Learning Representations and Algorithms for Human-Robot Interaction

Nicholas Roy

July 5, 2016

AERO ASTRO

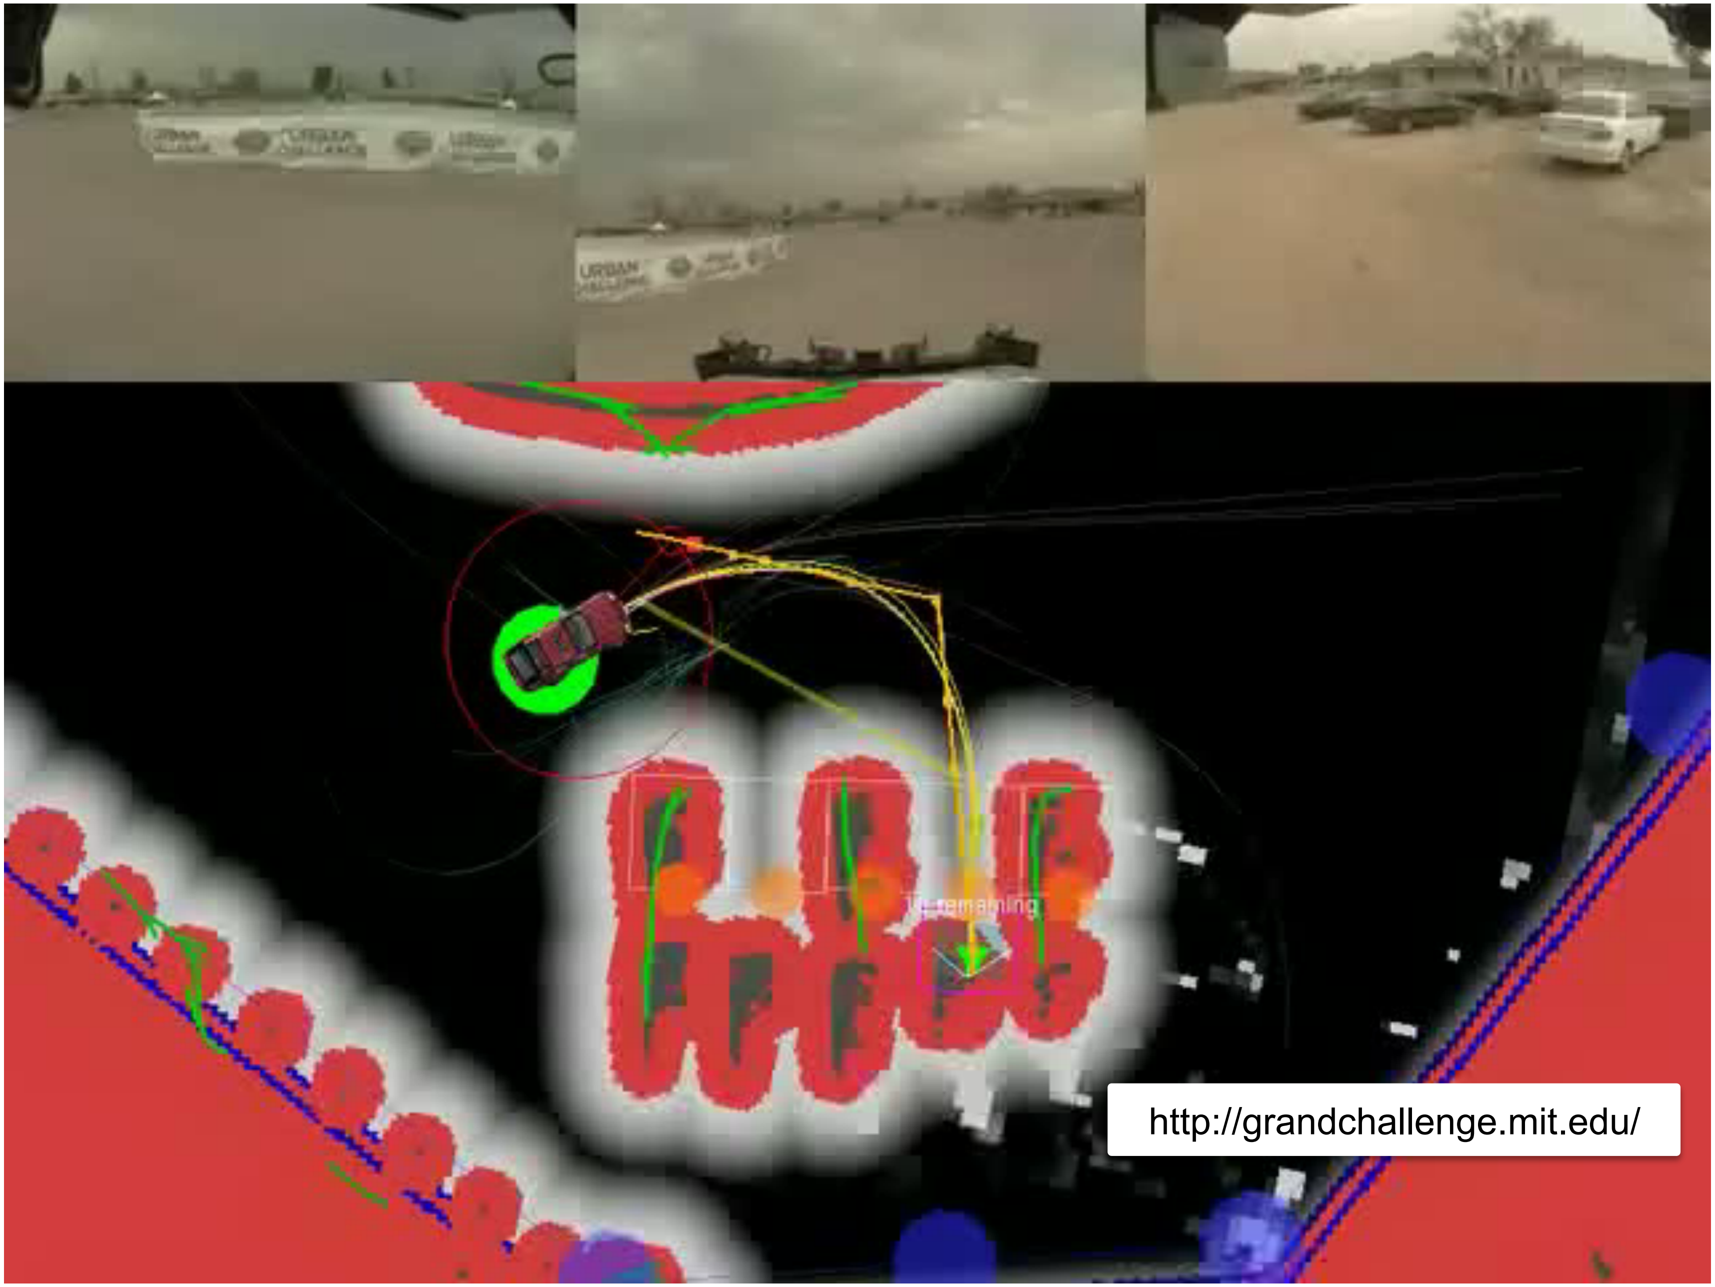
http://grandchallenge.mit.edu/

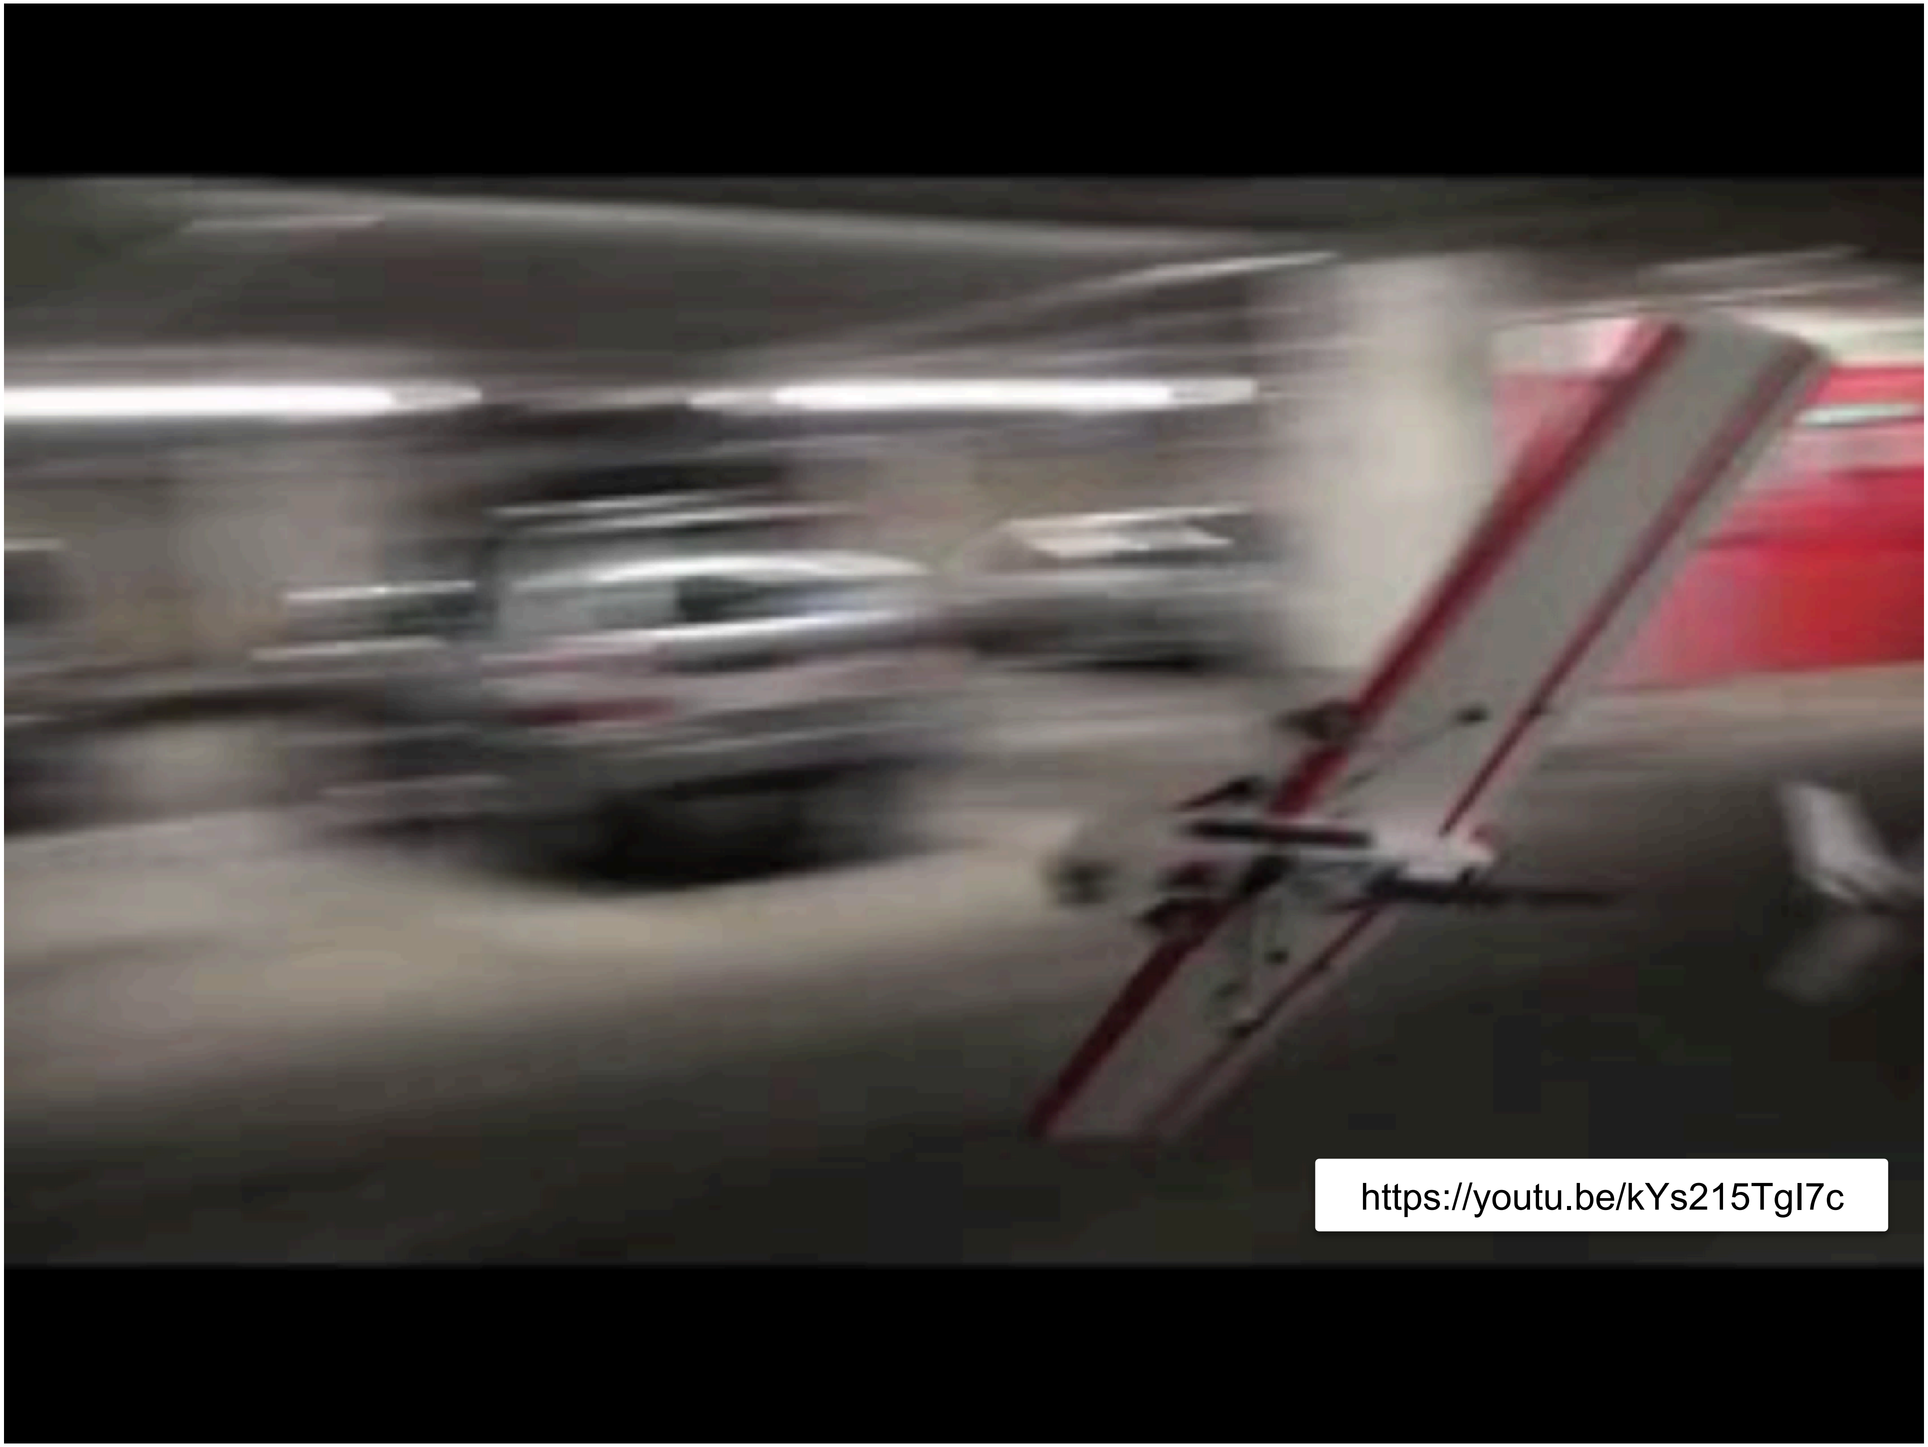

https://youtu.be/kYs215TgI7c

From the RCTA 2016 Program Review Meeting
Intelligence Introduction
Nicholas Roy (MIT) & Stuart Young (Gov. lead)

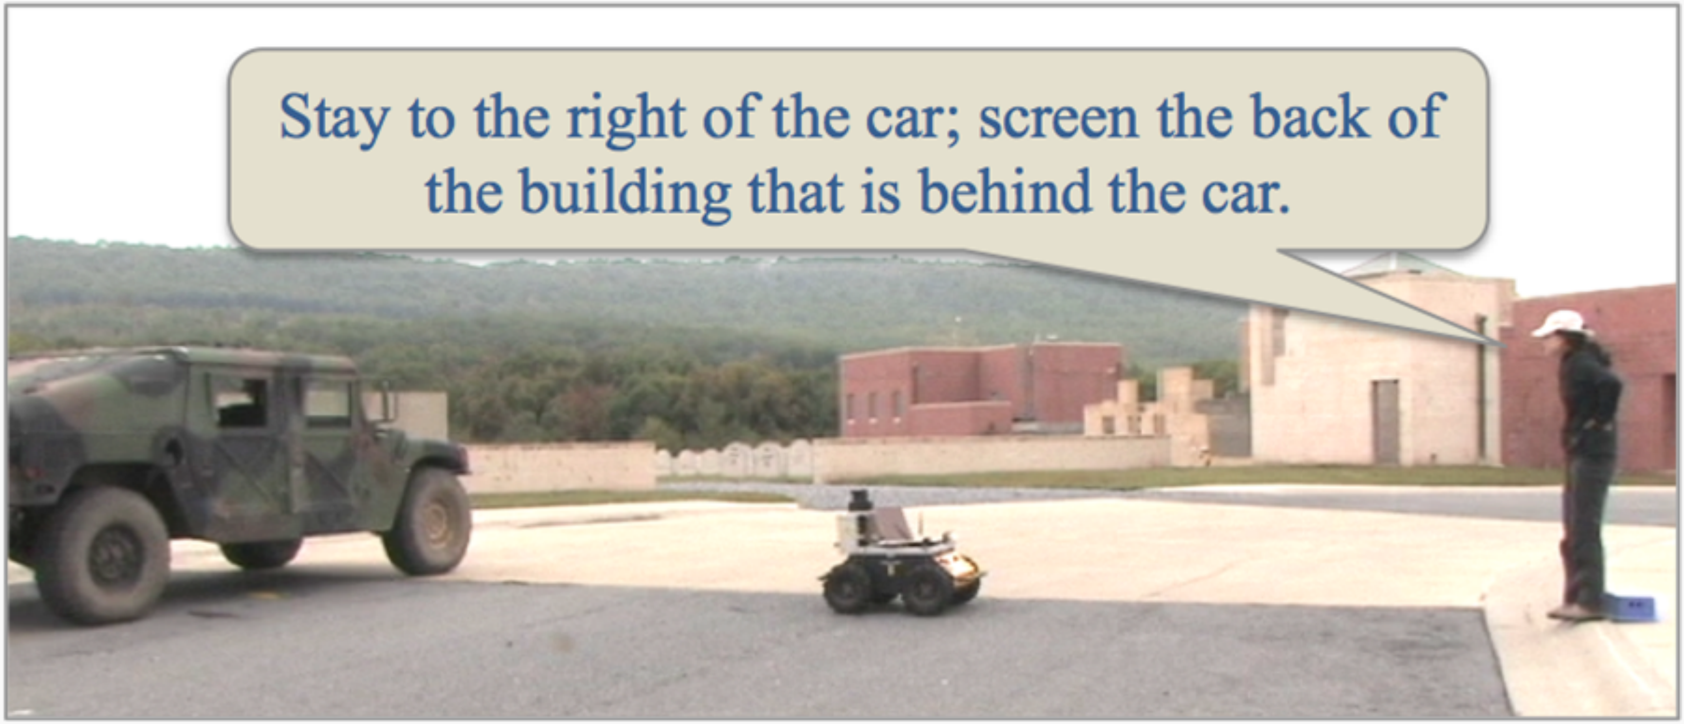

Intelligence and Teaming

Plan complex, temporally extended missions

Represent high-level properties of the world, e.g., the function of egress points, etc.

