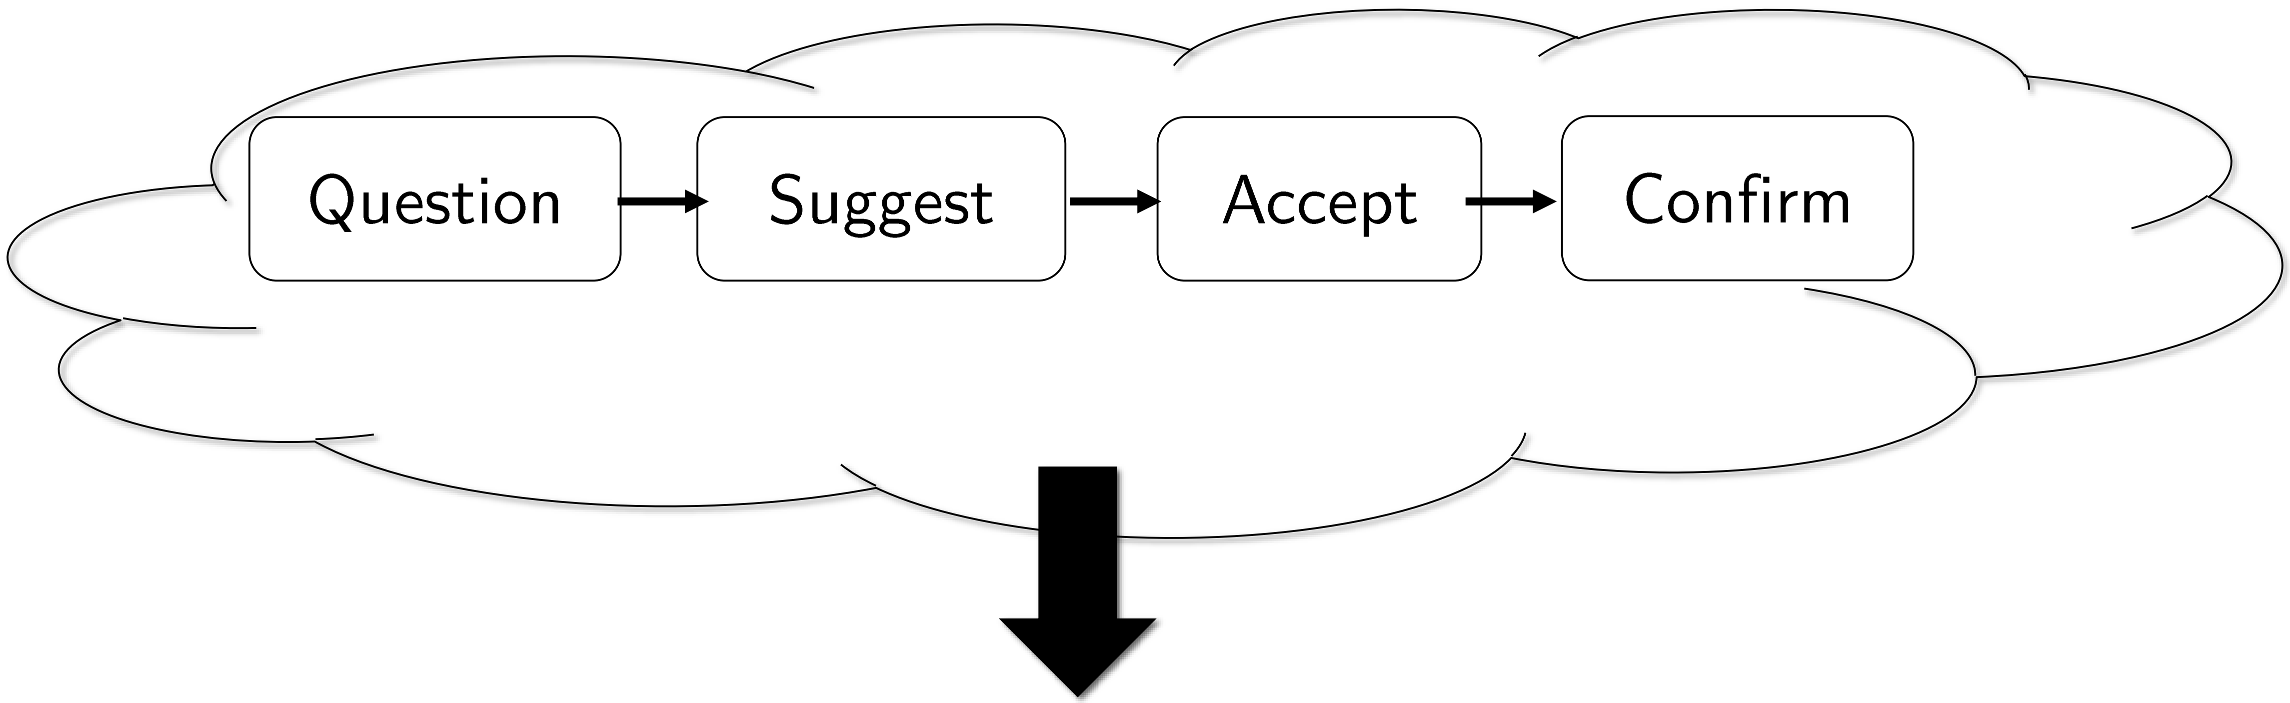

– Unendorsed option
– Partner decidable option
– Proposal
– Commit

B. Eugenio et al. (2000) The agreement process: An empirical investigation of human–human computer-mediated collaborative dialogs, International Journal on Human Computer Studies.

Machines that Learn from Listening to the Team

- **Approach:** learn conversational patterns from a large corpus of team meetings (AMI corpus ~100,000 utterances)
- **Features:** use of *dialogue acts* focused on capturing speakers' commitment process[1]
  - Unendorsed option: speaker "lays" an option with no subsequent actions from others
  - Partner decidable option: speaker presents option that requires further balancing of info.
  - Proposal: speaker presents an option to be accepted/rejected by the group
  - Commit: speaker indicates a full commitment towards an option

- Low-level classifier: Automatic tagging of dialogue acts (~80% accuracy)
- High-level classifier: HMM inference on group consensus[2] (~66% accuracy)

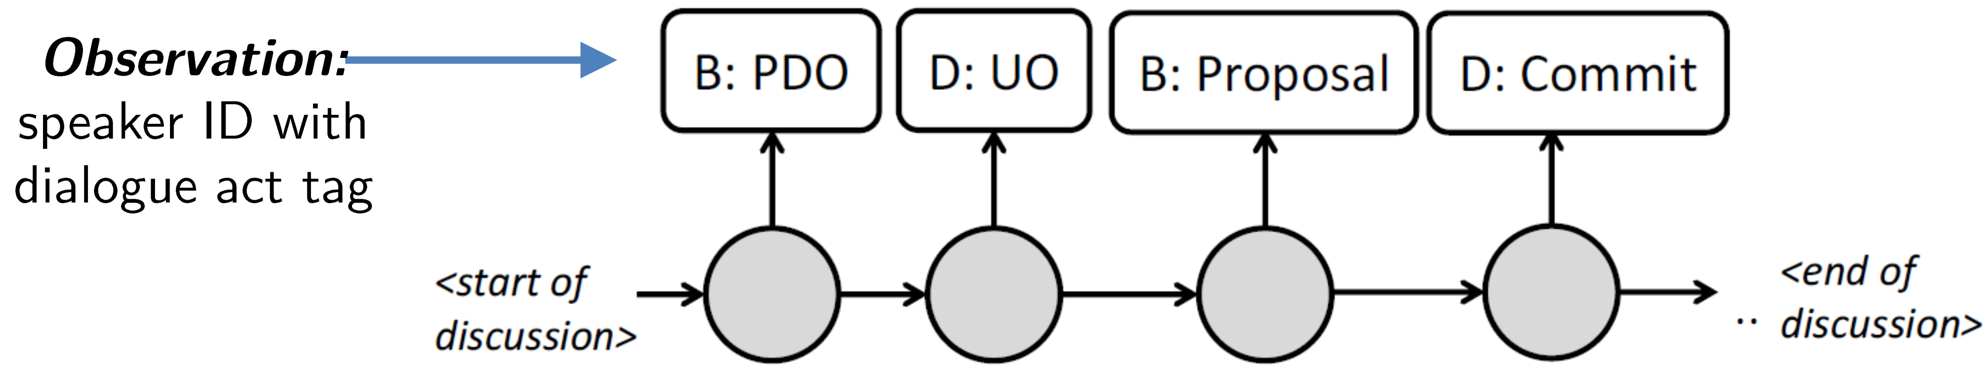

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

Machines that Learn from Listening to the Team

PREDICTION PERFORMANCE OF $HMM_{EUGENIO}$ AND BASELINES

	$\|O\|$	Acc. [%]	Rec. [%]	Prec. [%]	F1 [%]	FPR [%]
HMM_{DAs_full}	11	50.7	29.3	23.1	25.8	40.4
HMM_{DAs}	4	51.4	36.5	31.0	33.5	41.1
$HMM_{Eugenio}$	4	**62.1**	**44.7**	**43.8**	**44.2**	**29.5**

Meeting Phase	Discussion
Project kick-off	Getting acquainted with one another and discussing the project goals
Functional design	Setting user requirements, technical functionality and working design
Conceptual design	Determining conceptual specifications for components, properties and materials
Detailed design	Finalizing user interface and evaluating the final product

