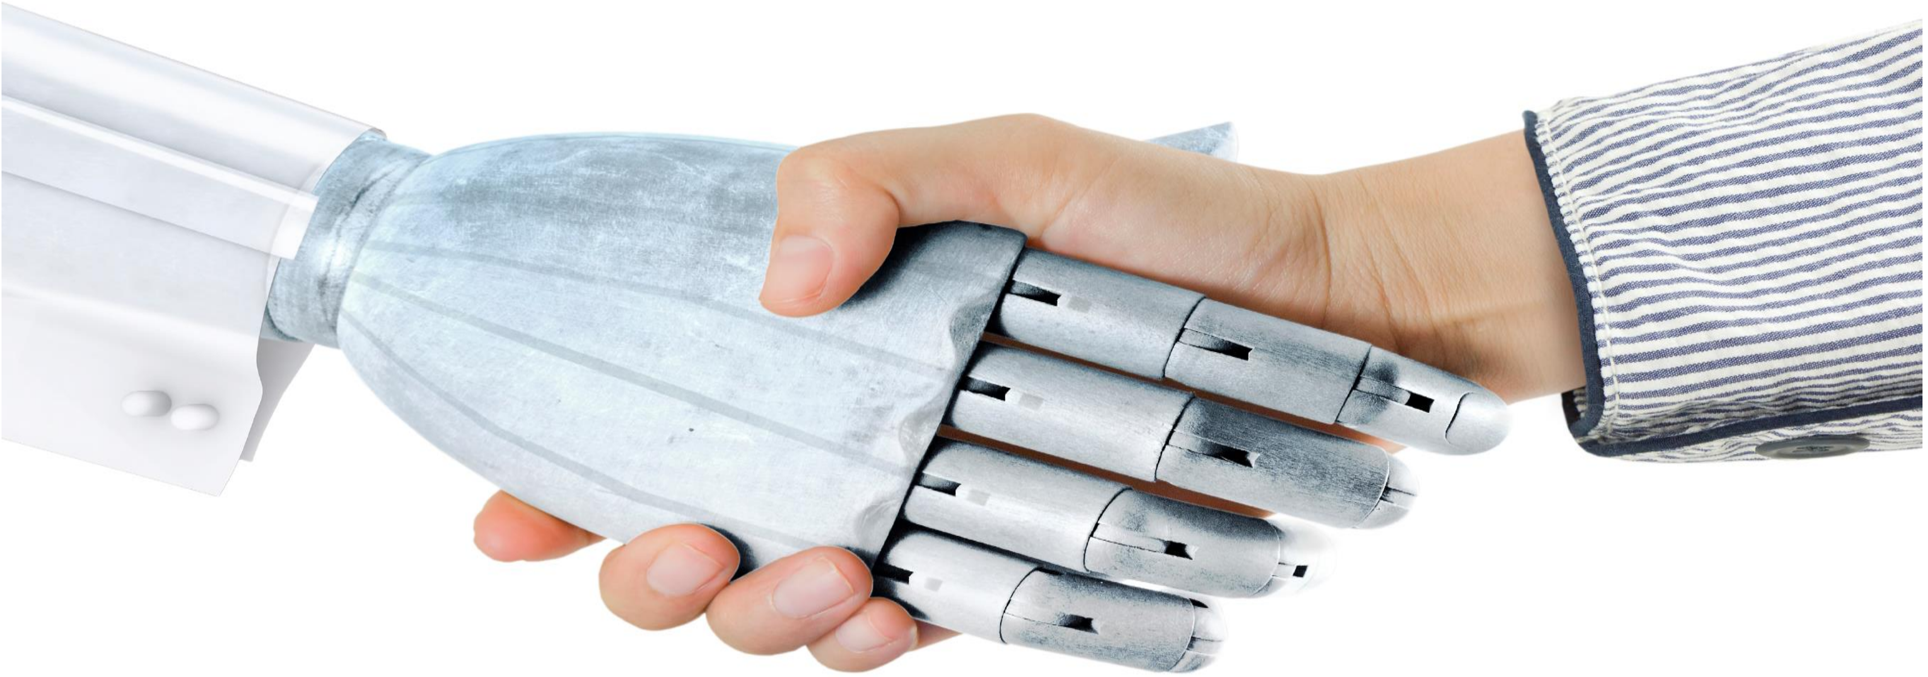

Enhancing Human Capability with Intelligent Machine Teammates

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

- 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

- 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

- 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.

Team Plan Formation

Pre-planning

Do we have consensus?

What have we agreed to?

Apprenticeship

What to do in this situation?

Dynamic Plan Execution

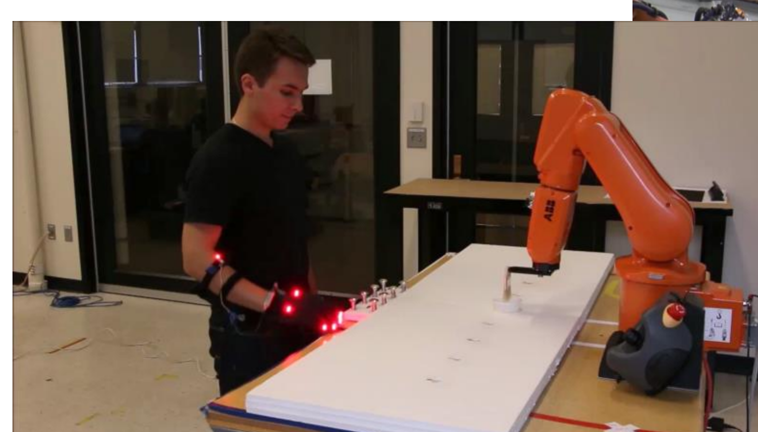

Inferring hidden mental states enables richer, flexible human-machine teaming

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

- [Hidden State] What is our teammate's context for 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? What have we agreed to?
- [Hidden State] What is our teammate's context for 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? What have we agreed to?
- [Hidden State] What is our teammate's context for 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

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? What have we agreed to?

- [Hidden State] What is our teammate's context for 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-Machine Collaboration to Improve Consistency of Shared Understanding

Interactive ROBOTICS @MIT

Machine observables ? Mental State

Input: Topic Discussion

A: We should send UAV to upper region. → Suggestion

B: What's the weather like up there? → Question-Info

…..

B: Okay, let's do that → Acceptance

Dialogue Acts

Suggestion

Question-Info

⋮

Acceptance

Predict

Output

Consistent

Inconsistent

"Could we review this **topic** once more?"

Human-Machine Collaboration to Improve Consistency of Shared Understanding

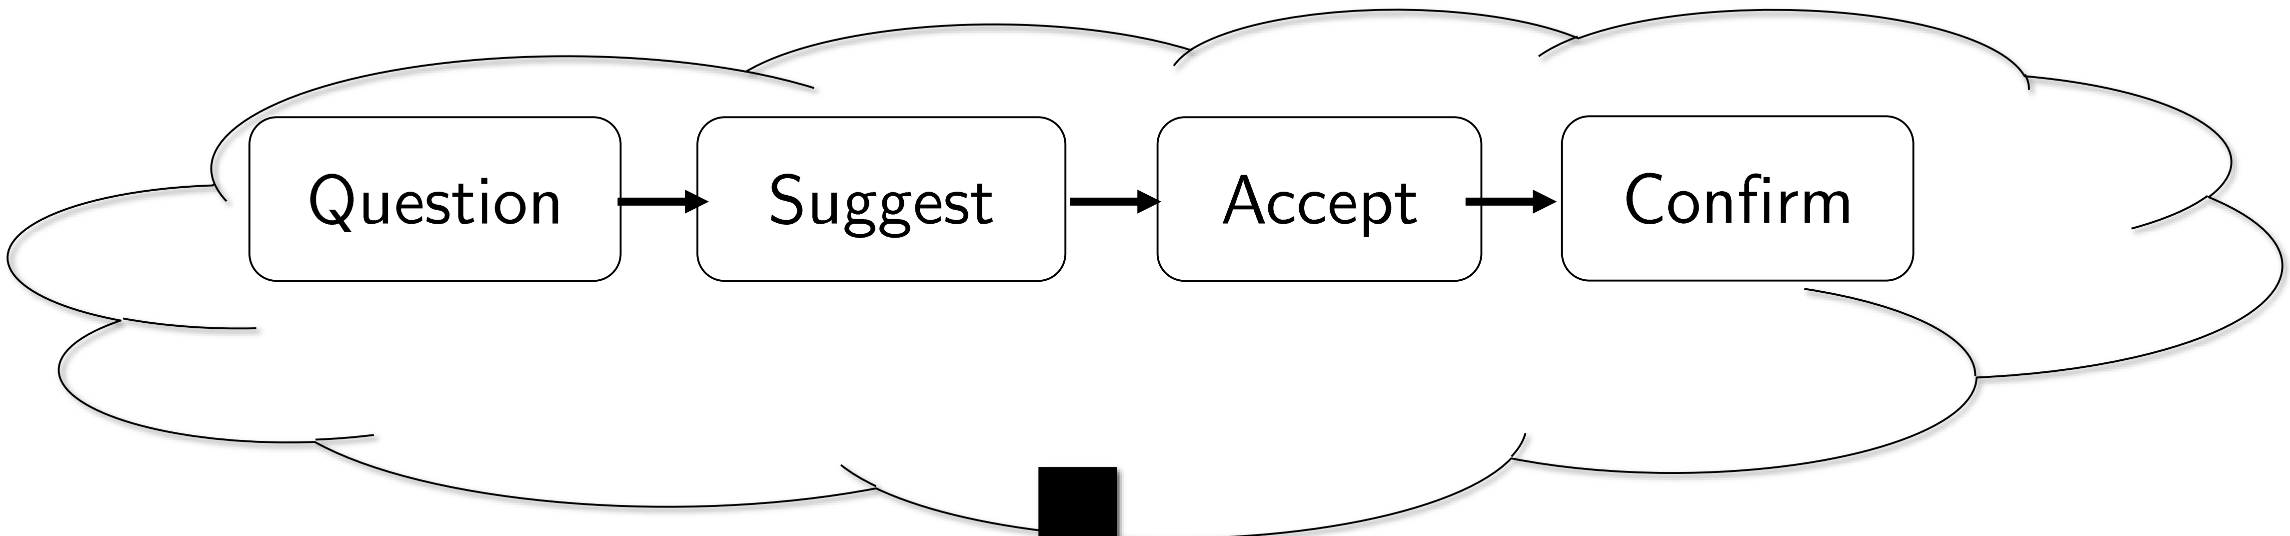

Four patterns of dialog acts to delineate strength of shared understanding

- 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.

Human-Machine Collaboration to Improve Consistency of Shared Understanding

Features: -- transformed observations that provide an informative basis for inference

- 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

Approach

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

Training Set: AMI corpus ~100,000 utterances

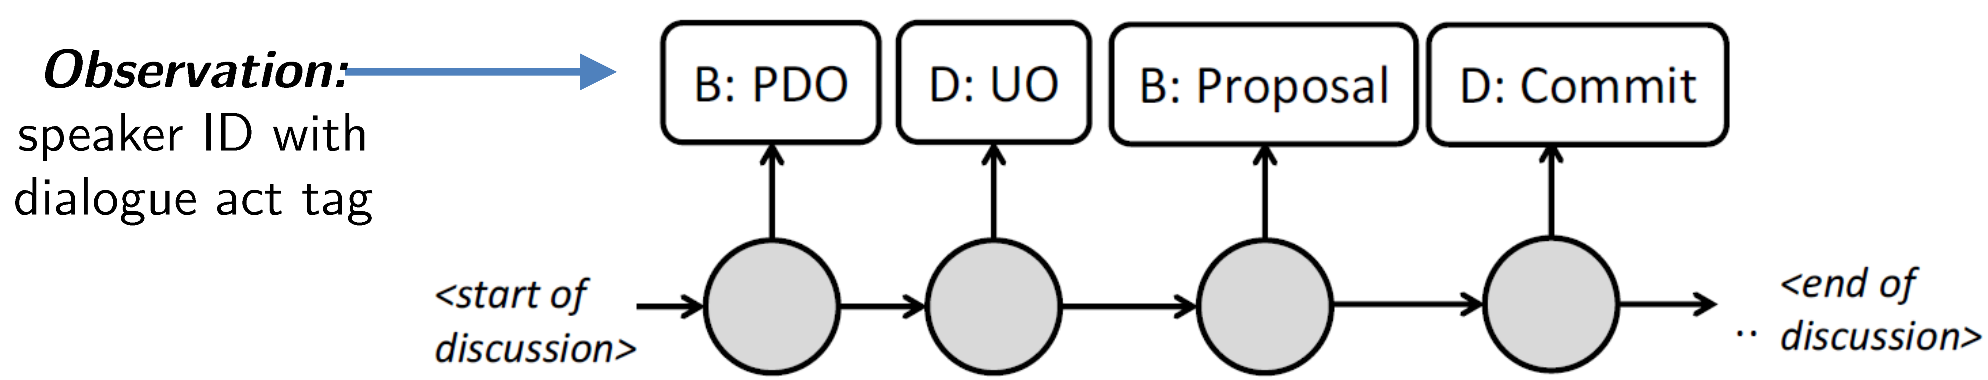

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

Human-Machine Collaboration to Improve Consistency of Shared Understanding

Successful translation of Eugenio's qualitative model into a predictive, statistical ML model

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

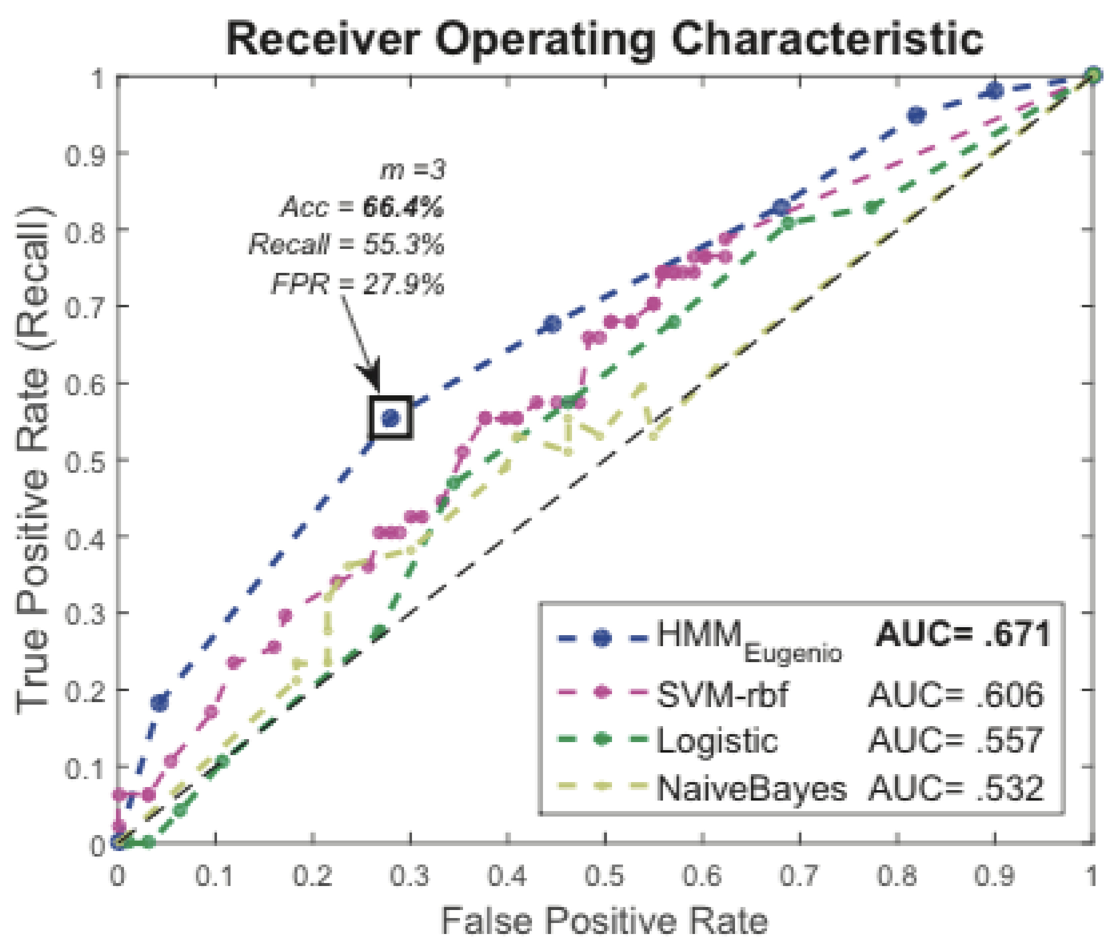

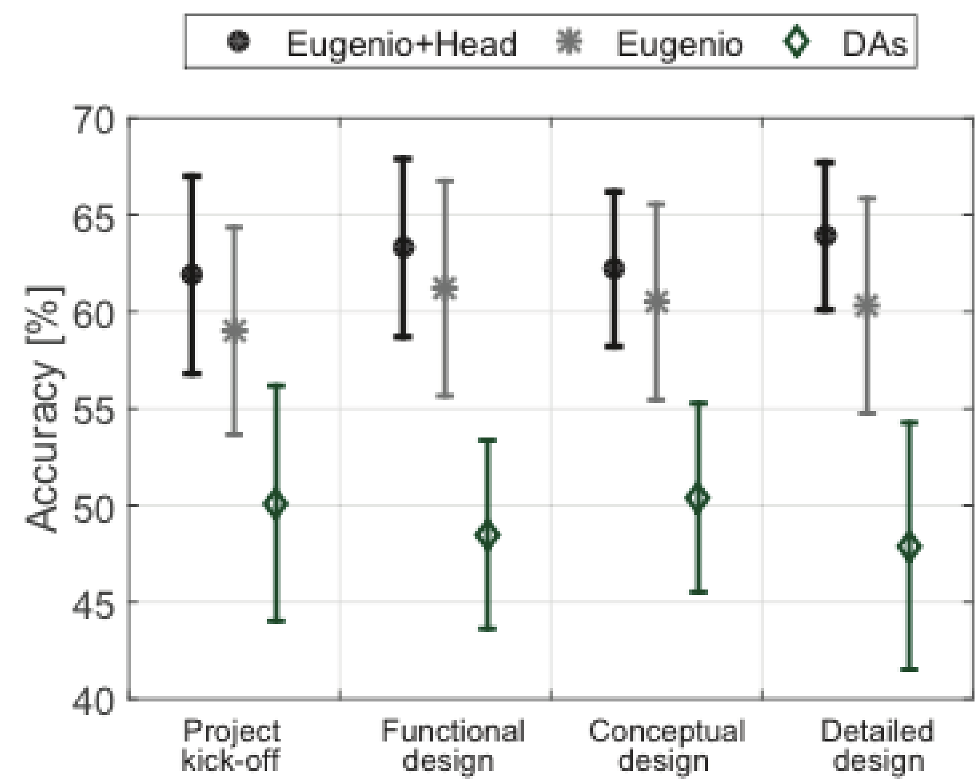

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

Human-Machine Collaboration to Improve Consistency of Shared Understanding

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

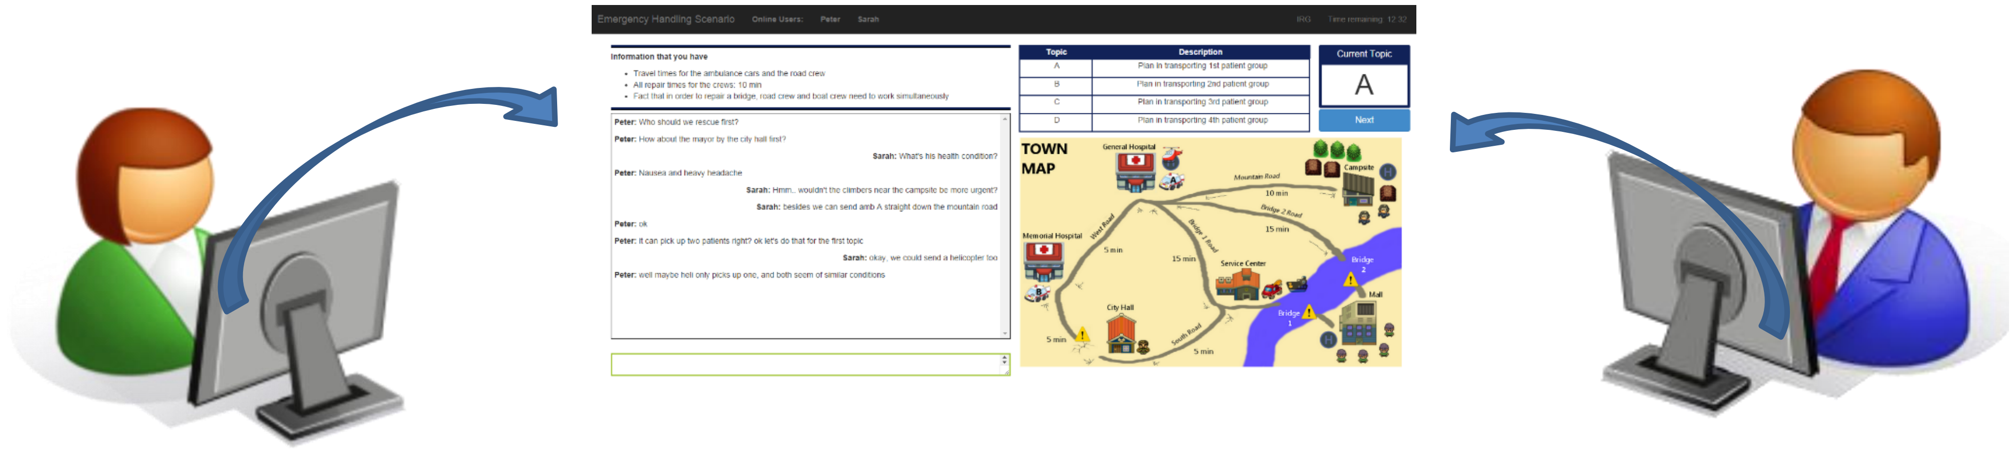

Treatment level	Definition
1. Adaptive review	System suggests review of the two topics with the **lowest** predicted c-scores (*weak* topics)
2. Maladaptive review	System suggests review of the two topics with the **highest** predicted c-scores (*strong* topics).