Reason about how the semantics of the world affect performance

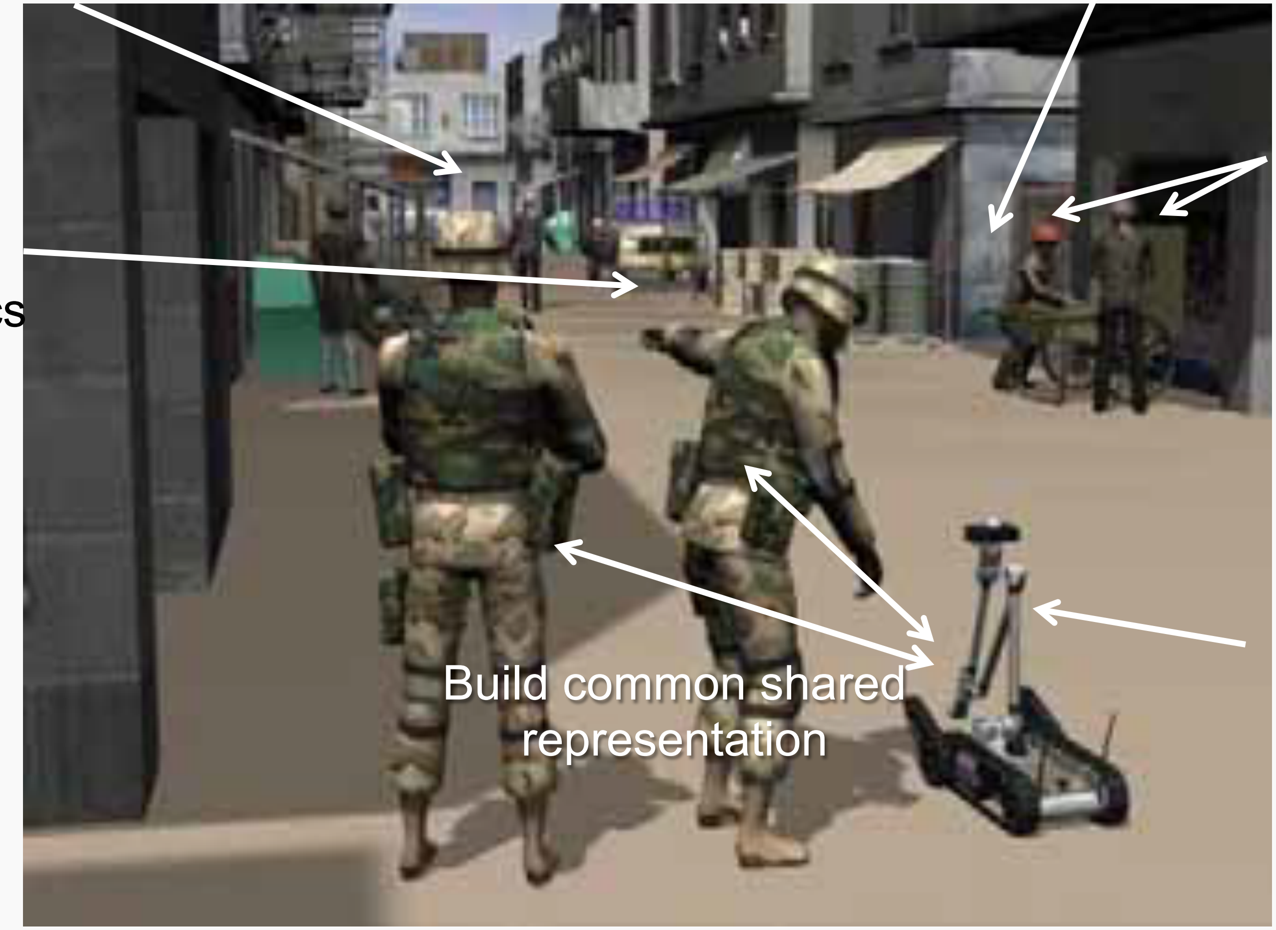

Create plans that are robust to incomplete or imperfect information

Learn complex dynamical models

Robotics Collaborative Technology Alliance

Representations for Robotics

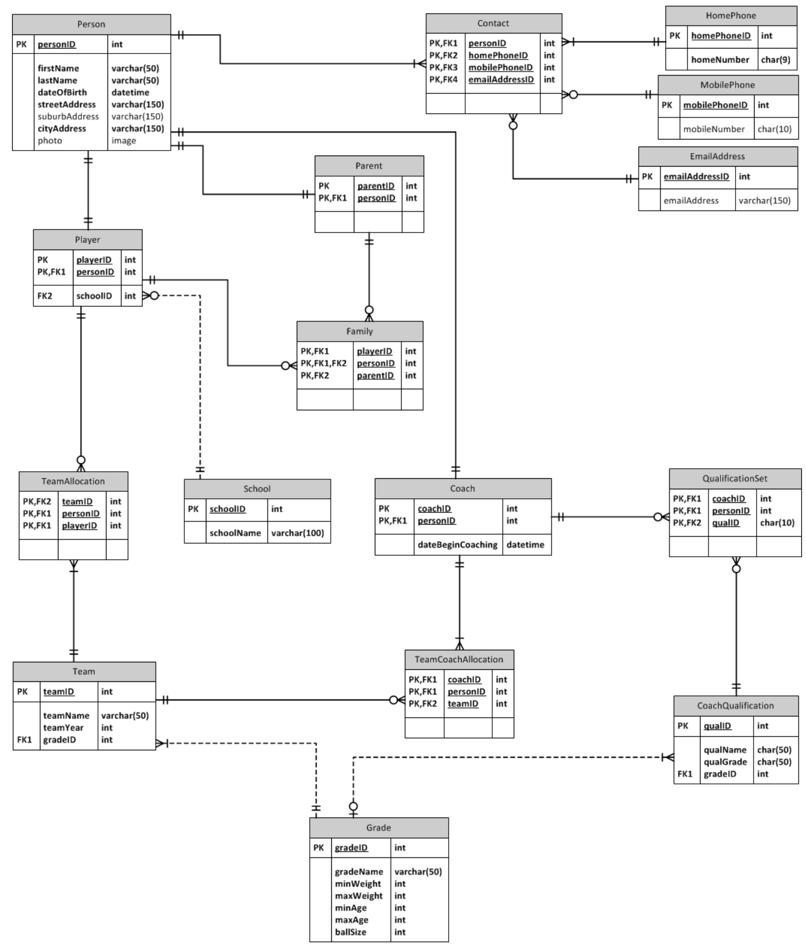

IN DEFENCE OF LOGIC

P.J. Hayes
Essex University
Colchester, U.K.

Introduction

Modern formal logic is the most successful precise language ever developed to express human thought and inference. Measured across any reasonably broad spectrum, including philosophy, linguistics, computer science, mathematics and artificial intelligence, no other formalism has been anything like so successful. And yet recent writers in the AI field have been almost unanimous in their condemnation of logic as a representational language, and other formalisms are in a state of rapid development.

I will argue that most of this criticism

performs inferences: some of its processes are the making of inferences.

But two different systems may be based on the same notion of inference and the same representational language. The *inference structure* of the language used by a system does not depend on the *process structure*. In particular, a system may have a logical inference structure - may be making deductively valid inferences - without being a classical uniform theorem-prover which just "grinds lists of clauses together".

Representations for Robotics

Portrait used of Bayes in a 1936 book,[1] but it is doubtful whether the portrait is actually of him.[2] No earlier portrait or claimed portrait survives.

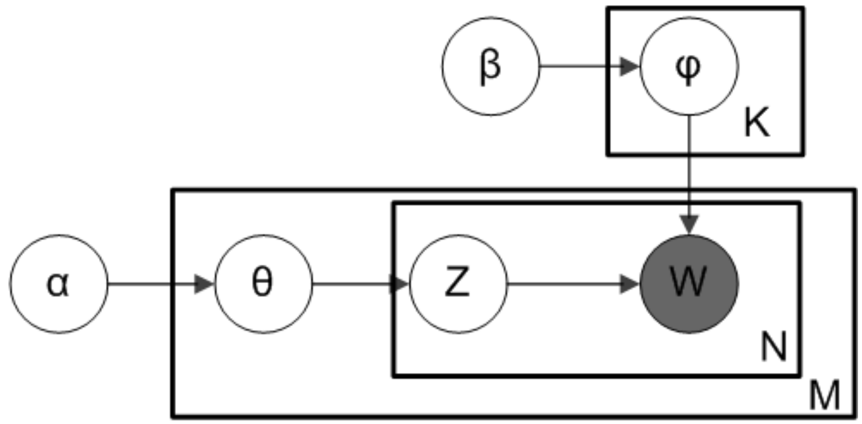

In Defense of Probability

Peter Cheeseman
SRI International
333 Ravenswood Ave., Menlo Park, California 94025

Abstract

In this paper, it is argued that probability theory, when used correctly, is suffrcient for the task of reasoning under uncertainty. Since numerous authors have rejected probability as inadequate for various reasons, the bulk of the paper is aimed at refuting these claims and indicating the scources of error. In particular, the definition of probability as a measure of belief rather than a frequency ratio is advocated, since a frequency interpretation of probability dras-

ference is that in probabilistic inference all the relevant inference paths ("proofs") connecting the evidence to the hypothesis of interest must be examined and "combined", while in logic it is sufficient to establish a single path between the axioms and the theorem of interest. Also, the output is different, the former includes at least one numerical measure, the latter simply true or false.

Unfortunately, the logical style of reasoning is so prevalent in AI that many have attempted to force intrinsically probabilistic situations into a logical straight-jacket with

Graphical Models

$$p(x_t, m \mid z_{1:t}, u_{1:t})$$

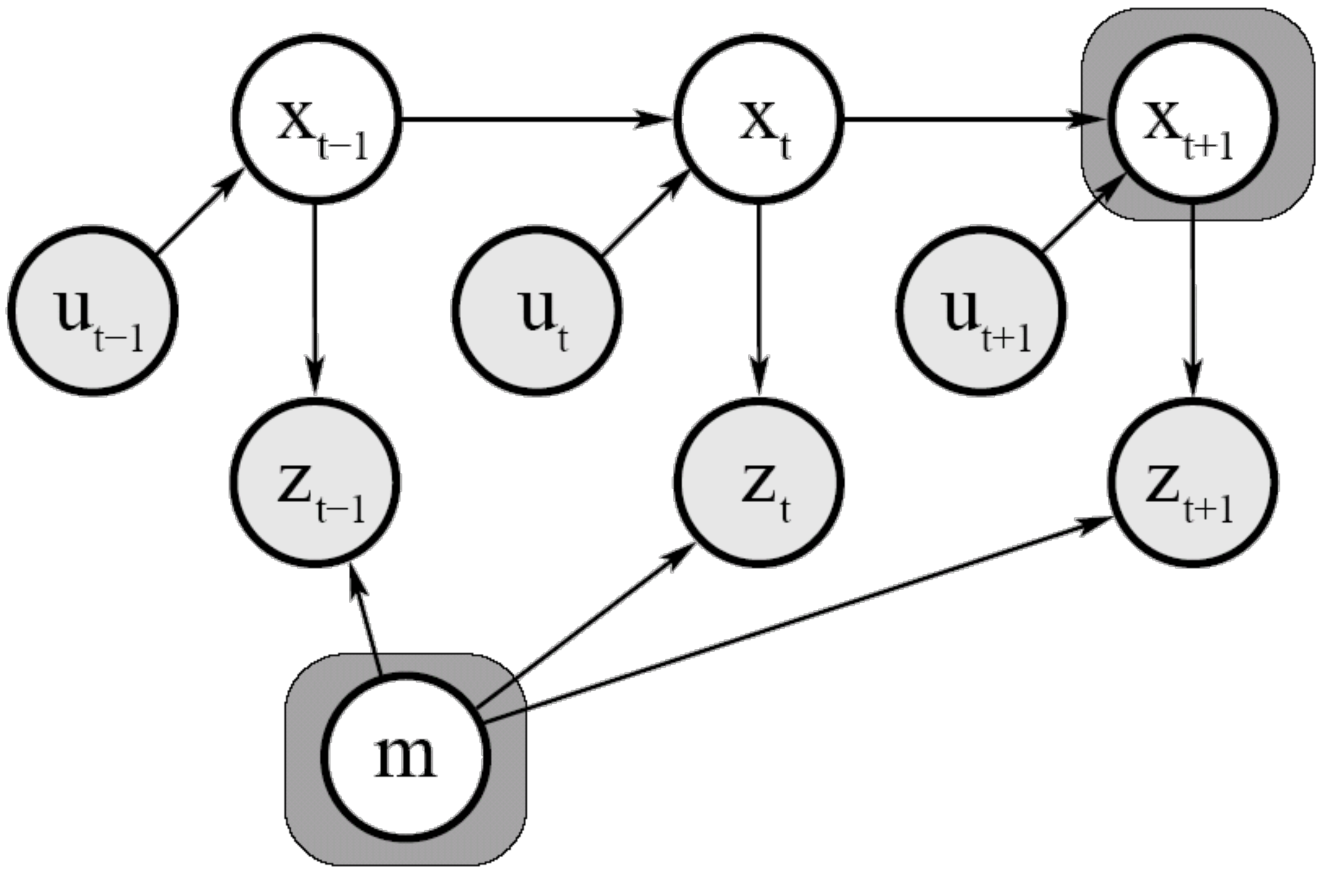

Mathematical Basis of Mapping and Navigation

$$p(x_t, m \mid z_{1:t}, u_{1:t}) = \int\int \ldots \int p(x_{1:t}, m \mid z_{1:t}, u_{1:t})\, dx_1 dx_2 \ldots dx_{t-1}$$

