Reasoning about Actions for Planning in Robotics

Shiqi Zhang
SUNY Binghamton

10/28/2018

The SUNY System

- 64 campuses
- Four PhD-granting "University Centers"
  - Albany
  - **Binghamton – most selective SUNY**
  - Buffalo
  - Stony Brook

Community — Greater Binghamton

- Located in New York state
- Birthplace of IBM (Endicott, NY)
- Home to several hi-tech companies.
- One of the safest U.S. midsized cities
- Low cost of living (12% below U.S. average)
- Close to major cities

Rankings

Binghamton University--SUNY

Binghamton, NY

#38 in Top Public Schools

#87 in National Universities (tie)

$24,403 (out-of-state), $9,523 (in-state) Tuition and Fees

13,632 Undergraduate Enrollment

2018 Rank	School
#25 Tie	Virginia Tech Blacksburg, VA
#29 Tie	University of Massachusetts—Amherst Amherst, MA
#33 Tie	Florida State University Tallahassee, FL
#33 Tie	Michigan State University East Lansing, MI
#38 Tie	Binghamton University—SUNY Binghamton, NY
#39 Tie	University of Colorado—Boulder Boulder, CO
#41 Tie	Stony Brook University—SUNY Stony Brook, NY
#41	University at Buffalo—SUNY Buffalo, NY

Computer Science Faculty

- 33 full-time faculty
  - 8 full professors
  - 8 associate professors
  - 11 assistant professors
  - 6 lecturers
- 4 adjunct lecturers
- 3 new faculty members in Robotics/AI, Computer Vision/Machine Learning, and Computer Architecture will join in Fall 2018.

Department also has close to 40 teaching Assistants

Reasoning about Actions for Planning in Robotics

Shiqi Zhang
SUNY Binghamton

10/28/2018

Why ***planning*** in robotics?

- Complex tasks in the real world require more than one action
- Robot actions (perception and actuation) are unreliable, and sometimes costly

Robots need to plan actions
to accomplish goals under uncertainty

Why ***reasoning*** in robotics?

- Robot faces many objects (locations, people, tools, etc) and their properties
- World state estimation with incomplete (qualitative and quantitative) knowledge

Robots need to reason
to understand the current world state

Reasoning (declarative) and planning (probabilistic)

Declarative knowledge representation & reasoning

Correct and natural
Incomplete knowledge
Explanation (good for HRI)
Goal-independent
Transferability

Strengths

Non-deterministic action outcomes
Imperfect perception
Unspecified, long horizon
Learning from experience (RL)

Robo`tics
decision-making

Probabilistic Planning & Reinforcement learning (RL)

Logical-probabilistic reasoning for probabilistic planning,

as illustrated in human-robot dialog

Robot needs to identify <Coffee, Office 1, Bob>, through spoken dialog

Time:	9:00am
Rooms:	Office 1, Office 2, …
Persons:	Alice, Bob, Carol, …
Items:	Coffee, Sandwich, ...

<Coffee, Office 1, Bob>

“I am a shopping robot, what item do you want?”

“Coffee, please”