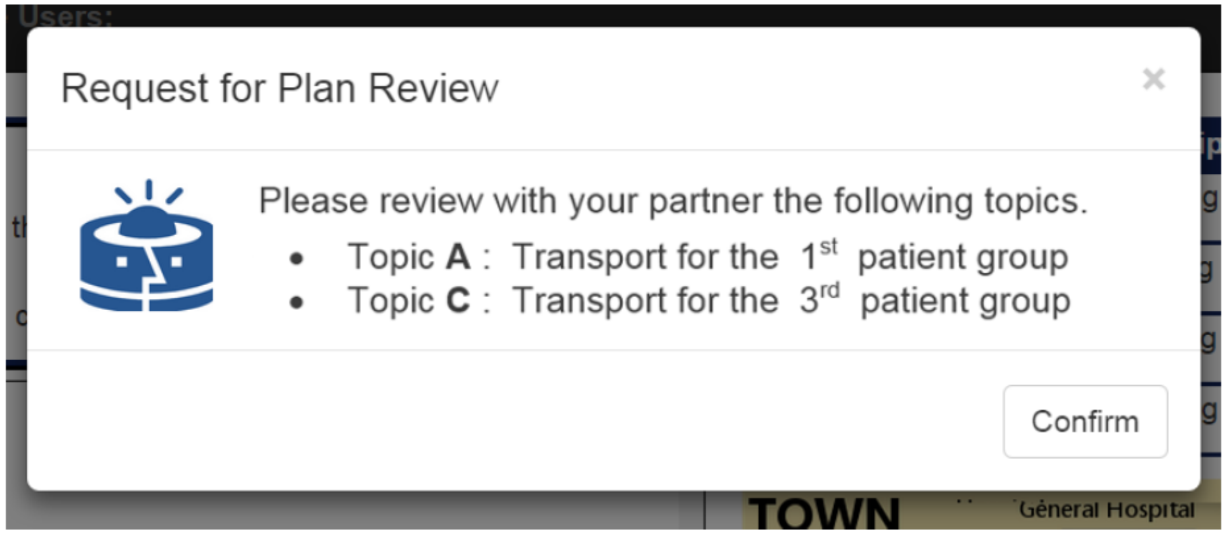

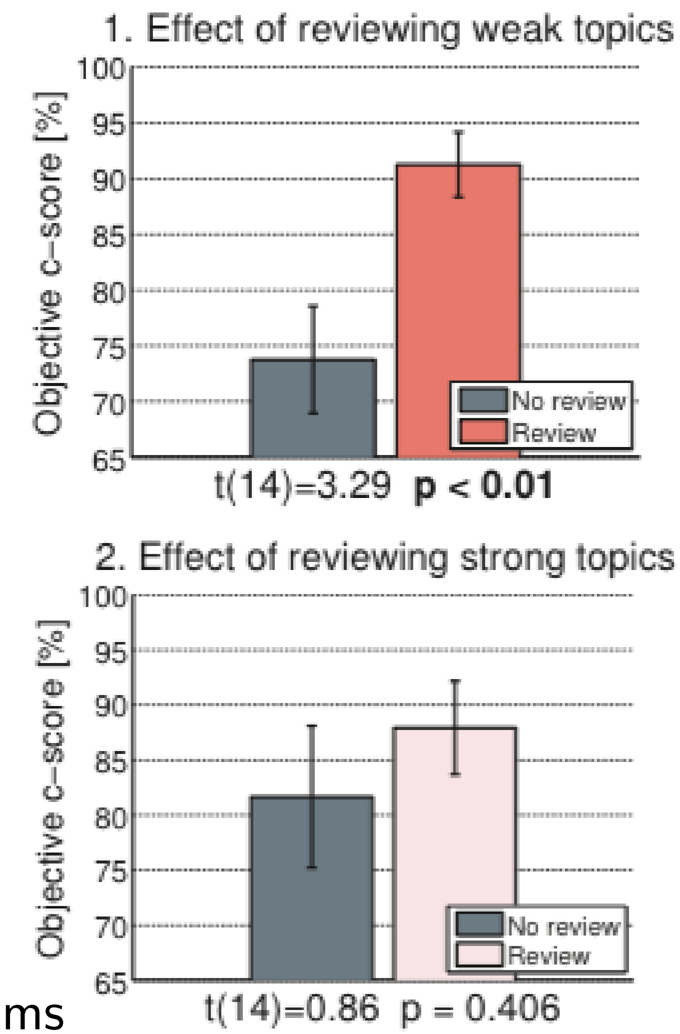

J. Kim and J. Shah (2016) IEEE Transactions on Human-Machine Systems

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

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

- [Hidden State] What is our teammate's context for decision-making?

Machine observables ? Mental State

Well-established cognitive models

Meaningful features that relate to mental state

Model structure to process complex information efficiently

What did we agree to?

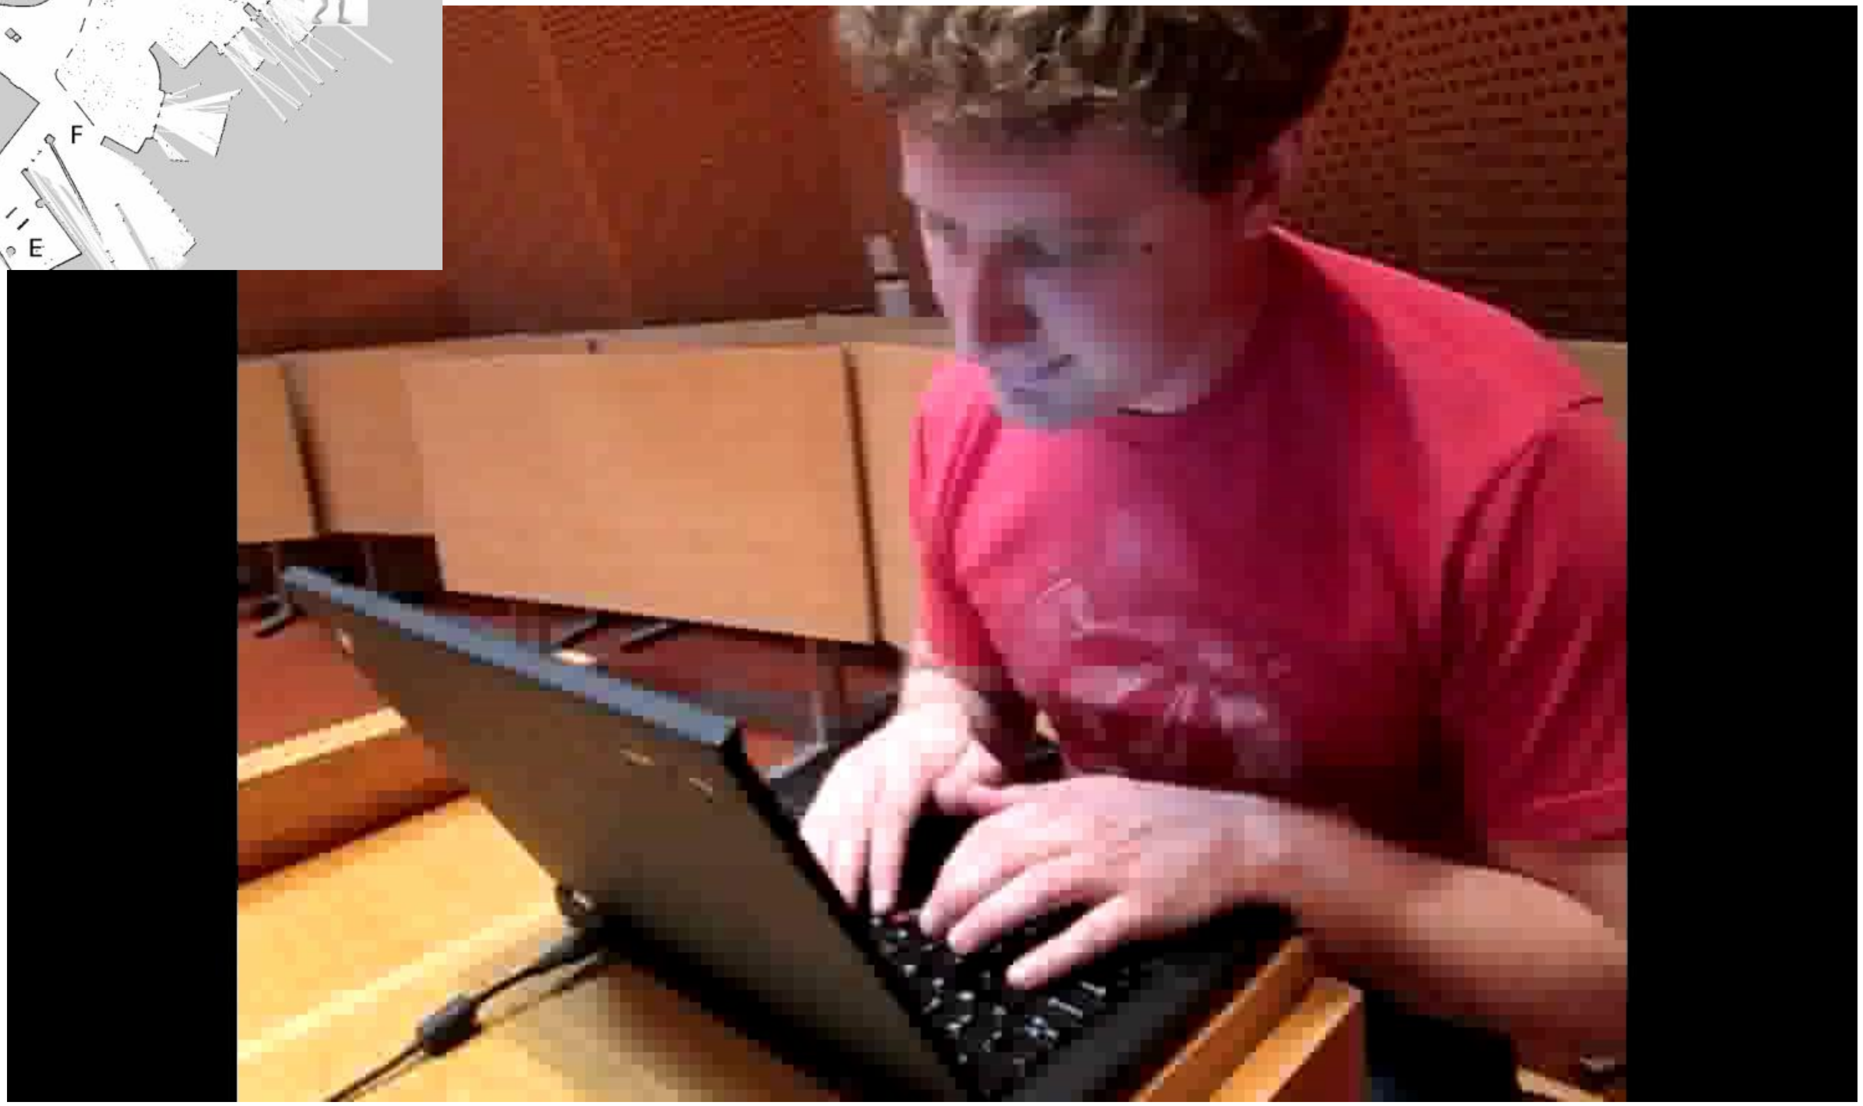

Scenario:

- 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.

What did we agree to?

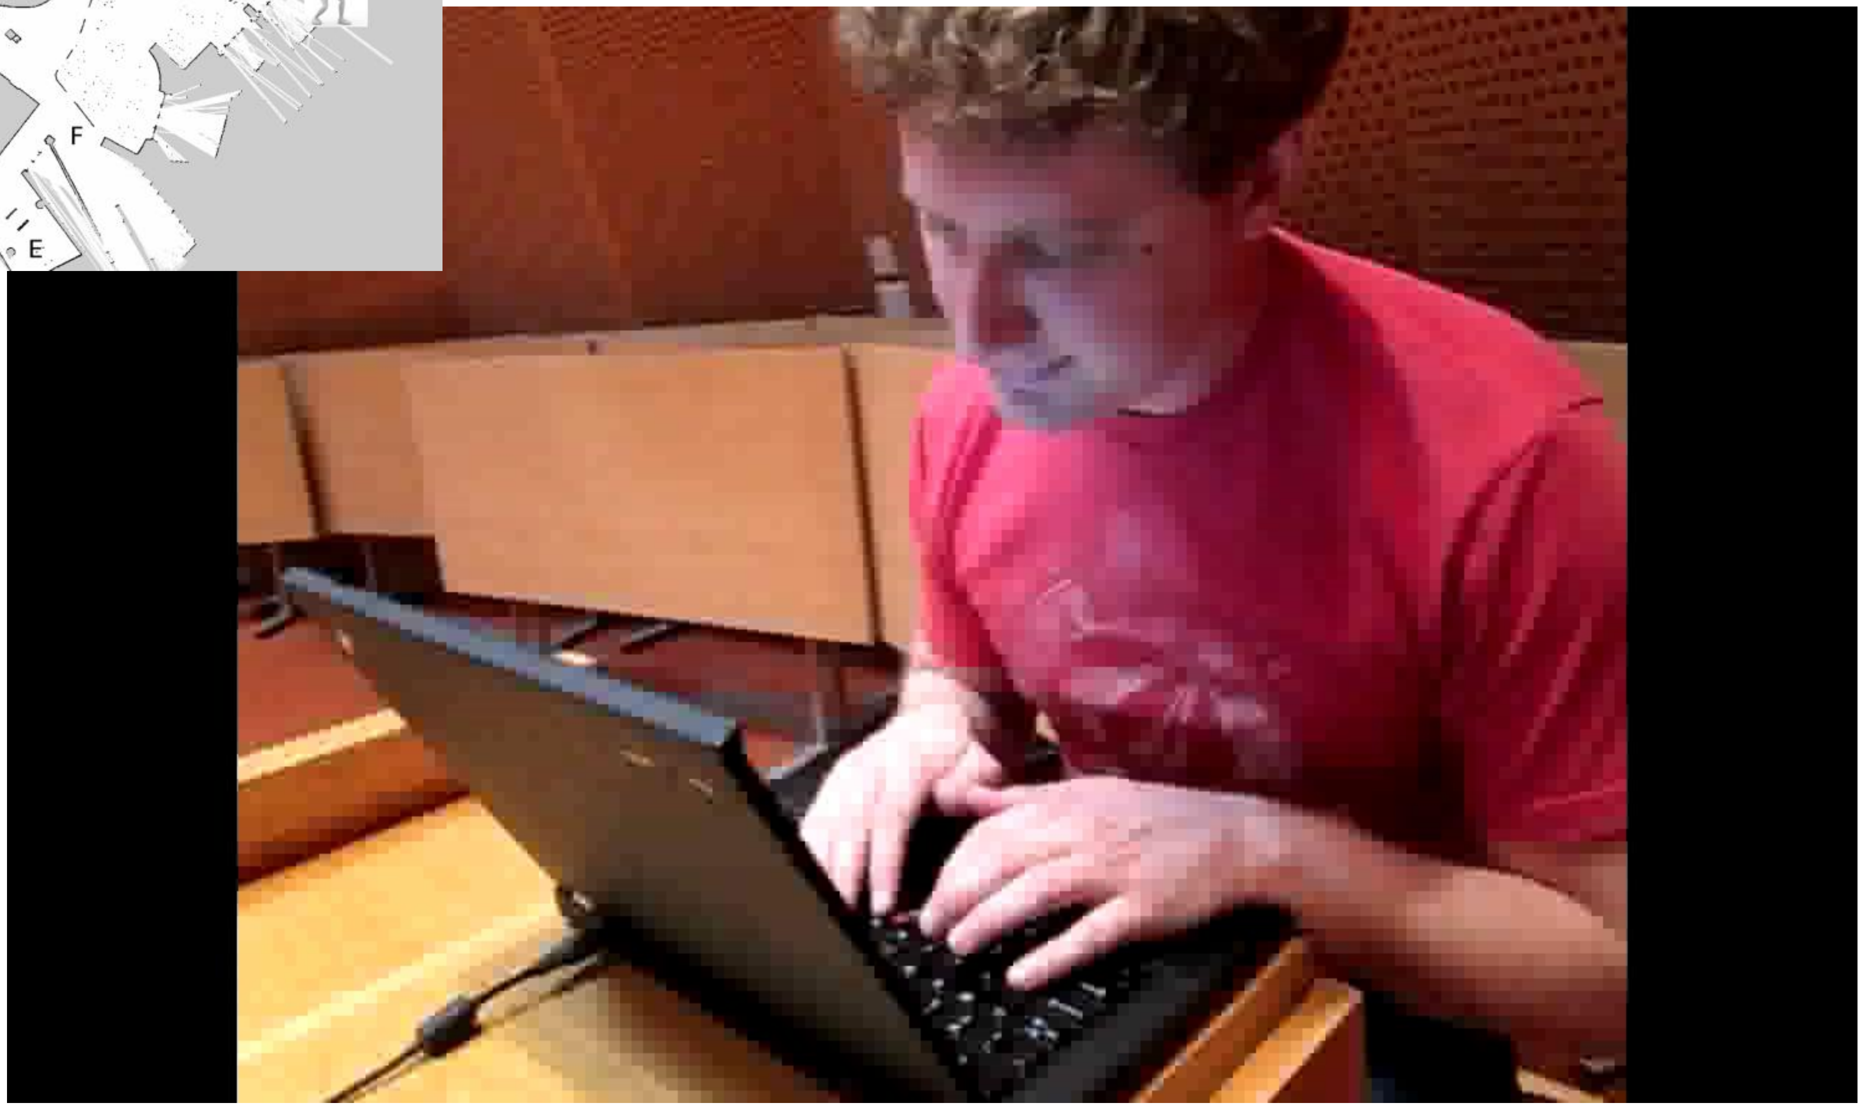

Trillions of possible solutions ….

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.

What did we agree to?