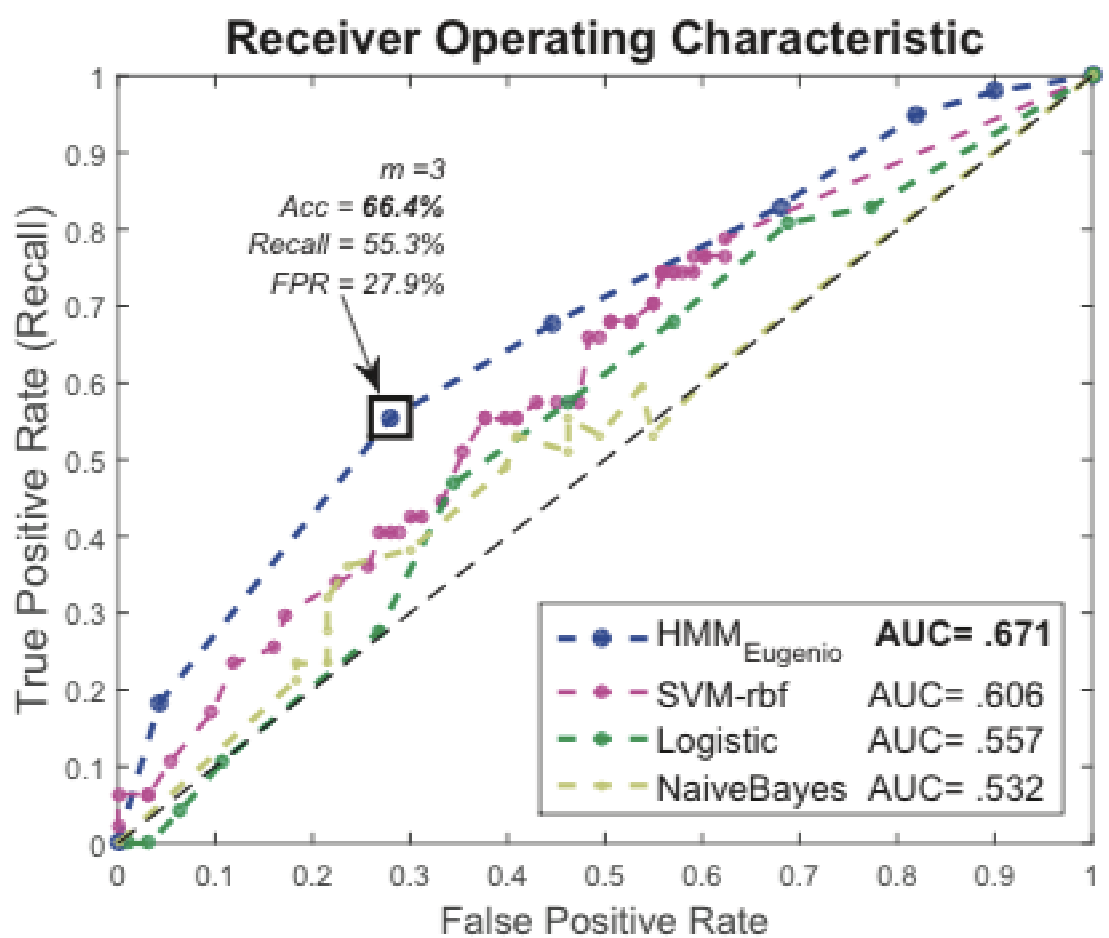

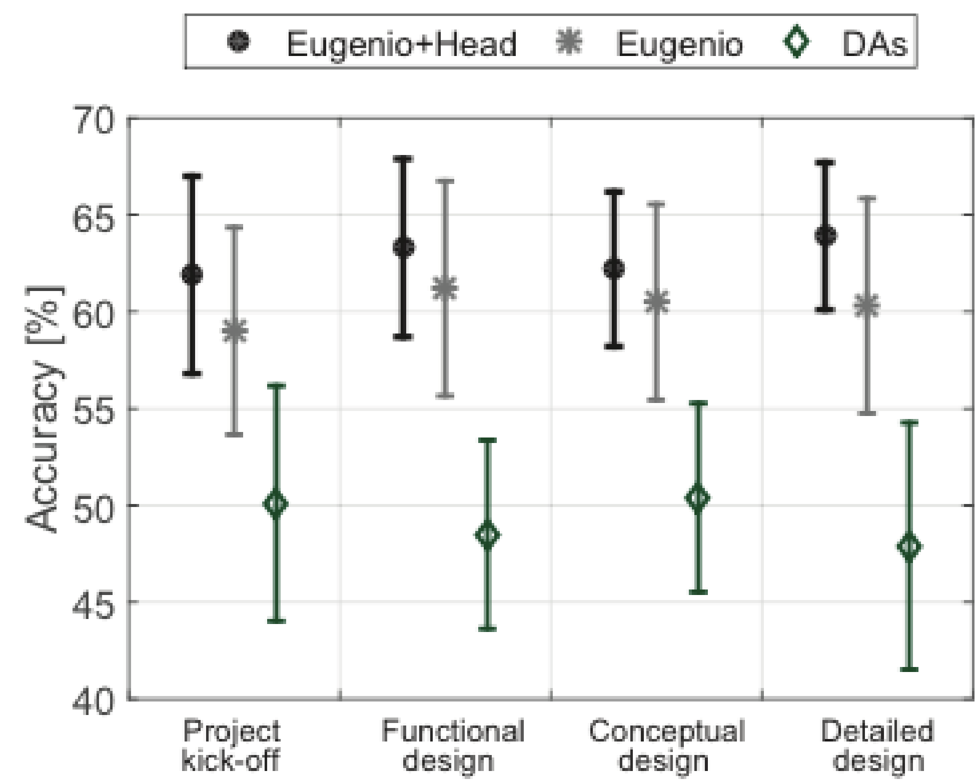

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

Evaluation with Live Teams

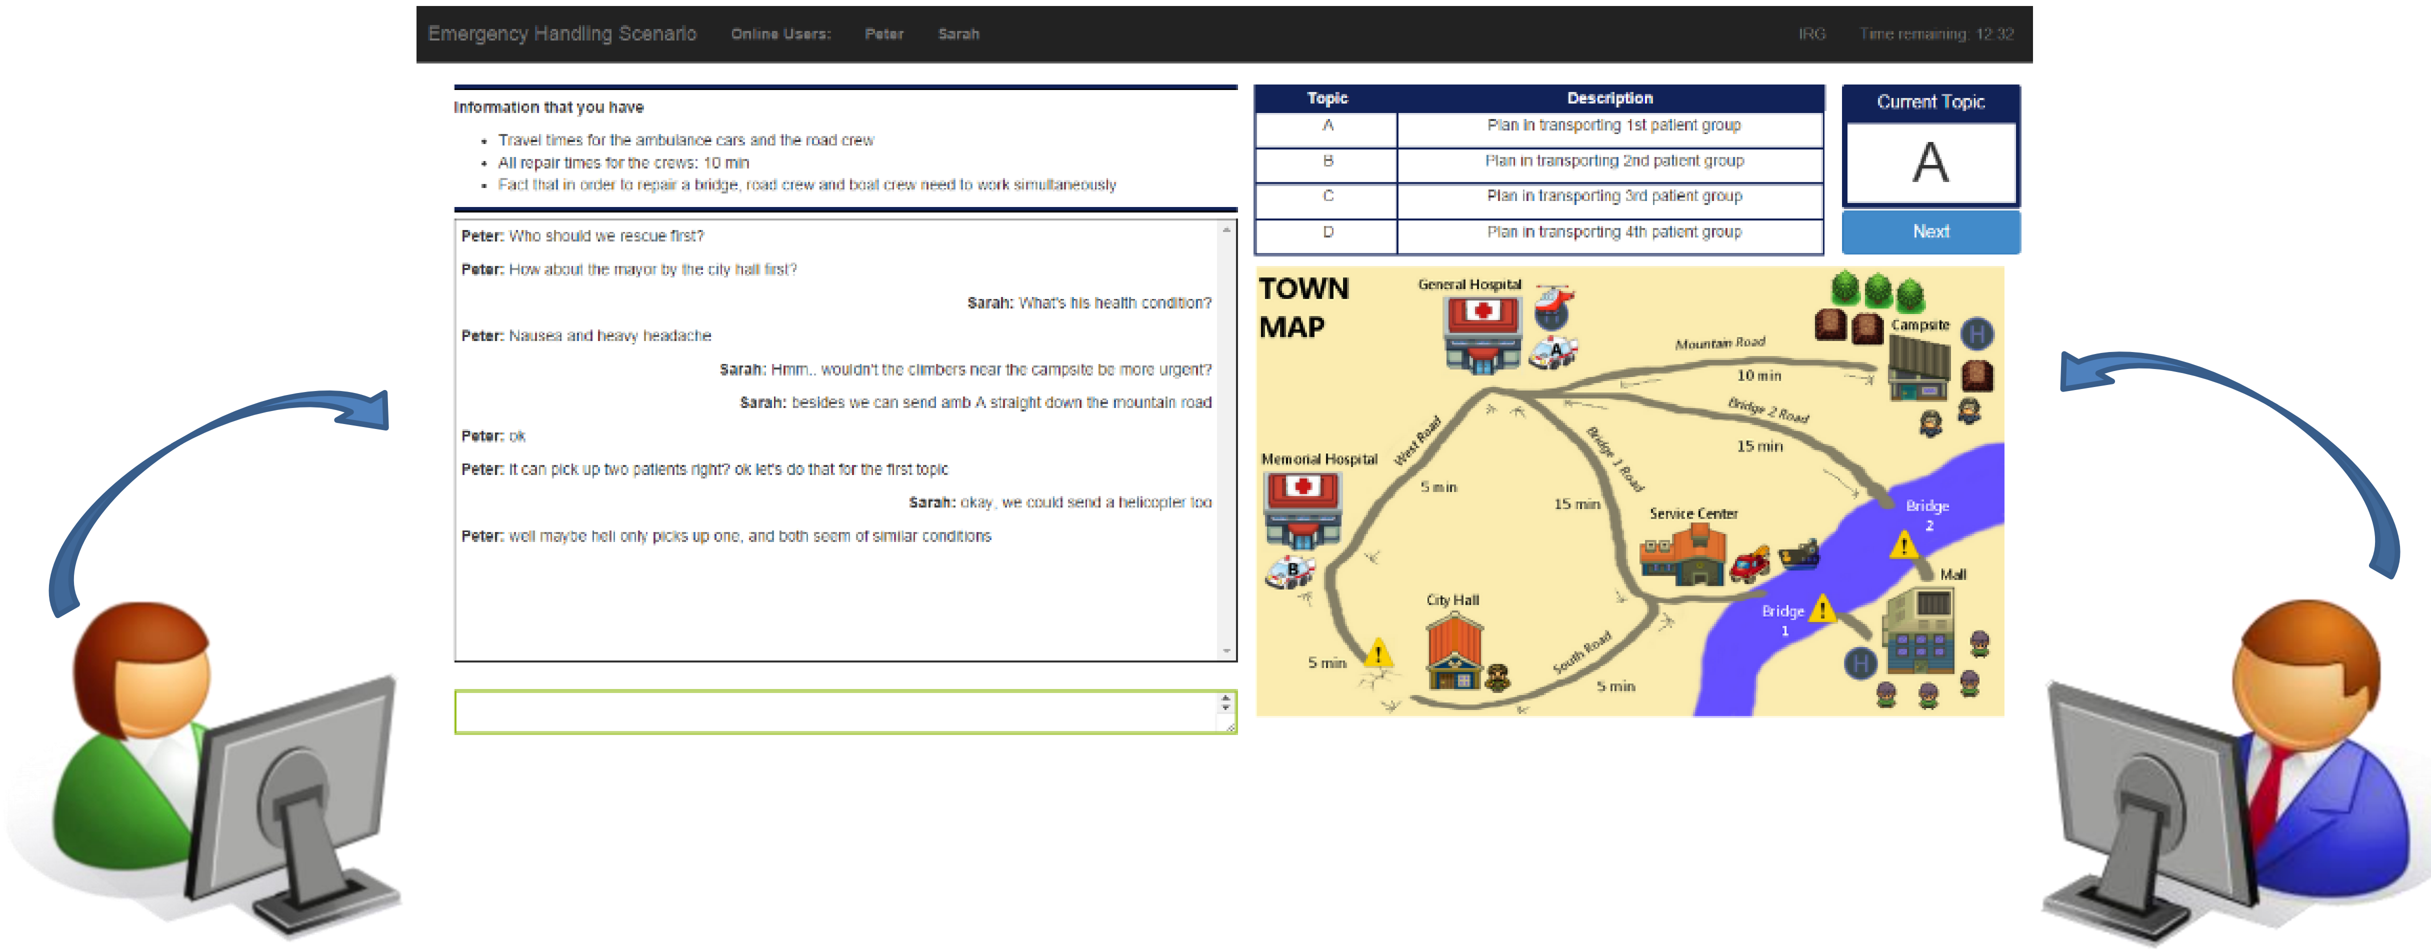

- **Evaluation:** assess utility of review in human team planning[3]
  - Simulated disaster response scenario
  - Communication through web chat interface (n=15 teams of two)
  - System identifies which topics most ideal to review

Evaluation with Live Teams

- **Findings:** statistically significant improvement (~18%) in objective measures of teams' consistency of understanding with intelligent review system

Treatment level	Definition
1. Adaptive review	System suggests review of the two topics with the **lowest** predicted c-scores (*weak* topics)
2. Maladaptive review	System suggests review of the two topics with the **highest** predicted c-scores (*strong* topics).

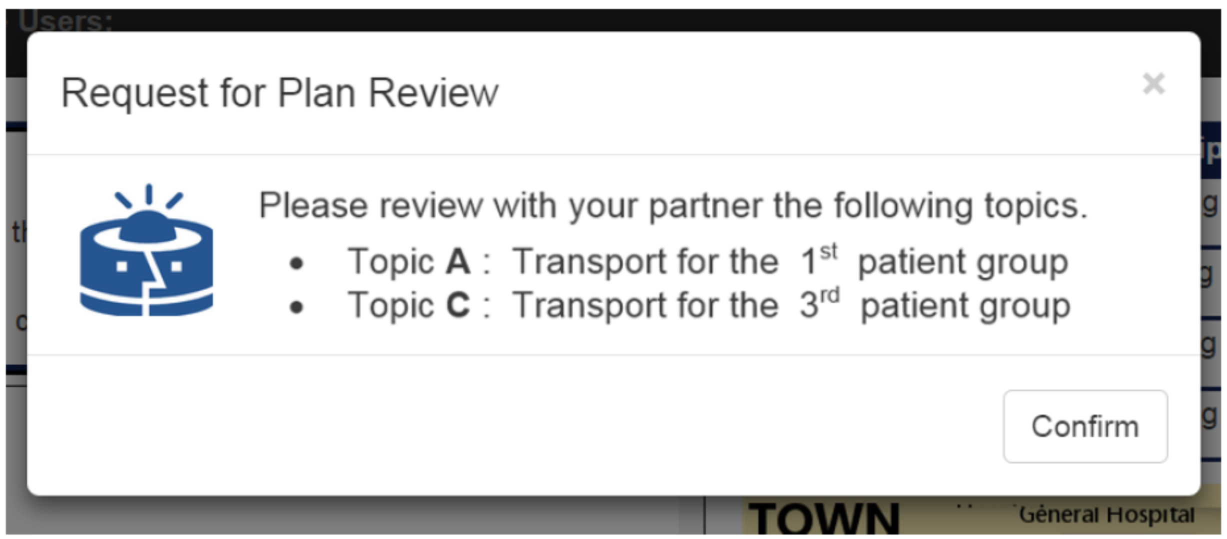

Measure	Questionnaire Items