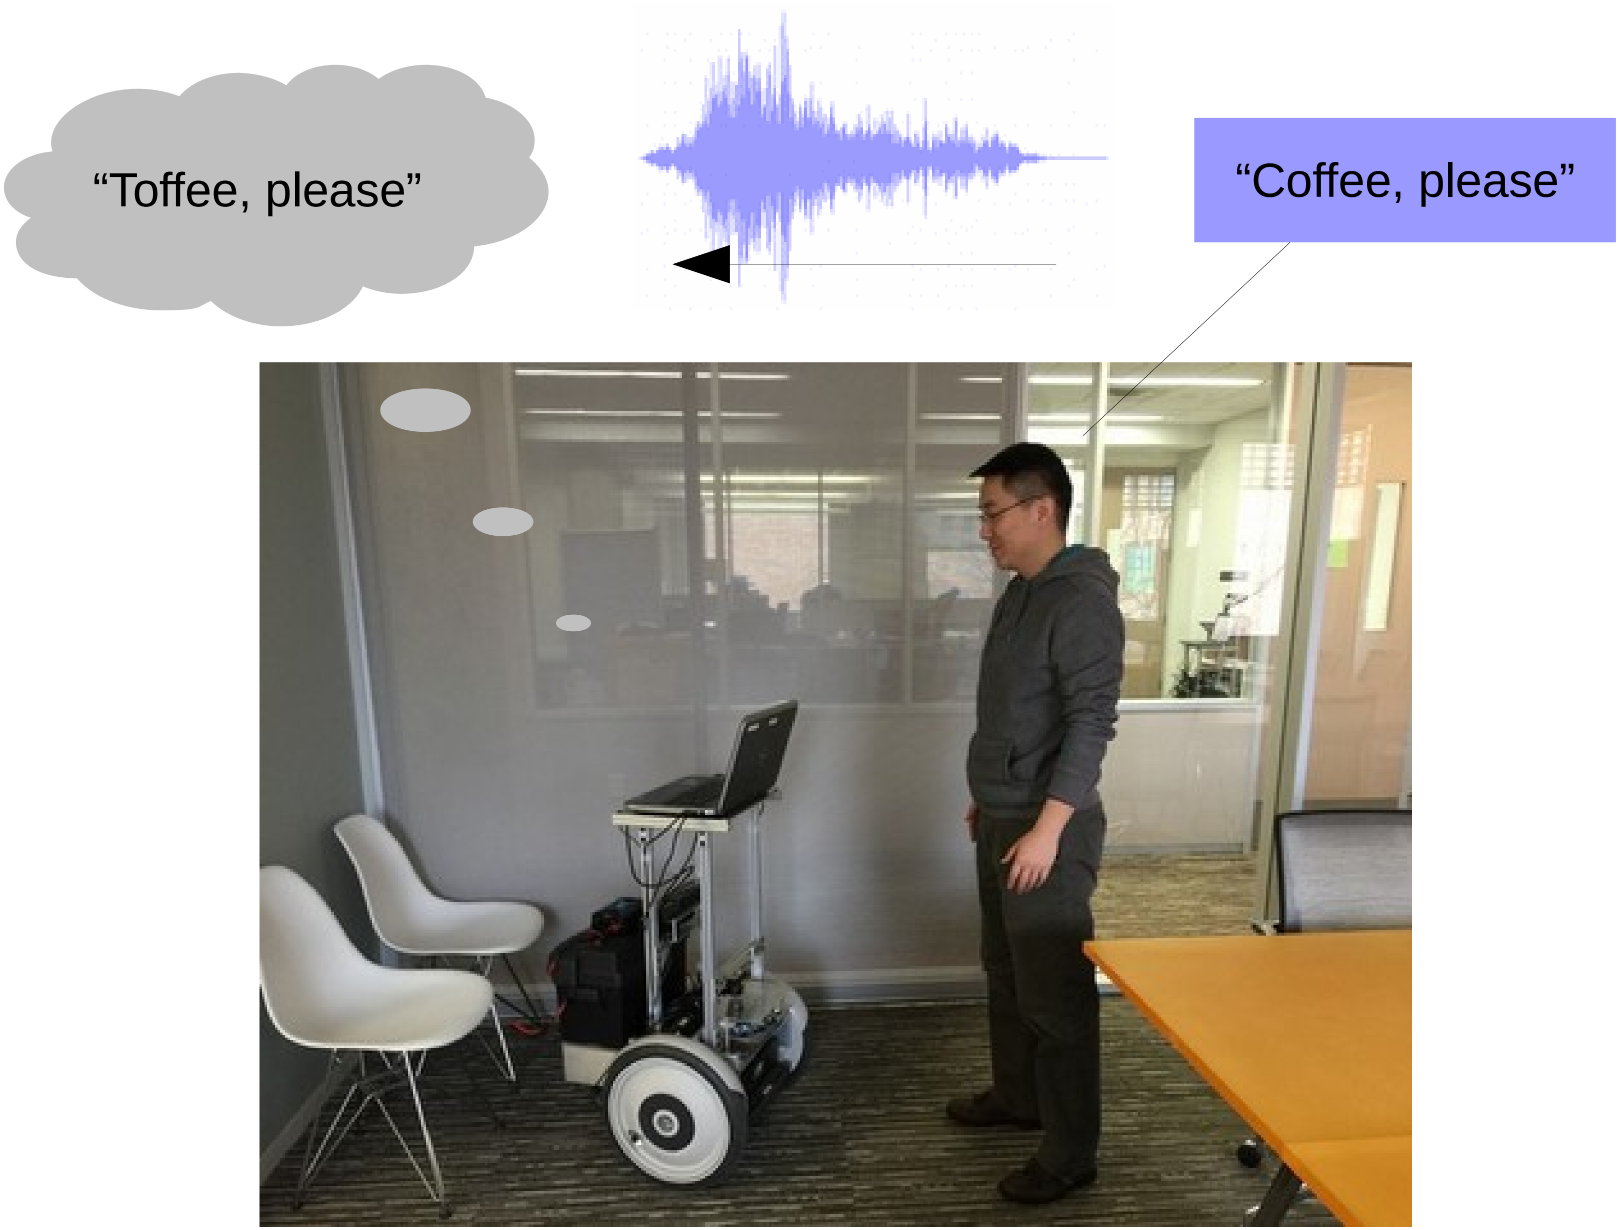
“Toffee, please”
“Coffee, please”

“Do you want me to buy toffee?”

“Coffee, please”

[***Demo video: integrated P-log and POMDP***]

[Zhang, Stone, AAAI 2015]