Interactive ROBOTICS @MIT

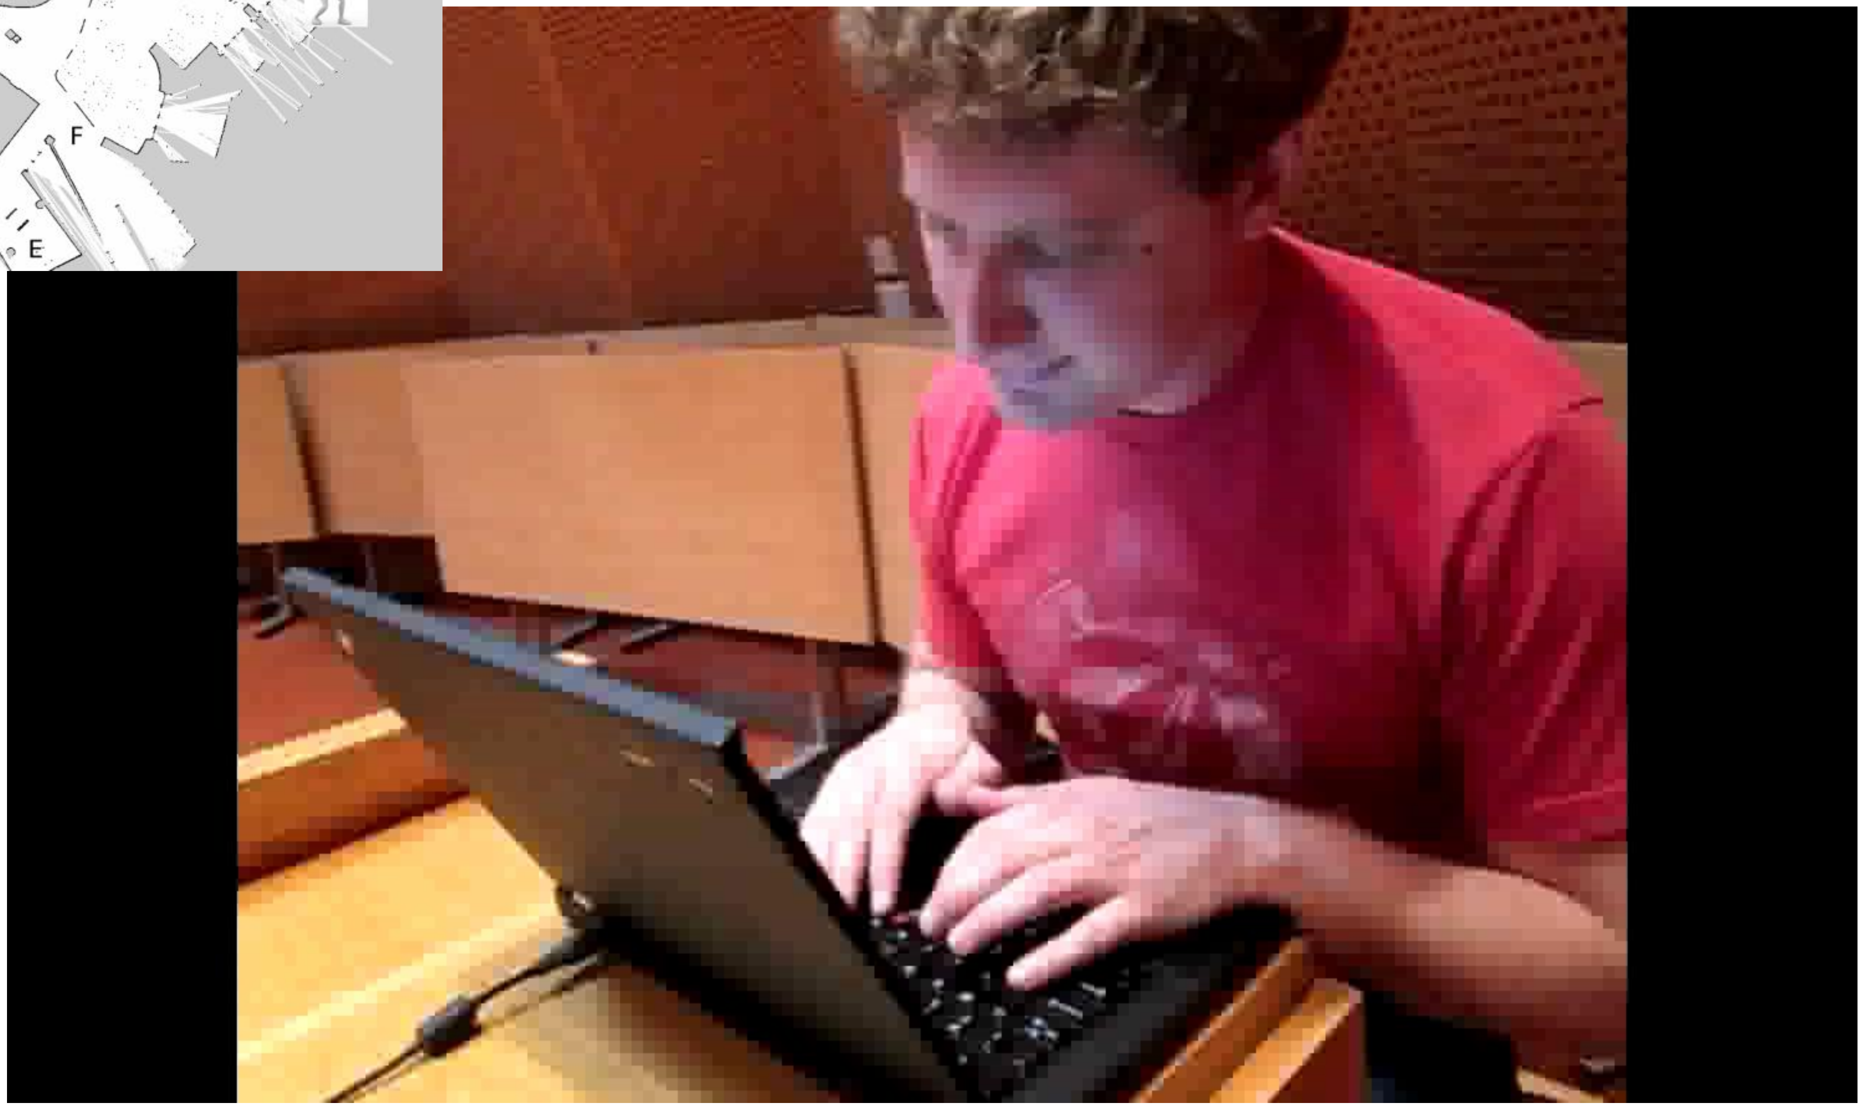

Using logical structure of planning problem, the inference task becomes almost as easy for a machine as for a person!

Trillions of possible solutions ….

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

Ordered tuple of sets of grounded predicates

nth predicate that appears in the tth utterance

plan step index assigned to each n predicates in utterance t

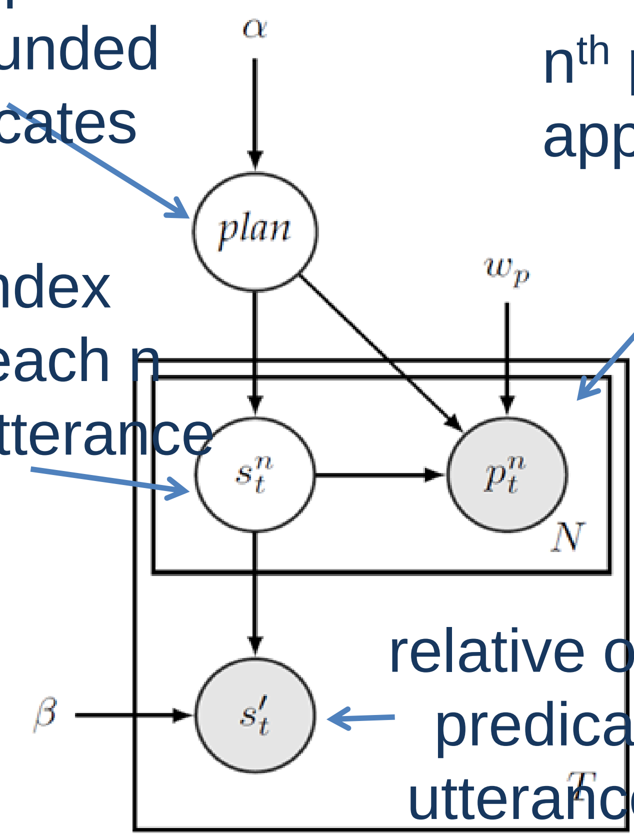

relative ordering of predicates in tth utterance as they appear in final plan

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

1. Perform inference through sampling (Gibbs & MH)

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

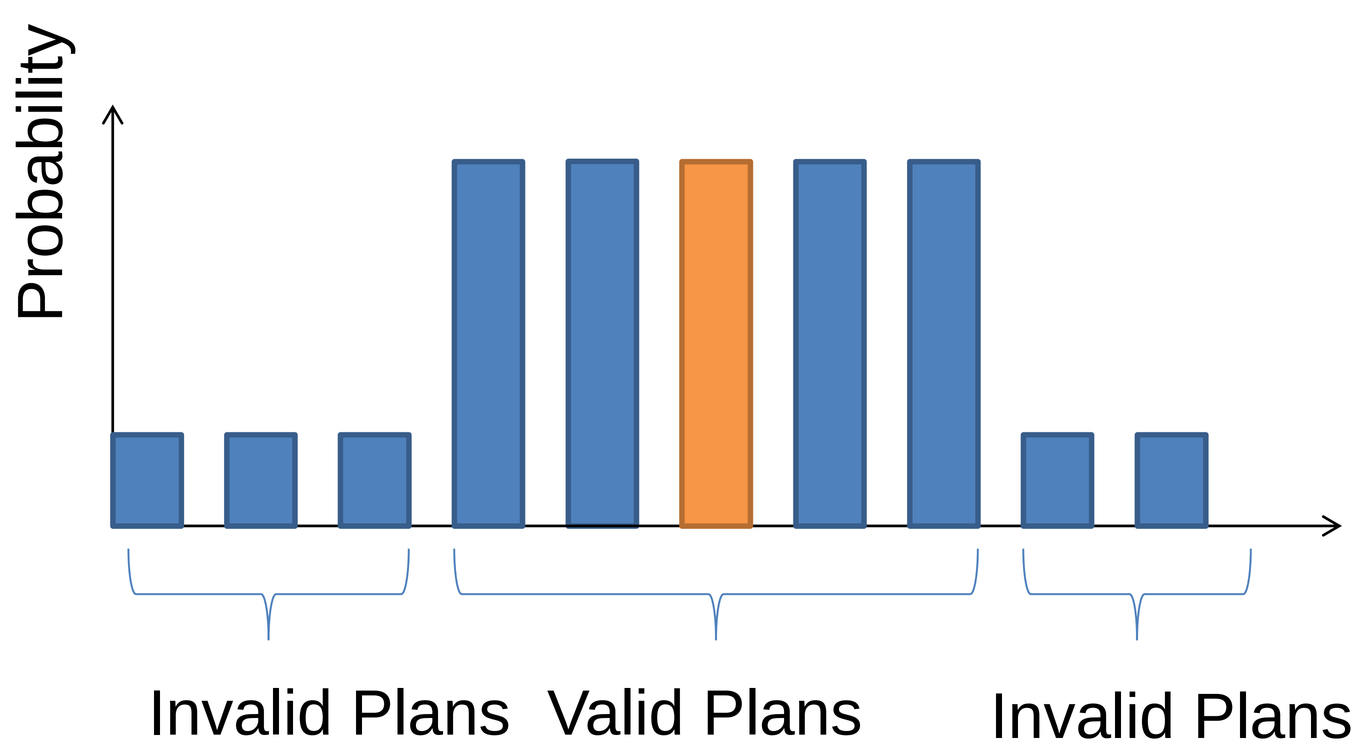

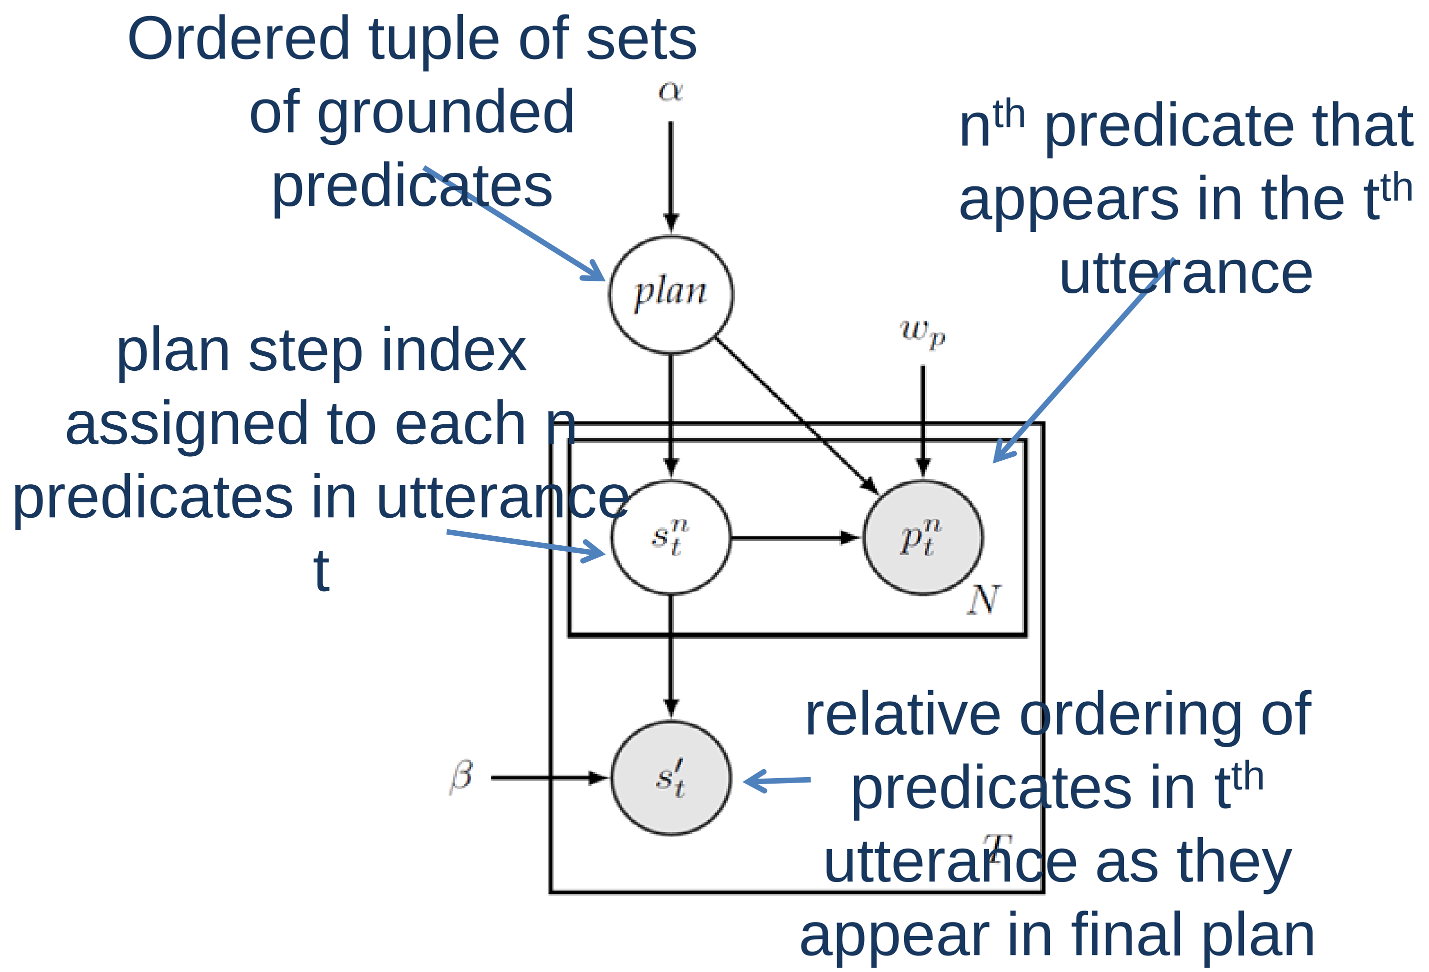

B. Kim, C. Chacha, J. Shah (2015) Journal of Artificial Intelligence Research

Successful Automatic Extraction of Final Agreed-Upon Plan

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

Successful Automatic Extraction of Final Agreed-Upon Plan

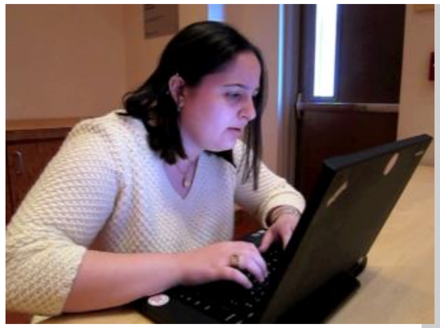

Scenario:

- 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

Successful Automatic Extraction of Final Agreed-Upon Plan

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

Pentagon touts "Loyal Wingman" for combat jets

30 MARCH, 2016 | BY: JAMES DREW | WASHINGTON DC

US Air Force plans to insert a “brain” into current-generation fighter jets to create autonomous flying wingmen paired with the Lockheed Martin F-35 were given a bump today, with the Pentagon’s second-in-charge saying he expects to see “unmanned wingmen in the air” before convoys of driverless Humvees.

Deputy defence secretary Robert Work touted the long-considered “loyal wingman” concept at a forum hosted by the Washington Post in Washington DC on 30 March, where he explained that the air force will pair unmanned Lockheed F-16s with F-35s in future battles.

Projects with AFRL, Lockheed, and NavAir

Aim:

Demonstrate human-machine collaborative **mission planning, pre-mission brief** and **after-action review**

for **Multi-Platform Air Operations**

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] What is our teammate's context for 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 Complex Strategies for Decision-Making from Apprenticeship

Gombolay et al. IJCAI’16, RSS’16

- How to learn complex strategies just by watching?

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.

- Anti-ship missile defense (with MIT LL)

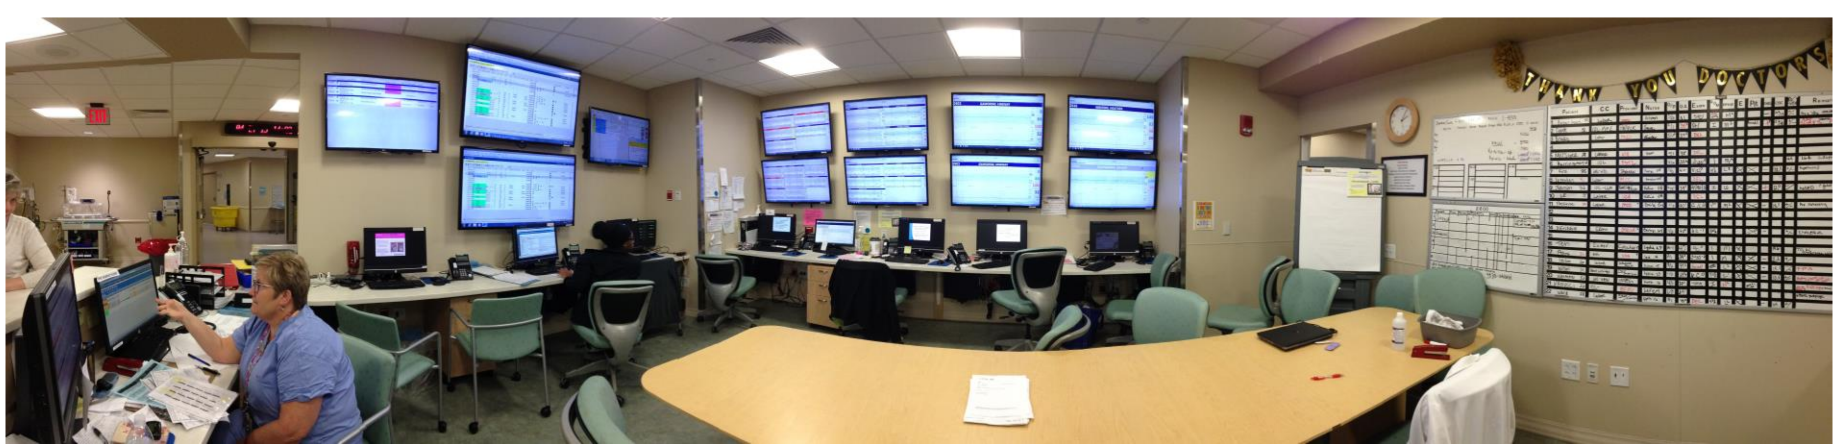

- Coordination of patient care in a hospital

Machines that Learn Complex Strategies for Decision-Making from Apprenticeship

Gombolay et al. 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*

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

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

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

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

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

Gombolay et al. 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.

- 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, $p < 0.011$)

Machines that Learn Complex Strategies for Decision-Making from Apprenticeship

Gombolay et al. IJCAI'16

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

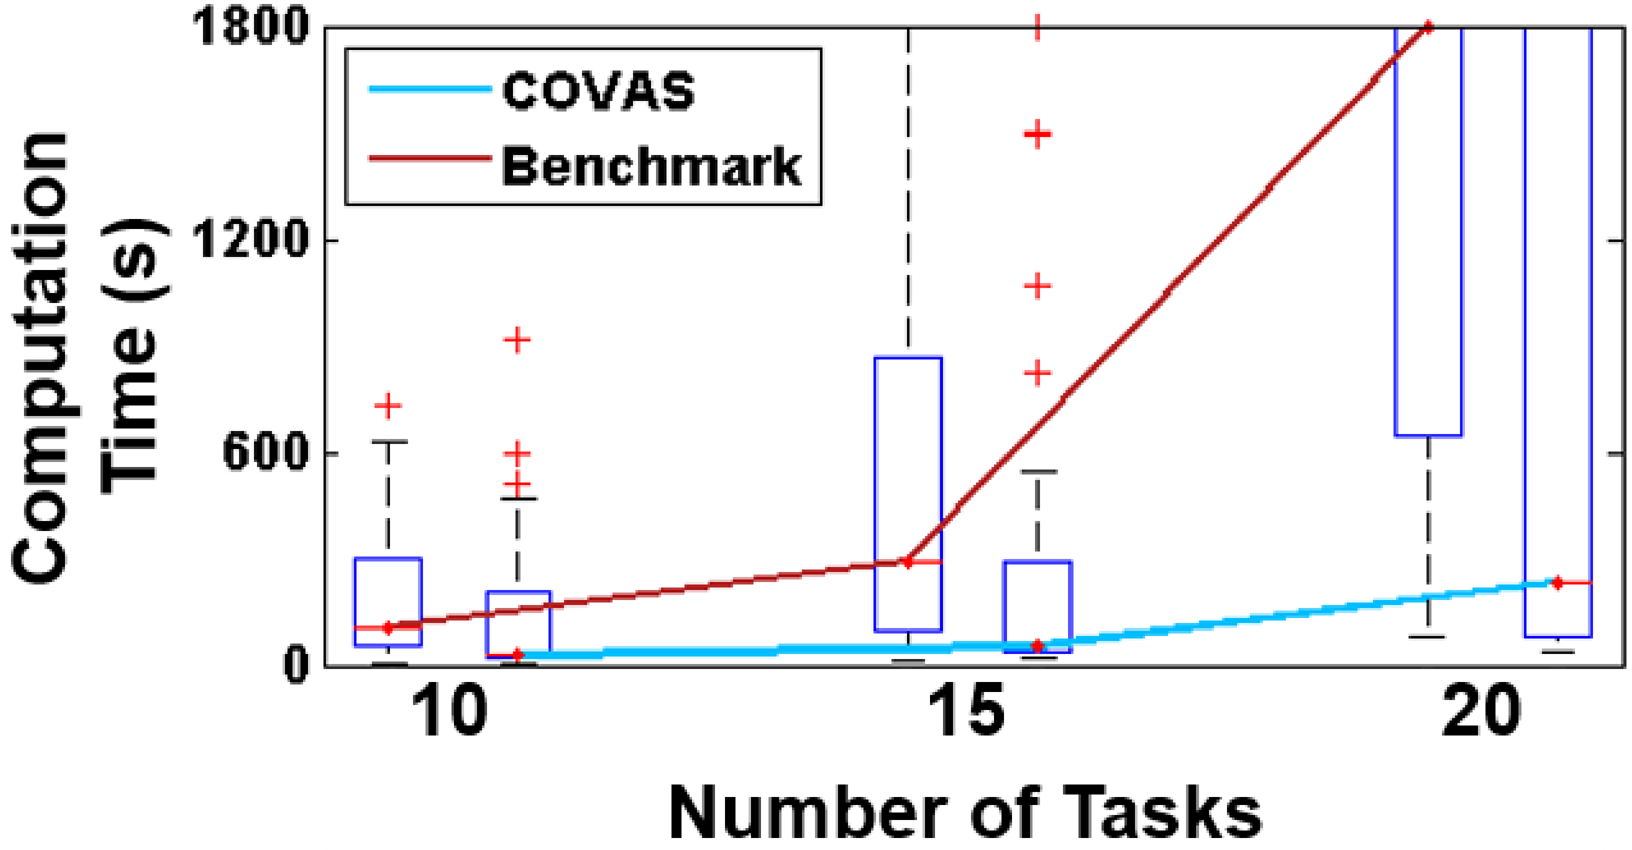

- Strategies learned from humans for small problems are used by the machine to quickly solve problems that are too large for the human or machine alone.