Perceived utility	*"The review phase of topics suggested by the system helped my teammate and I reach a stronger understanding over those topics."*
Perceived recall	*"The system suggested the two topics where there was potential for lack of understanding between my teammate and I."*

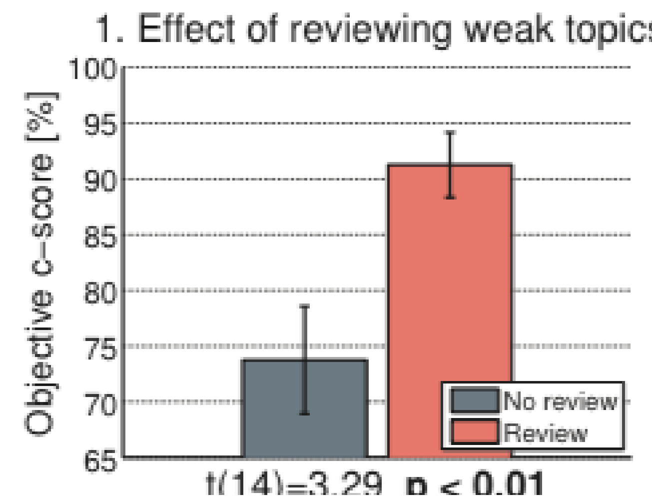

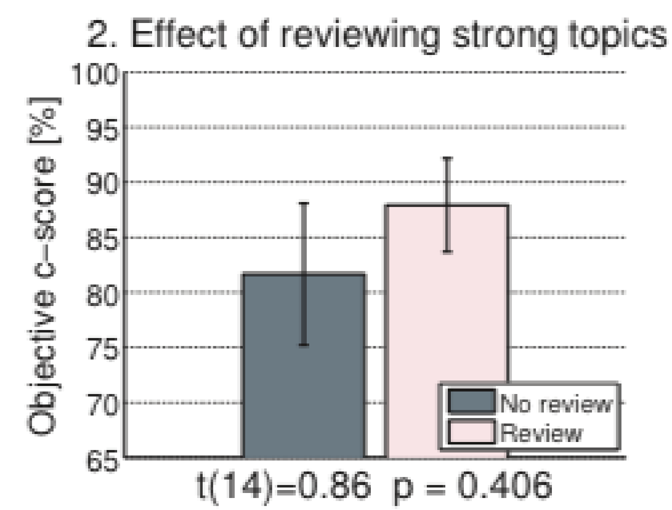

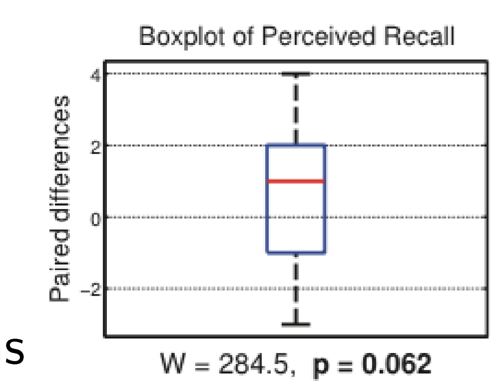

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

Inferring hidden mental states enables richer flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision
- [Hidden State] How does each teammate's prior experience inform their decision-making?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision
- [Hidden State] How does each teammate's prior experience inform their decision-making? – How do we aggregate data from prior experience?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

How does prior experience inform decision-making?