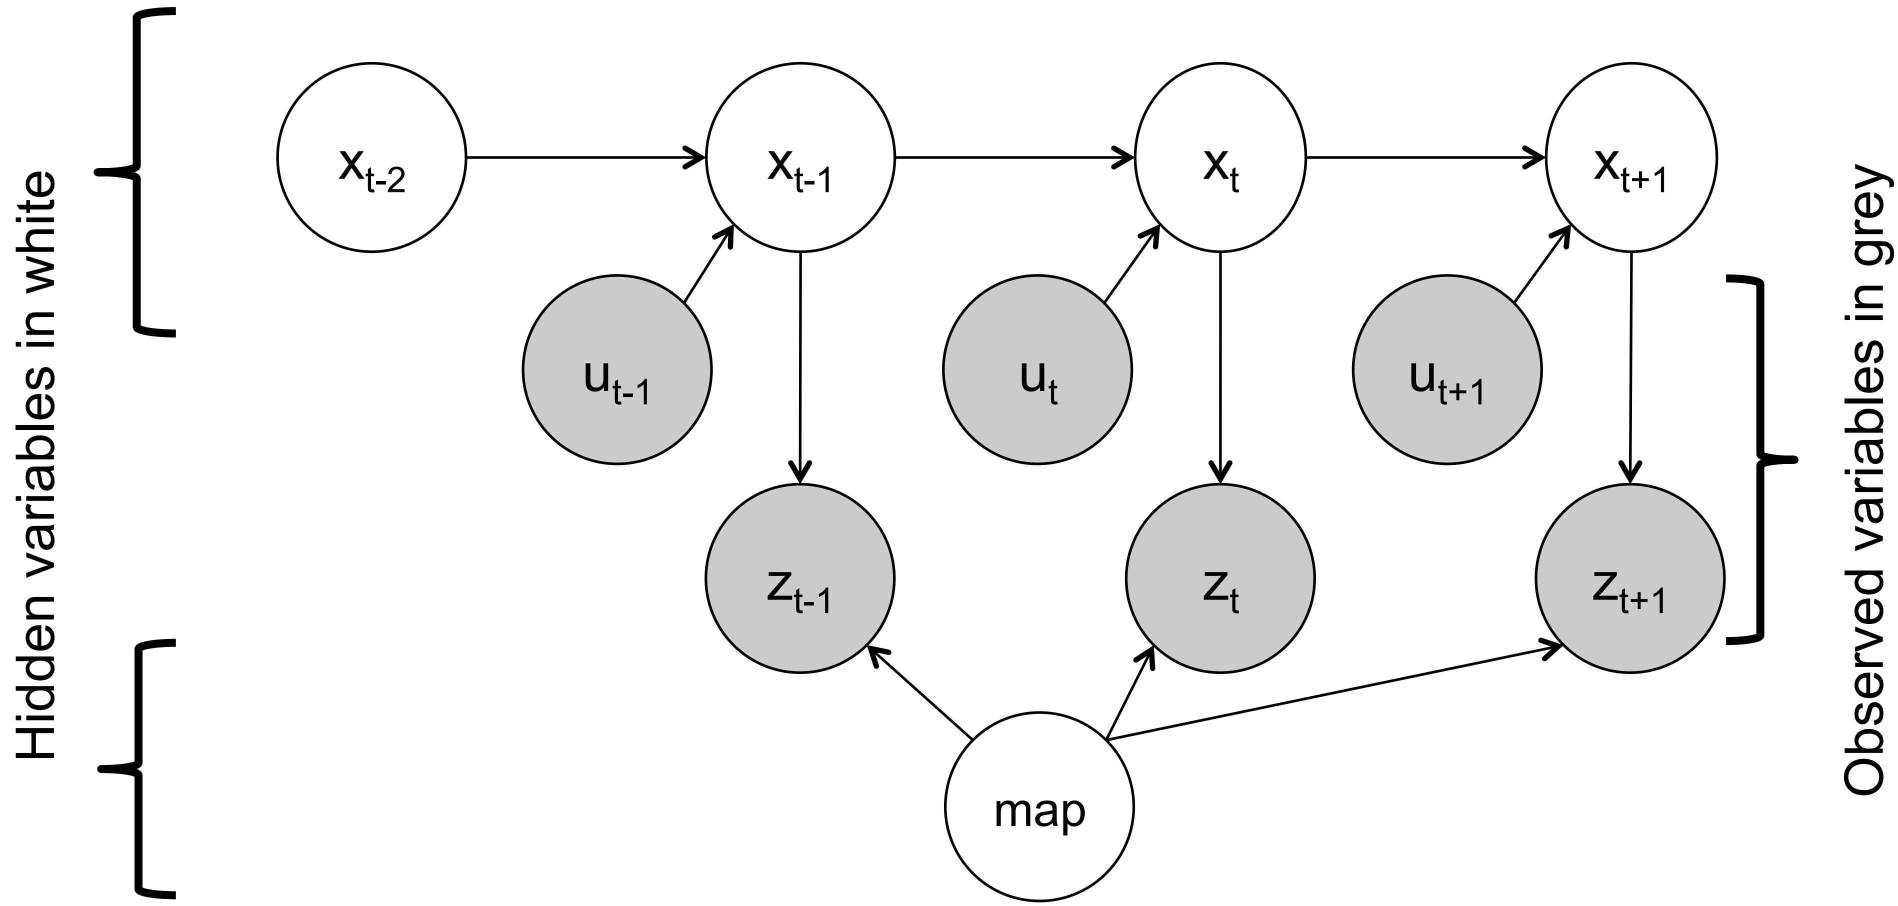

Filtering is Weighting by the Present and Marginalizing out the Past

$$p(x_t, m \mid z_{1:t}, u_{1:t}) = \int\int\ldots\int p(x_{1:t}, m \mid z_{1:t}, u_{1:t})\, dx_1 dx_2 \ldots dx_{t-1}$$

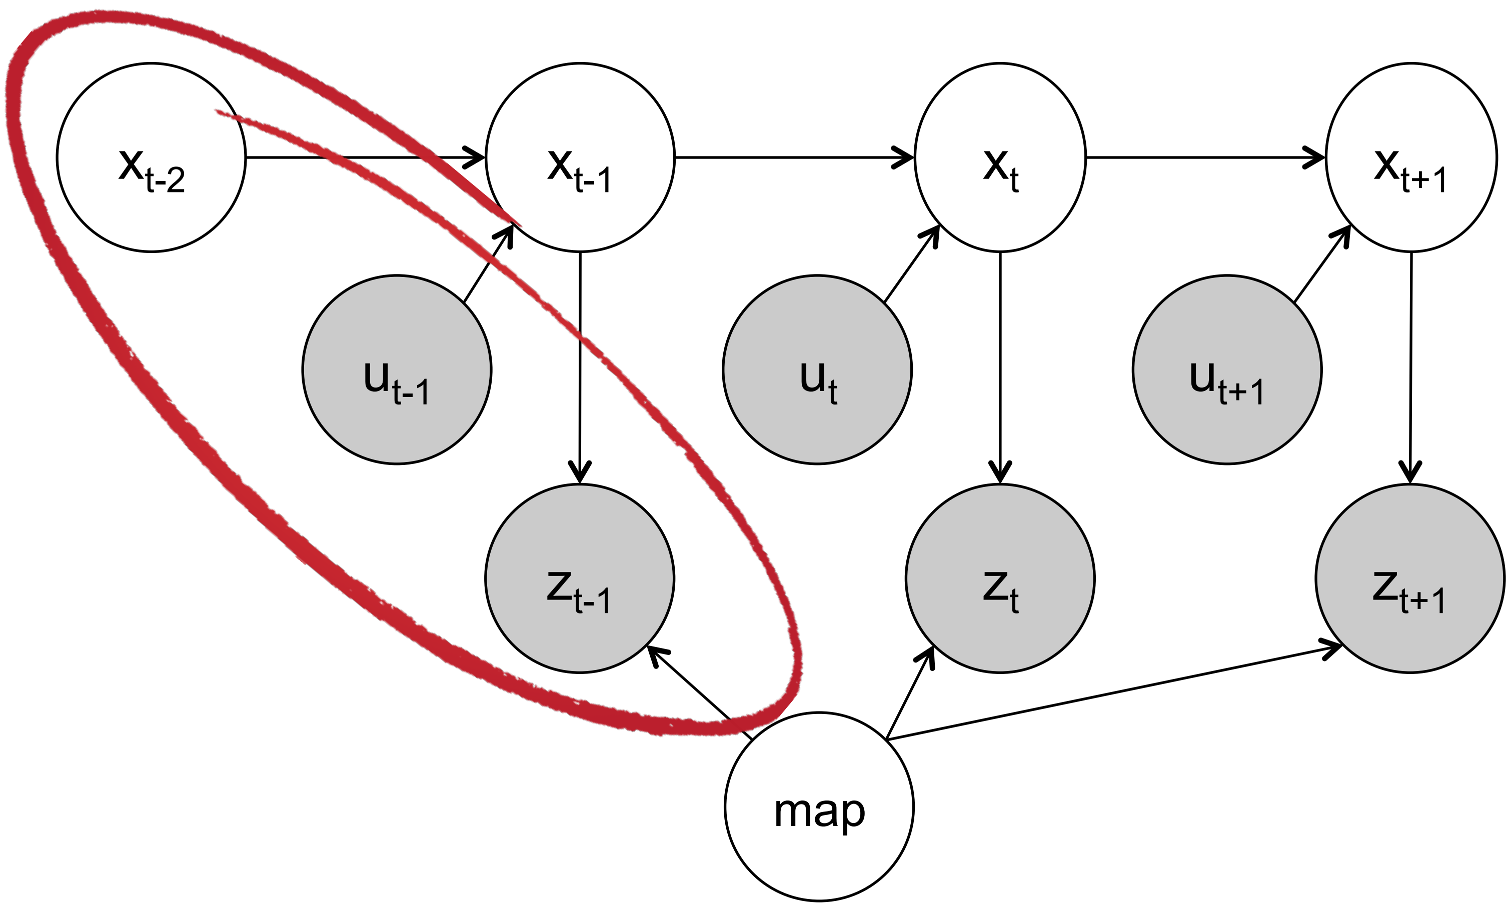

Change in Representation: Marginalize out Measurements Instead

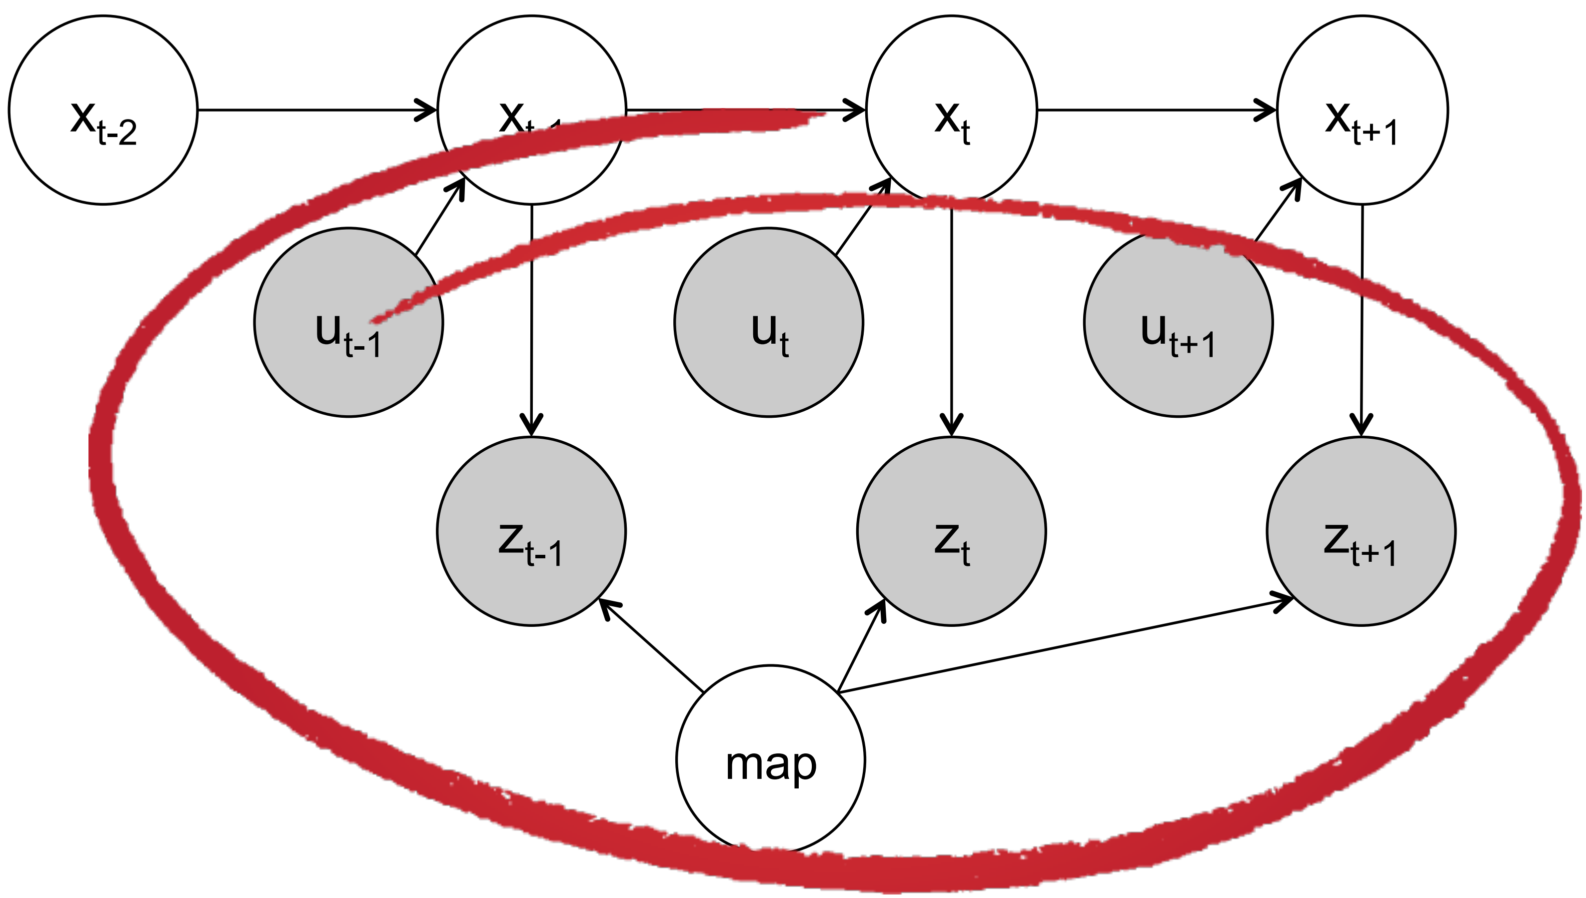

Change in Representation: Marginalize out Measurements

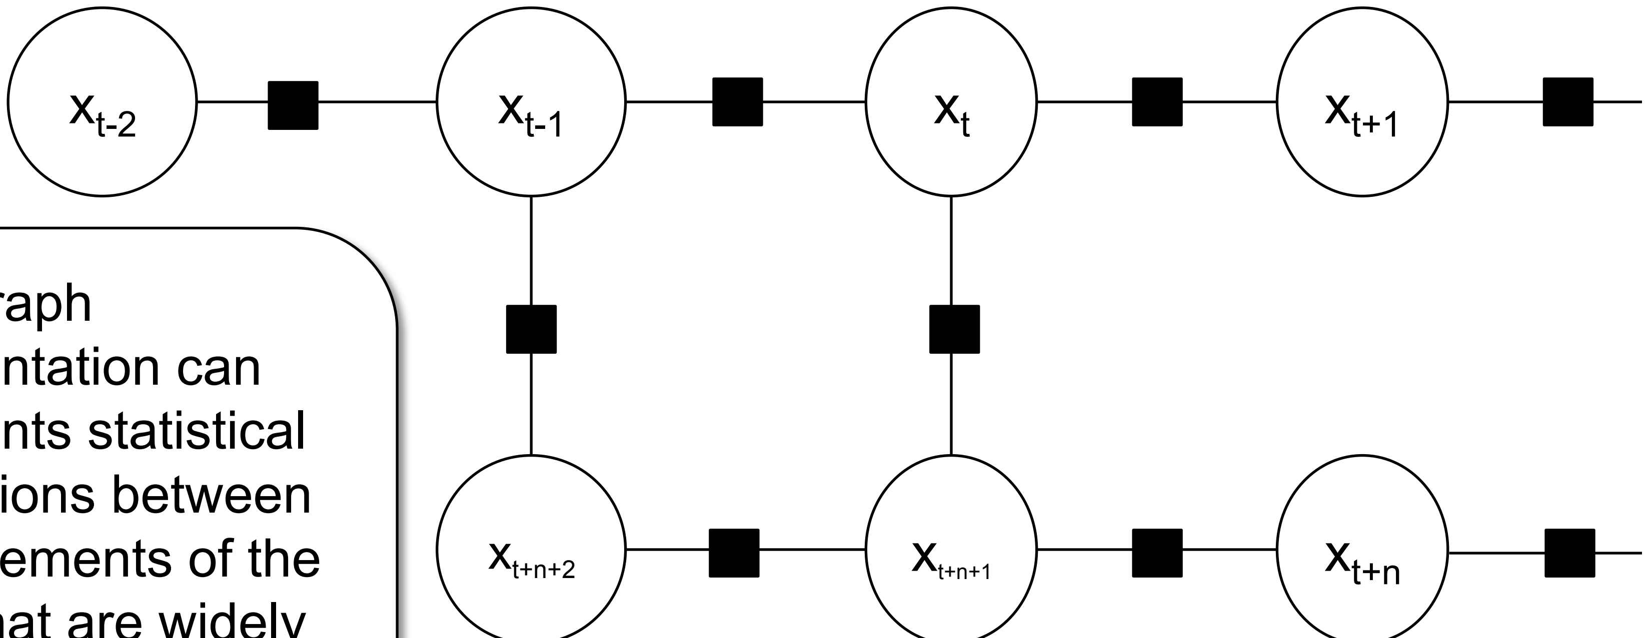

Pose graph representation can represents statistical correlations between measurements of the world that are widely separated in time but not distance.

Motion Planning

Build a graph or grid in configuration space that captures the collision-free space and search for the shortest path.

As the dimensionality of the c-space grows, building and maintaining this representation becomes painful.