Next Steps - From Drones to Teammates

Gombolay et al. RSS'16

- Coordination of patient care in a hospital

Next Steps - From Drones to Teammates

Gombolay et al. RSS’16

- Coordination of patient care in a hospital

From Planning to Flexible Execution

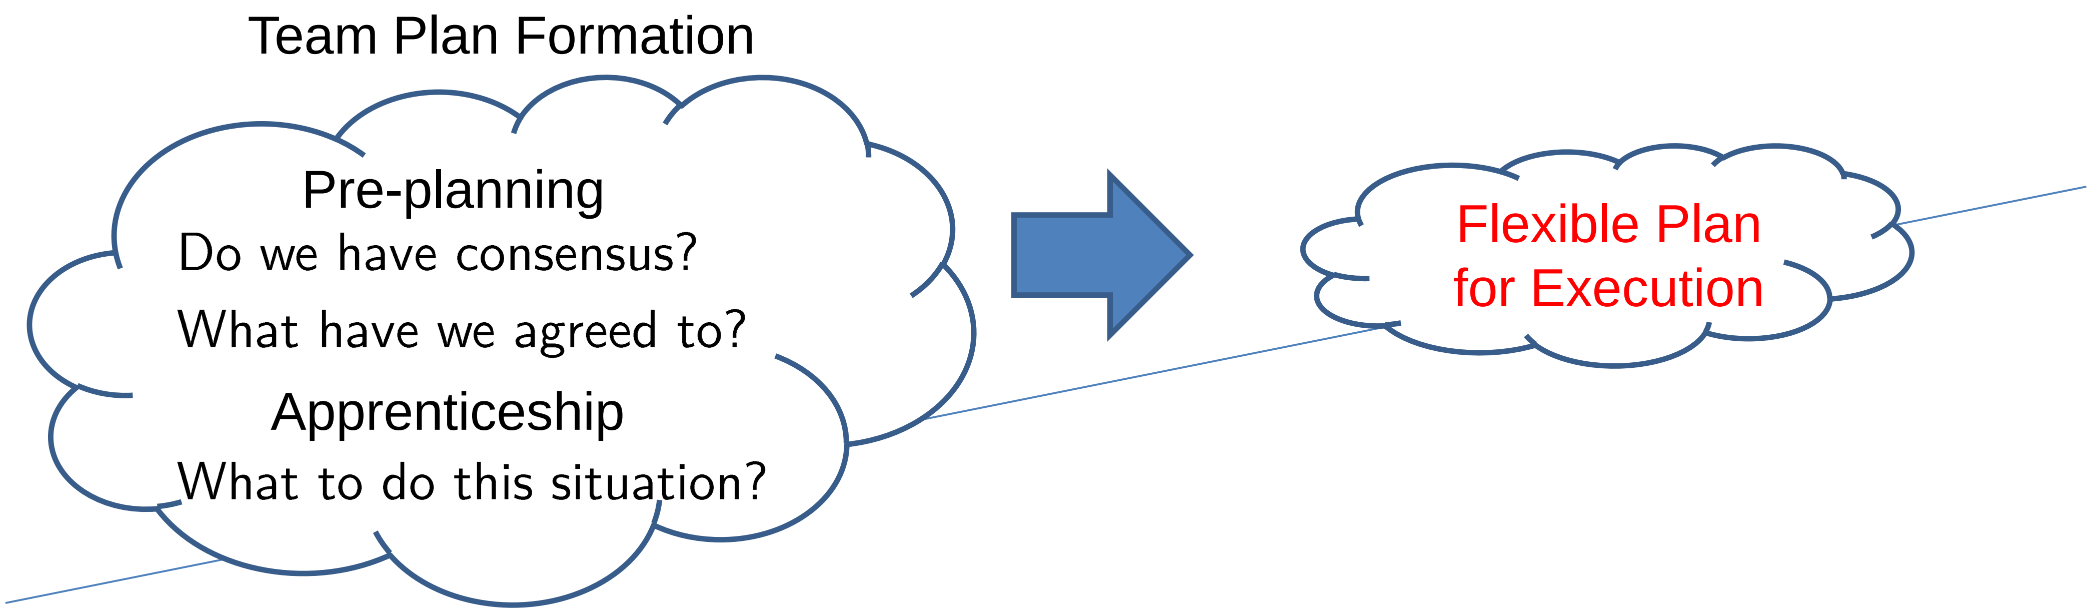

From Planning to Flexible Execution

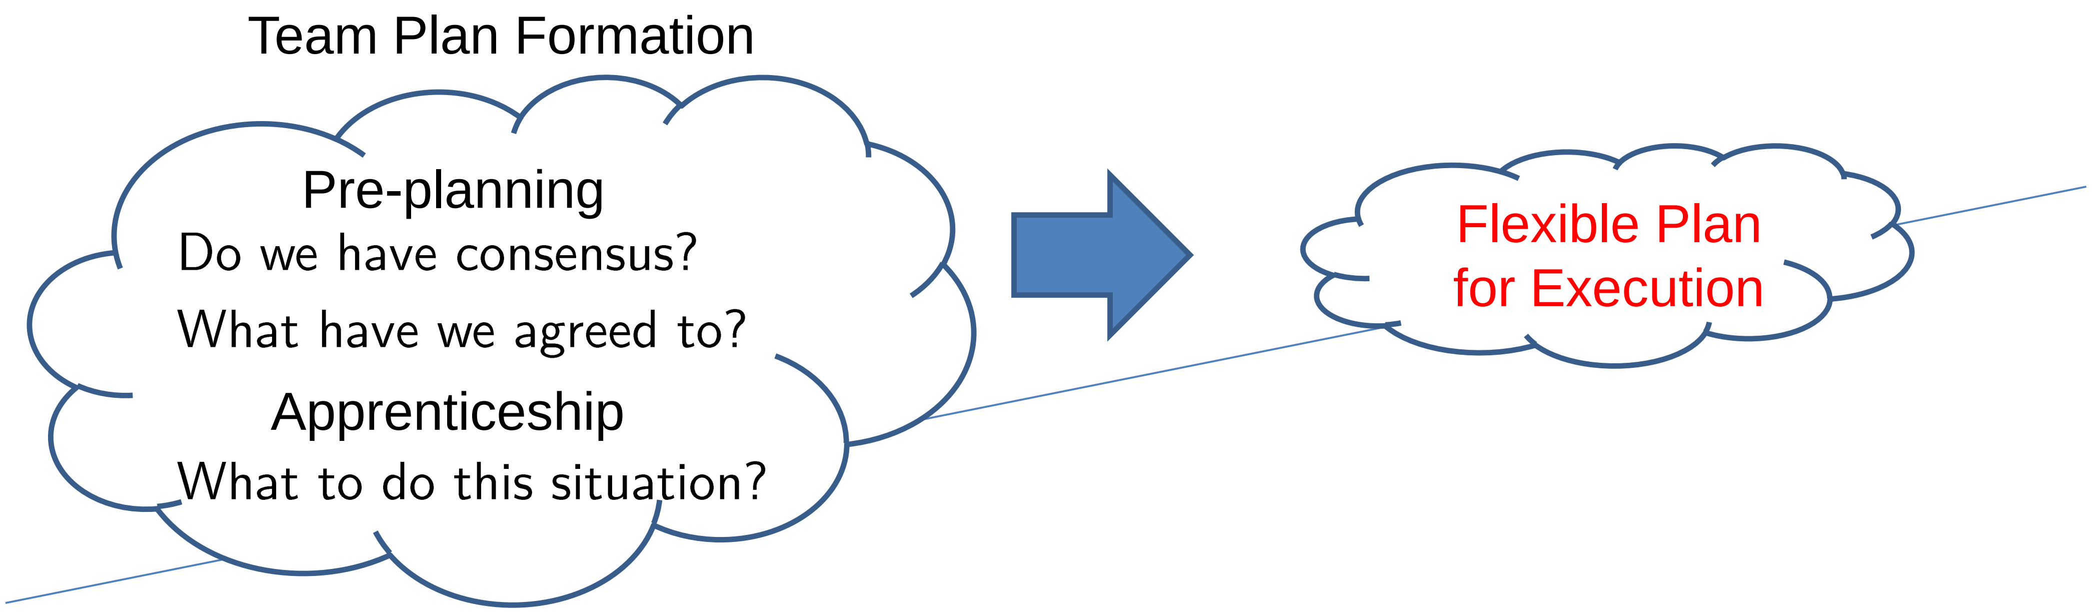

Methods for flexible planning and scheduling:

- Multi-Robot Task Allocation & Scheduling at Scale through Schedulability Analysis [RSS'13, JAIS'14]
- Computational techniques for "fair" allocation of resources under uncertainty [NIPS'14, AAAI'15]
- Fast computation of flexible schedules with preferences that accommodate disturbance [RSS'12]
- Multi-level optimization of coordination strategy, allocation and schedule [IJCAI'16]

From Planning to Flexible Execution

Team Plan Formation

Pre-planning
Do we have consensus?
What have we agreed to?
Apprenticeship
What to do this situation?

Flexible Plan
for Execution

100% Automation

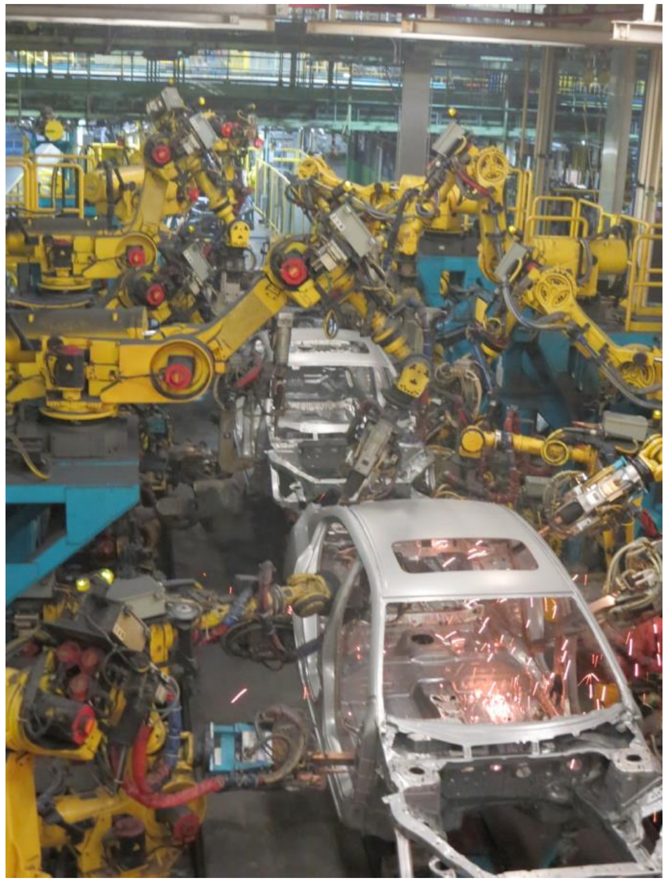

100% Manual work

Why is Collaborative Robotics Important?

From Planning to Flexible Execution

Team Plan Formation

Pre-planning

Do we have consensus?

What have we agreed to?

Apprenticeship

What to do this situation?

Flexible Plan for Execution

Why is Collaborative Robotics Important?

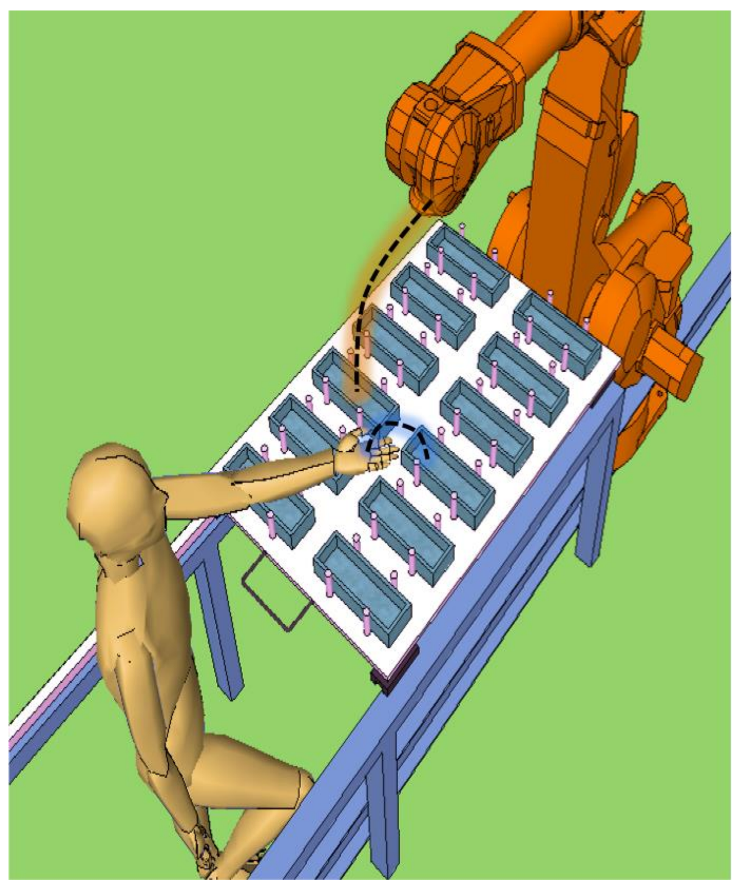

Reduction of NET-B

Process No.:		Process Name:	METER INSTALL	
Zone :	IL	Sub-Zone:	% Time	Category
1	WALK FROM TABLE TO PARTS 1-2 STEPS	60400	14.11%	Net B
2	GET METER	78100	18.81%	Net A
3	GET AND ALIGN COUPLER	78100	18.81%	
4	SET COUPLER	78100	14.11%	
5	PULL CHECK COUPLER	78100	4.70%	
6	PLACE METER TO IN PANEL	78100	18.81%	Net A
7	PACKAGING (A1 30 PER)	78100	10.66%	Net B
	Total		100.00%	