Machine observables ? Mental State

Feature set that characterizes a person's prior experiences

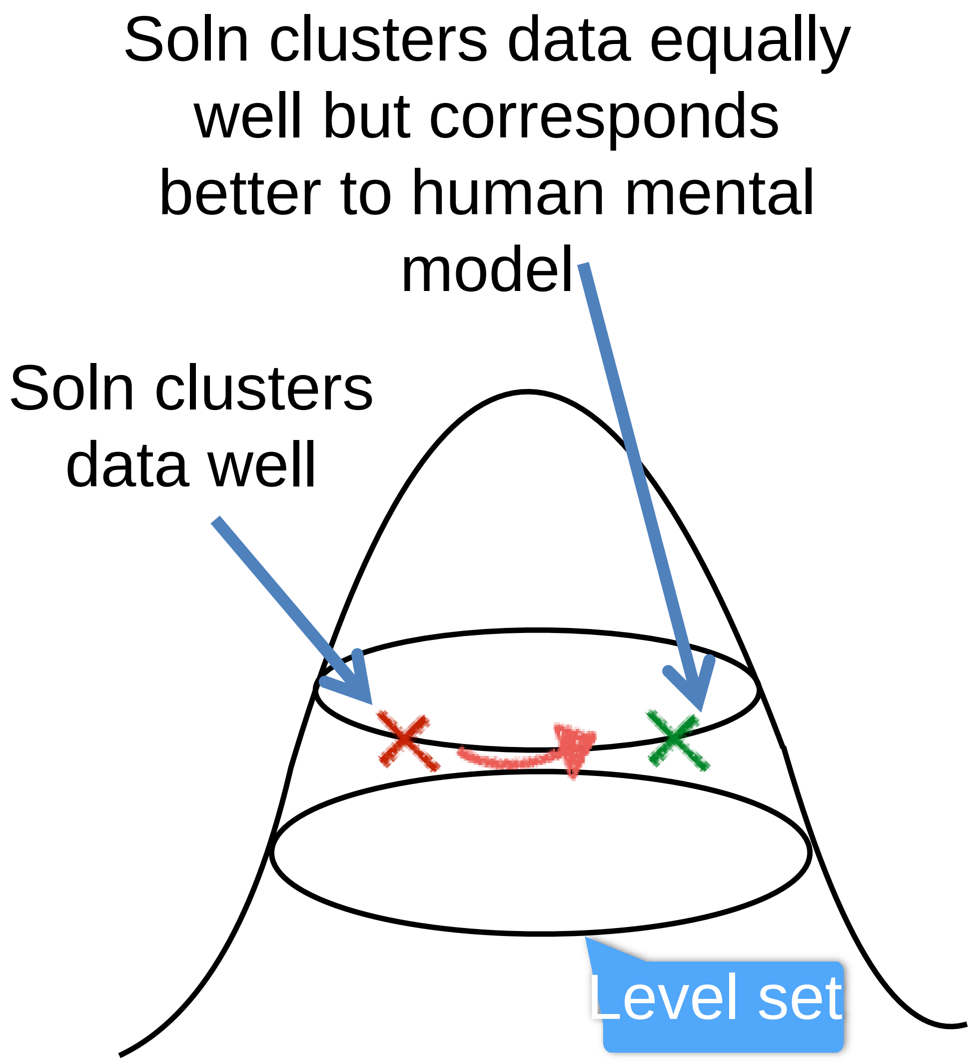

Posterior distribution

How does prior experience inform decision-making?

- Human's tactical decision is based on exemplar-based reasoning (matching and prototyping)
- Skilled fire fighters use recognition-primed decision making — a situation is matched to typical cases

Machine observables → Mental State

examples (prototypes) and subspaces (important features)

[1] M.S. Cohen, J.T. Freeman, and S. Wolf. Metarecognition in time-stressed decision making: Recognizing, critiquing, and correcting. *Human Factors*, 1996. .
[2] A. Newell and H.A. Simon. *Human problem solving*. Prentice-Hall Englewood Cliffs, 1972
[3] G.A. Klein. Do decision biases explain too much. *HFES*, 1989.

How does prior experience inform decision-making?

Machine observables ? Mental State

Flour
Vanilla
Egg
Salt
Blueberry
Sugar
Baking Powder
Soy Sauce
Chicken
Sesame Seeds
Oil
Salt

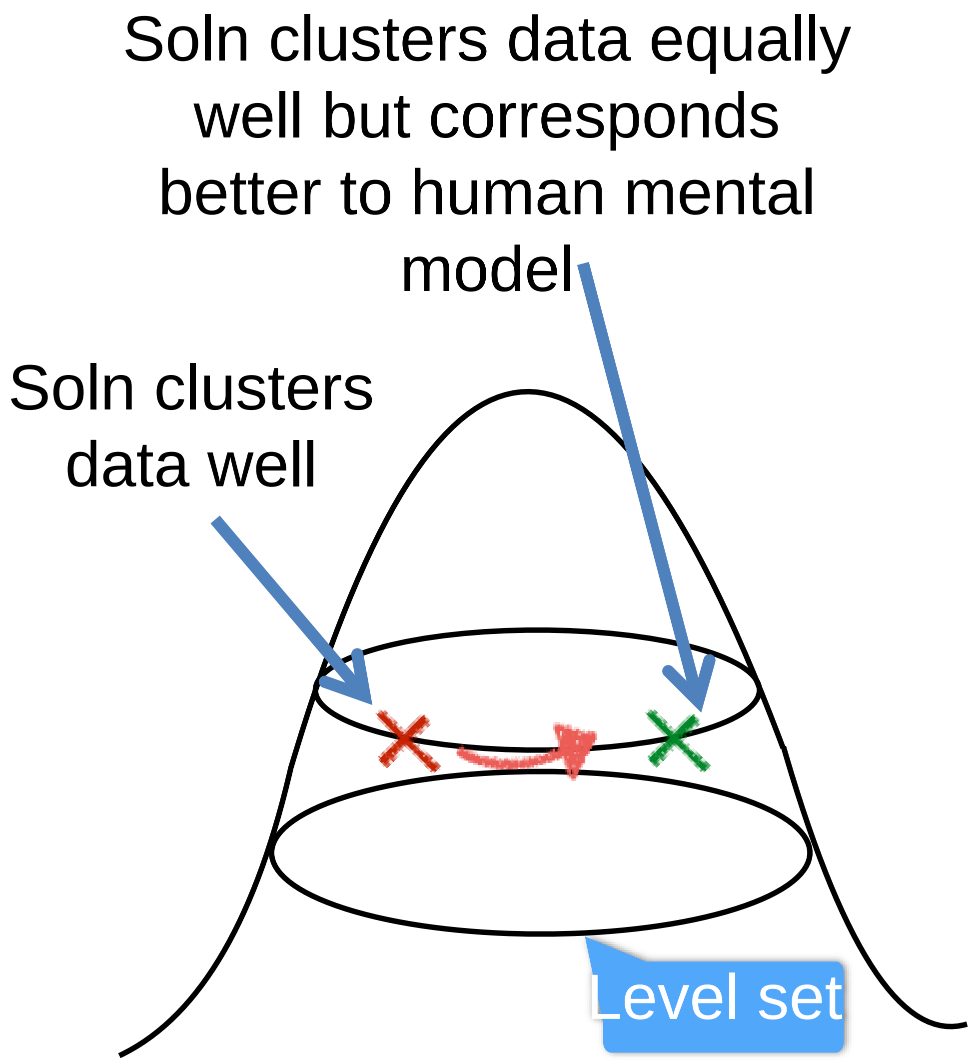

Posterior distribution

How does prior experience inform decision-making?