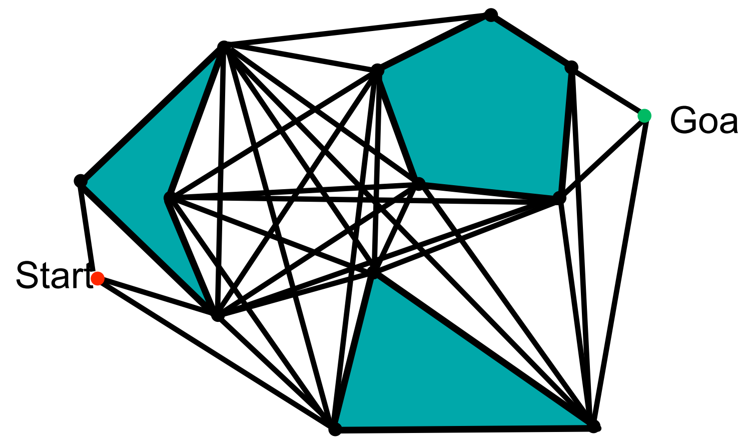

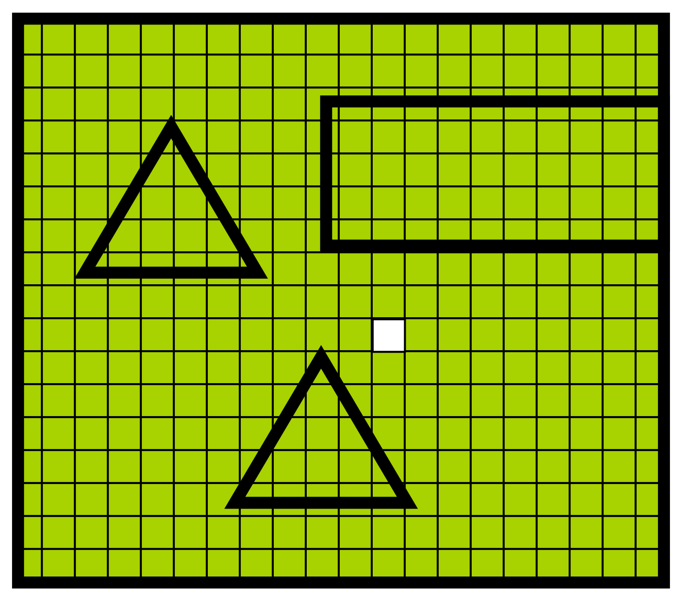

Randomized Motion Planning

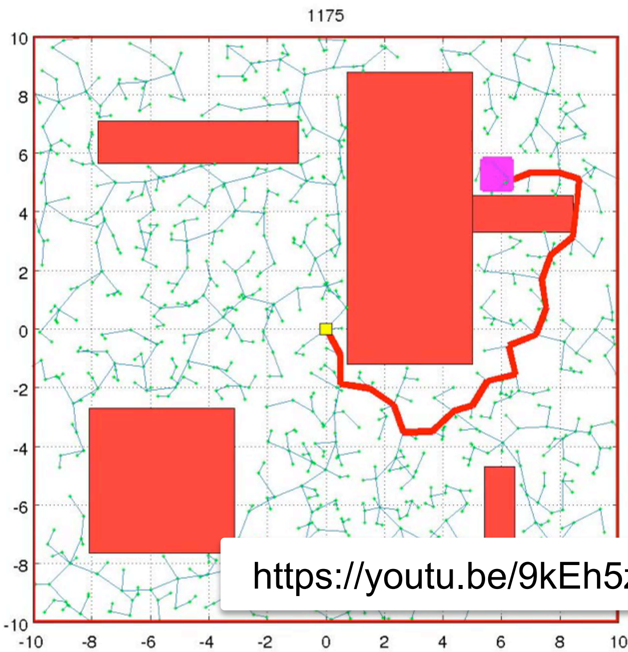

https://youtu.be/9kEh5zm1F_A

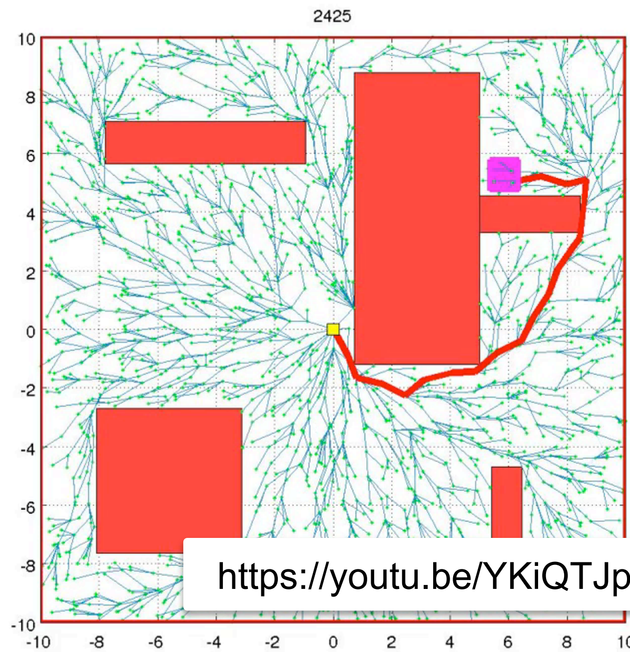

https://youtu.be/YKiQTJpPFkA

Major change in motion planning: represent the world as a randomly generated graph in the free space.

Videos courtesy of S. Karaman, MIT

What about ~~deep learning~~?

Hyperparametric Function Approximation!*

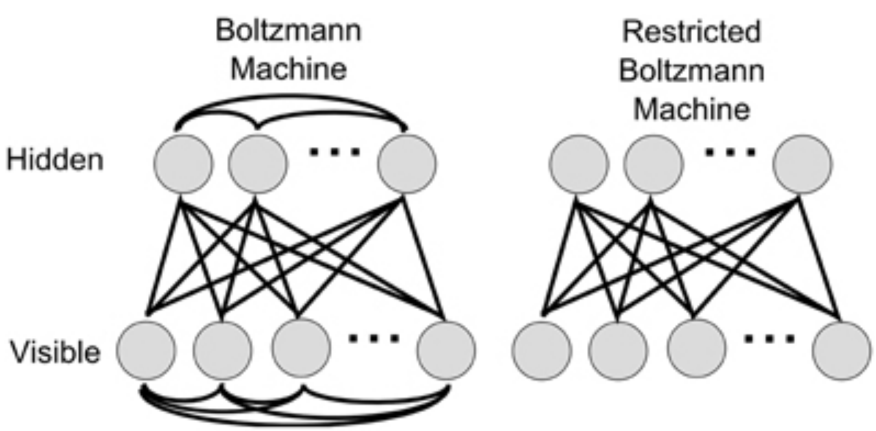

Convolutional networks and applications in vision
Y LeCun, K Kavukcuoglu… - Circuits and Systems (…, 2010 - ieeexplore.ieee.org
... Applications to visual object recognition and vision navigation for off-road mobile **robots** are described. **...** While the issue of **learning** features has been a topic of interest for many years **...** been achieved in the last few years with the development of so-called **deep learning** methods. **...**
Cited by 223 Related articles All 21 versions Cite Save

Playing atari with **deep** reinforcement **learning**
V Mnih, K Kavukcuoglu, D Silver, A Graves… - arXiv preprint arXiv: …, 2013 - arxiv.org
... [12] Sascha Lange and Martin Riedmiller. **Deep** auto-encoder neural networks in reinforcement **learning**. In Neural Networks (IJCNN), The 2010 International Joint Conference on, pages 1–8. IEEE, 2010. [13] Long-Ji Lin. Reinforcement **learning** for **robots** using neural networks. **...**
Cited by 105 Related articles All 24 versions Cite Save

Reinforcement **learning** for **robots** using neural networks
LJ Lin - 1993 - DTIC Document
Page 1. AD-A261 434 Reinforcement **Learning** for **Robots** Using Neural Networks Long-Ji Lin January 6, 1993 CMU-CS-93-103 DTIC **...** By a .. Reinforcement **Learning** for **Robots** DItbfi,'* o Using Neural Networks Availability Codes LONG-JI LIN Dist Avail and jor / pca **...**
Cited by 357 Related articles All 4 versions Cite Save More

**** Inspired by Ken Goldberg***

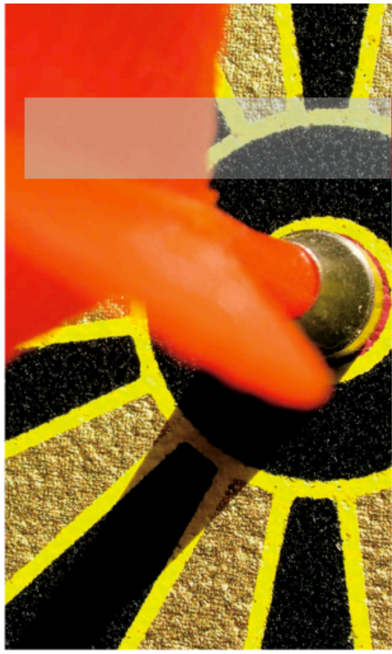

EXPERT OPINION

Contact Editor: **Brian Brannon,** bbrannon@computer.org

The Unreasonable Effectiveness of Data

Alon Halevy, Peter Norvig, and Fernando Pereira, *Google*

Eugene Wigner's article "The Unreasonable Effectiveness of Mathematics in the Natural Sciences"[1] examines why so much of physics can be

behavior. So, this corpus could serve as the basis of a complete model for certain tasks—if only we knew how to extract the model from the data.

Exact Models vs Approximate Models

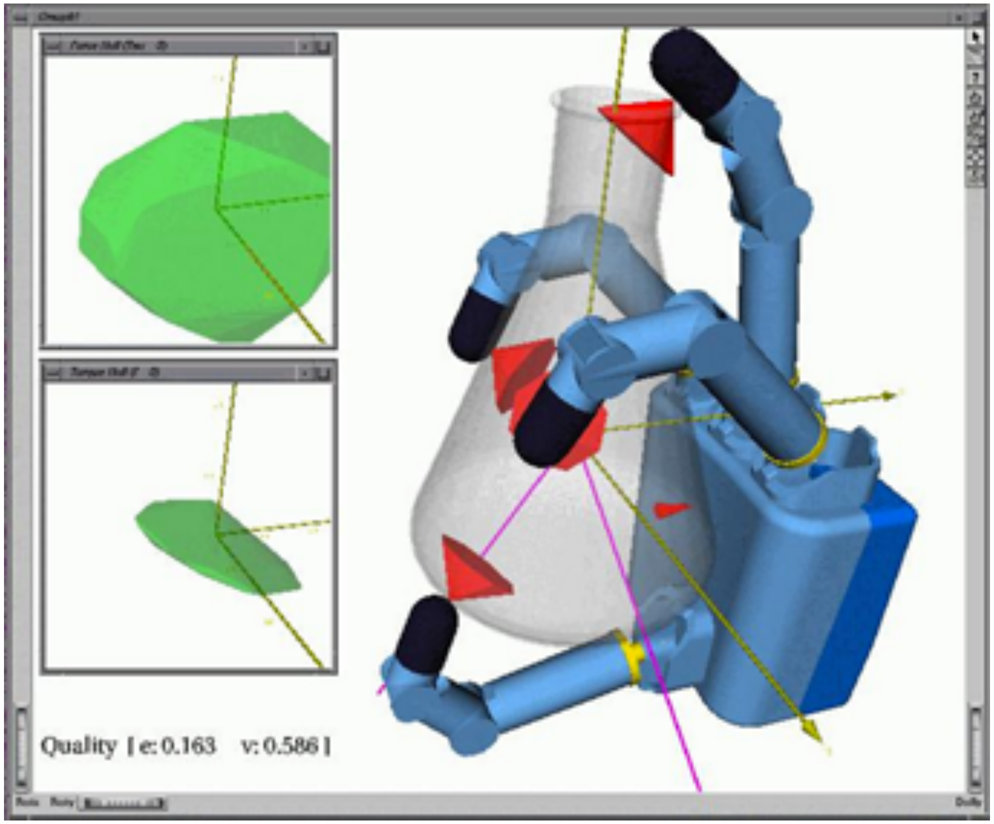

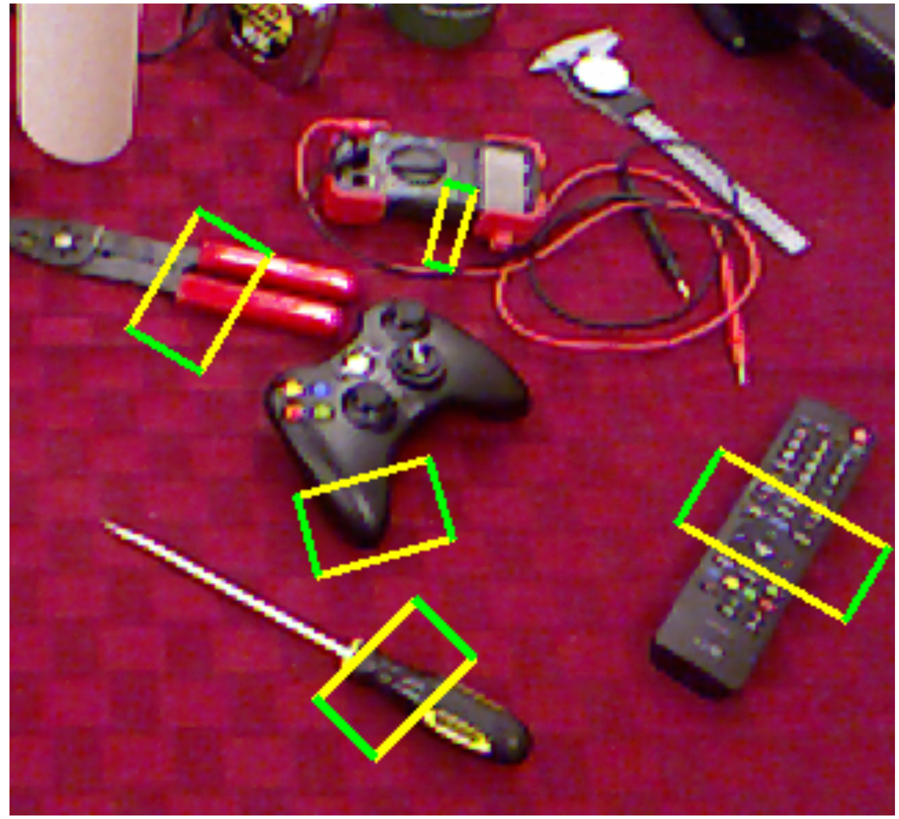

Left image, Miller and Allen, IMG 2004 Right image, Lenz et al, IJRR 2014

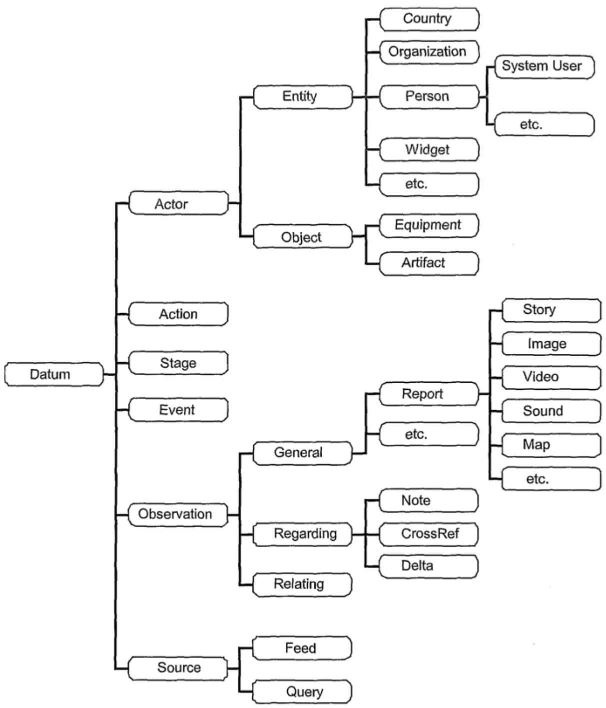
Datum
Actor
Entity
Country
Organization
Person
System User
etc.
Widget
etc.
Object
Equipment
Artifact
Action
Stage
Event
Observation
General
Report
Story
Image
Video
Sound
Map
etc.
etc.
Regarding
Note
CrossRef
Delta
Relating
Source
Feed
Query

?

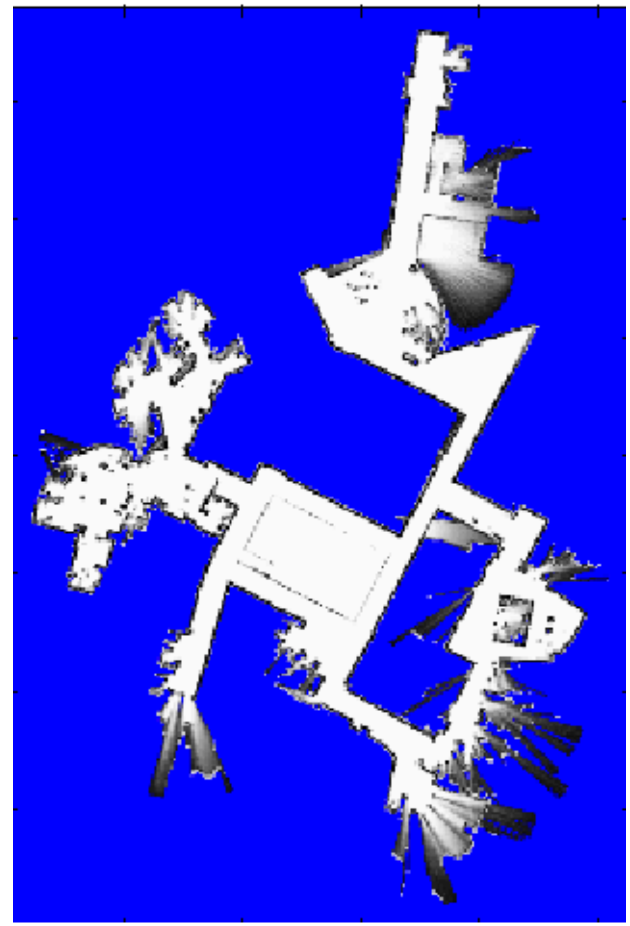

Inferring Robot Actions

- Given some language $z_{1:T}$
  - *"Go past the elevators, through the door, down the hallway, the conference room is on your left."*
- A set of actions $r_{1:M}$
- And a map m,
- Find lowest cost path through the map

$$\underset{\pi_{1:T}}{\operatorname{argmin}}\, c\left(\pi_{1:T} \mid z_{1:T}, m\right)$$

- Where $\pi_i \in \{r_{1:M}\}$

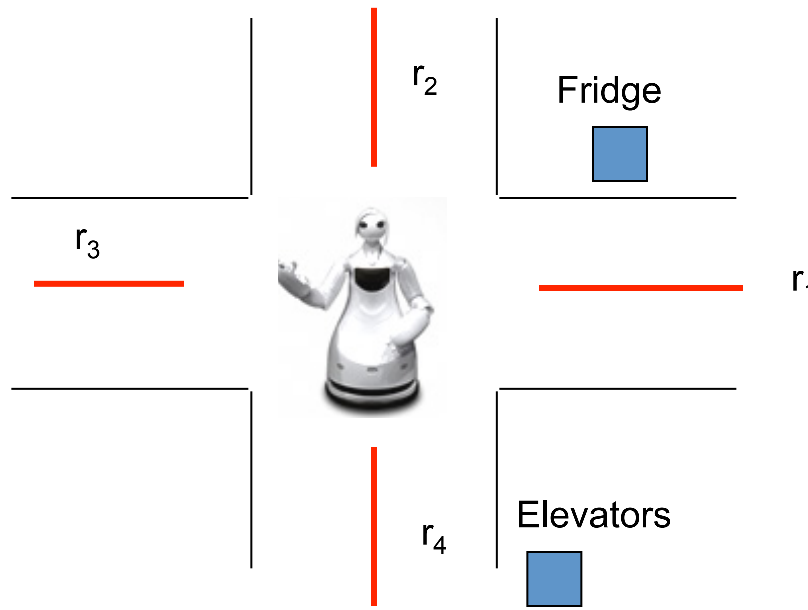

The Problem of Generalization

Not obvious how to turn this problem statement:

Given
- some language $z_{1:T}$
- a set of actions $r_{1:M}$
- and a map m…

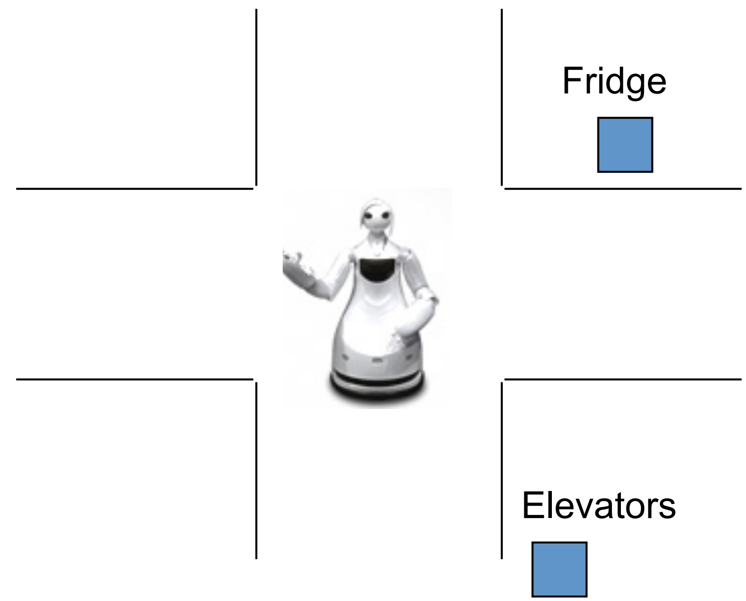

into a more general system for understanding natural language.