Net B reduction with collaborative robot (For only meter install part): 24.76%

ISO/TS 15066:2016
Robots and robotics devices – Collaborative robots


KUKA

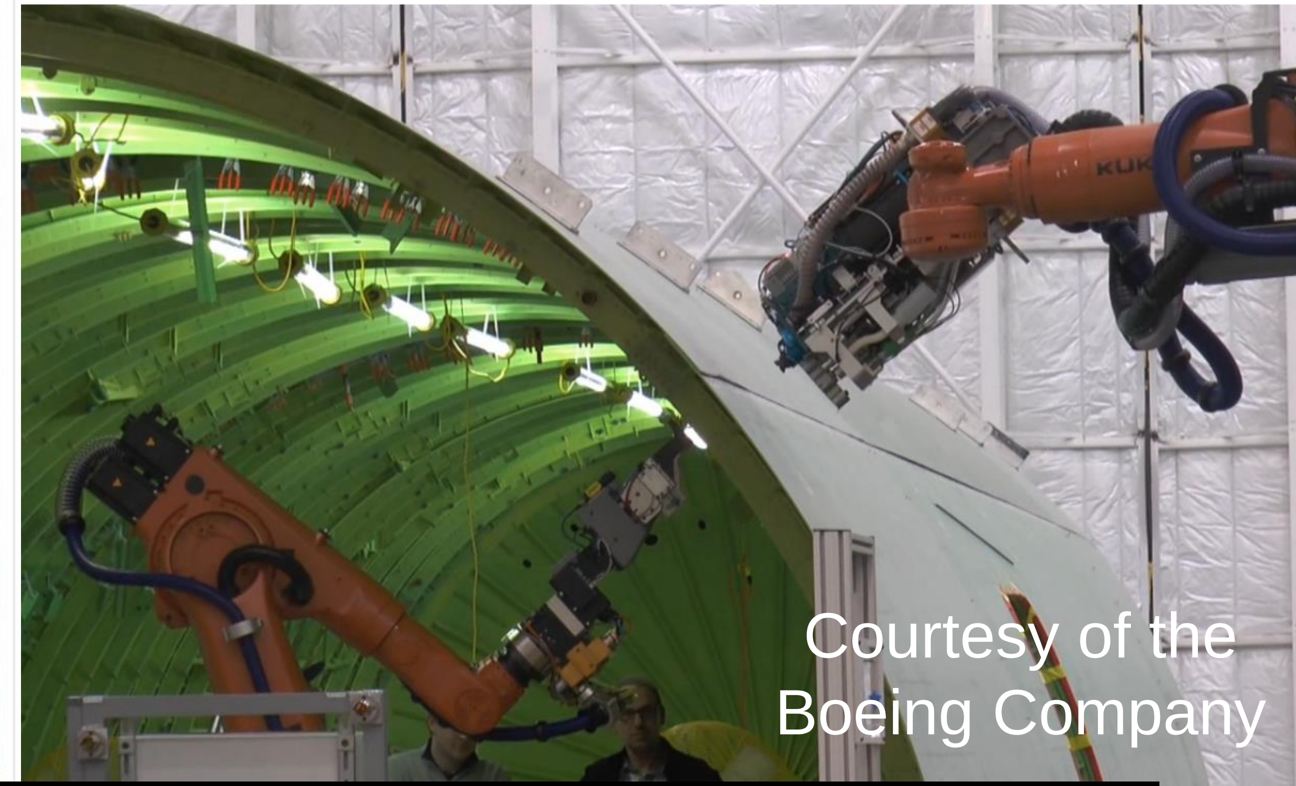
Courtesy of the
Boeing Company

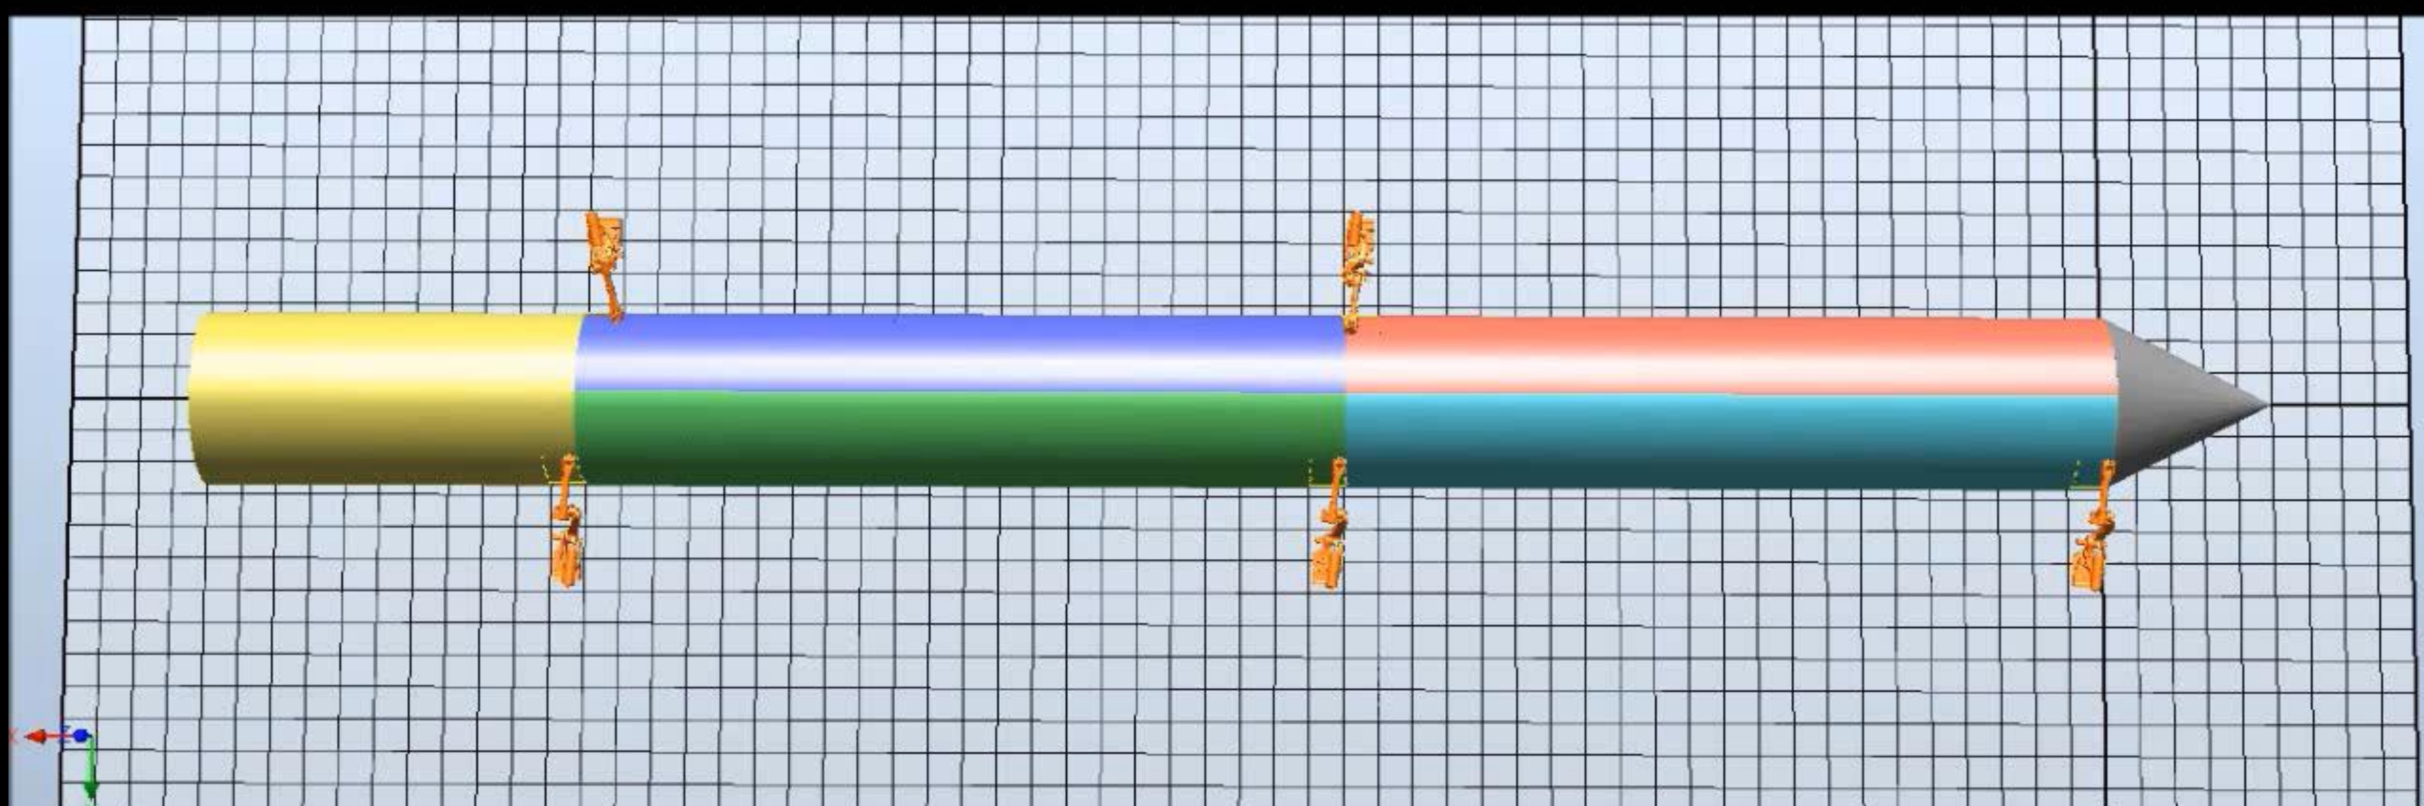

Robots that Plan to Work Flexibly with People

- Real-time processor techniques enable efficient pruning of search space for multi-robot task allocation & sequencing

Computation Time & Schedule Quality for 10-robot Task Allocation & Scheduling

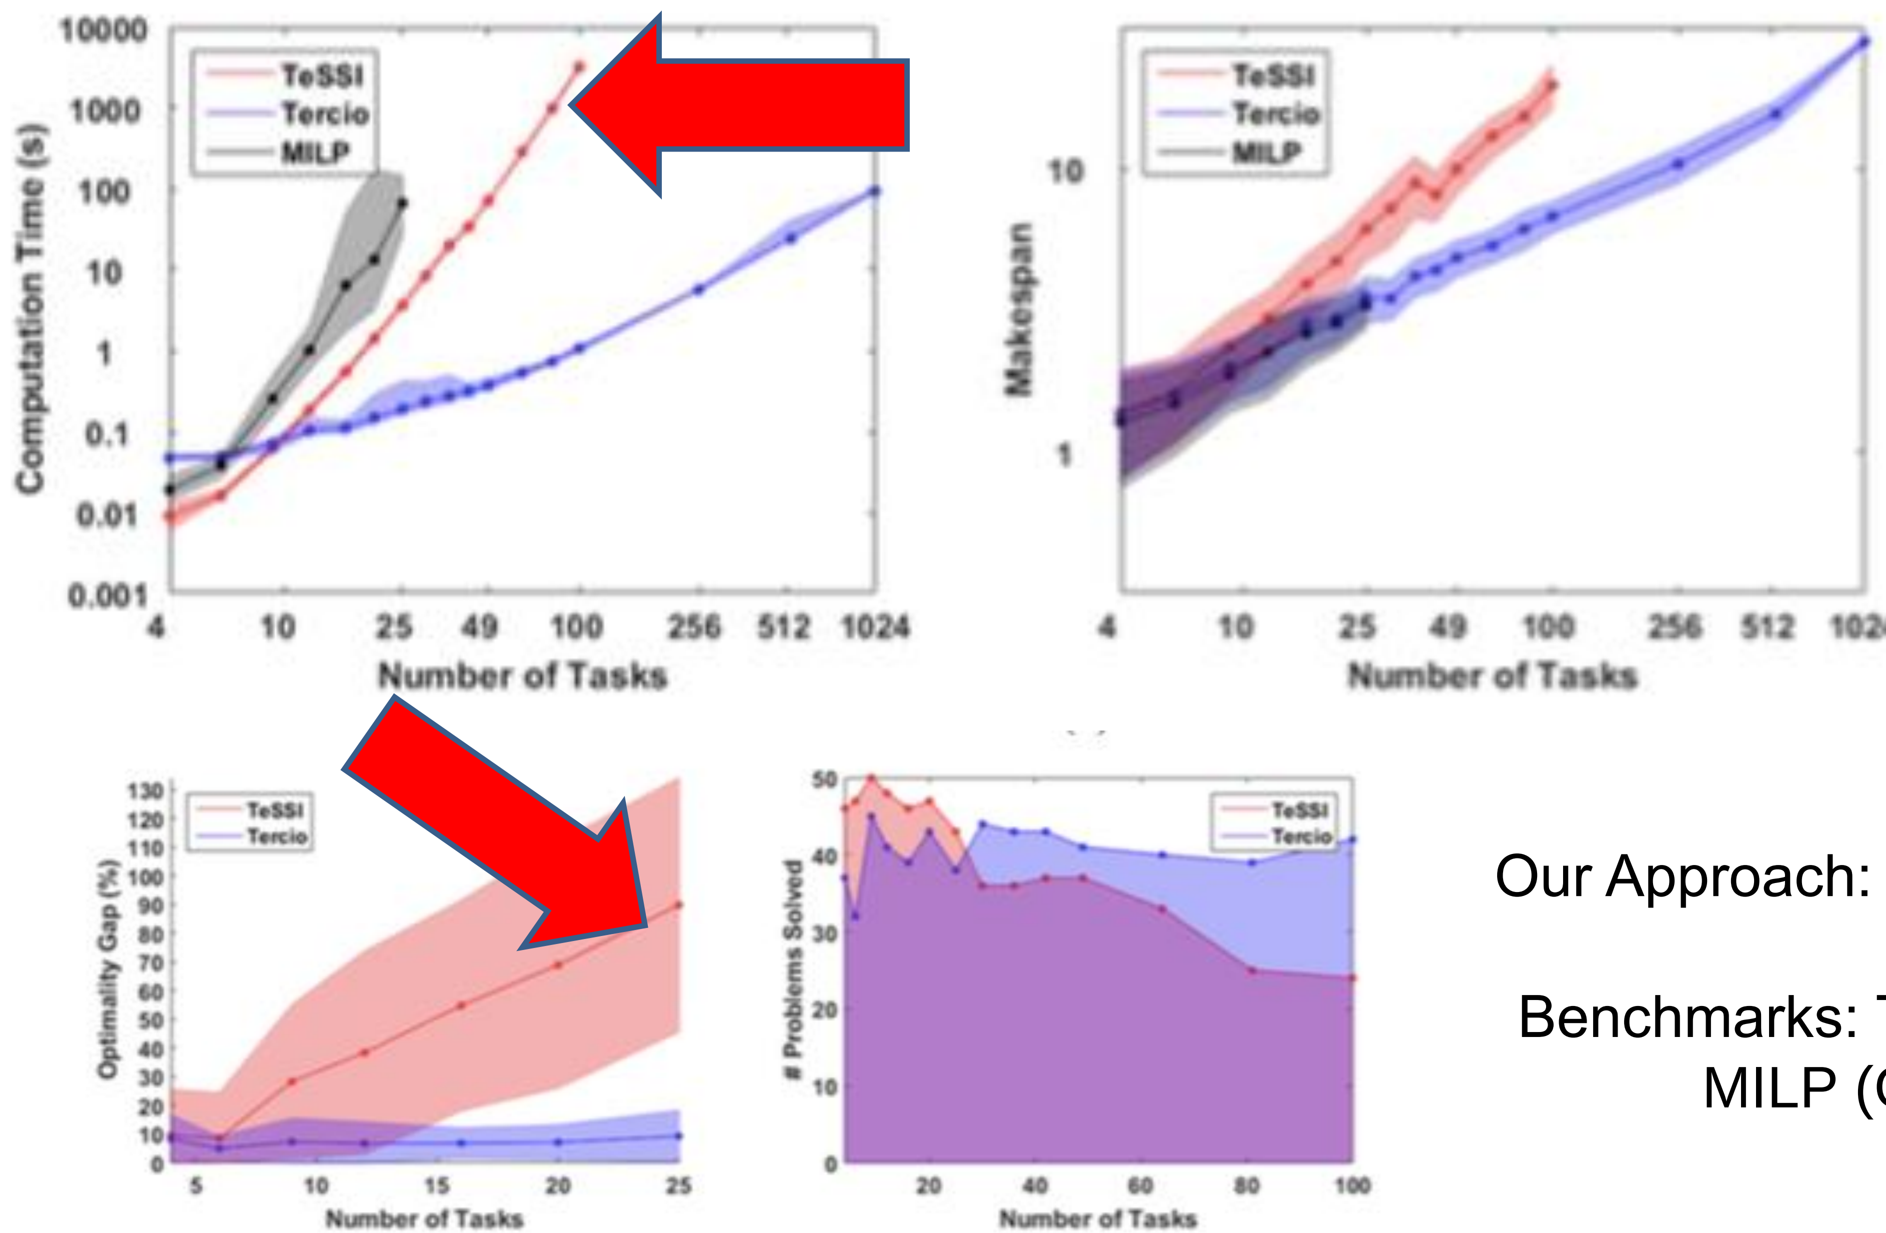

Our Approach: Tercio RSS '13

Benchmarks: TeSSi AAAI'15
MILP (Gurobi)

Robots that Plan to Work Flexibly with People

- Efficient computation through Real-time Processor Schedulability Analysis

Our scheduling problem:

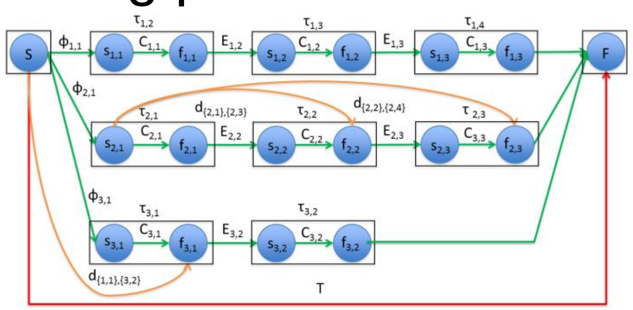

Prior approaches provide schedulability tests for restricted problem structure:

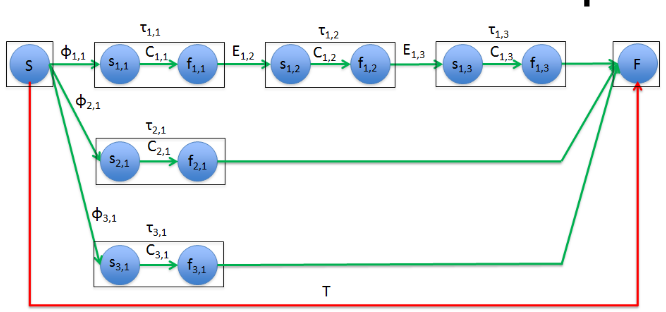

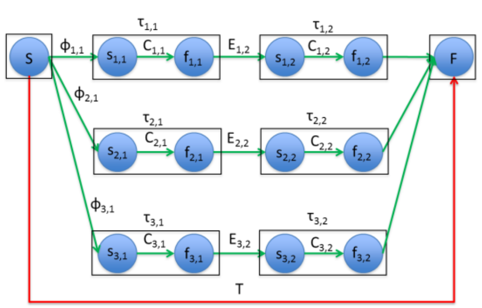

Out test relies on worst case analysis of task orderings and resulting idle times

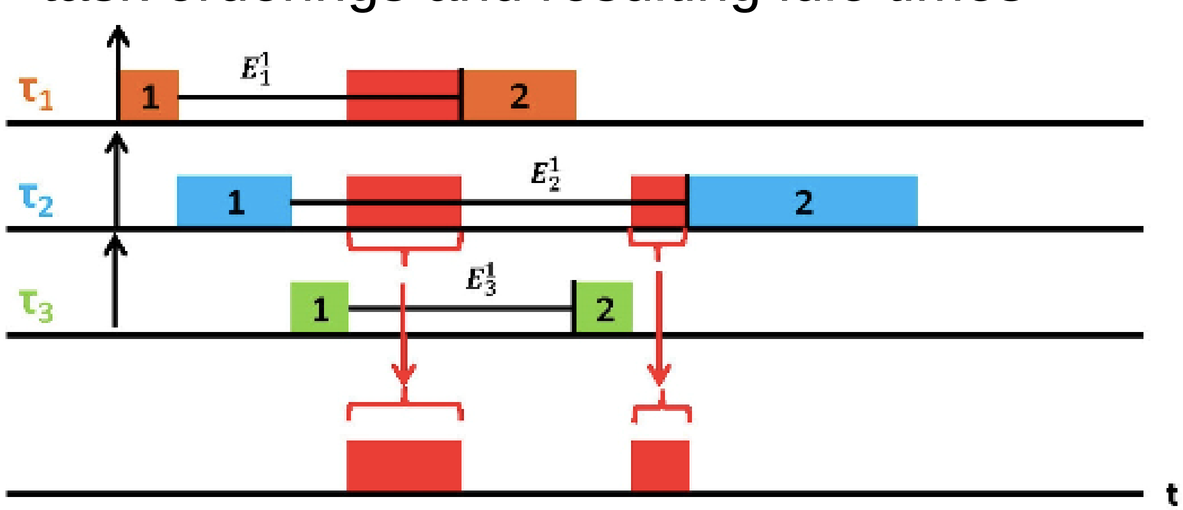

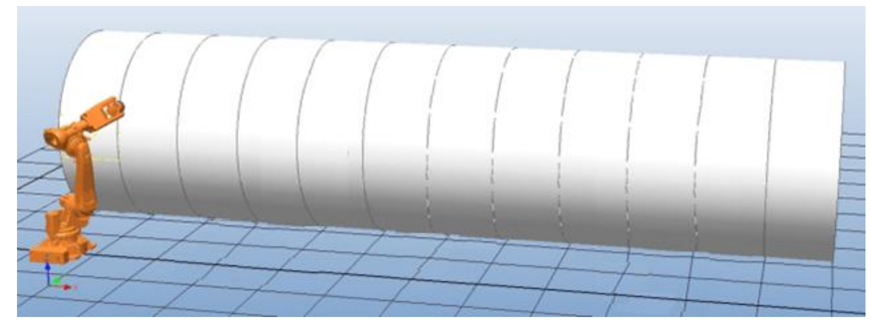

First closed-form polynomial time schedulability test for task sets with upper- and lowerbound temporal constraints

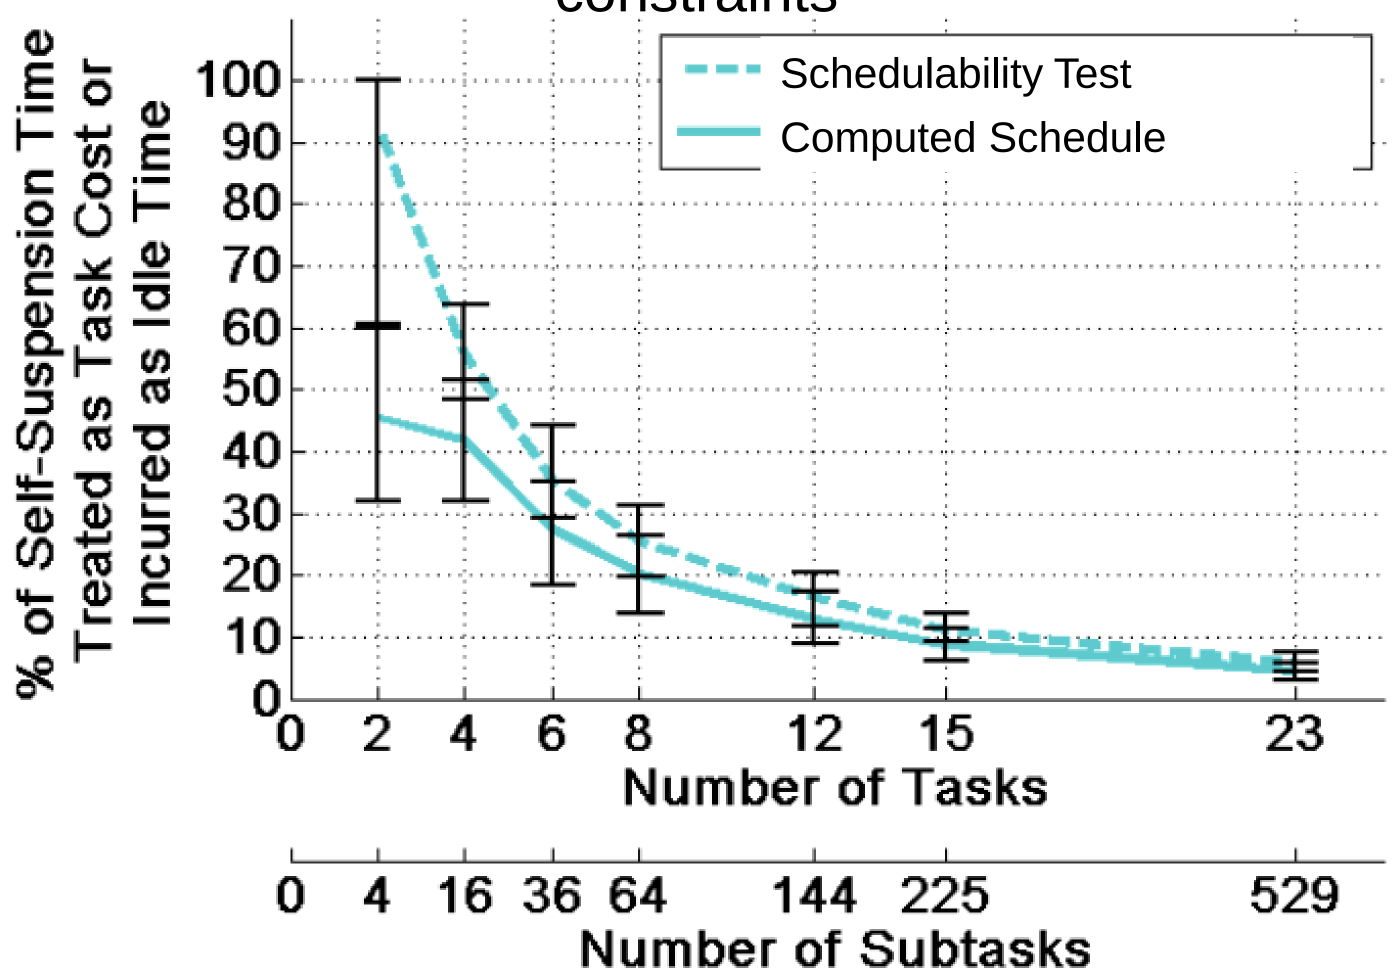

Robots that Plan to Work Flexibly with People

- Real-time processor techniques enable efficient pruning of search space for multi-robot task allocation & sequencing