CORPP: commonsense reasoning and probabilistic planning, a complete example

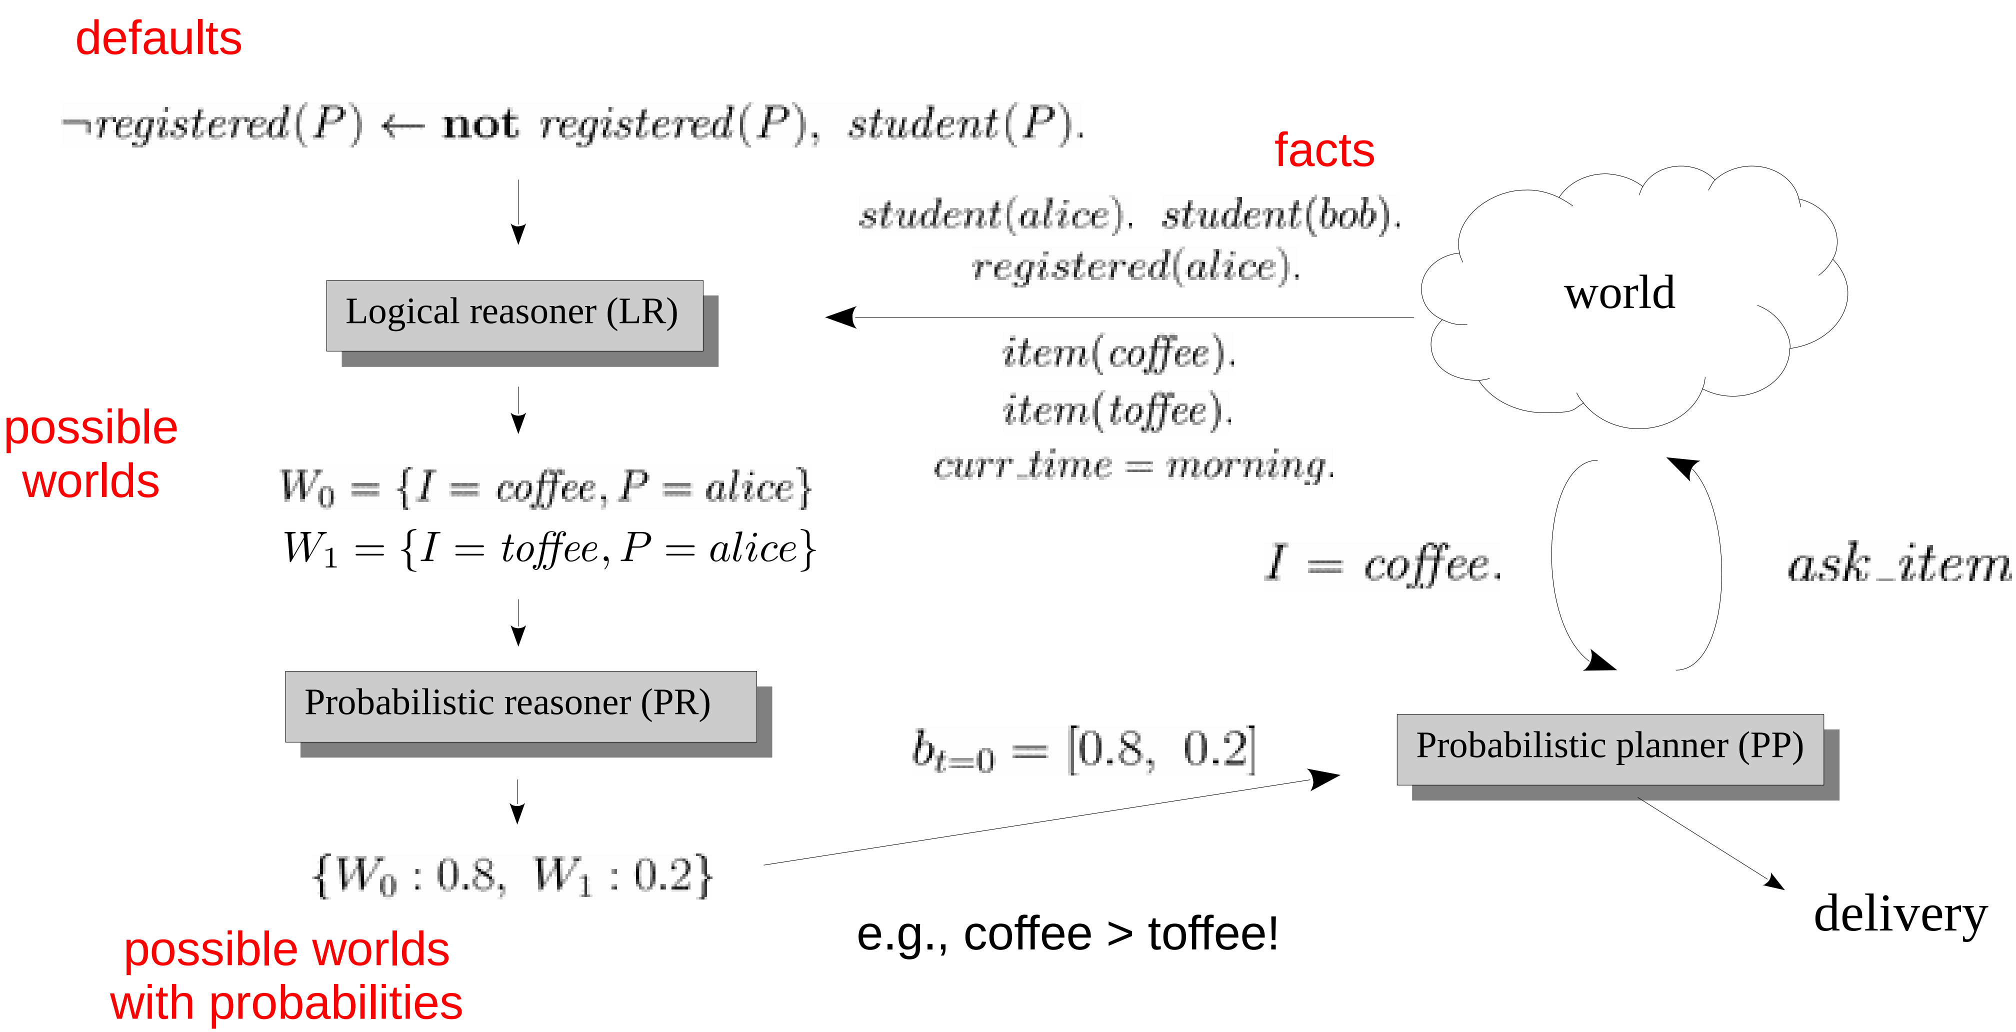

[Zhang, Stone, AAAI 2015]