Grounding Graphs

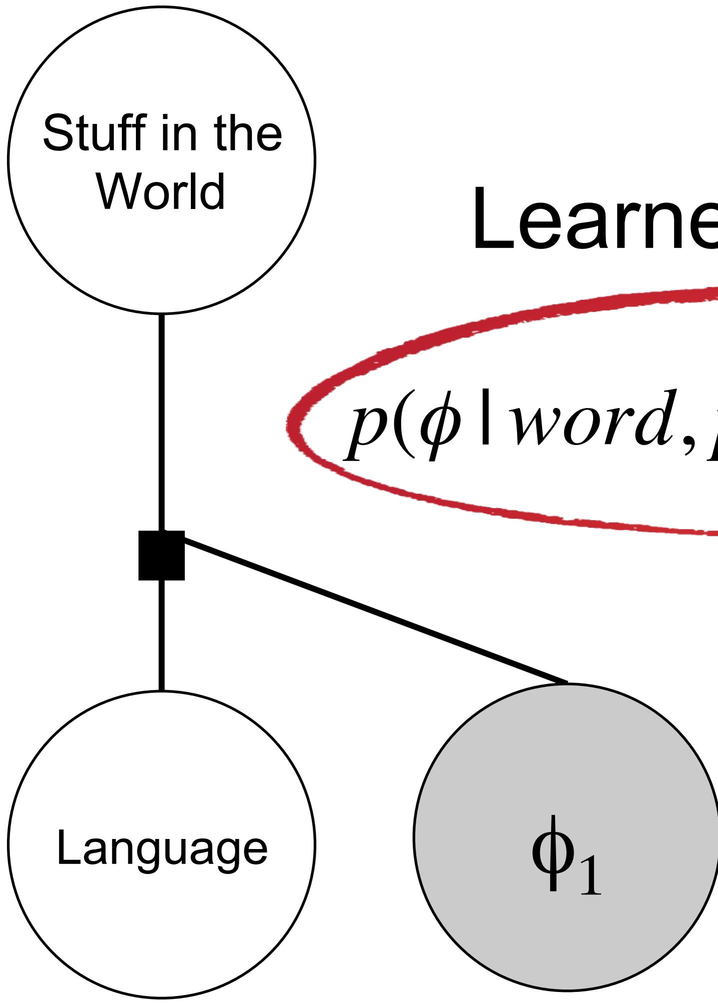

Learned from data

$$p(\phi \mid word, piece\ of\ the\ world)$$

Grounding Graphs

Go pas

“Go past the elevators.”

Groun h for Go past:

- Ac t
- Ol ark

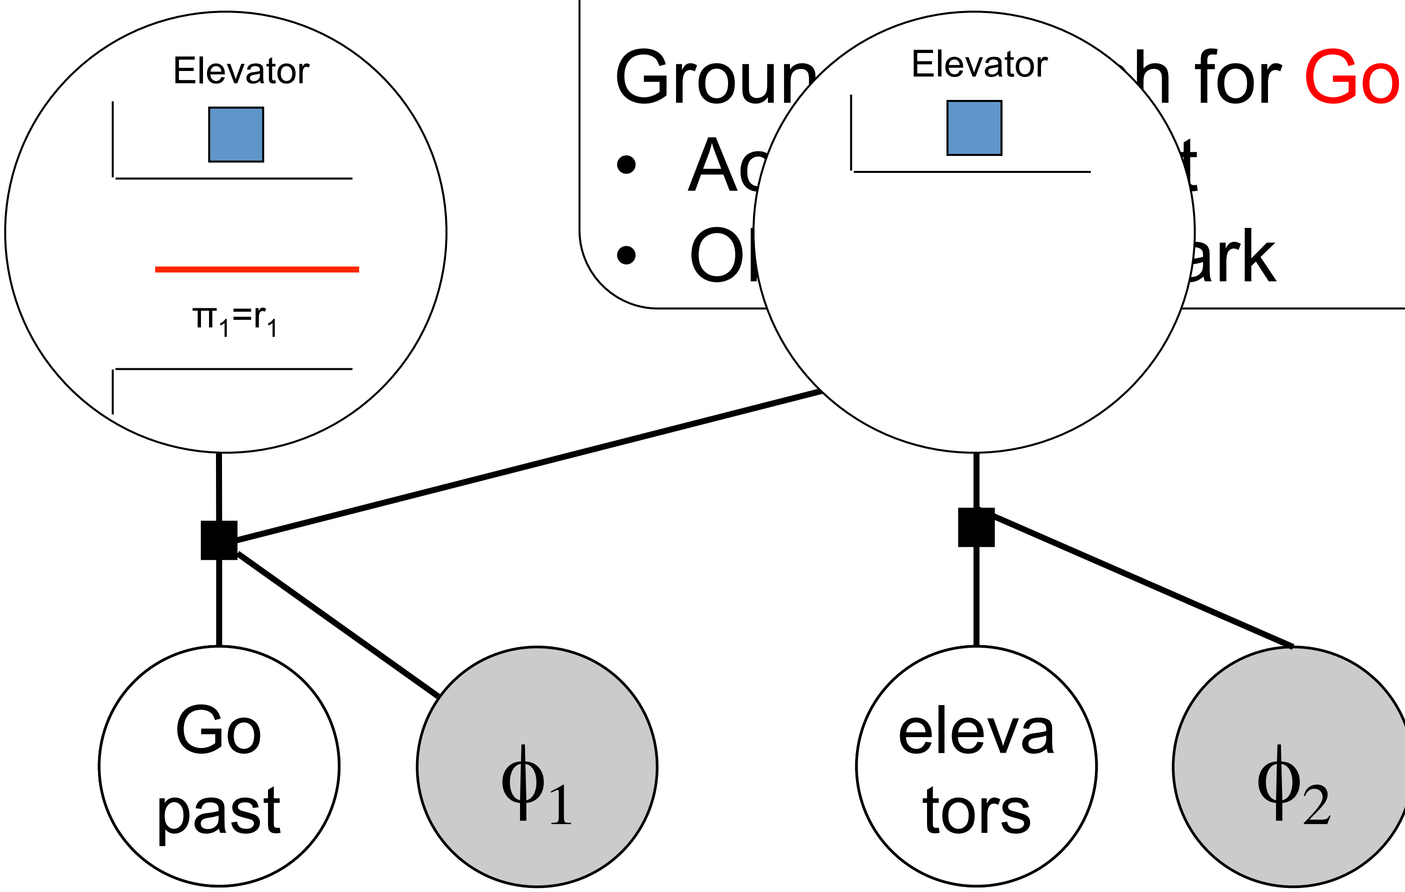

Where does the structure come from?

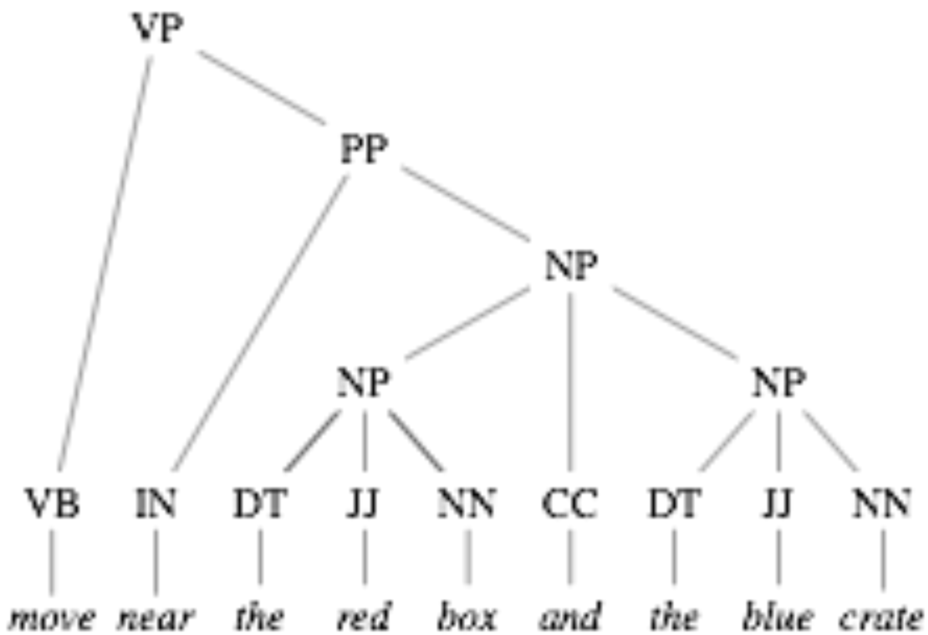

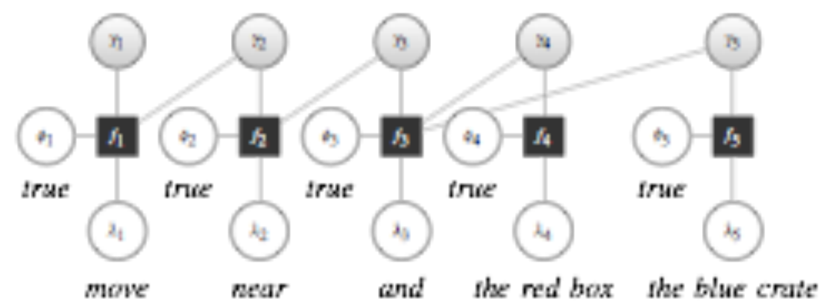

“Put the pallet on the truck.”

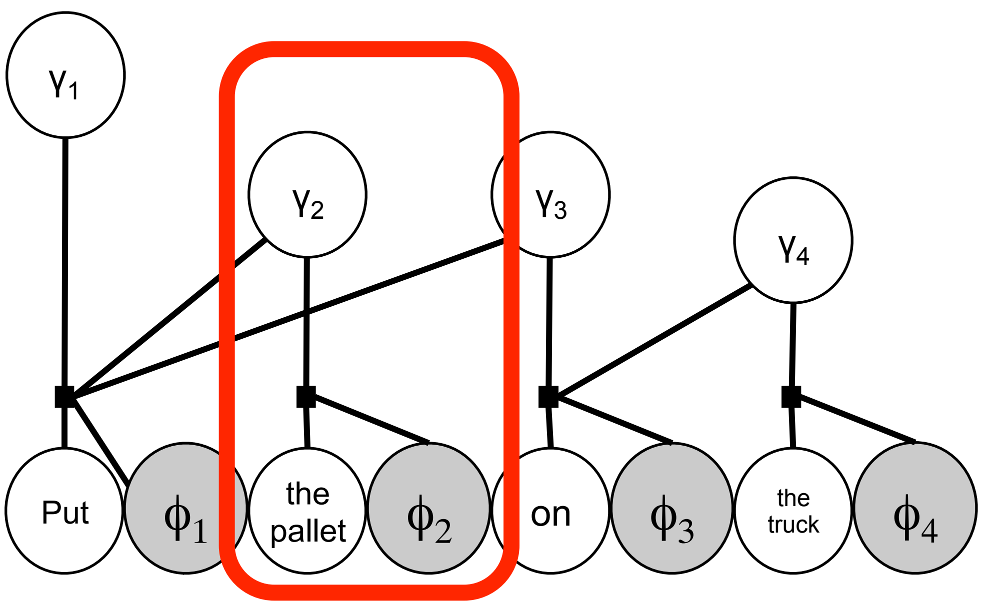

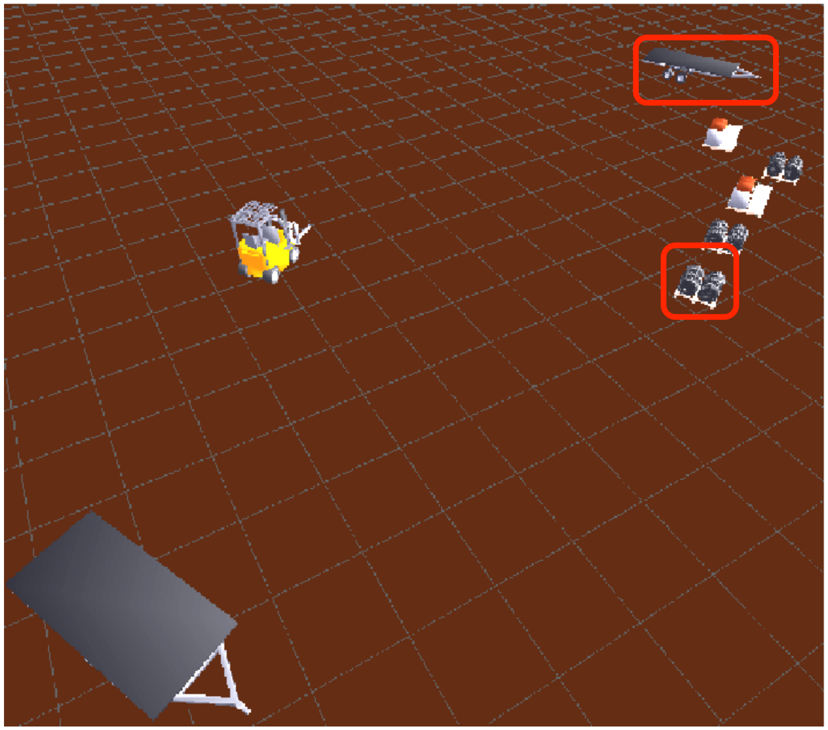

"Put the pallet on the truck."

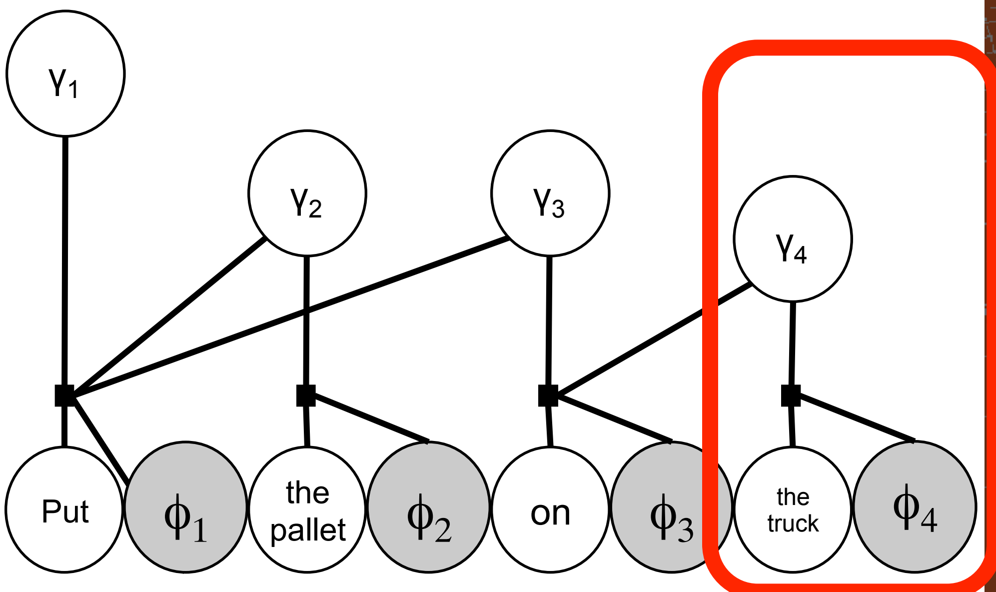

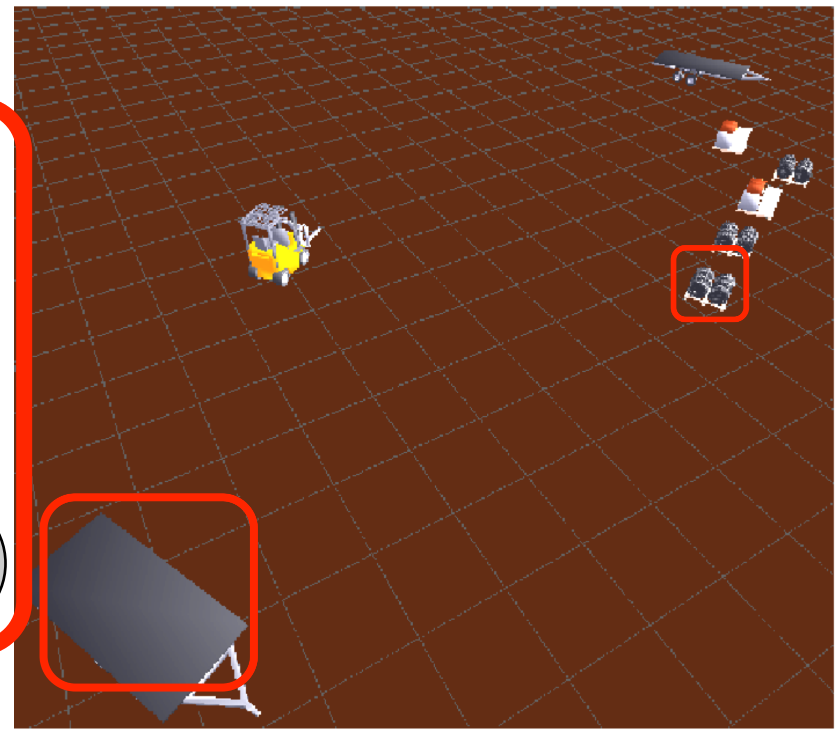

Candidate for γ_4

"Put the pallet on the truck."