Computation Time & Schedule Quality for 10-robot Task Allocation & Scheduling

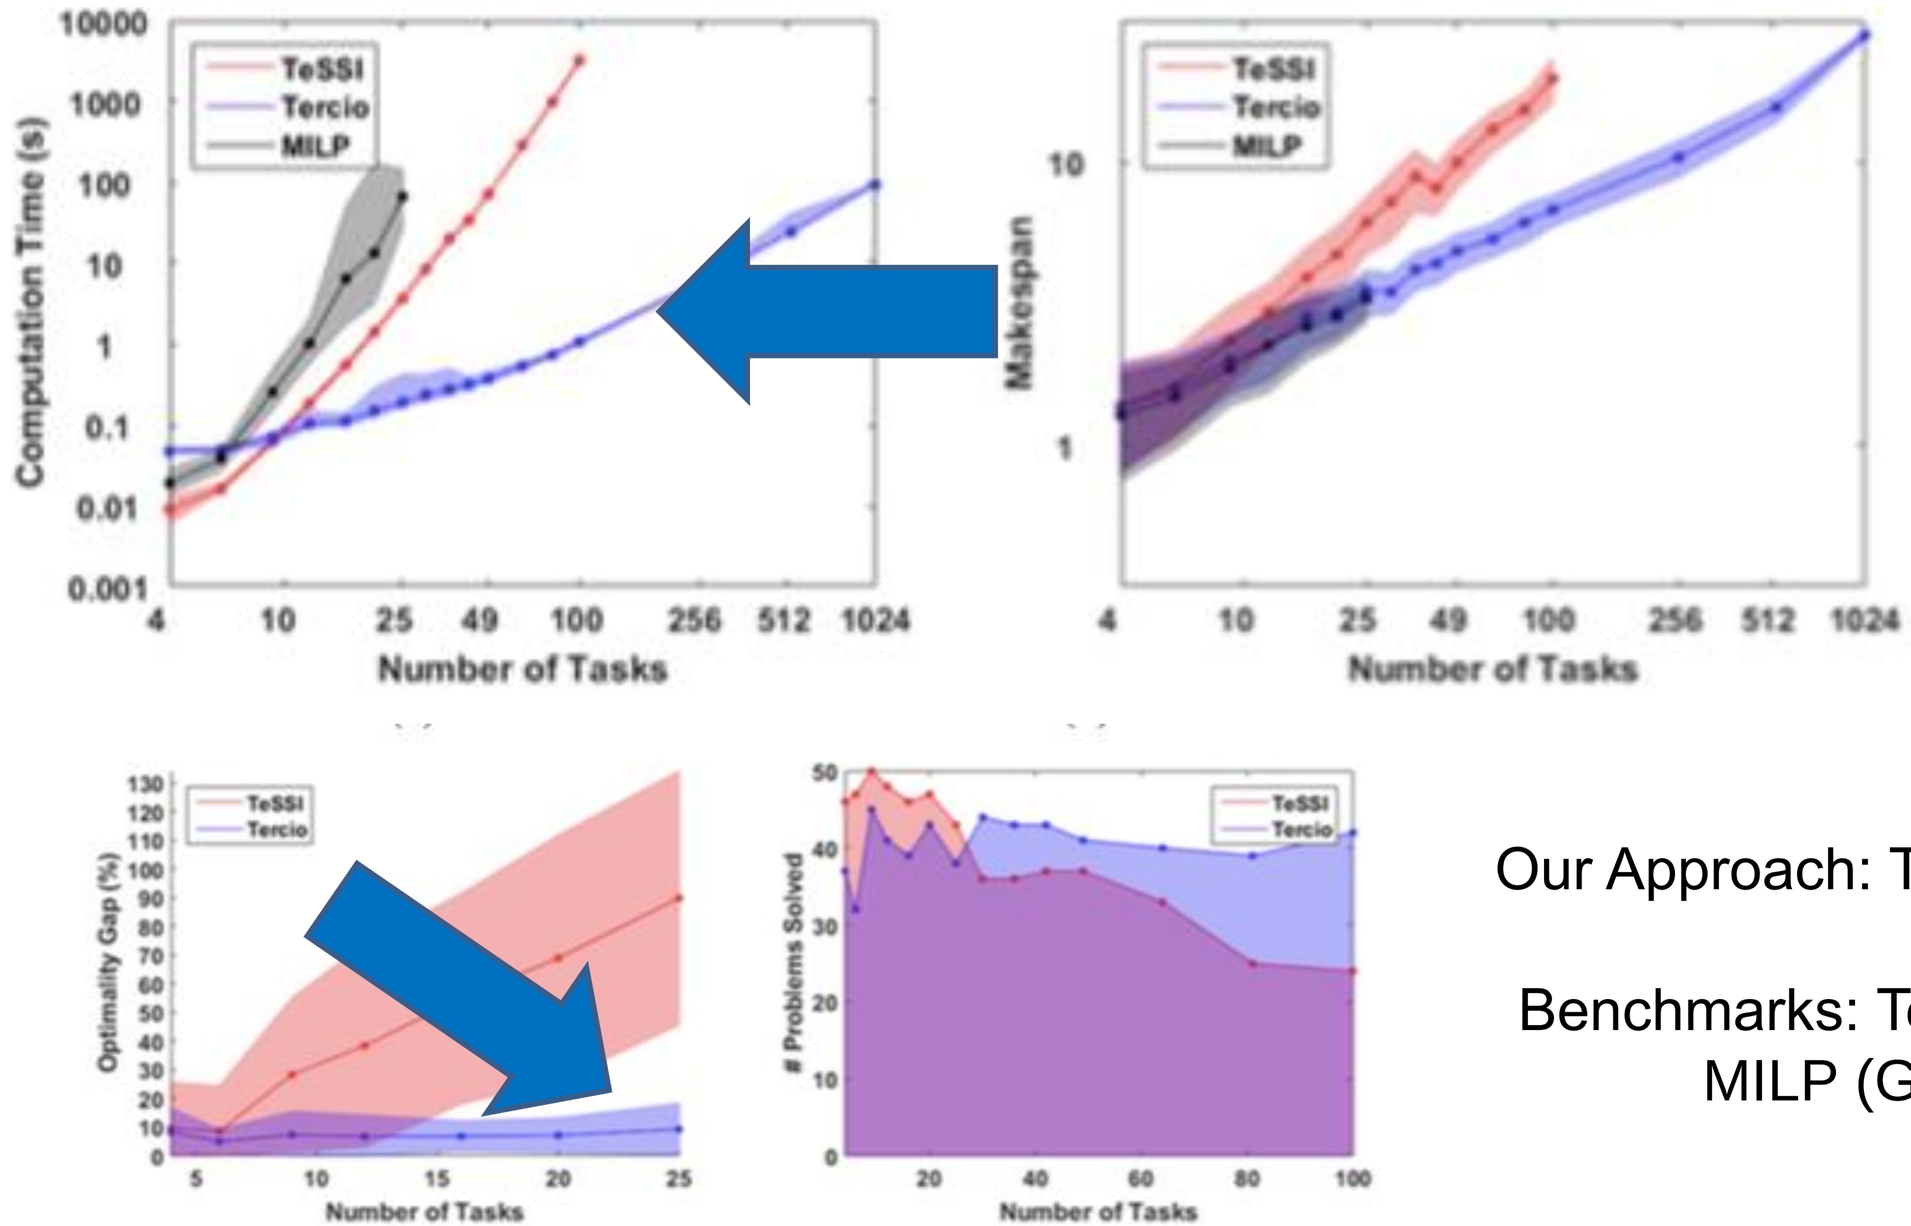

Our Approach: Tercio RSS '13

Benchmarks: TeSSi AAAI'15
MILP (Gurobi)