- Joint inference on **prototypes**, **subspaces** and cluster labels

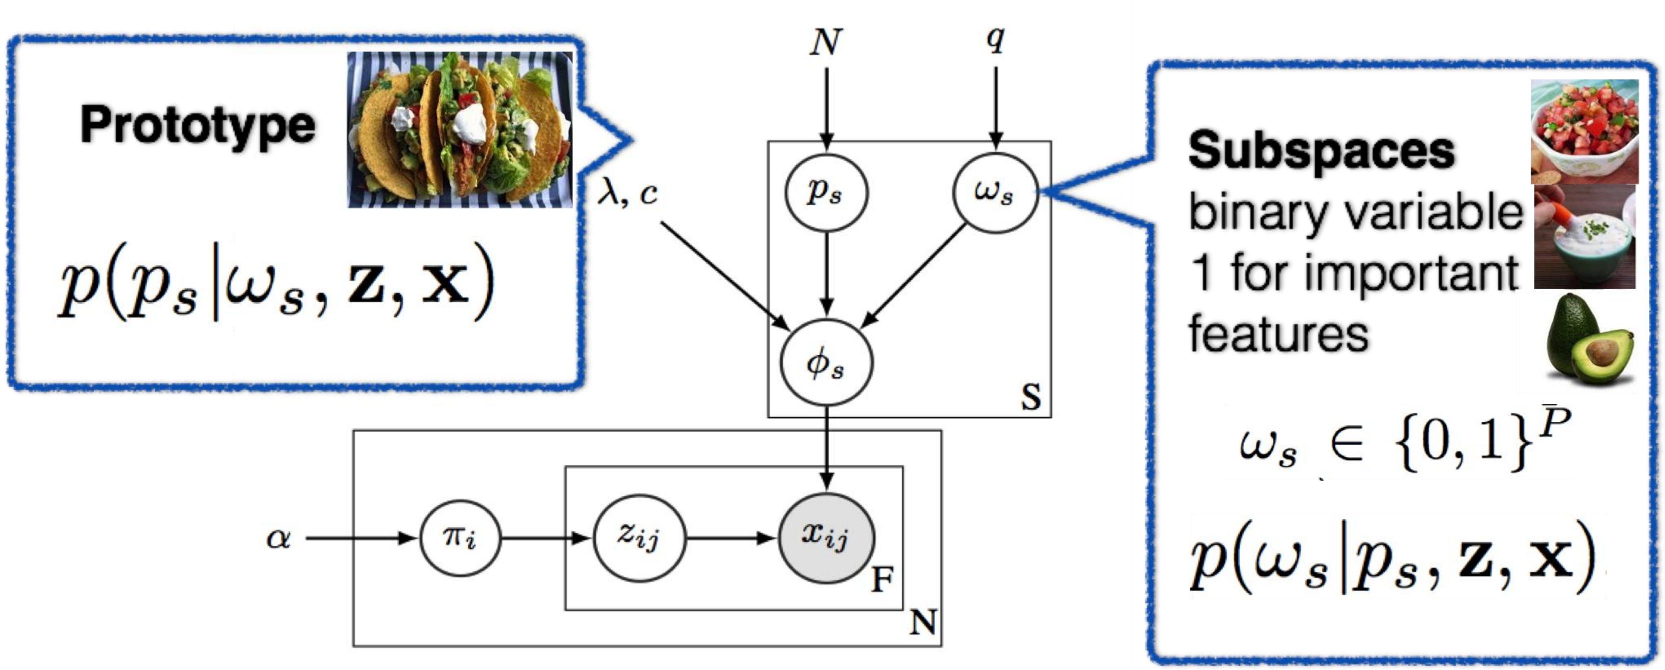

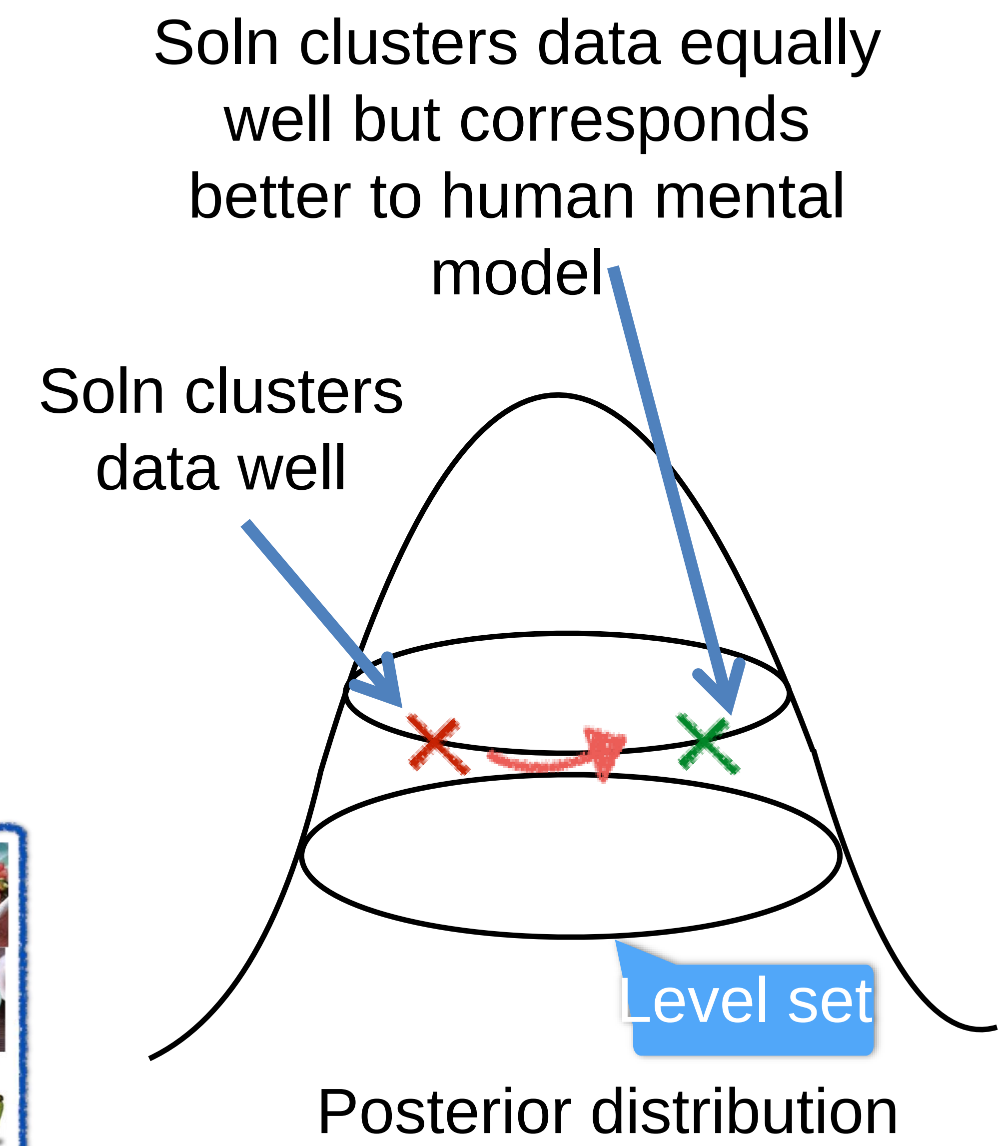

Kim, Rudin & Shah NIPS'14

Classification Performance on Standard Datasets

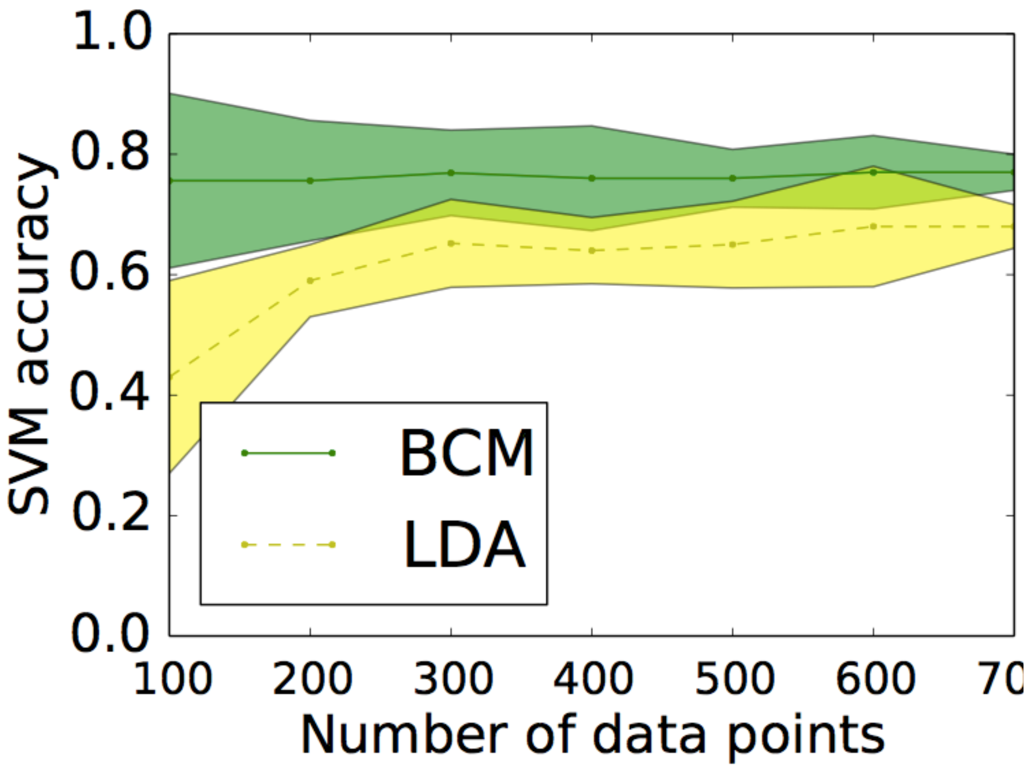

- Handdigit
- dataset

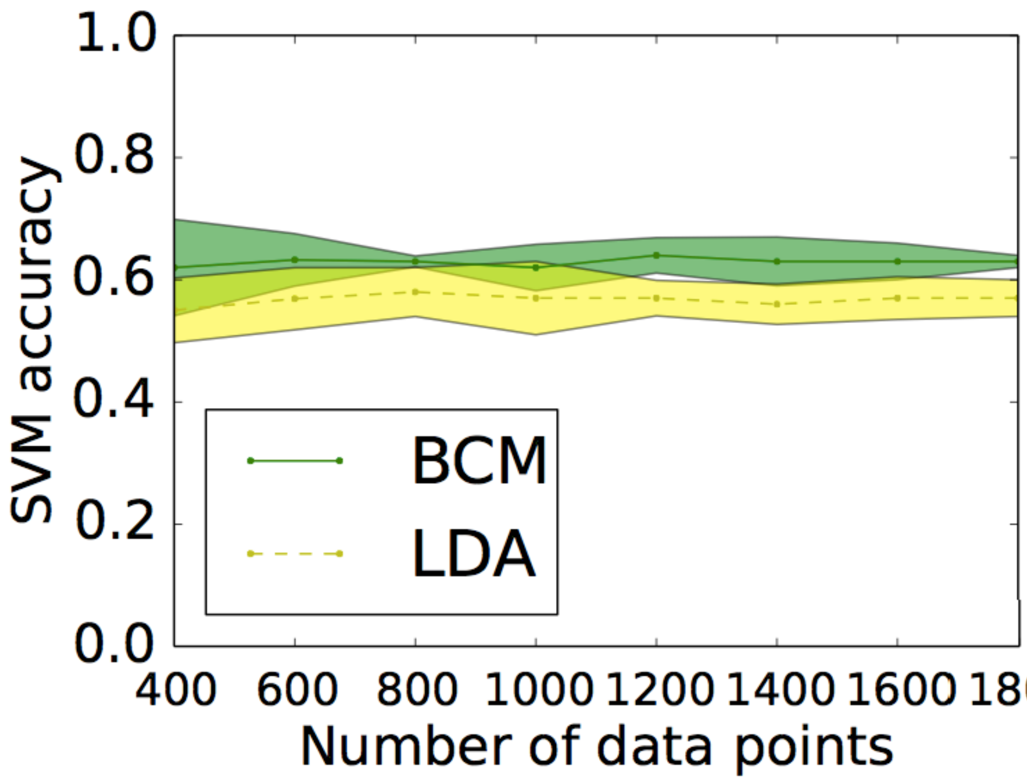

20Newsgroups dataset

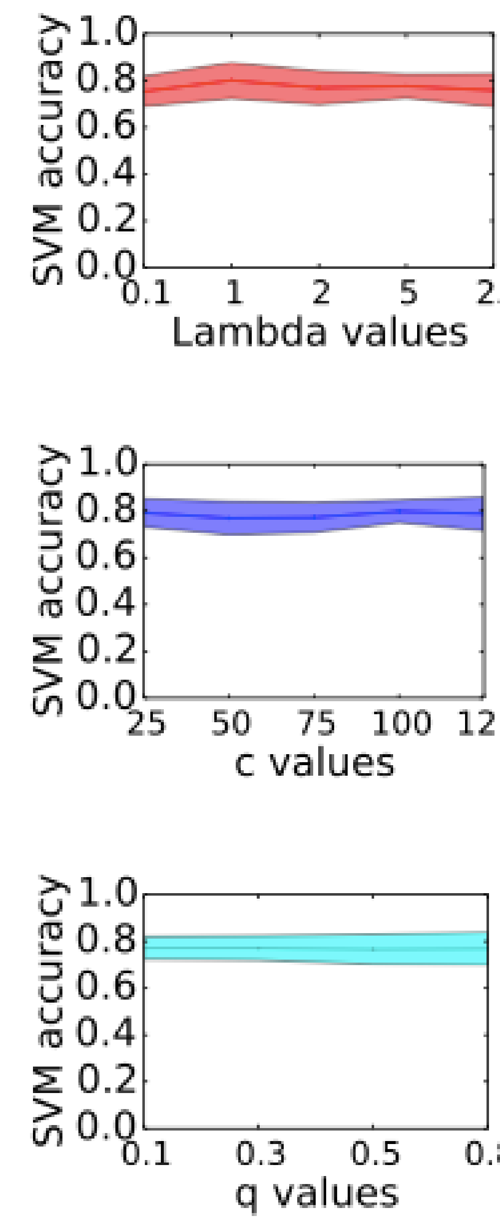

Kim, Rudin & Shah NIPS’14

Gibbs sampling iteration

Assessing Compatibility with Human Decision-Making

a new data point to be classified

- Participant's task is to assign the ingredients of a specific dish (a new data point) to a cluster
- Each cluster is explained using either BCM or LDA.