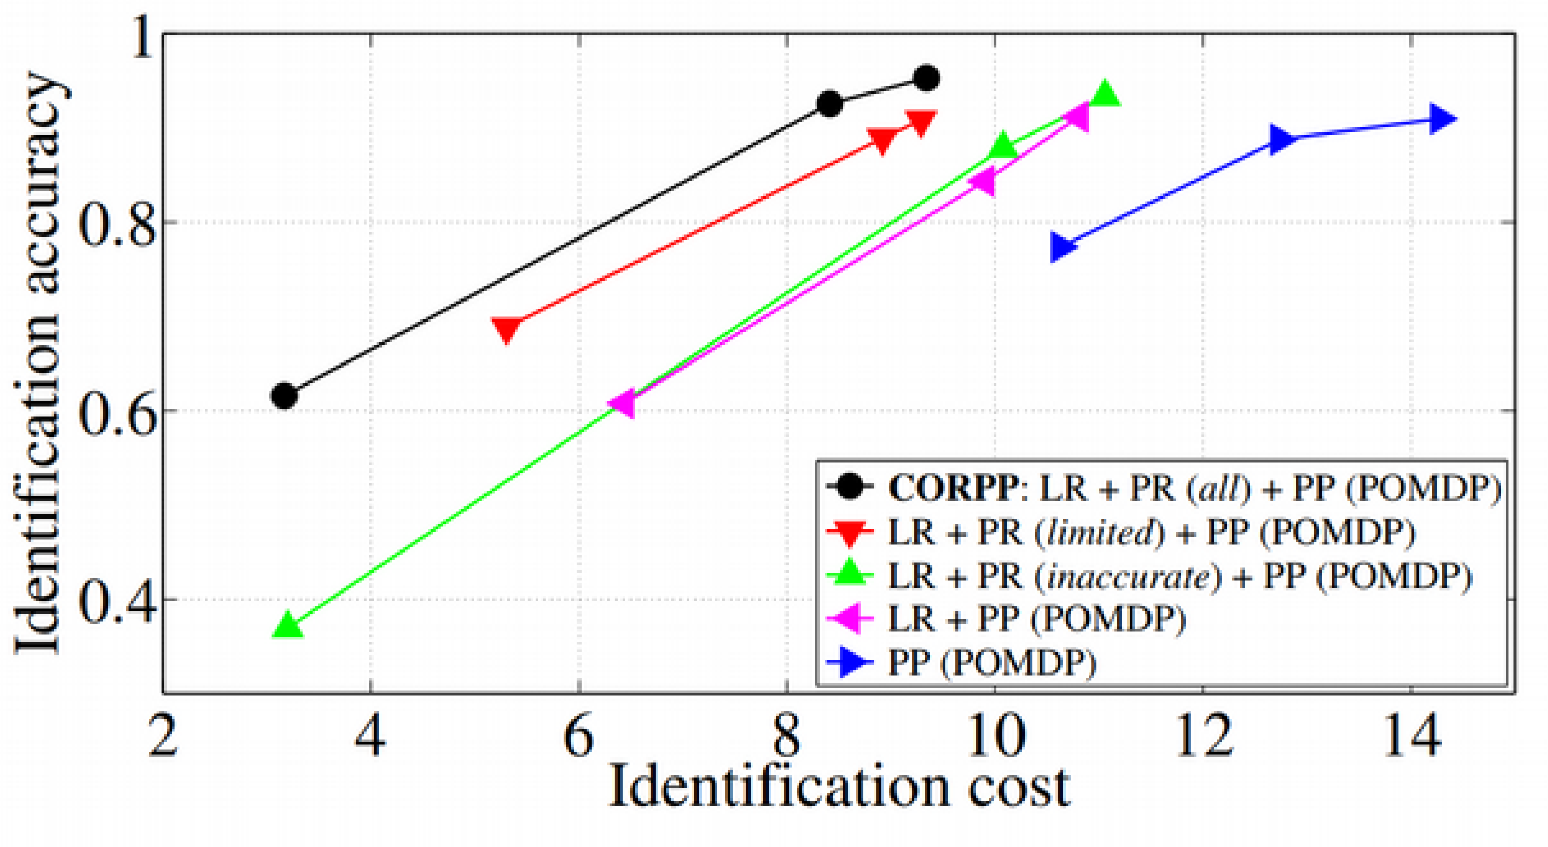

CORPP reasons with logical and probabilistic knowledge, improving robot behaviors

Interleaved CORPP (iCORPP)