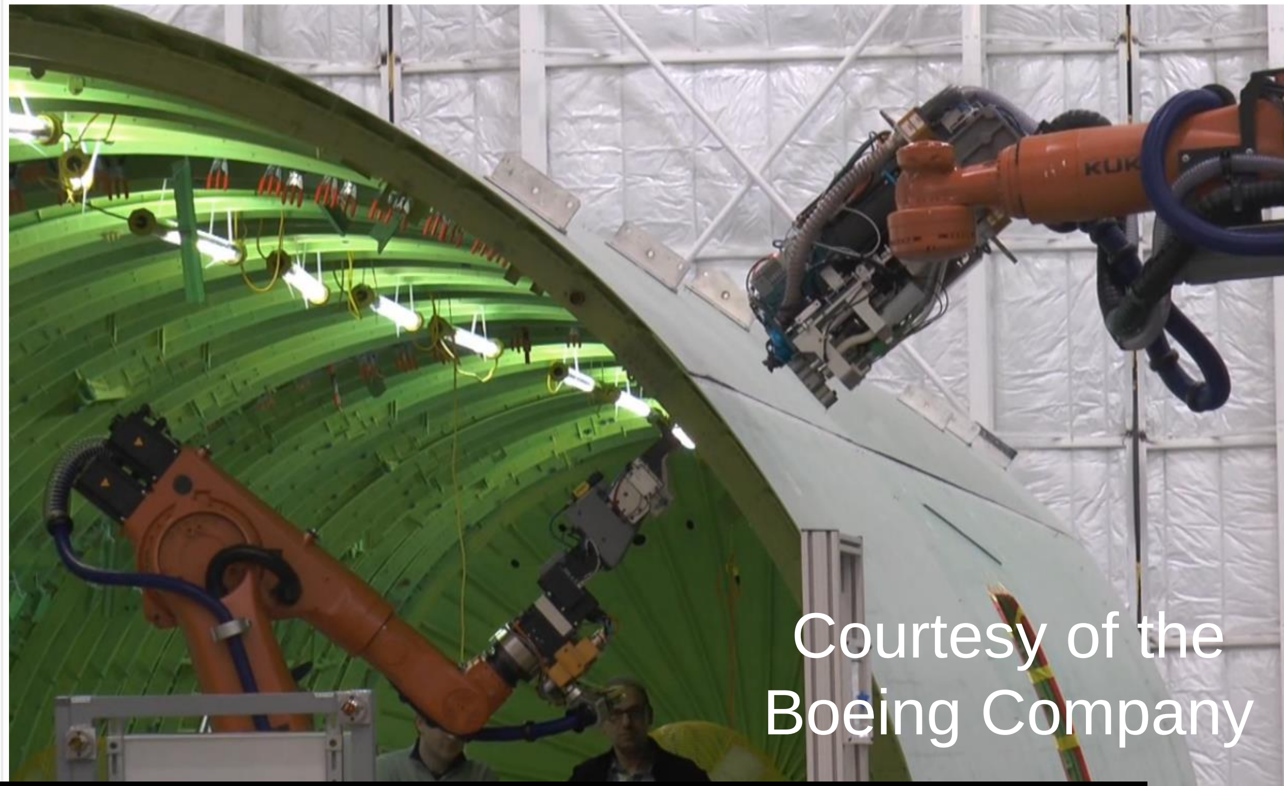
Courtesy of the
Boeing Company

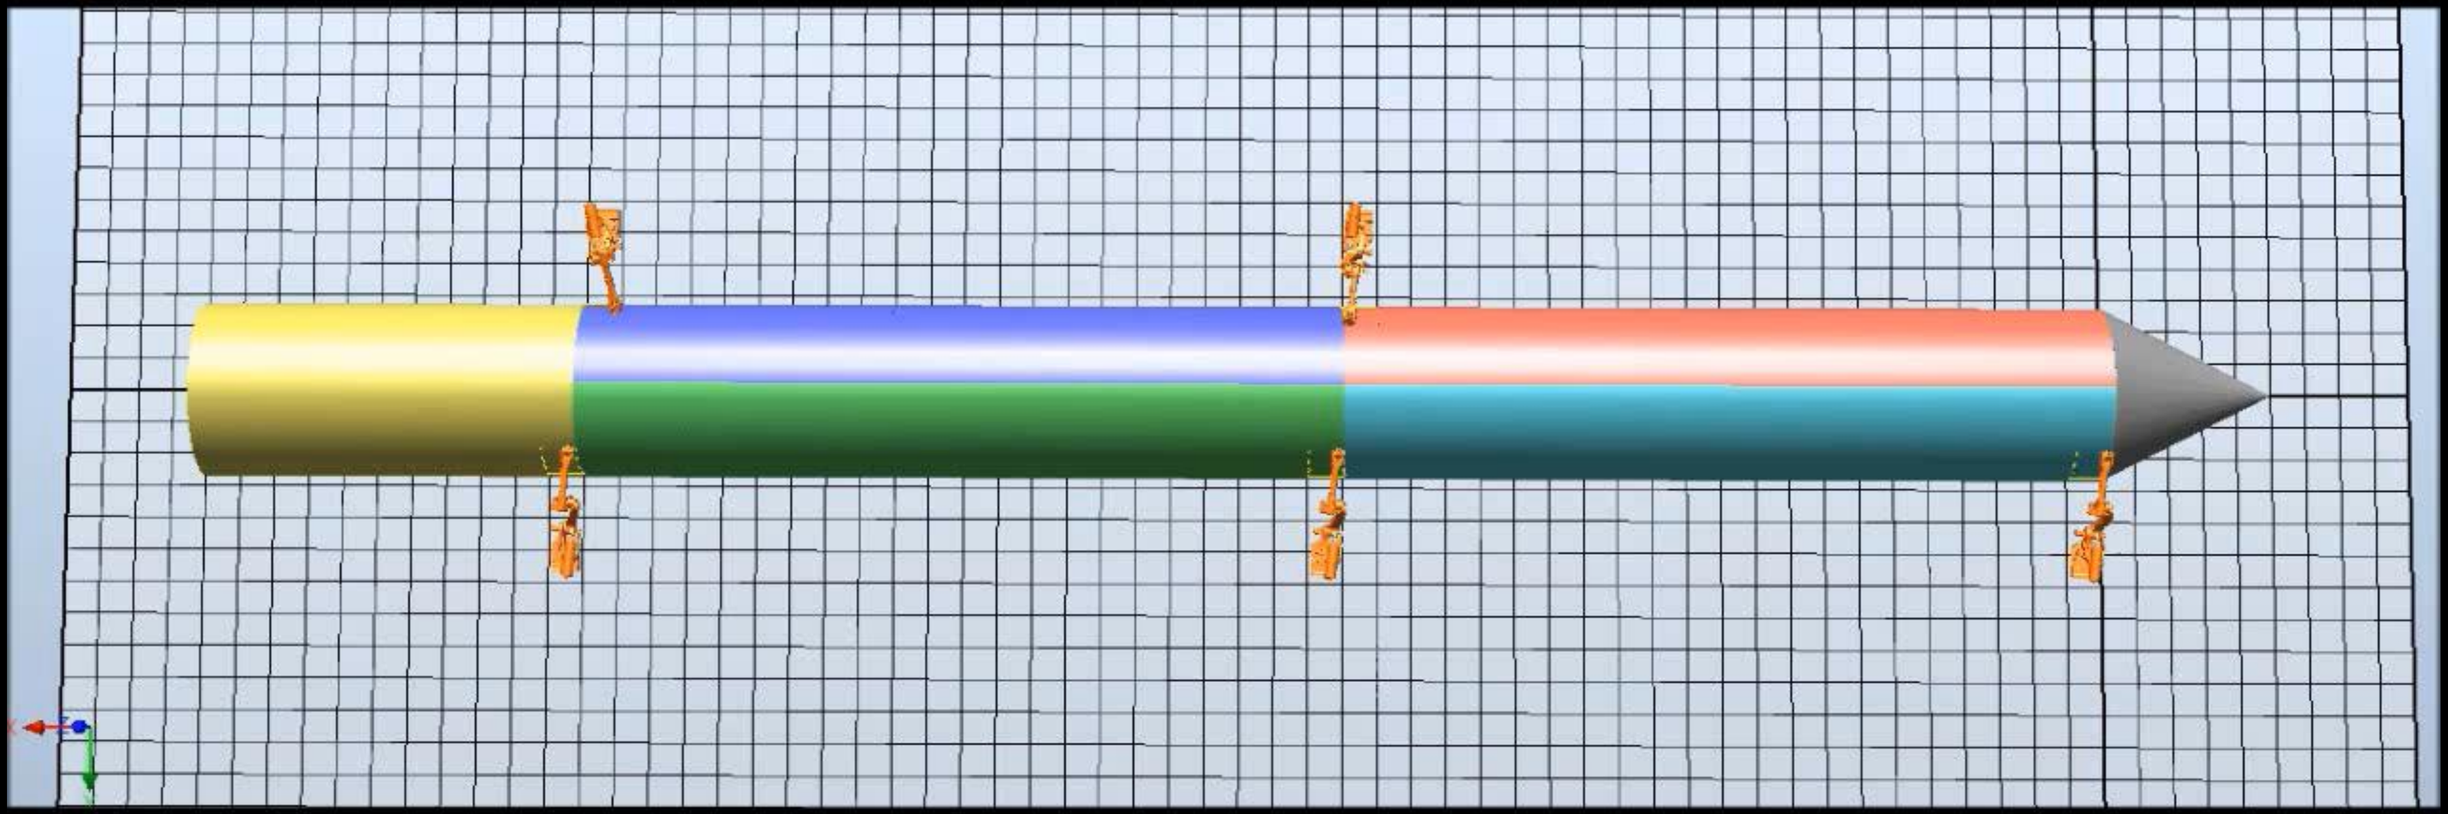

Prediction enables close-proximity collaborative robotics

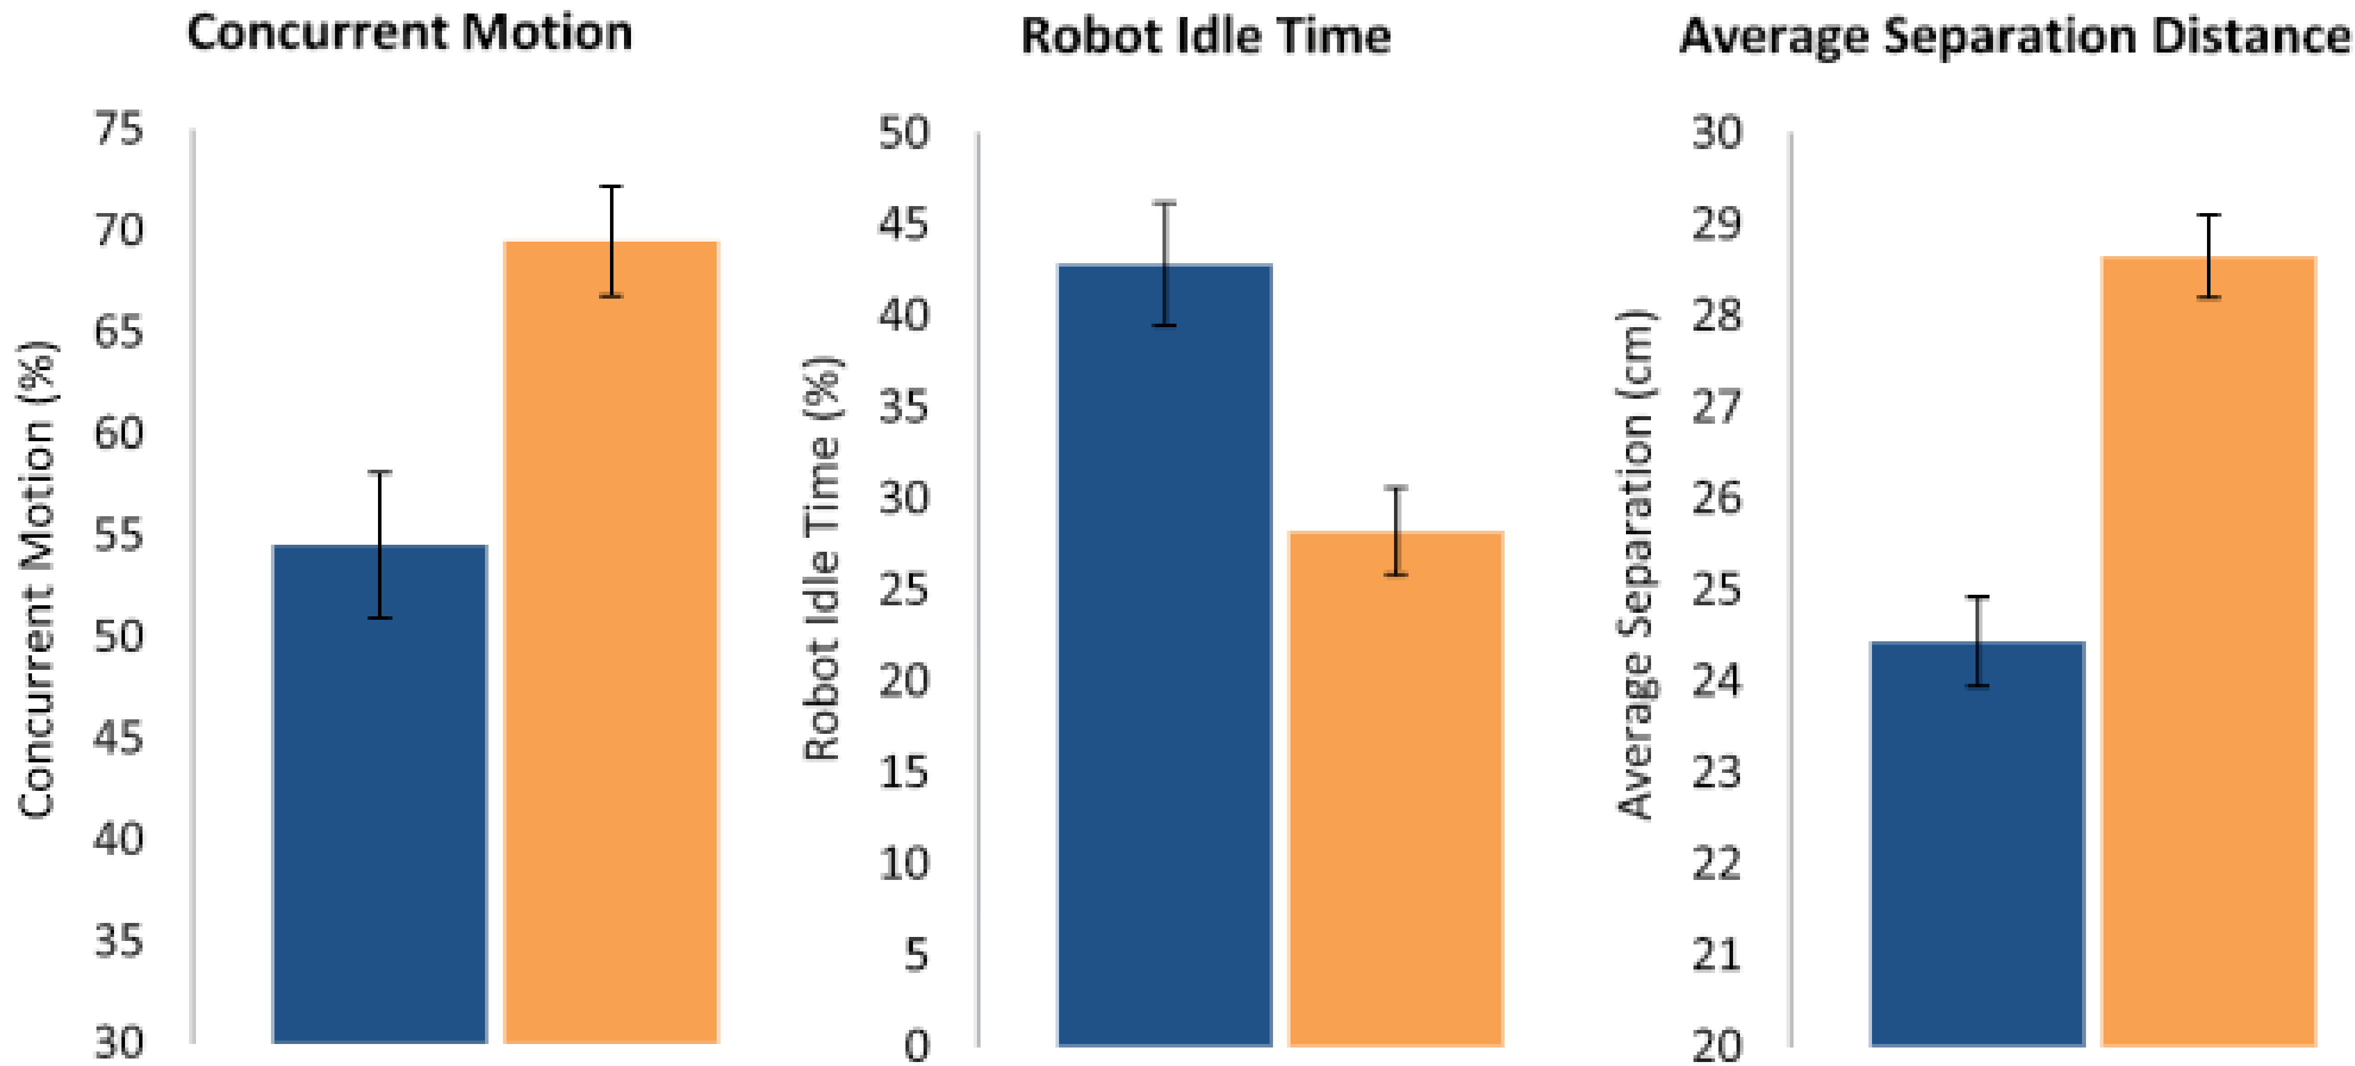

Improved team fluency metrics

Lasota, P., Shah, J. A.: Analyzing the Effects of Human-Aware Motion Planning on Close-Proximity Human–Robot Collaboration. Human Factors: The Journal of the Human Factors and Ergonomics Society, 2015.

Leveraging the Biomechanical Model for Fast Target Prediction of Human Motion

Perez D'Arpino et al. ICRA'15, Lasota et al. ICRA'17, Hayes et al. ICRA'17

Building a library of motions (offline, using human motion data)

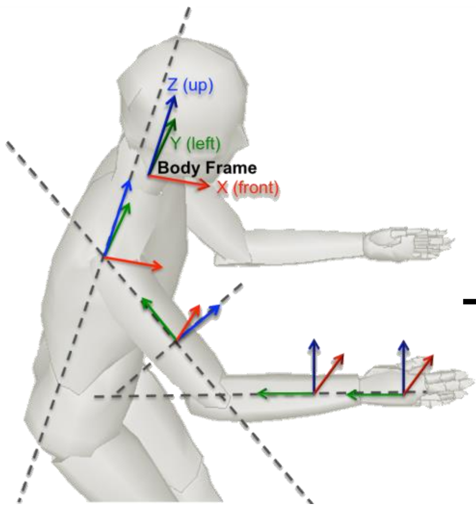

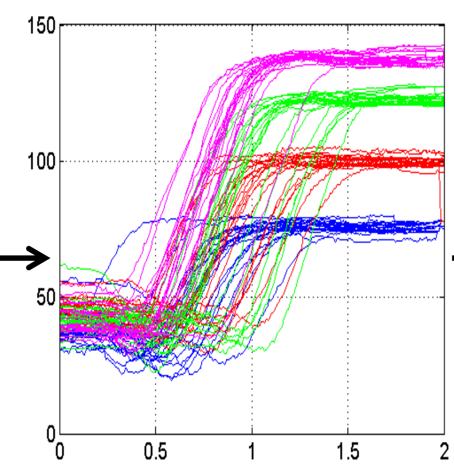

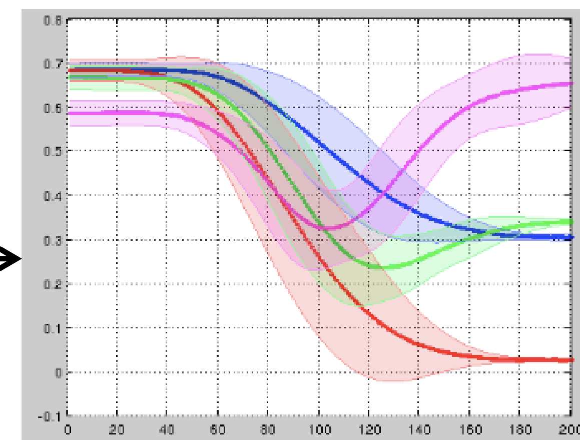

Average prediction accuracy after observing 20% and 43% of the trajectory

	20%	43%
This method	73.26%	89.55%
GMM*	57.08%	85.83%

7 demonstrations, 12 motion classes

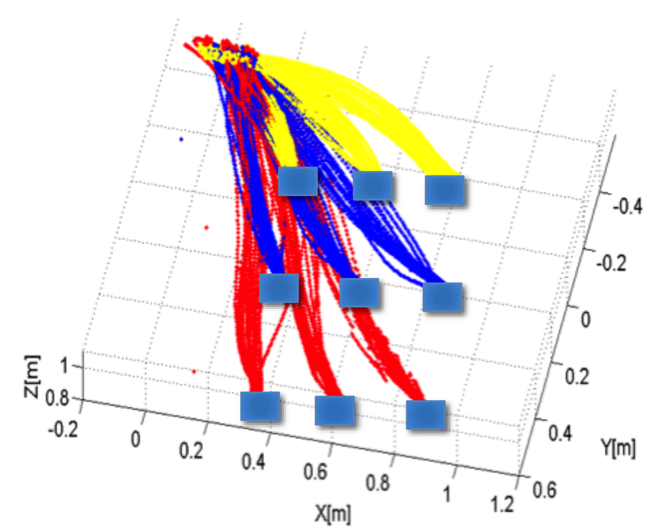

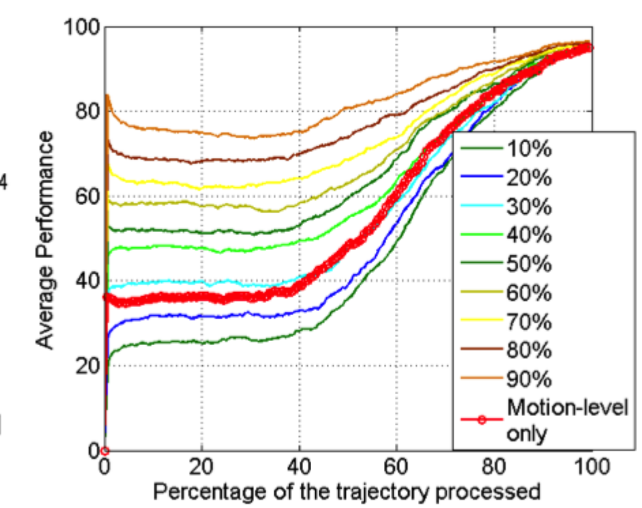

Pérez-D'Arpino, C., and J. Shah, "Fast Target Prediction of Human Reaching Motion for Cooperative Human-Robot Manipulation Tasks Using Time Series Classification", ICRA 2015.

* Mainprice, Jim and Berenson, Dmitry: Human-robot collaborative manipulation planning using early prediction of human motion. IROS, 2013

Application to Target Prediction of Human Walking

Prediction without Anticipatory Indicators

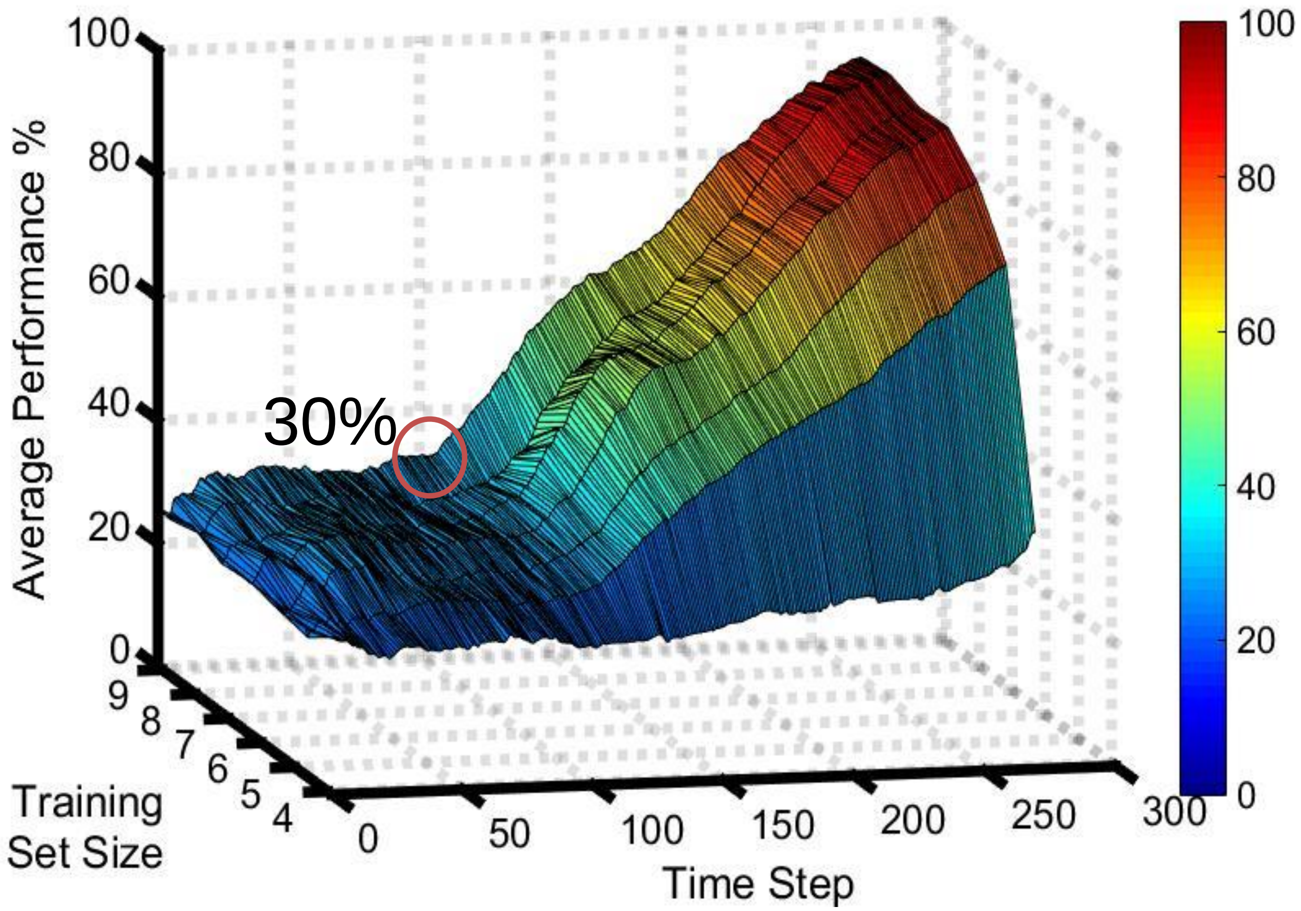

Prediction with Anticipatory Indicators

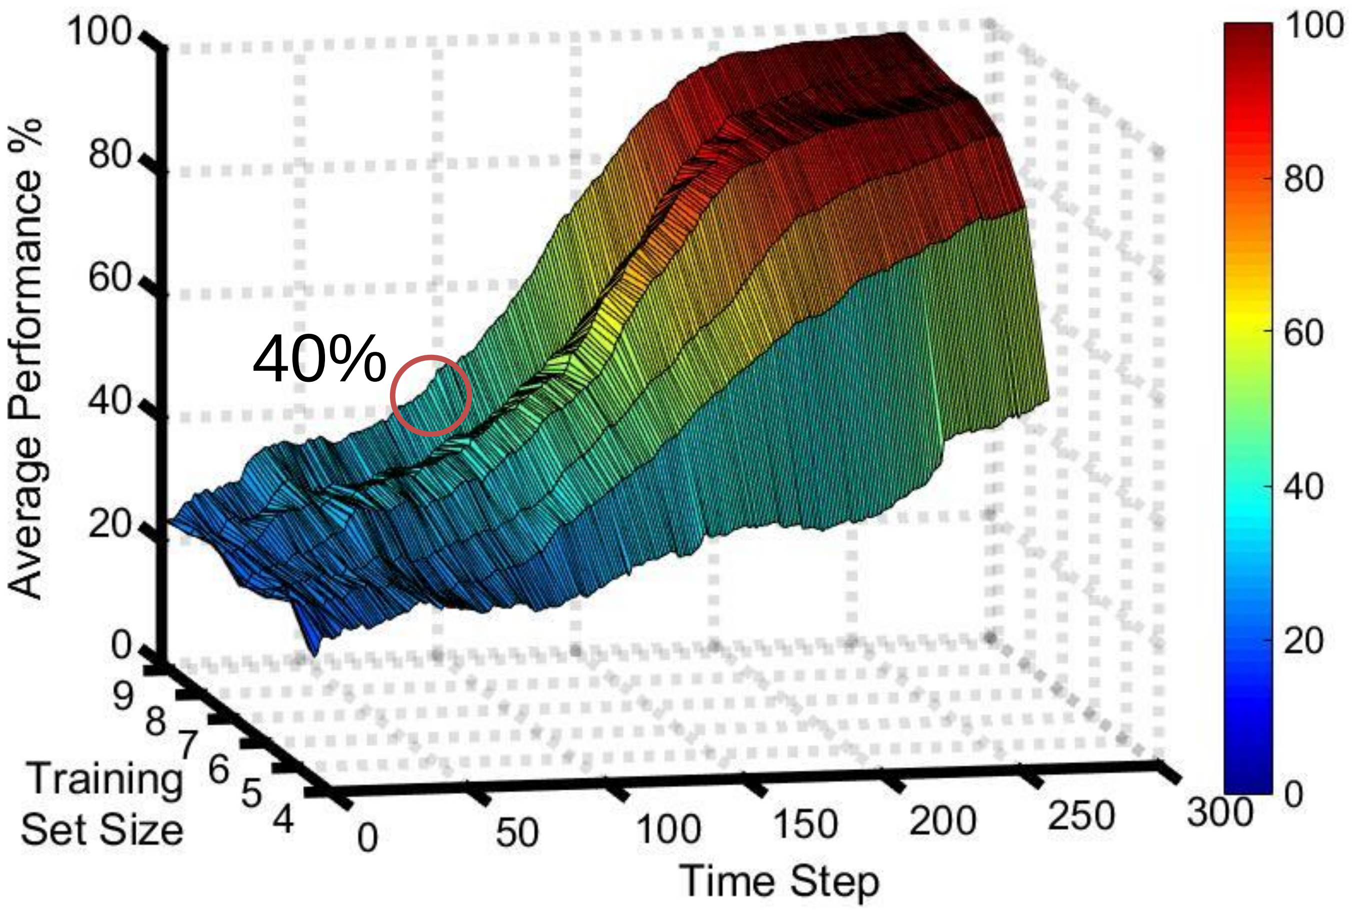

- Data from 6 participants, 25 demonstrations collected per participant.
  - Data set aggregates demonstrations from all participants.
    - Random subsampling validation within that data set.

Vaibhav V. Unhelkar**, Claudia Pérez-D'Arpino**, Leia Stirling, Julie A. Shah:
Human-Robot Co-Navigation Using Anticipatory Indicators of Human Walking Motion. ICRA 2015.
** These authors contributed equally to this work.