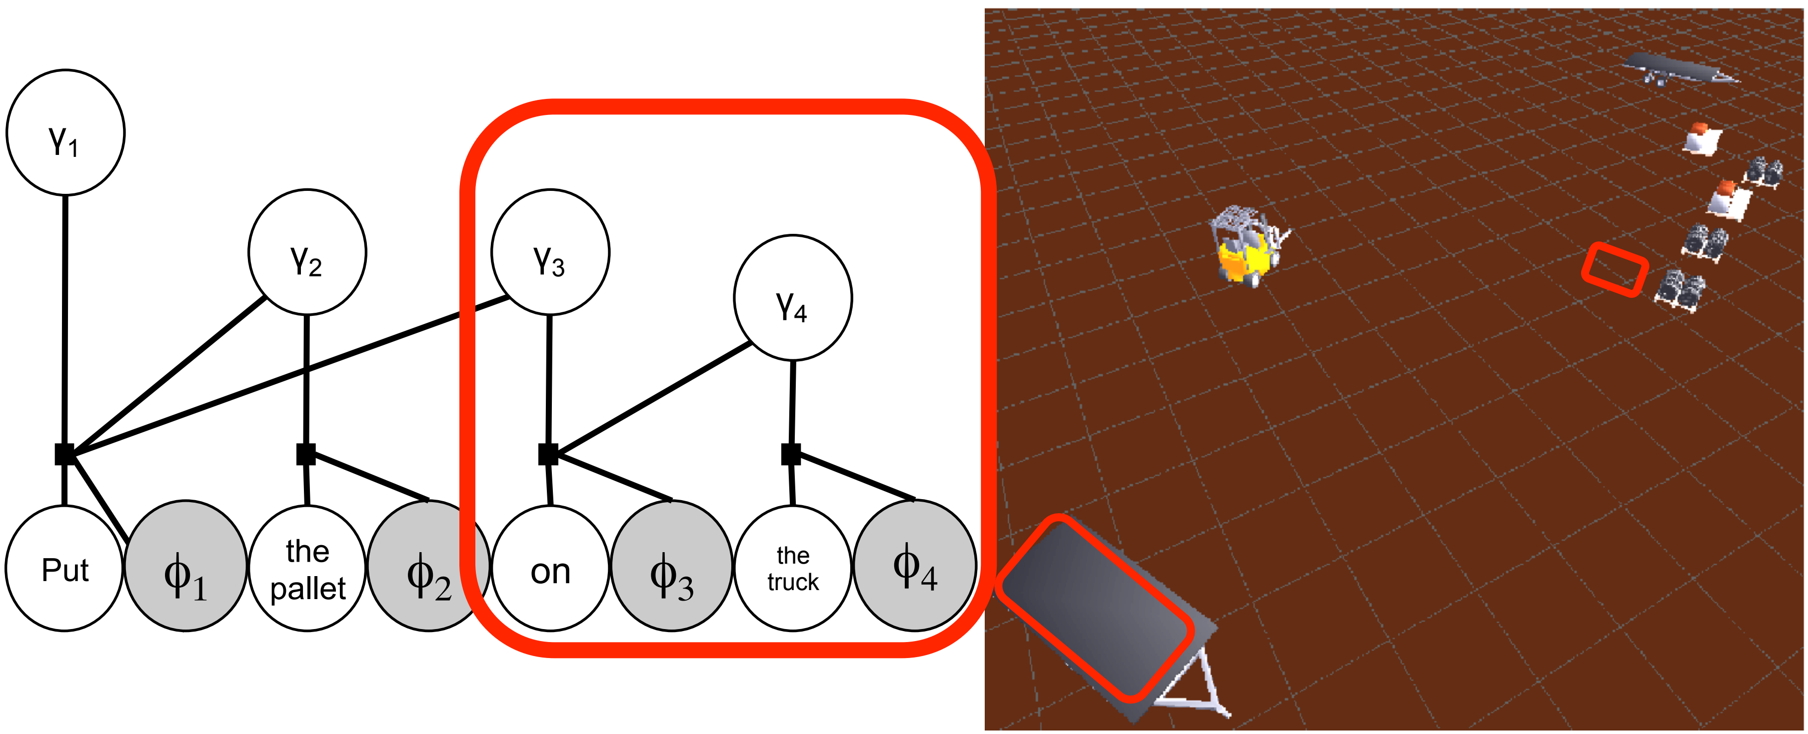

Candidate for γ_3

Place the lifted tyre pallet, next to another tyre pallet on the trolley.
Lift the tire pallet in the air, then proceed to deposit it to the right of the tire pallet already on the table right in front of you.
Place the pallet of tires on the left side of the trailer.
Please lift the set of six tires up and set them on the trailer, to the right of the set of tires already on it.
Place a second pallet of tires on the trailer.
lift the tire pallet you are carrying and set on the truck in front of you
Place the pallet of tires that is on the forklift next to the pallet of tires .
Lift t on on truck. Lower tire pallet.
Revers End.
lift t
Arrang
Place of the truck.
Lift t set it on the platform directly
ahead, eady there.
Put th

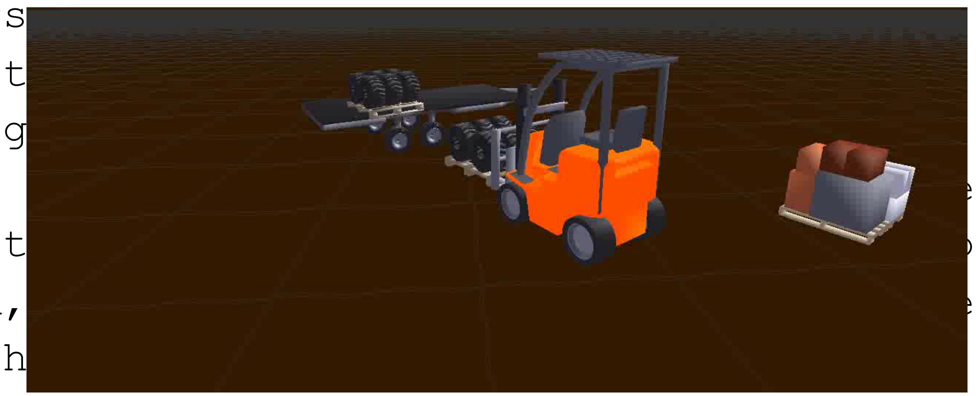

Load the skid right next to the other skid of tires on the trailer.
Put the tire pallet on the trailer to the right of the other tire pallet.
Lift pallet up and place the pallet beside the other pallet on the truck bed. Reverse the forklift slowly from the truck bed.
Raise tire pallet. Move forward to unoccupied location on truck. Lower tire pallet. Reverse to starting position. Lower forks.
Move the pallet on the ground to the platform; place it to the right of the pallet that is already on the platform.

S. Tellex et al, AAAI 2011, ISER 2012

Generalized Grounding Graphs

https://youtu.be/wzRp4BY0U1g

A problem with generalization: speed

This grounding variable is a motion plan.