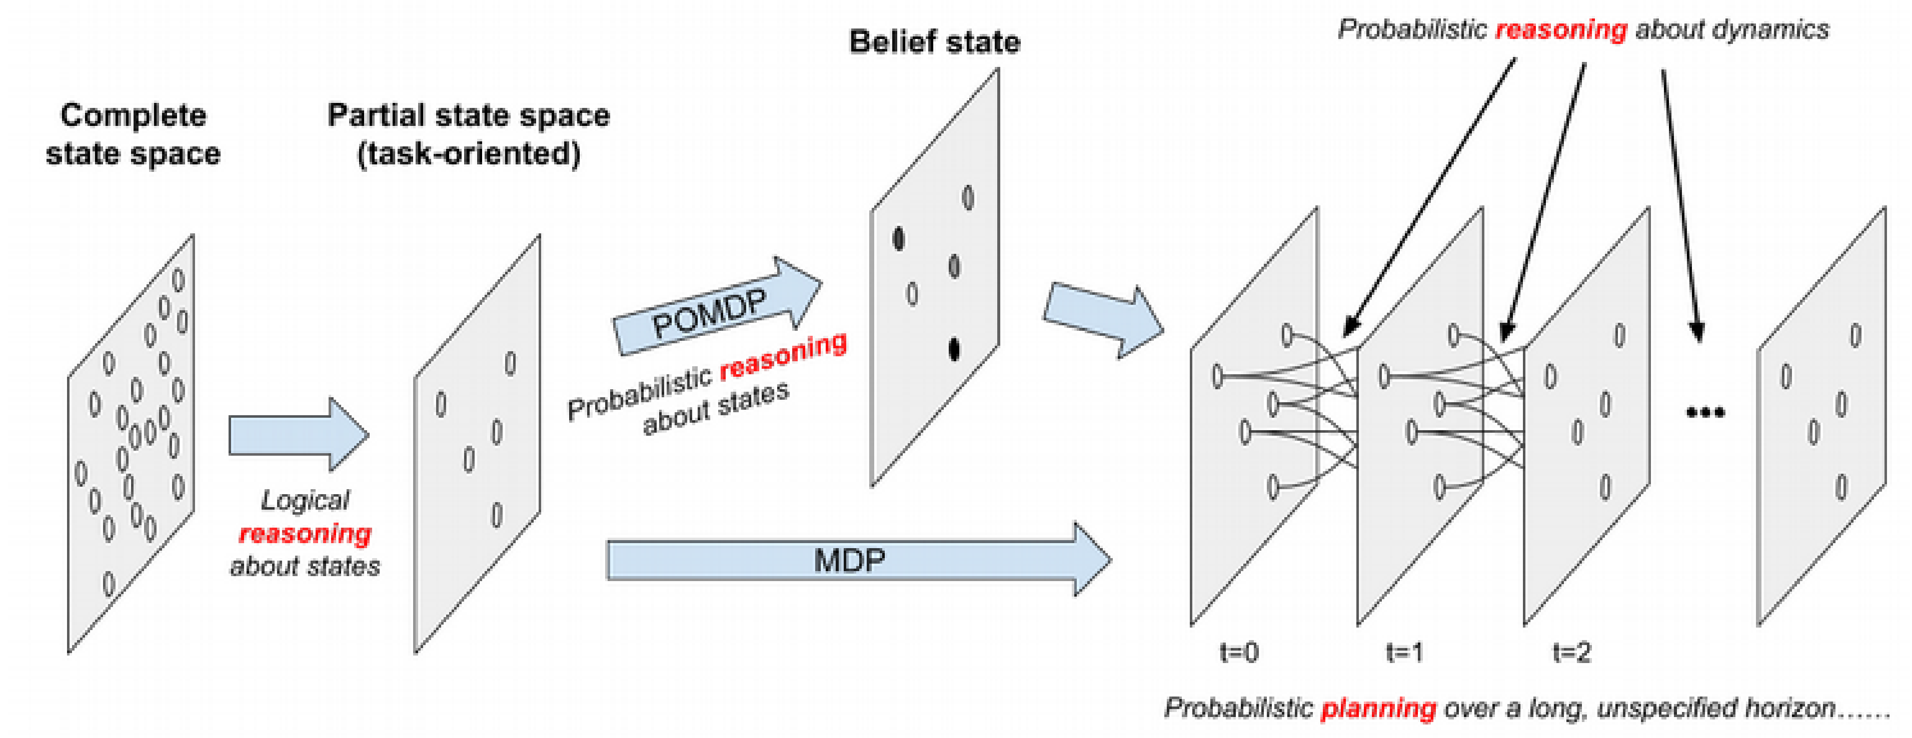

[Zhang, Khandelwal, Stone, AAAI 2017]

Interleaved CORPP (iCORPP): Dynamically Constructed (PO)MDPs for Adaptive Robot Planning