Assessing Compatibility with Human Decision-Making

Specific dish

flour
egg
cranberry
brown sugar
pumpkin
baking powder
oil

a new data point to be classified

Examples of types of dishes

Dish 1 ingredients	Dish 2 ingredients
flour	soy sauce
vanila	chicken
egg	sugar
salt	semame seeds
sugar	rice
blue berry	oil
baking powder	salt

- 384 classification questions asked to 24 people

Clusters explained using

1. **BCM** : ingredients of the prototype recipe for each cluster without recipe name nor subspaces for fairness
2. **LDA**: representative ingredients of each cluster

- Statistically significantly better performance with Bayesian Case Model for clustering (85.9% v.s. 71.3%)

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – Do we have consensus?
- [Hidden State] How does each teammate's prior experience inform their decision-making? – How do we aggregate data from prior experience?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – ~~Do we have consensus?~~
- [Hidden State] How does each teammate's prior experience inform their decision-making? – ~~How do we aggregate data from prior experience?~~

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – What did we agree to?

- [Hidden State] How does each teammate's prior experience inform their decision-making? – What will we do next?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Human Team Planning

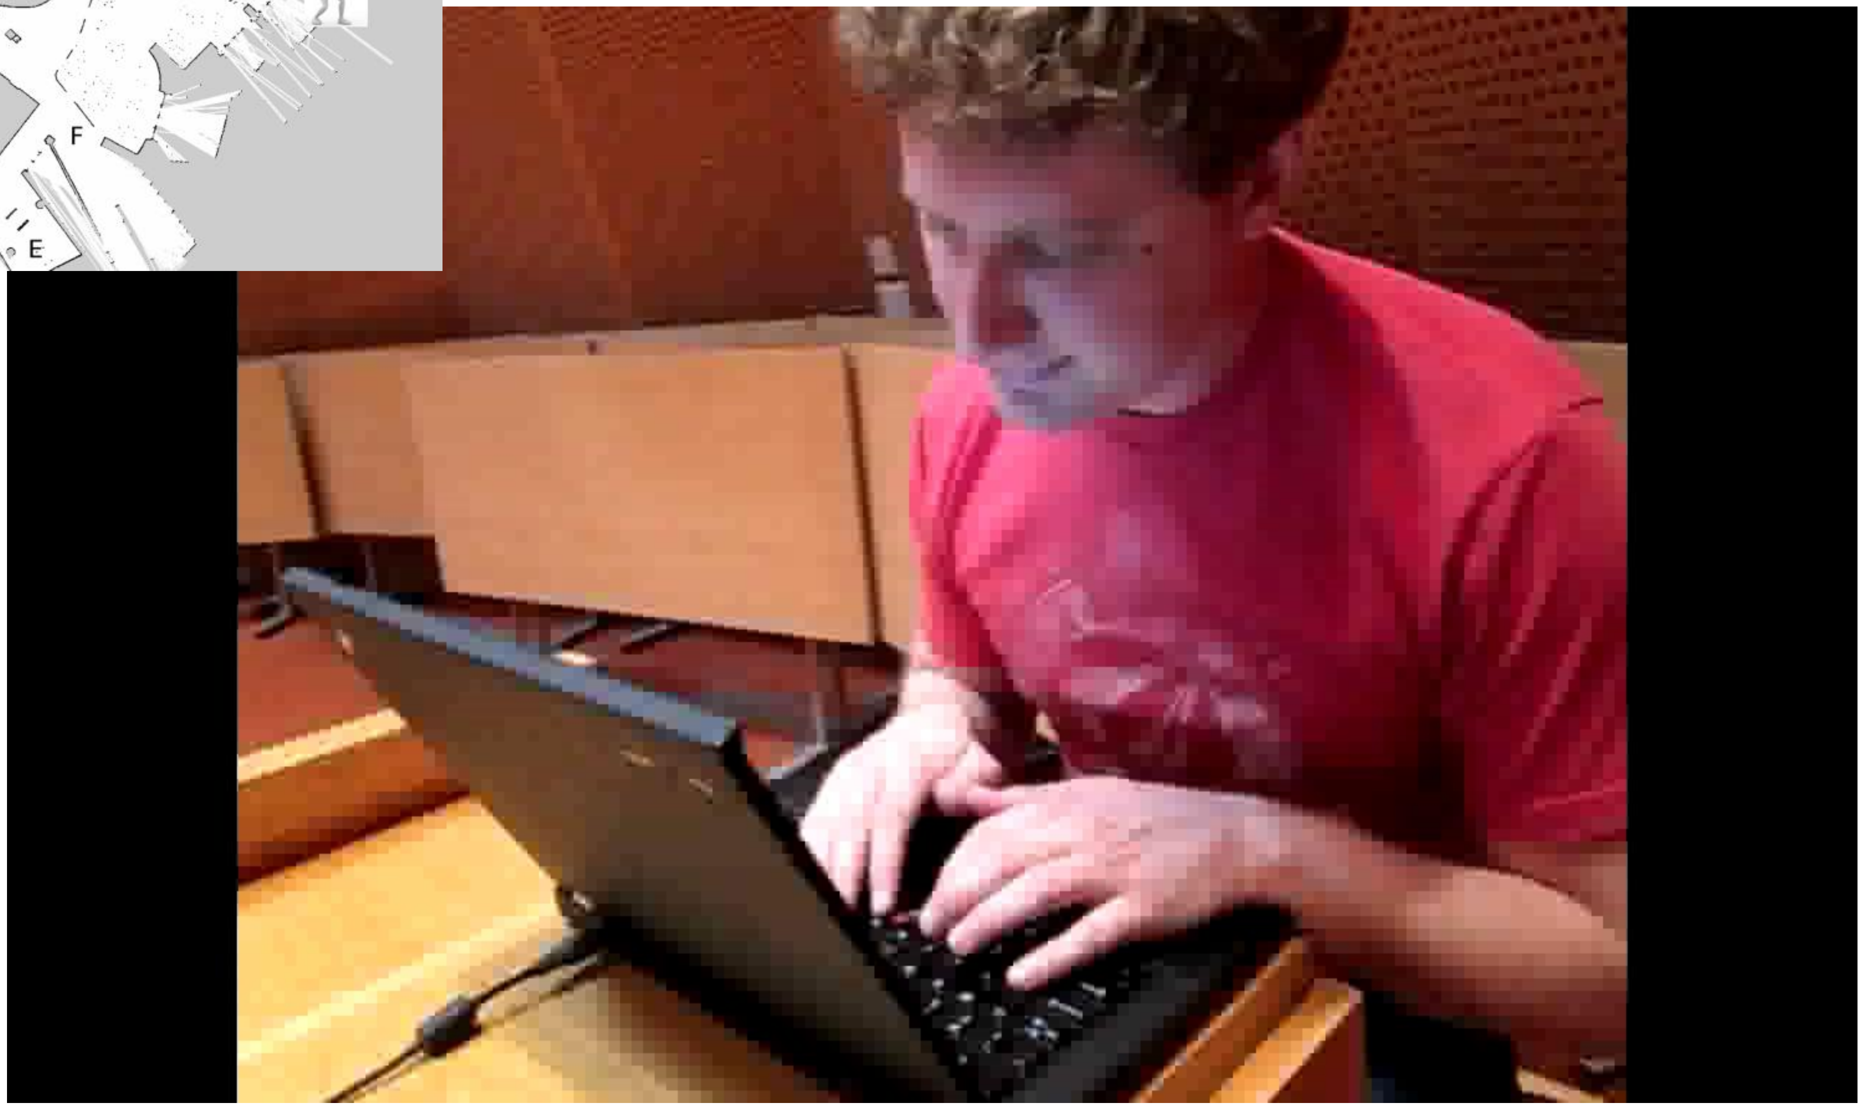

Scenario:

- 8 rooms
- B, D, G rooms have patients that need to be rescued
- C, F rooms have leaking valves that need to be fixed
- Robots must inspect the rooms before human crews enter.