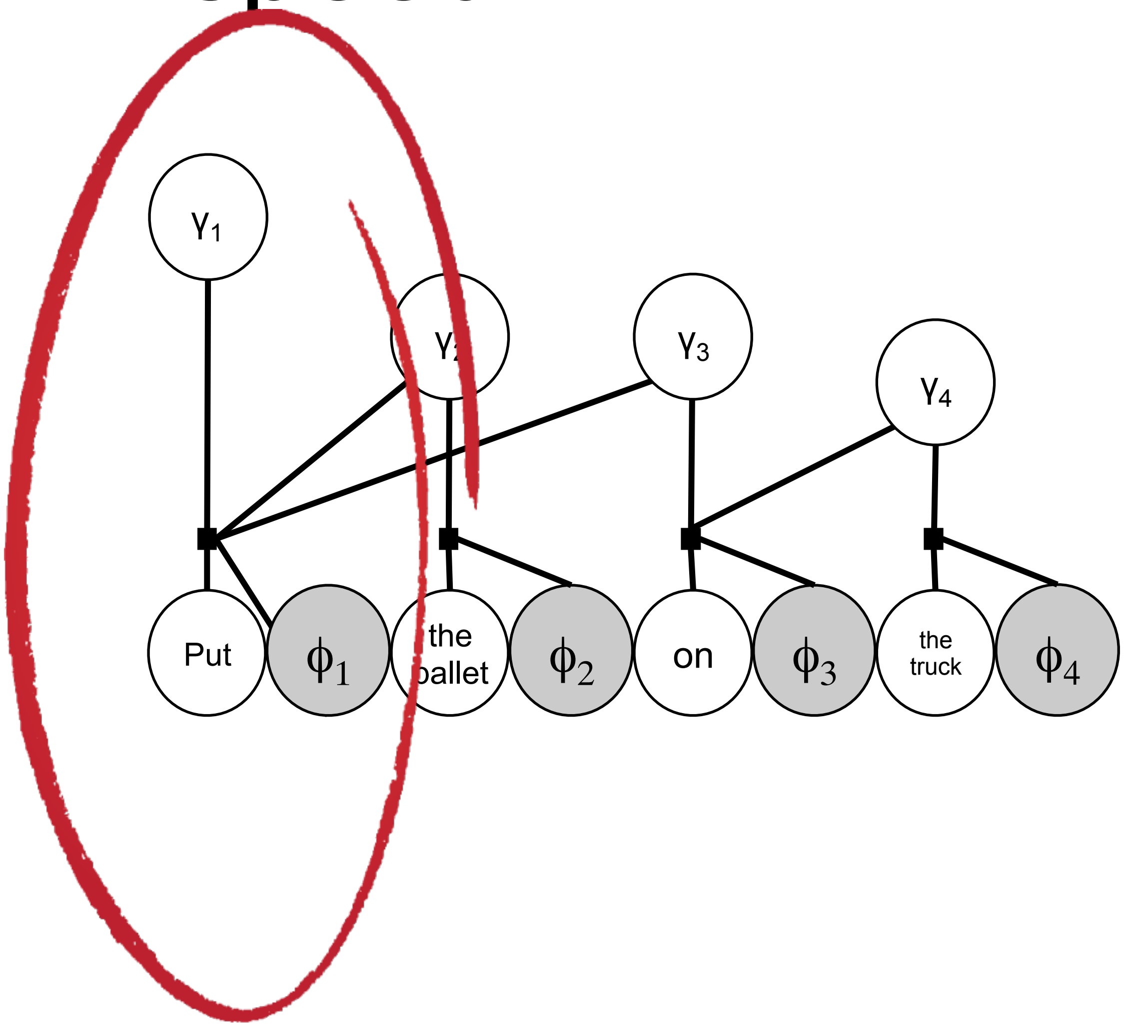

T. Howard et al, ICRA 2014

Solve the motion planning problem separately

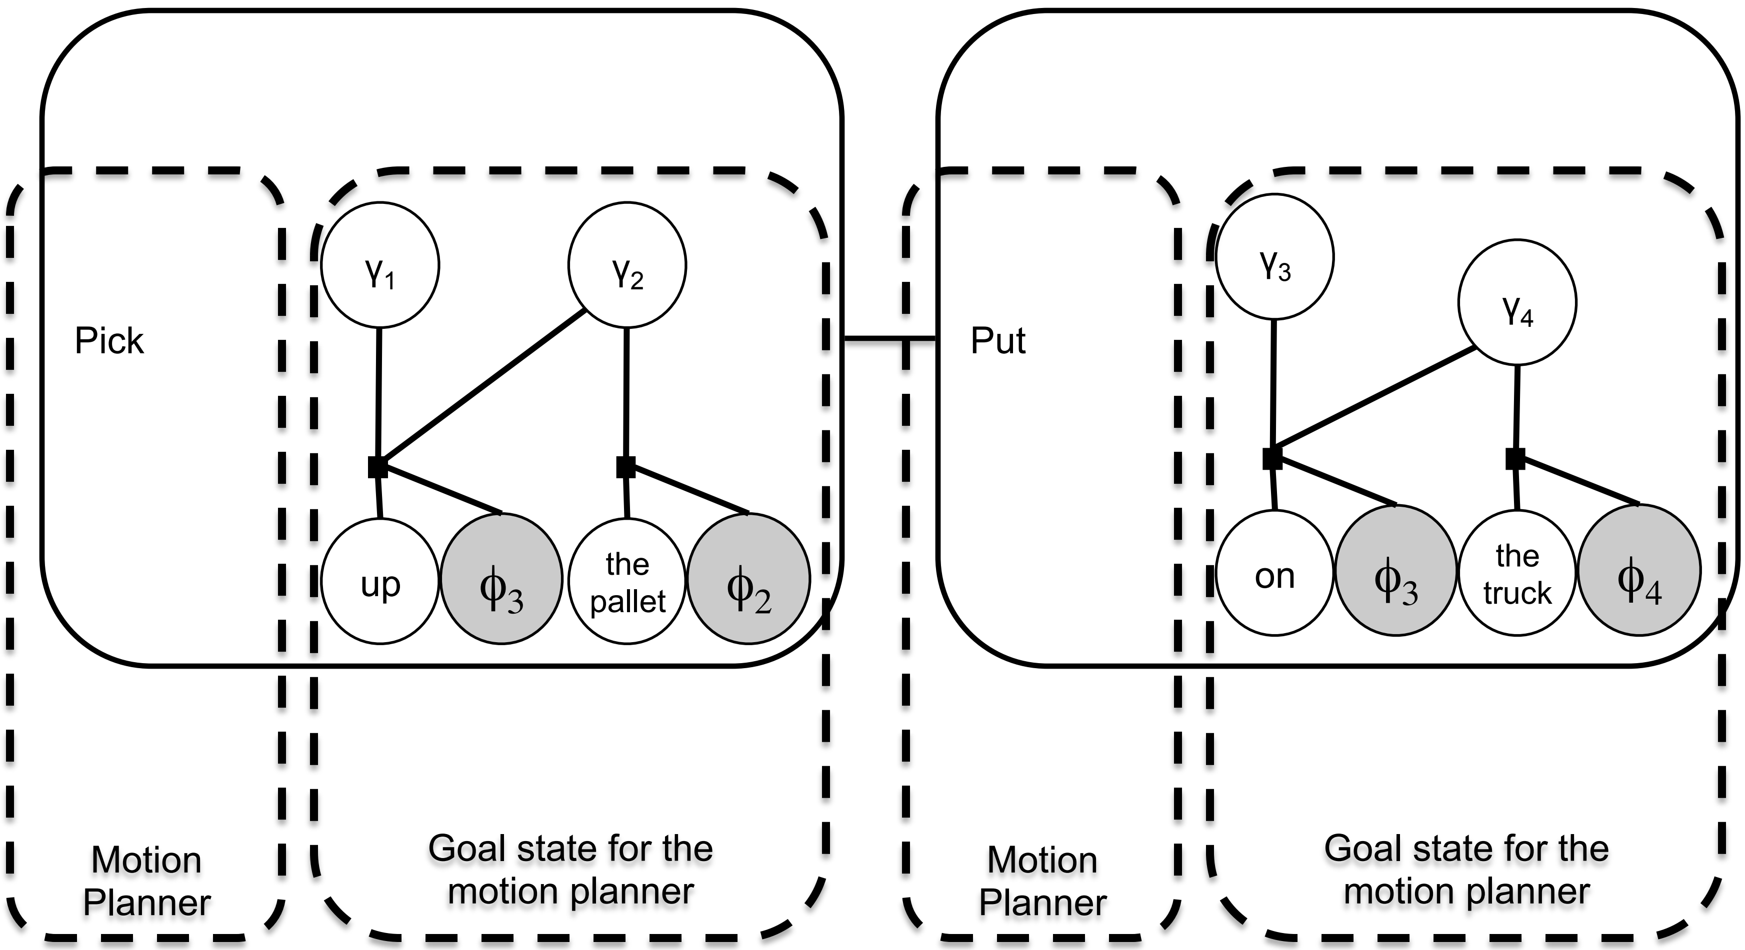

T. Howard et al, ISER 2014

Adding Hierarchy

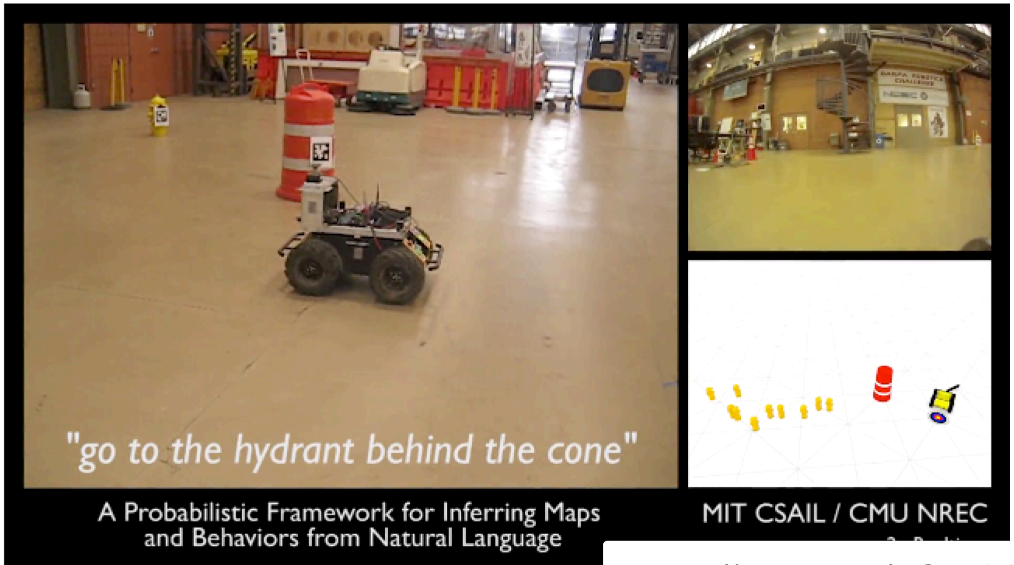

https://youtu.be/nGIA818ozBY

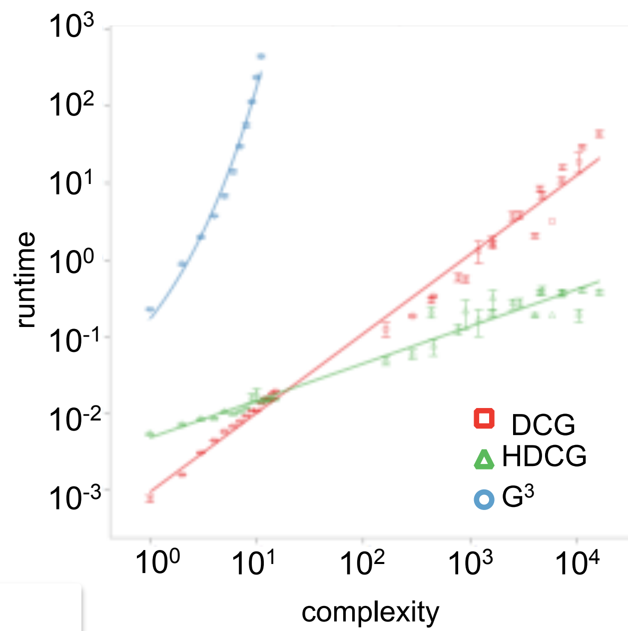

Another problem

- No notion of abstract concepts:
  - “pick up the first block in the row of blocks”
  - “grasp the nearest block in the group”
  - “place the tool in the middle of the circle”

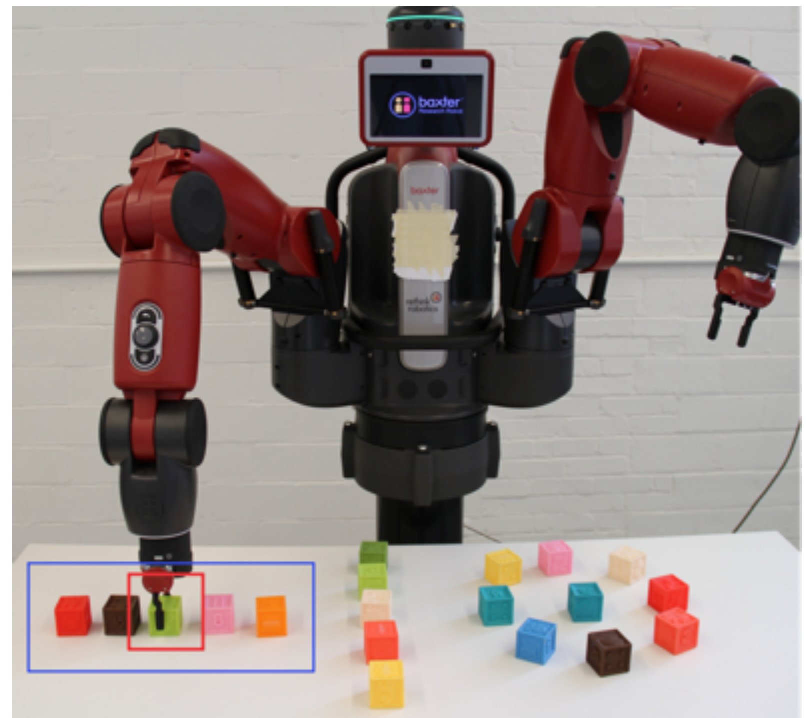

Abstract concepts require reasoning about large sets of objects

What exactly gets grounded here?
What is even the domain of
these variables?

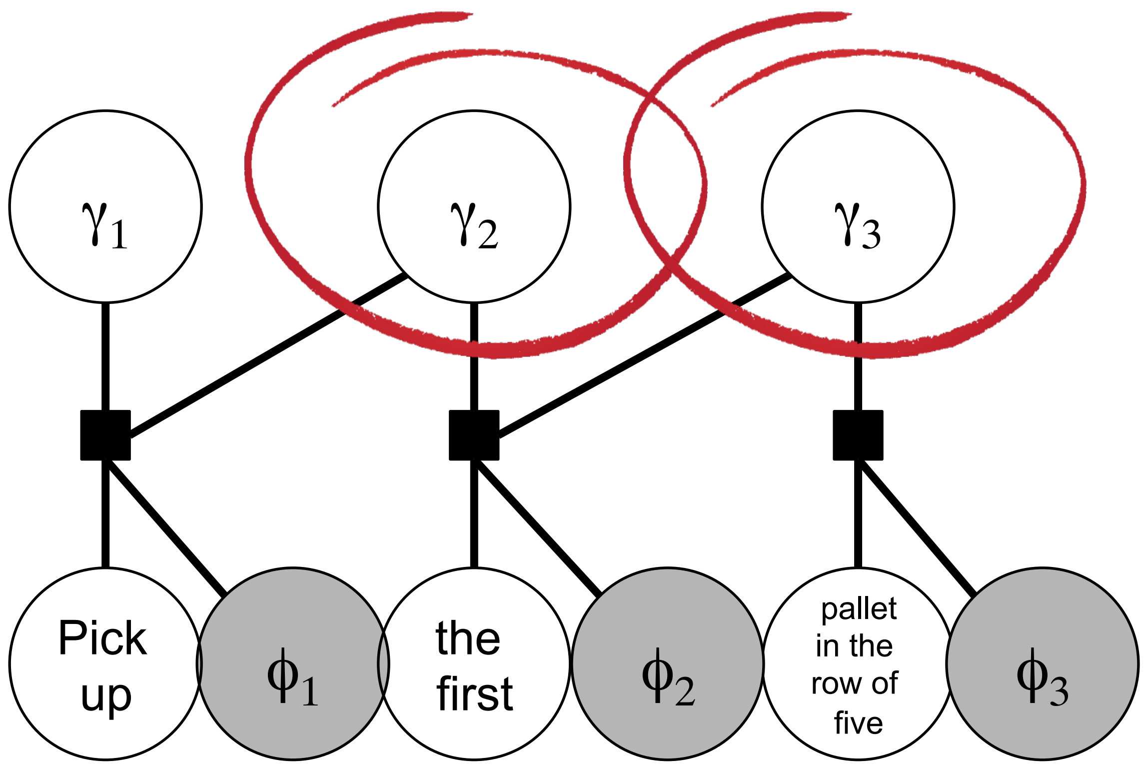

Reasoning about large sets of objects is painful

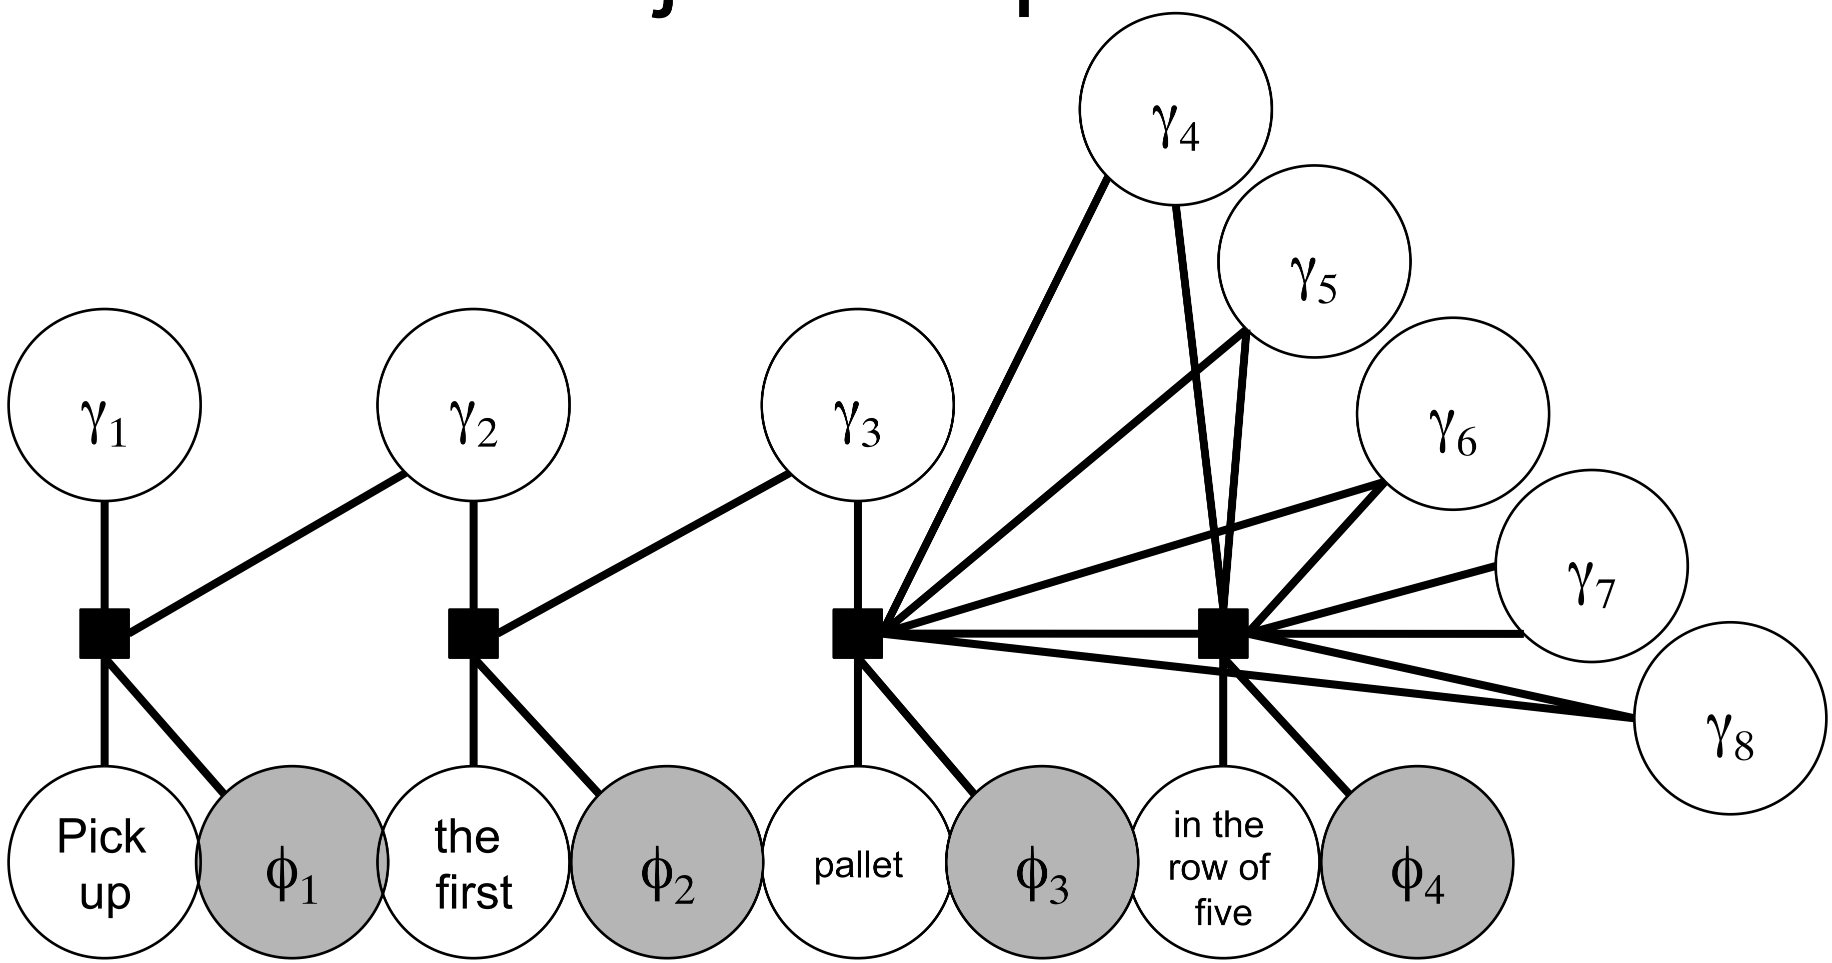

Abstract concepts require abstraction

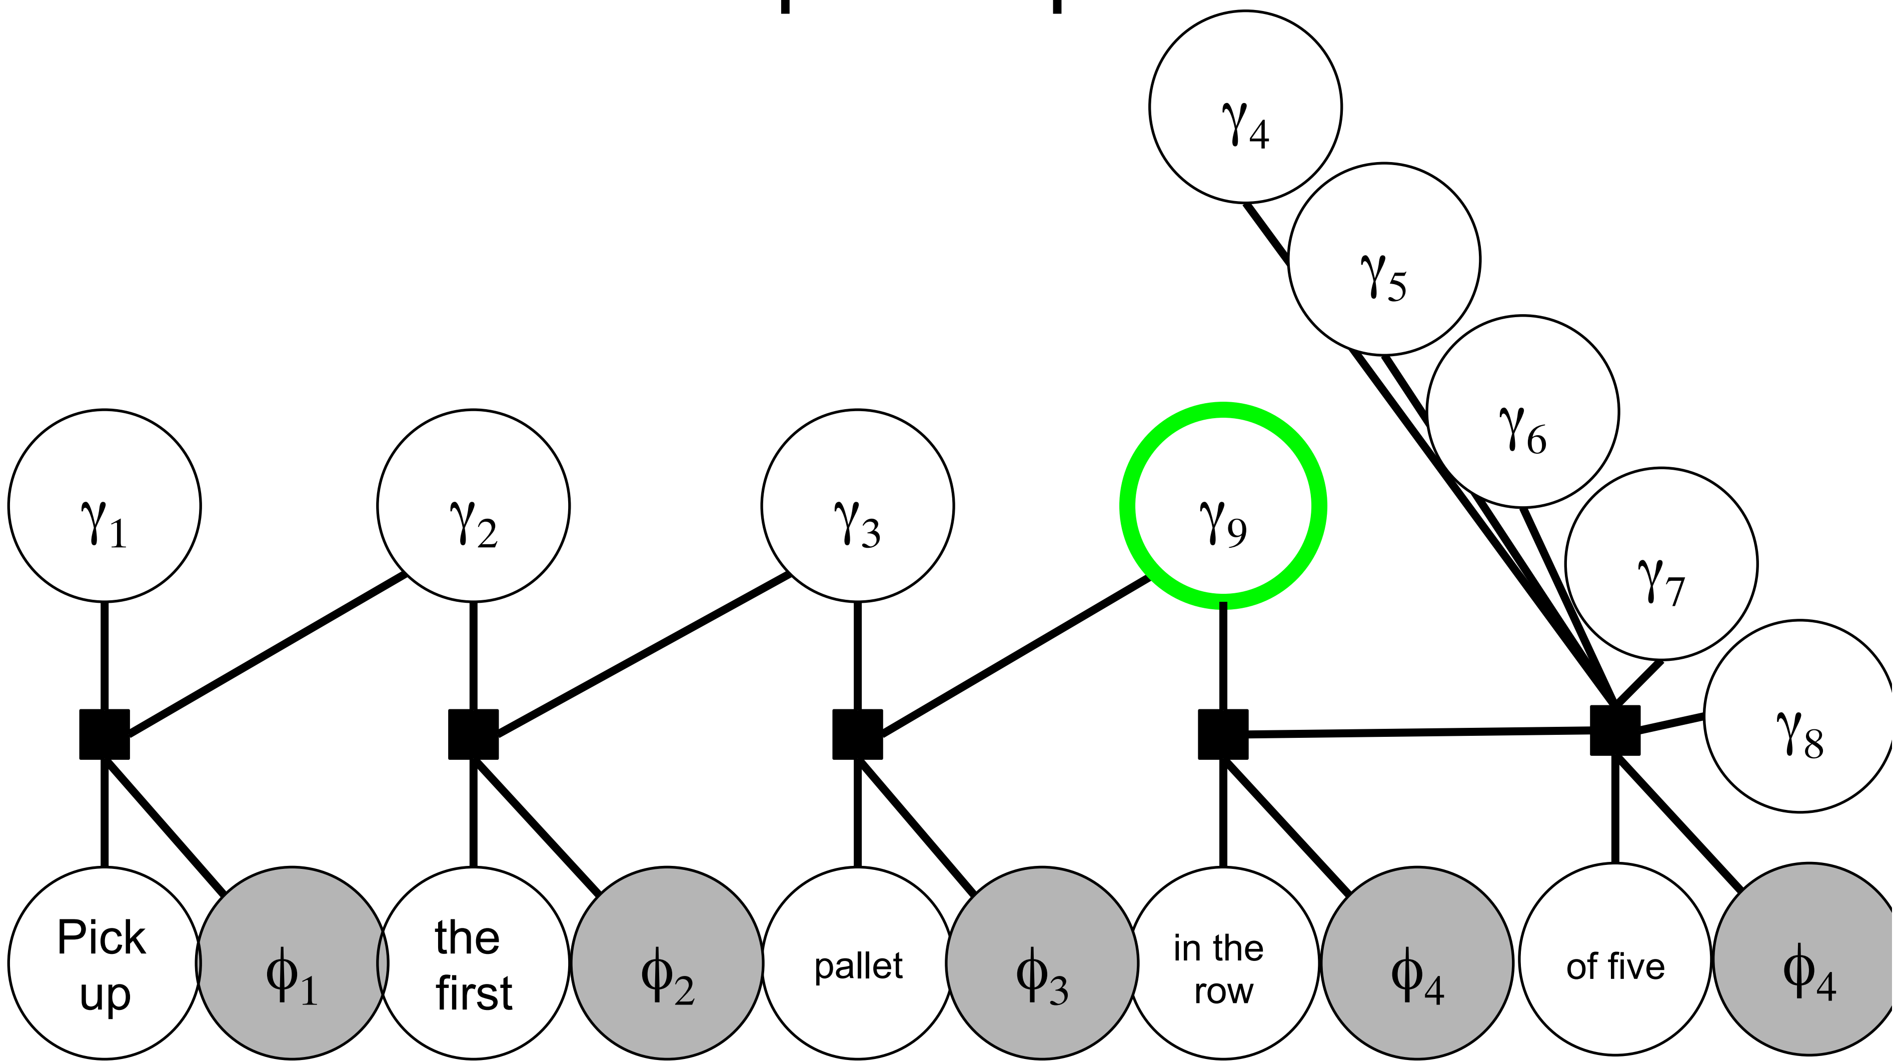

R. Paul et al, RSS 2016

Inference Speed

	Runtime	
Objects	DCG Baseline	Proposed Model
4	0.14 ±0.0031	0.0070 ±0.00023
5	0.21 ±0.0092	0.0091 ±0.00057
7	0.47 ±0.033	0.010 ±0.00079
10	2.96 ±0.18	0.010 ±0.00010
12	14.25 ±0.51	0.011 ±0.00072
Total	1.89	0.062