Example domain: ***robot navigation***

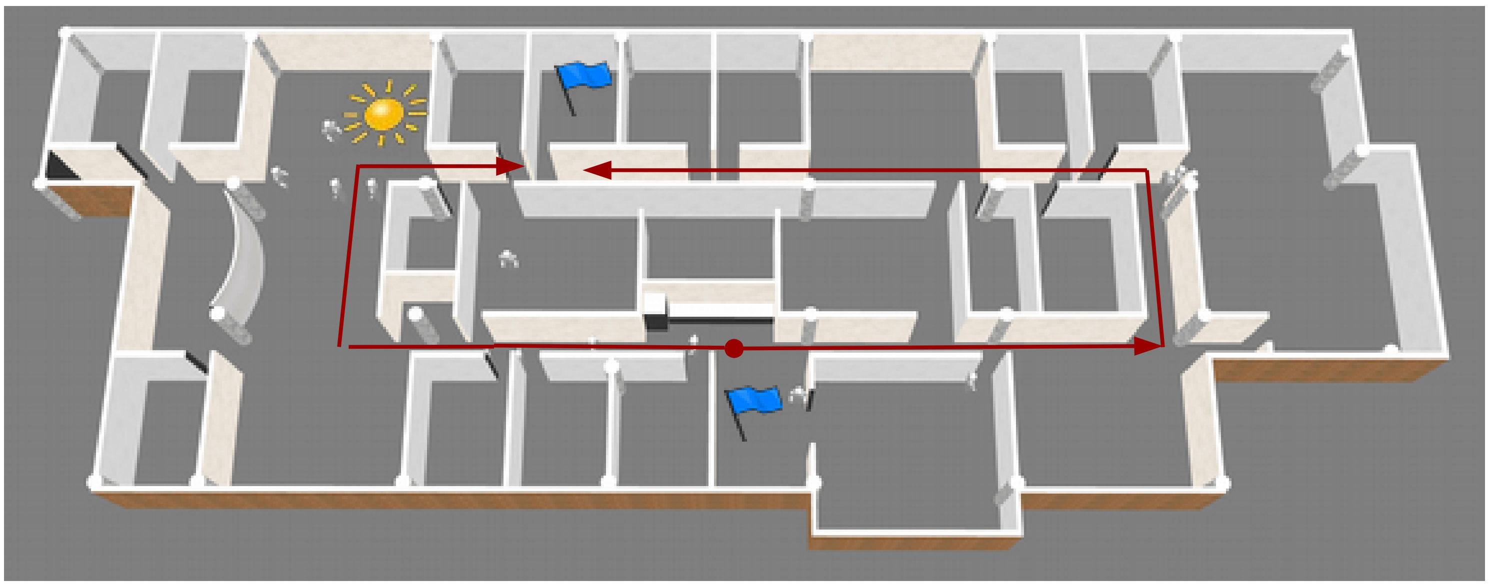

Robot locations	Areas under sunlight	Areas blocked	Weather	Time
10	2^10	2^10	5	3

More than **2^27** states!

Interleaved CORPP (iCORPP):
Dynamically Constructed (PO)MDPs for Adaptive Robot Planning