Generative model with logic-based prior improves efficiency of inference process

1. Sample a plan from

$$p(plan) \propto \begin{cases} e^{\alpha} & \text{if plan is valid} \\ 1 & \text{if plan is invalid} \end{cases}$$

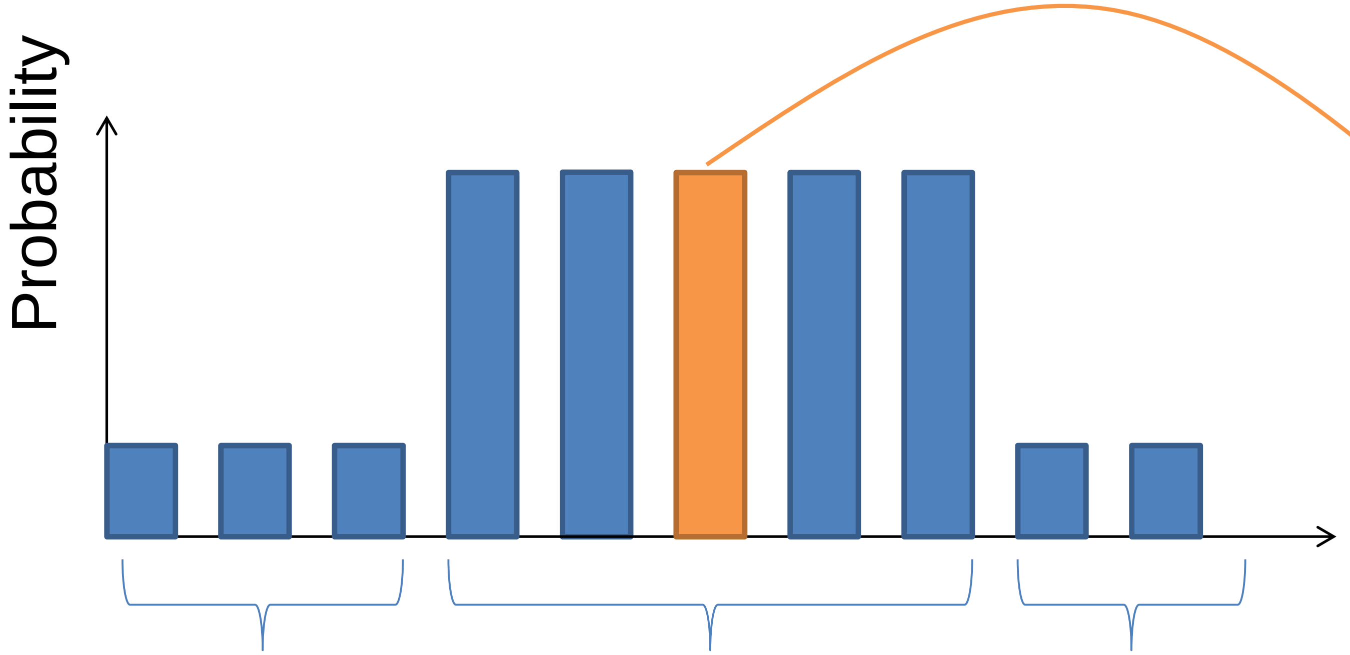

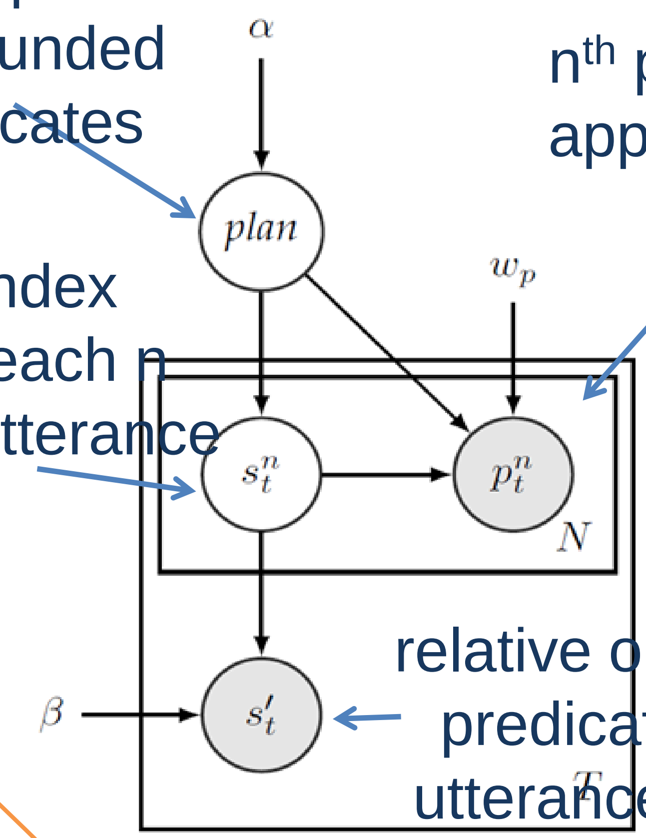

Step 1. Do A and B

Step 2. Do C,D and E

Step 3. Do F and G

B. Kim, C. Chacha, J. Shah (2015) Journal of Artificial Intelligence Research

Generative model with logic-based prior improves efficiency of inference process

Scenario:

- 8 rooms
- B, D, G rooms have patients that need to be rescued
- C, F rooms have leaking valves that need to be fixed
- Robots must inspect the rooms before human crews enter.

	Task Allocation		% Seq	Avg.
	% Inferred	% Noise Rej		
PDDL	84	100	91	91
PDDL with missing goals and constants	100	54	75	76
PDDL with missing a constraint	88	77	84	83
No PDDL	85	75	87	82

Technique correctly infers 80-90% of plan, on average.

N=48 distinct plans

Inferring hidden mental states enables richer, flexible human-machine teaming

- [Hidden State] What is the current state of our commitment to each decision? – What did we agree to?
- [Hidden State] How does each teammate's prior experience inform their decision-making? – What will we do next?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

Machines that Learn from Watching People Solve Resource Allocation & Scheduling Problems

Interactive ROBOTICS @MIT

IJCAI'16

- Goal: Emulate problem solving capability of human domain experts.
- Approach: Pairwise rank formulation used to train a machine learning model
  - Define a set of scheduling-relevant features for the problem
    - E.g. deadline, duration of task, earliest time task is available, resources required by task
  - Each observation of expert commitment is described by the feature vector
  - Positive and negative training examples computed through pairwise comparison
    - Differences computed for scheduled versus unscheduled tasks
  - Classifiers trained to predict highest priority next action to take, and whether to take action at time *t*

$$\begin{aligned} &{}^{rank}\theta^{m}_{\langle \tau_i, \tau_x \rangle} := \left[\xi_{\boldsymbol{\tau}}, \gamma_{\tau_i} - \gamma_{\tau_x} \right], y^{m}_{\langle \tau_i, \tau_x \rangle} = 1, \\ &\forall \tau_x \in \boldsymbol{\tau} \backslash \tau_i, \forall O_m \in \boldsymbol{O} | \tau_i \text{ scheduled in } O_m \end{aligned} \tag{1}$$

$$\begin{aligned} &{}^{rank}\theta^{m}_{\langle \tau_x, \tau_i \rangle} := \left[\xi_{\boldsymbol{\tau}}, \gamma_{\tau_x} - \gamma_{\tau_i} \right], y^{m}_{\langle \tau_x, \tau_i \rangle} = 0, \\ &\forall \tau_x \in \boldsymbol{\tau} \backslash \tau_i, \forall O_m \in \boldsymbol{O} | \tau_i \text{ scheduled in } O_m \end{aligned} \tag{2}$$

$$\widehat{\tau_i^*} = \underset{\tau_i \in \boldsymbol{\tau}}{\operatorname{argmax}} \sum_{\tau_x \in \boldsymbol{\tau}} f_{priority}\left(\tau_i, \tau_x\right) \tag{3}$$

$$\begin{aligned} &{}^{act}\phi^{m}_{\tau_i} := \left[\xi_{\boldsymbol{\tau}}, \gamma_{\tau_i} \right], \\ &y^{m}_{\tau_i} = \begin{cases} 1 : \tau_i \text{ scheduled in } O_m \wedge \\ \qquad \tau_i \text{ scheduled in } O_{m+1} \\ 0 : \tau_\emptyset \text{ scheduled in } O_m \end{cases} \end{aligned} \tag{4}$$

Machines that Learn from Watching People Solve Resource Allocation & Scheduling Problems

Interactive ROBOTICS @MIT

IJCAI'16, RSS'16

- Successful application of technique to anti-ship missile defense (with MIT LL)

ONR makes a serious game of missile defense, electronic warfare

BY KEVIN MCCANEY • FEB 04, 2015

A Navy ship being under missile fire is no game, but making a game of that scenario can help sailors prepare for it.

- Successful application to coordination of patient care in a hospital

Machines that Learn from Watching People Solve Resource Allocation & Scheduling Problems

IJCAI’16

- Successful application of technique to anti-ship missile defense (with MIT LL)

ONR makes a serious game of missile defense, electronic warfare

BY KEVIN MCCANEY • FEB 04, 2015

A Navy ship being under missile fire is no game, but making a game of that scenario can help sailors prepare for it.

  - Problem involved 5 decoys, 10 types of threats, 16 game configurations.
  - Dataset: 162 games played by 27 human experts with expertise in ASMD.
    - E.g. deadline, duration of task, earliest time task is available, resources required by task
  - Model trained on 16 demonstrations in which a player mitigated all enemy missiles

  - Average human player’s score: $74,728 \pm 26,824$
  - Learned model’s average score: $87,540 \pm 16,842$
  - Learned scheduling policy performed better than the human demonstrators on more scenarios than vice versa (12 vs. 4 scenarios, $p < 0.011$)