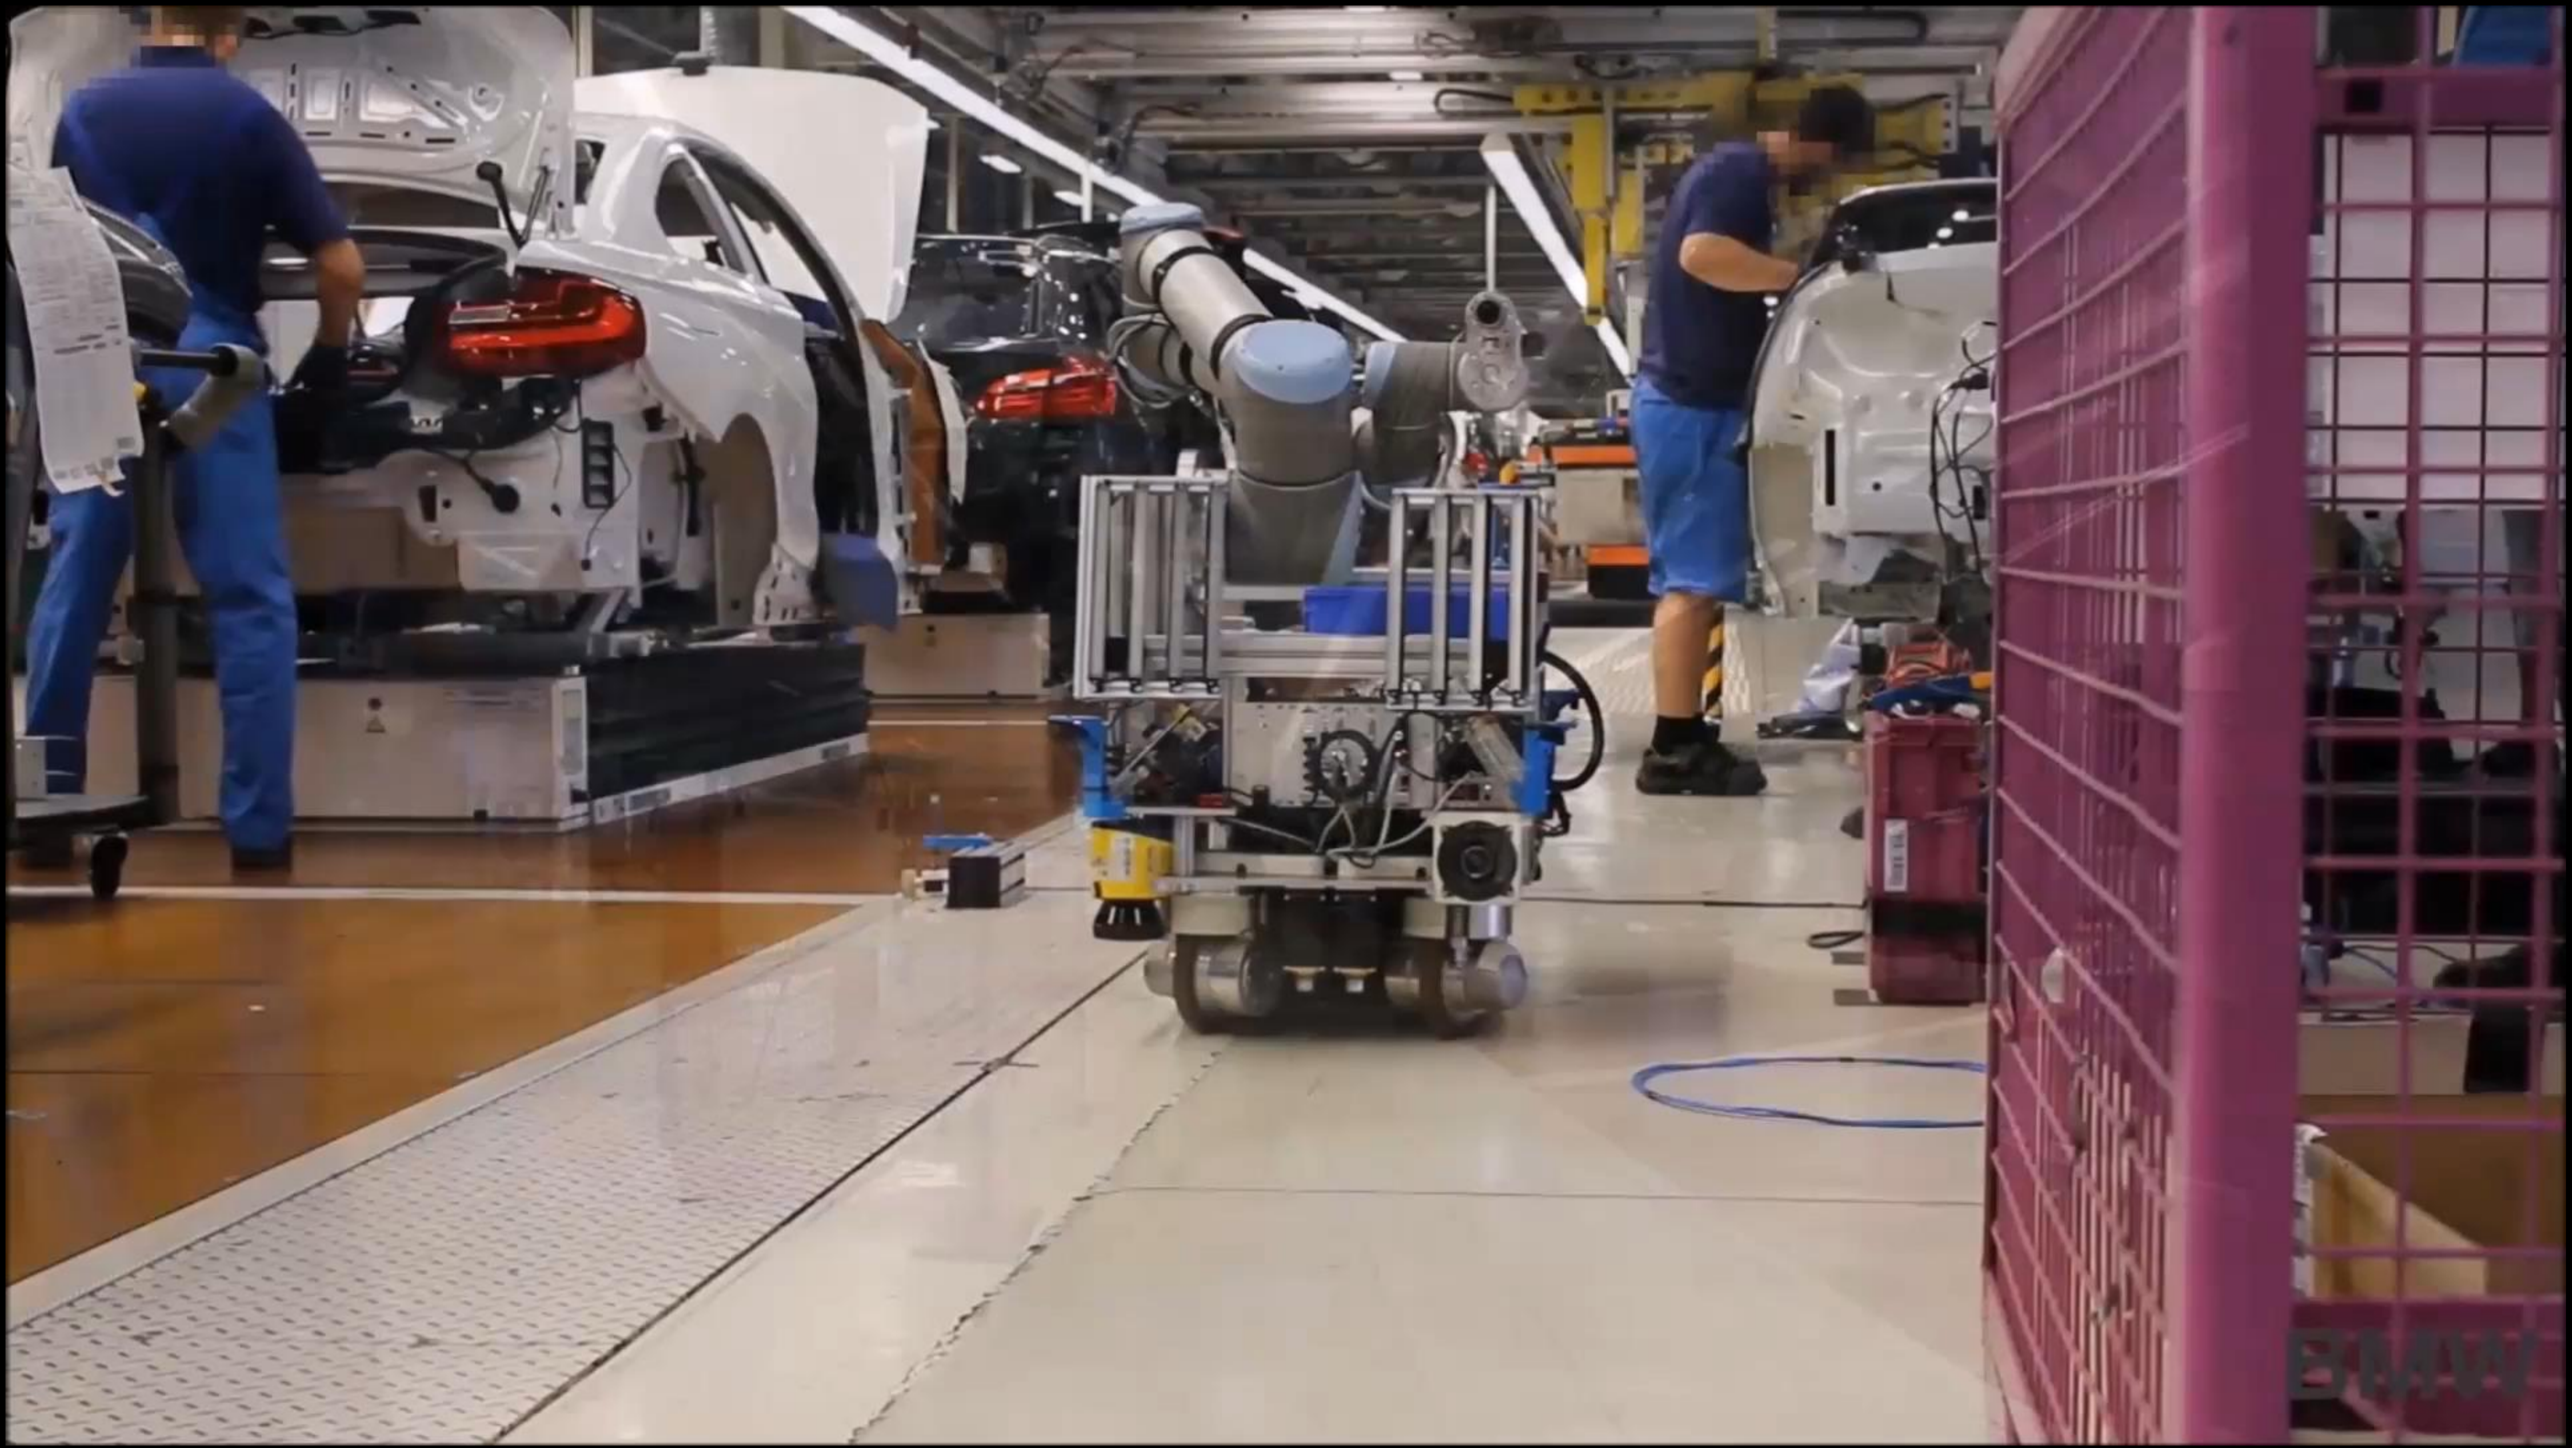

Courtesy of BMW

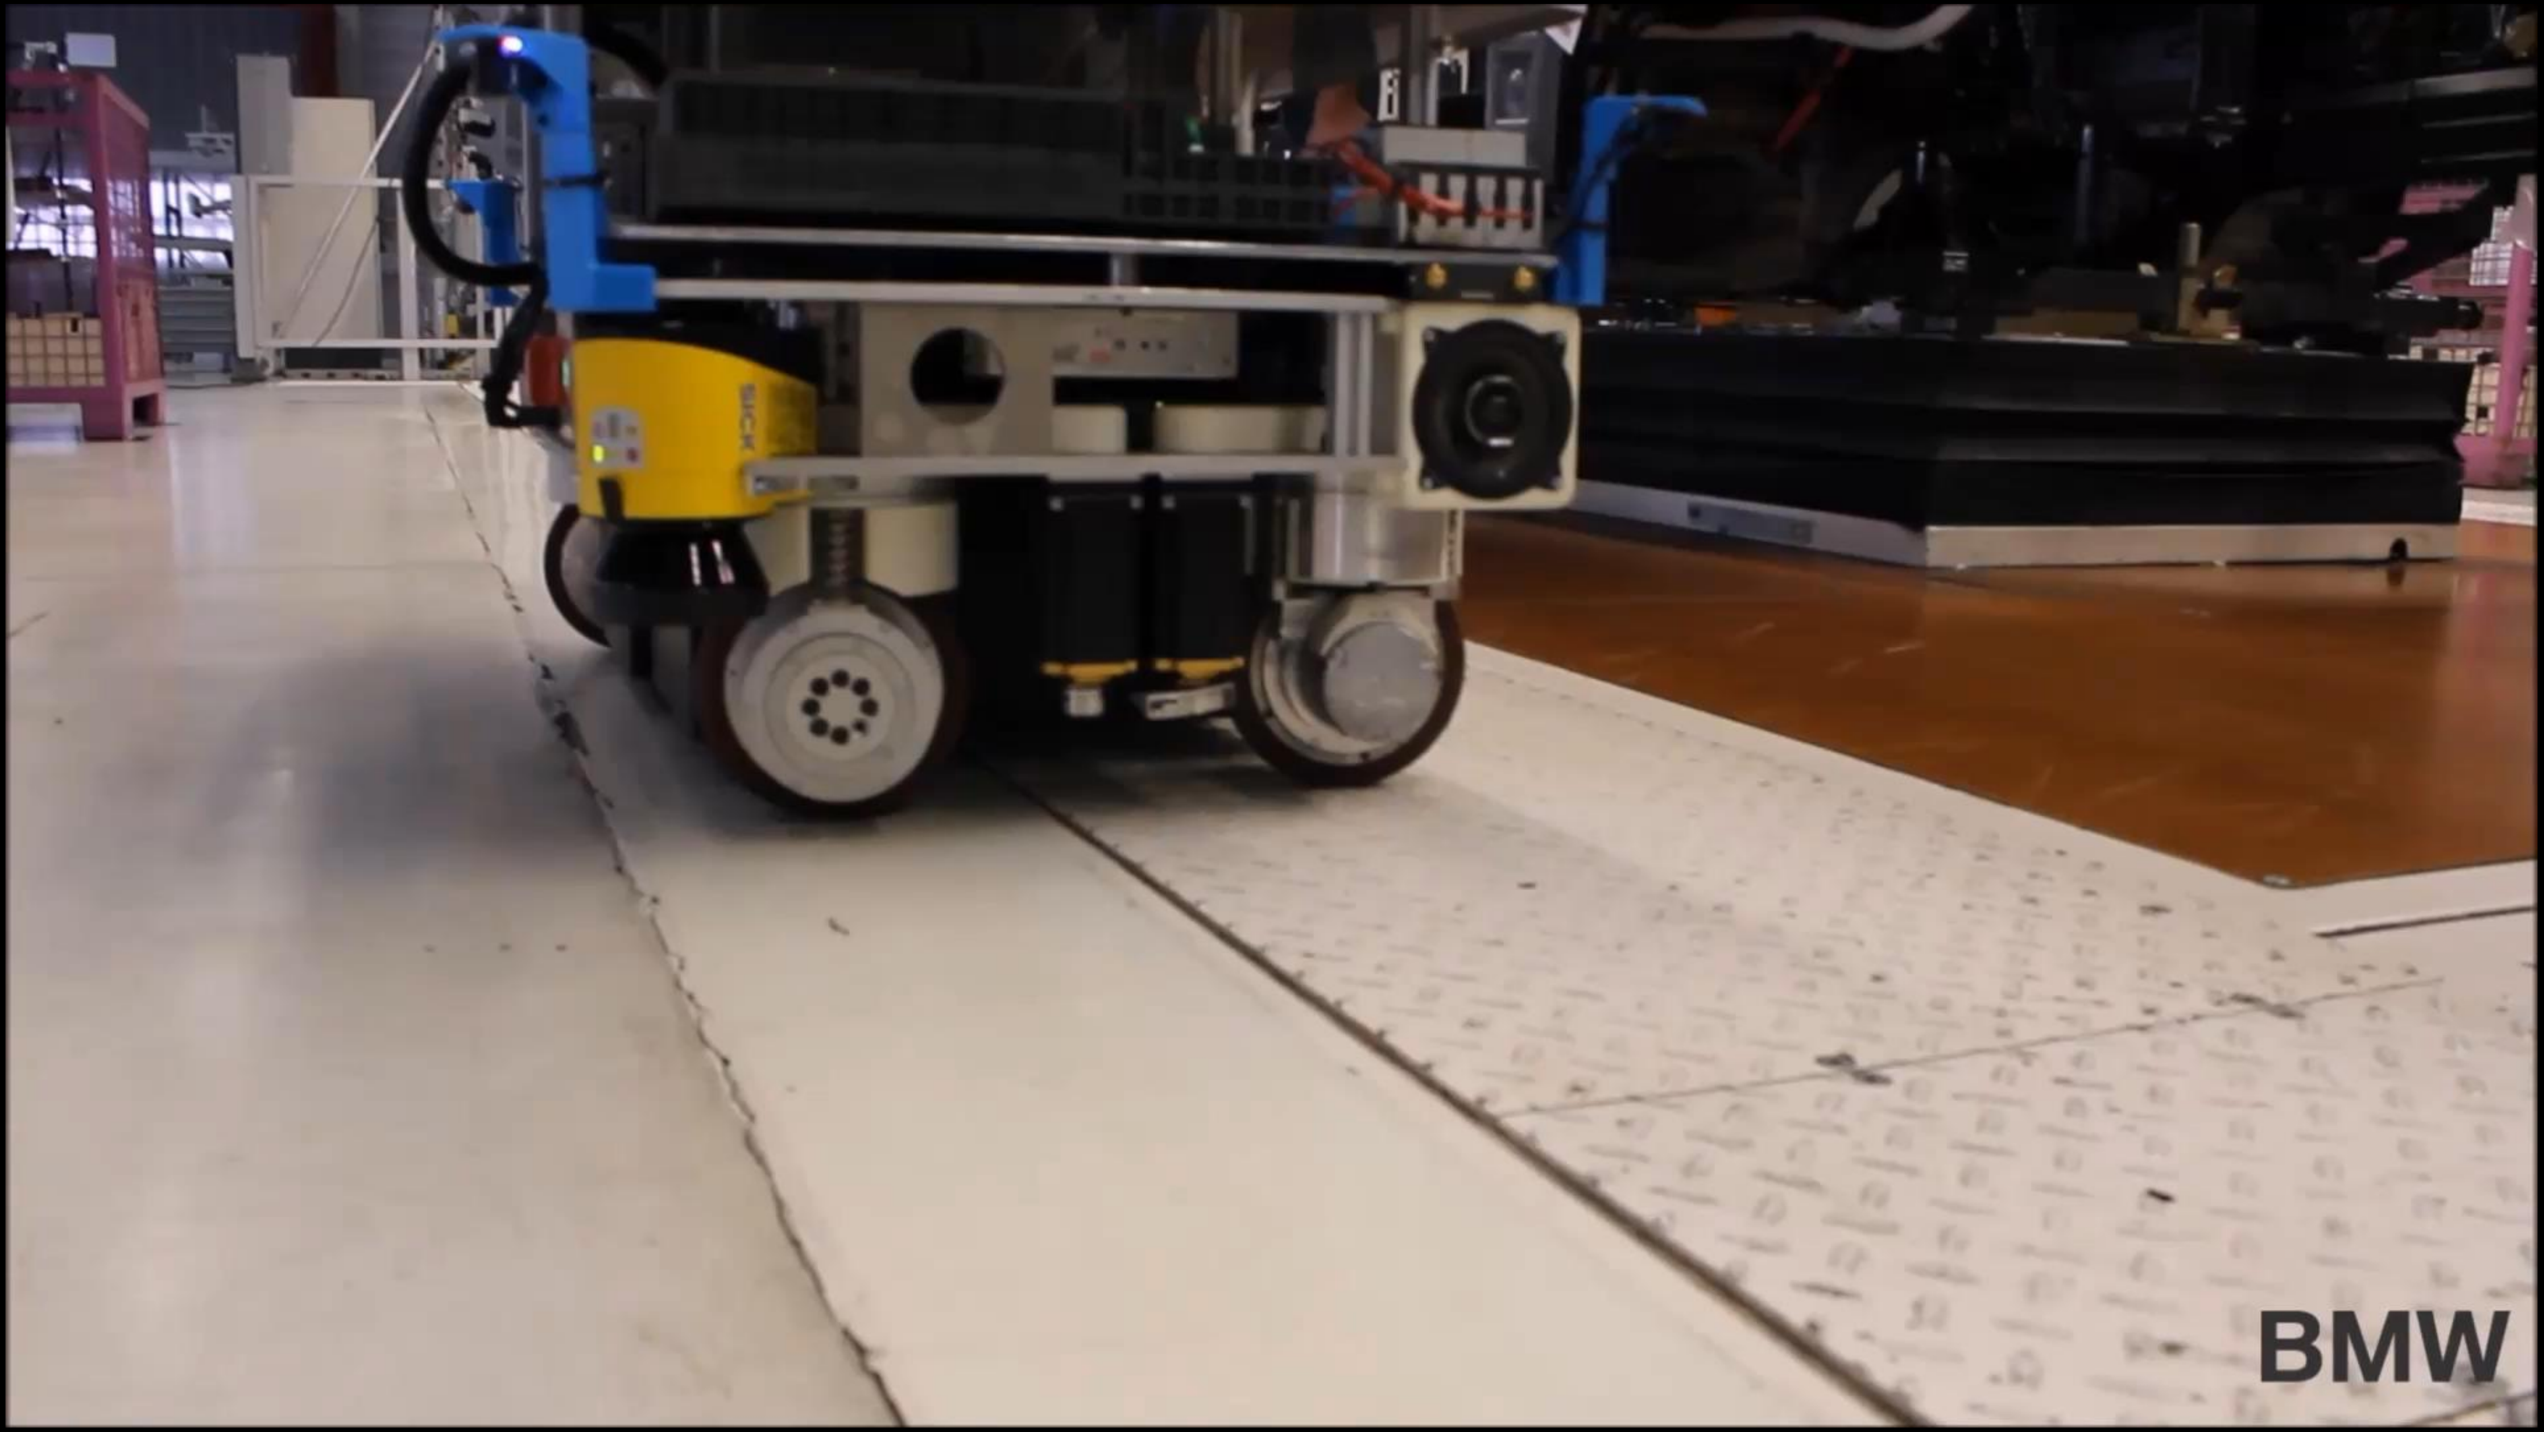

Courtesy of BMW

Contributions

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

Well-established cognitive models → Meaningful features that relate to mental state

Well-established cognitive models → Model structure to process complex information efficiently

- Experiments validate that these models yield richer, flexible human-machine interactions
  - making higher quality shared plans [IEEE THMS'16, RSS'16 JAIR'15]
  - learning complex shared plans from observation [IJCAI'16]
- Dynamic plan execution supports flexible, real-time teaming
  - multi-robot task allocation and scheduling at scale [RSS'12, RSS'13, NIPS'14, AAAI'15 IJCAI'16]
  - human-aware planning using human motion prediction [ICRA'15a, ICRA'15b, Human Factors '15]