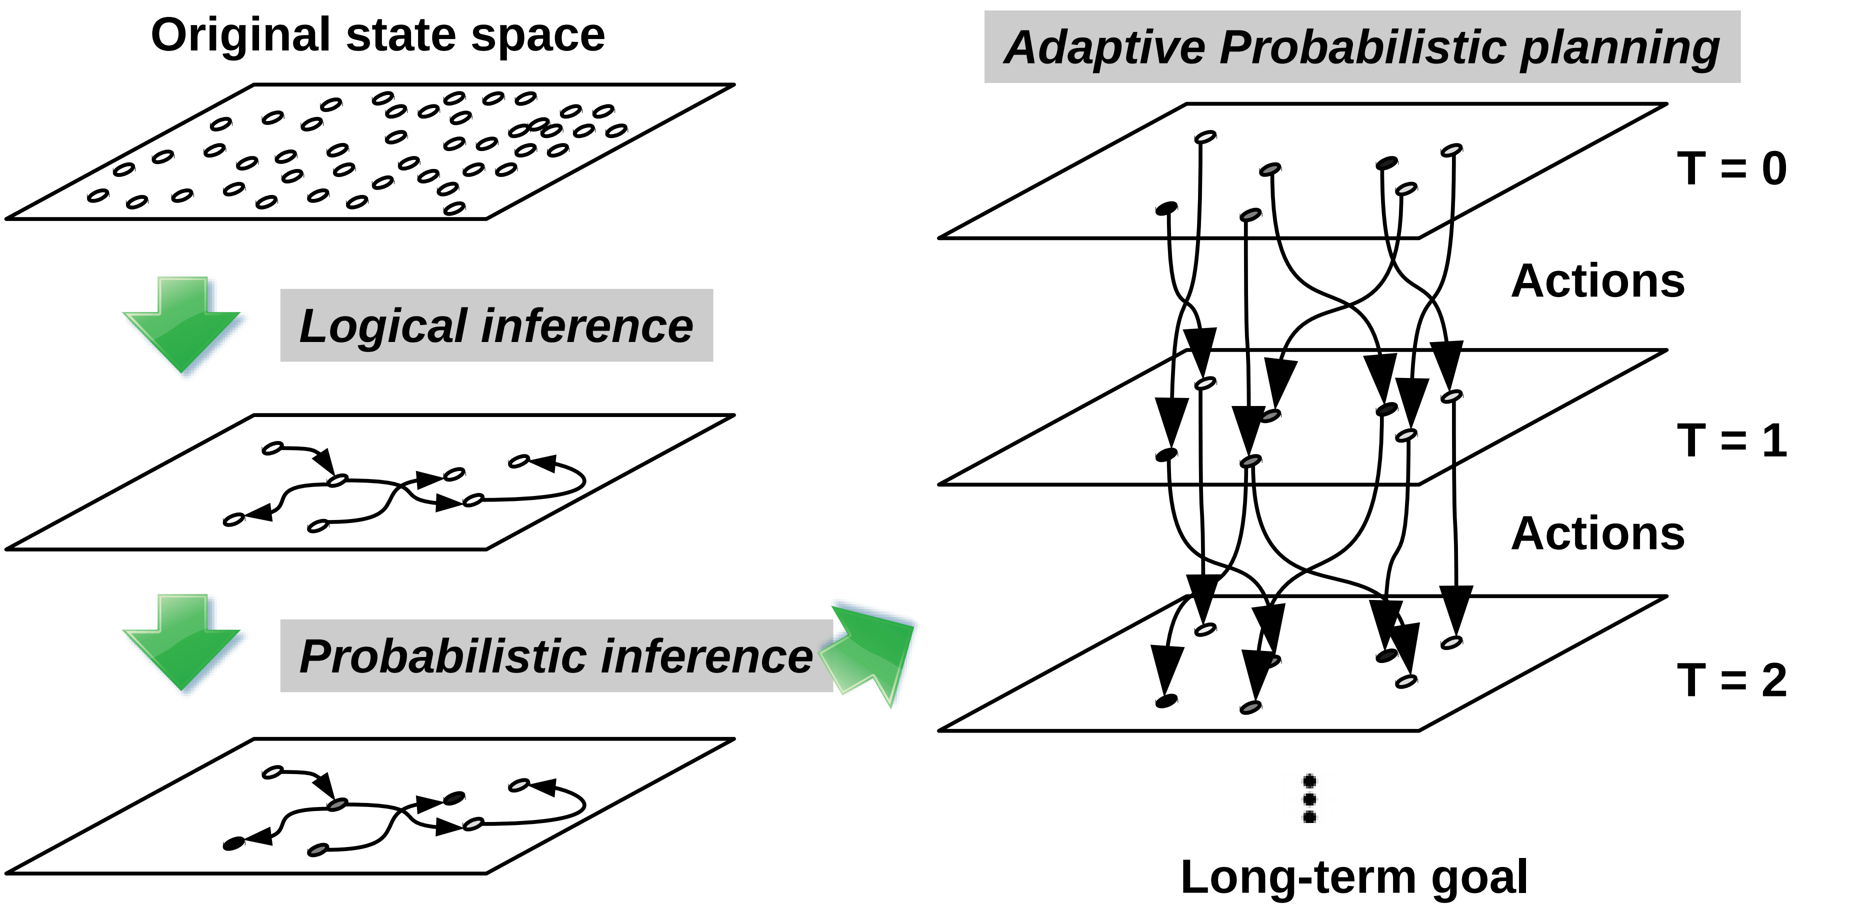

This work enables robot behaviors to **adapt** to exogenous domain changes **without including** these exogenous attributes in probabilistic planning models

Integrated *learning*, *reasoning*, and *planning* for robot sequential decision-making

Human Intention Estimation problem

Robot needs to identify human intention (e.g., interested to interact or not) as accurate and early as possible