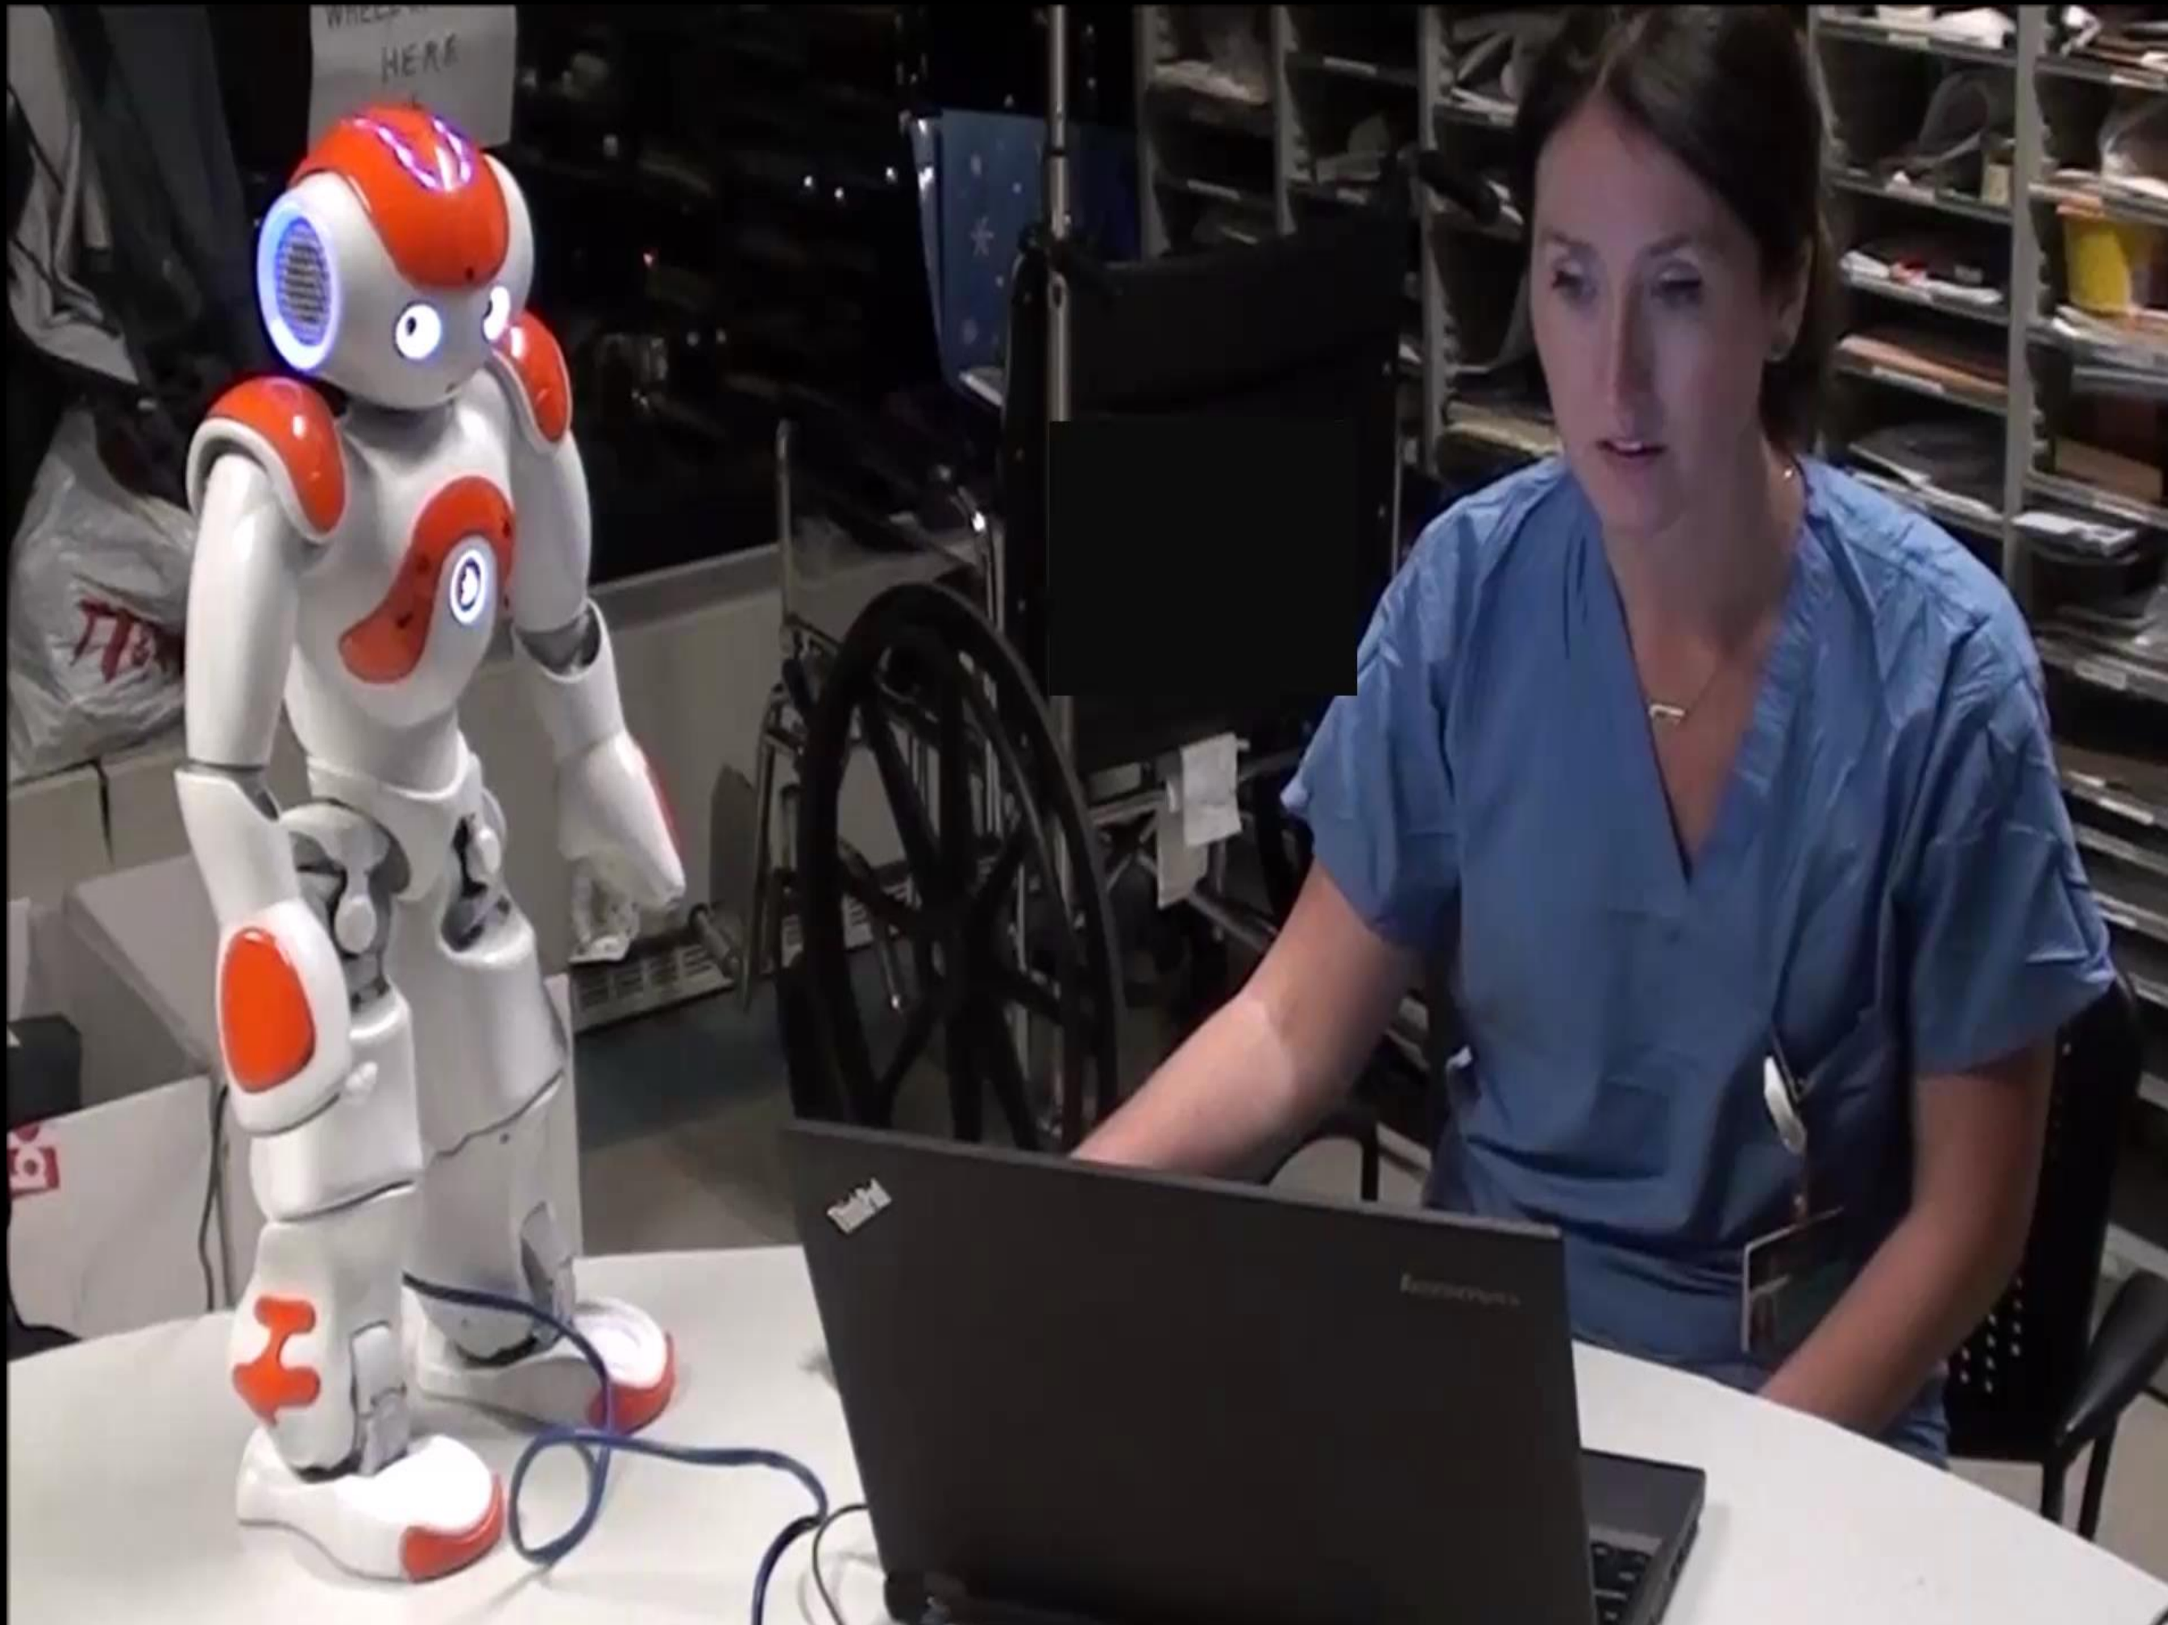
HERE
ThinkPad

Contributions

- Approach: translate well-established cognitive models into new computational models that allow machines to
  - infer our mental state
  - process complex information efficiently

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

- Experiments validate that these models yield richer, flexible human-machine teaming
  - making higher quality shared plans [IEEE THMS'16, RSS'16 JAIR'15]
  - making better sense of big data [NIPS'14]
  - learning complex shared plans from observation [IJCAI'16]