LSTM-CORPP

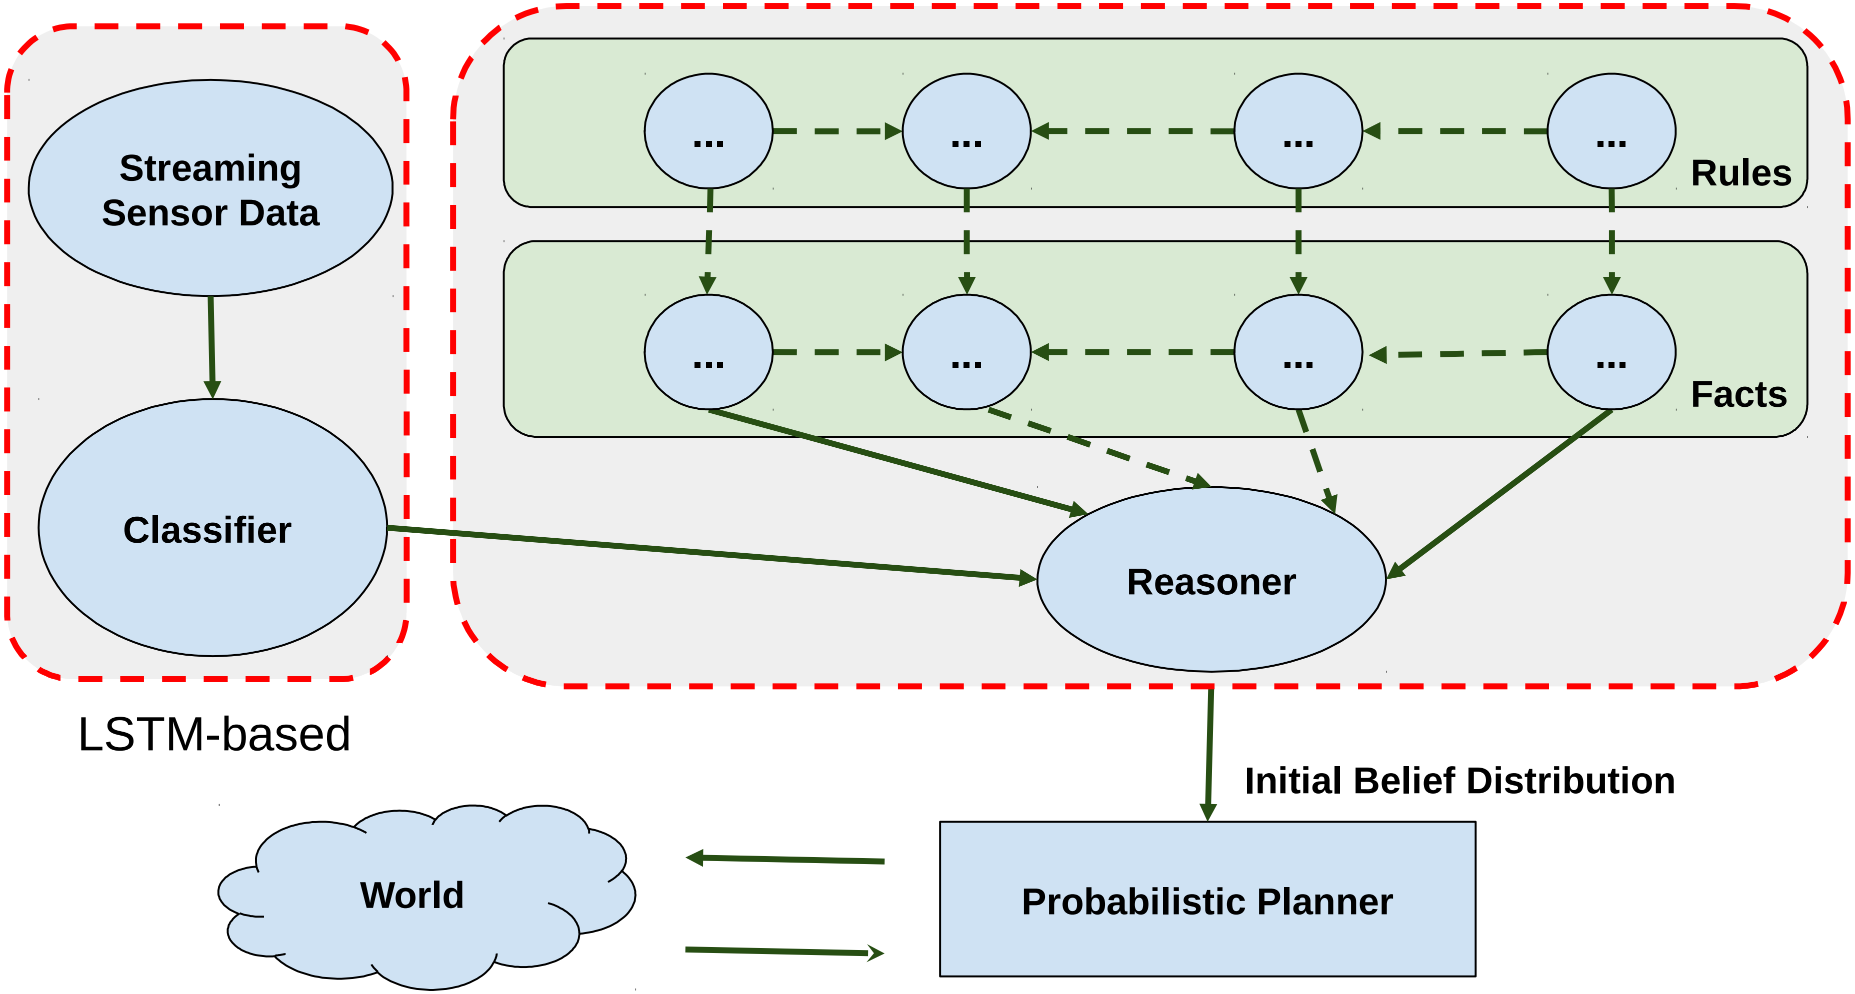

[Amiri, Shirazi, Zhang, R2K Workshop with KR, 2018]

LSTM-CORPP: preliminary results

	Accuracy	Precision	Recall	F1 Score	Cost
Learning	0.61	0.56	0.30	0.39	N/A
Reasoning	0.60	0.54	0.62	0.58	N/A
Learning + Reasoning	0.58	0.51	0.72	0.60	N/A
Reasoning + Planning (CORPP)	0.79	0.67	0.94	0.78	21.6
LSTM-CORPP (Ours)	**0.83**	**0.74**	0.86	**0.80**	**13.1**

AAAI'19 Tutorial

Knowledge-based Sequential Decision-Making under Uncertainty

(1/4 day tutorial)

Knowledge representation and reasoning (KRR)

Sequential decision-making (SDM)

Papers

- Shiqi Zhang and Peter Stone, **CORPP: Commonsense Reasoning and Probabilistic Planning, as Applied to Dialog with a Mobile Robot**, AAAI 2015
- Shiqi Zhang, Mohan Sridharan and Jeremy Wyatt, **Mixed Logical Inference and Probabilistic Planning for Robots in Unreliable Worlds**, IEEE Transactions on Robotics (TRO), 31 (3): 699-713, 2015
- Shiqi Zhang, Piyush Khandelwal and Peter Stone, **Dynamically Constructed (PO)MDPs for Adaptive Robot Planning**, AAAI 2017
- Saeid Amiri, Mohammad Shirazi, and Shiqi Zhang, **Leveraging Supervised Learning and Automated Reasoning for Robot Sequential Decision-Making**, KR'18 R2K Workshop, 2018
- Keting Lu, Shiqi Zhang, Peter Stone, and Xiaoping Chen, **Robot Representation and Reasoning with Knowledge from Reinforcement Learning**, arXiv preprint: 1809.11074, 2018

How to integrate?

Declarative knowledge representation & reasoning

Correct and natural
Incomplete knowledge
Explanation (good for HRI)
Goal-independent
Transferability

Non-deterministic action outcomes
Imperfect perception
Unspecified, long horizon
Learning from experience (RL)

Robotics decision-making

Probabilistic Planning & Reinforcement learning (RL)

Credits:
Mohan Sridharan, Peter Stone, Michael Gelfond, Jeremy Wyatt
Saeid Amiri, Piyush Khandelwal

Thank you!