(Inference speed in seconds)

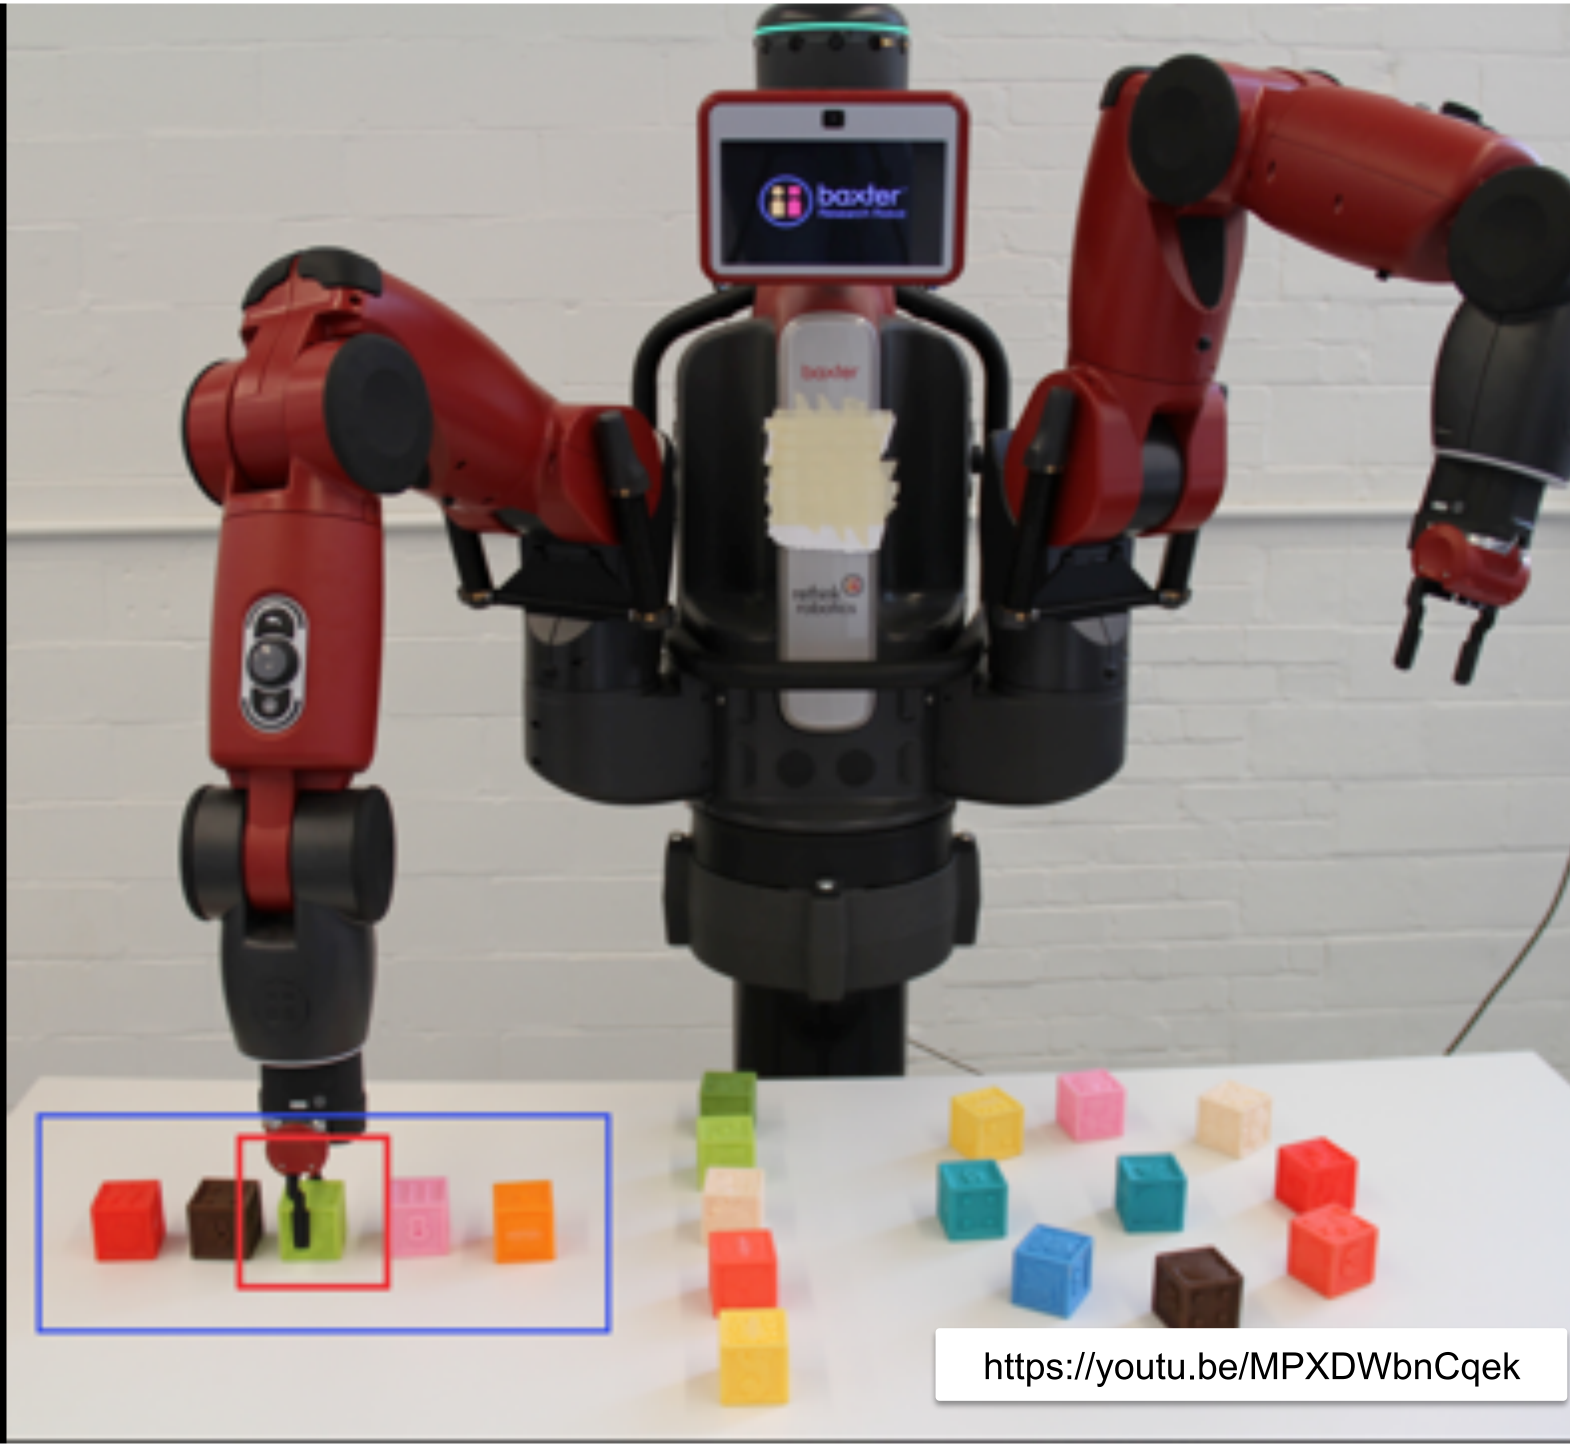
baxter
https://youtu.be/MPXDWbnCqek

What we have today	What we need

• Unstructured, flat world representations
  • Hand-coded concepts in the representation, hand-coded relations between perception and planning
  • Hand-coded behaviors and motion strategies.
	• Representations that are not hand-coded but learned from data, and support wide range of tasks.
	  • The representations must be able to use context to focus computation on relevant concepts, and capture higher-level (abstract) concepts.

• State is assumed to be fully observable and known perfectly	• Need strategies that can plan to avoid failures due to uncertainty, and plan to gather more information when needed • These strategies must be computationally efficient.

• Learning available only on targeted learning tasks (no end-to-end learning)
  • No requirement of real-time response or model of computational cost of inference and learning
	• Need the ability to carry out end-to-end learning, and adapt to changes in the world and the model over time.
	  • These learning algorithms must be efficient enough to run online and provide guarantees that performance will not be degraded by learning

Higher Level Autonomy

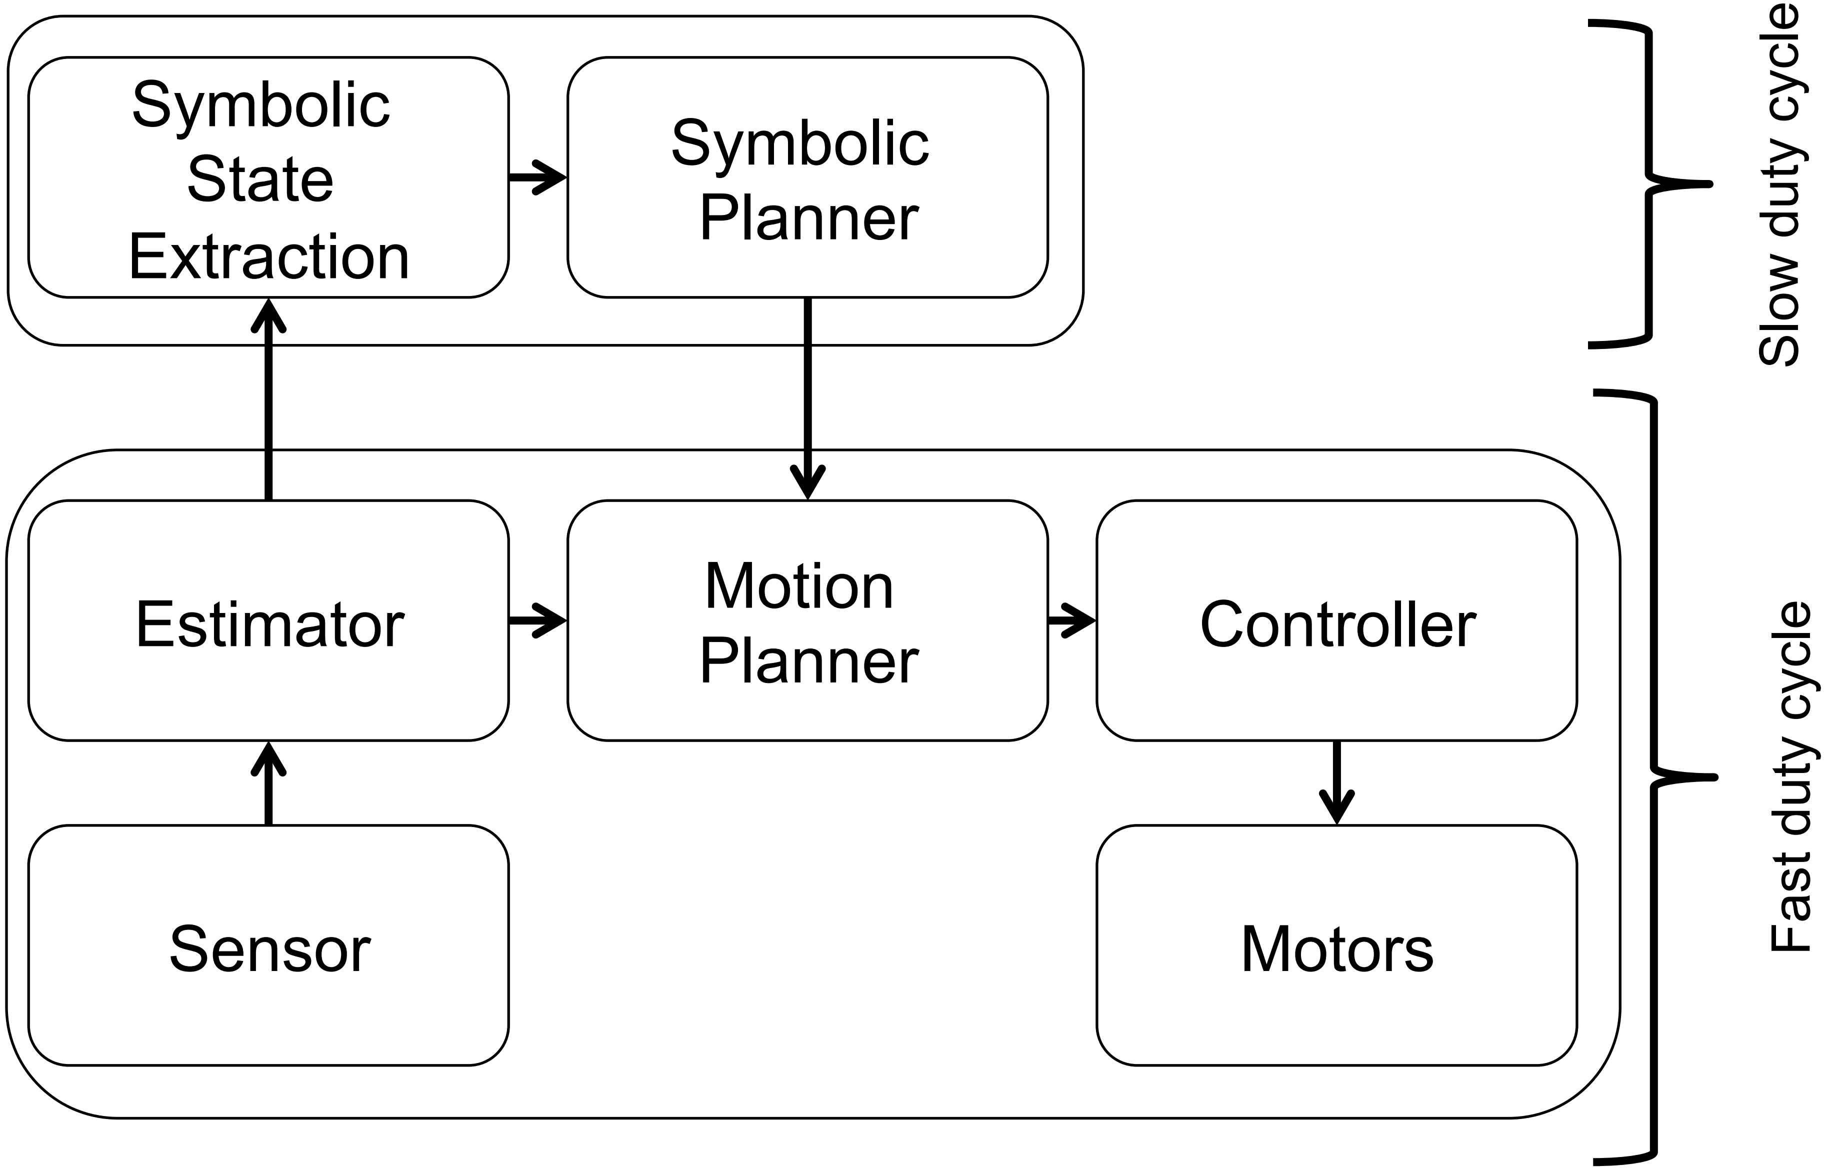

Task and Motion Planning

2011 IEEE International Conference on Robotics and Automation
Shanghai International Conference Center
May 9-13, 2011, Shanghai, China

Hierarchical Task and Motion Planning in the Now

Leslie Pack Kaelbling and Tomás Lozano-Pérez

***Abstract*— In this paper we outline an approach to the integration of task planning and motion planning that has the following key properties: It is aggressively hierarchical; it makes choices and commits to them in a top-down fashion in an attempt to limit the length of plans that need to be constructed, and thereby exponentially decrease the amount of search required. It operates on detailed, continuous geometric representations and does not require a-priori discretization of the state or action spaces.**

of a complete low-level plan with guaranteed soundness or optimality conditions. Our goal is to design a system that can work effectively with non-determinism in the environment or in the low-level controllers. In such cases, plan... detail far into the future will typically be wasted, due to the inability to predict exactly what will happen. For this re... we plan 'in the now': we construct a plan at an abstract l..., commit to it, and then recursively plan and execute actions to

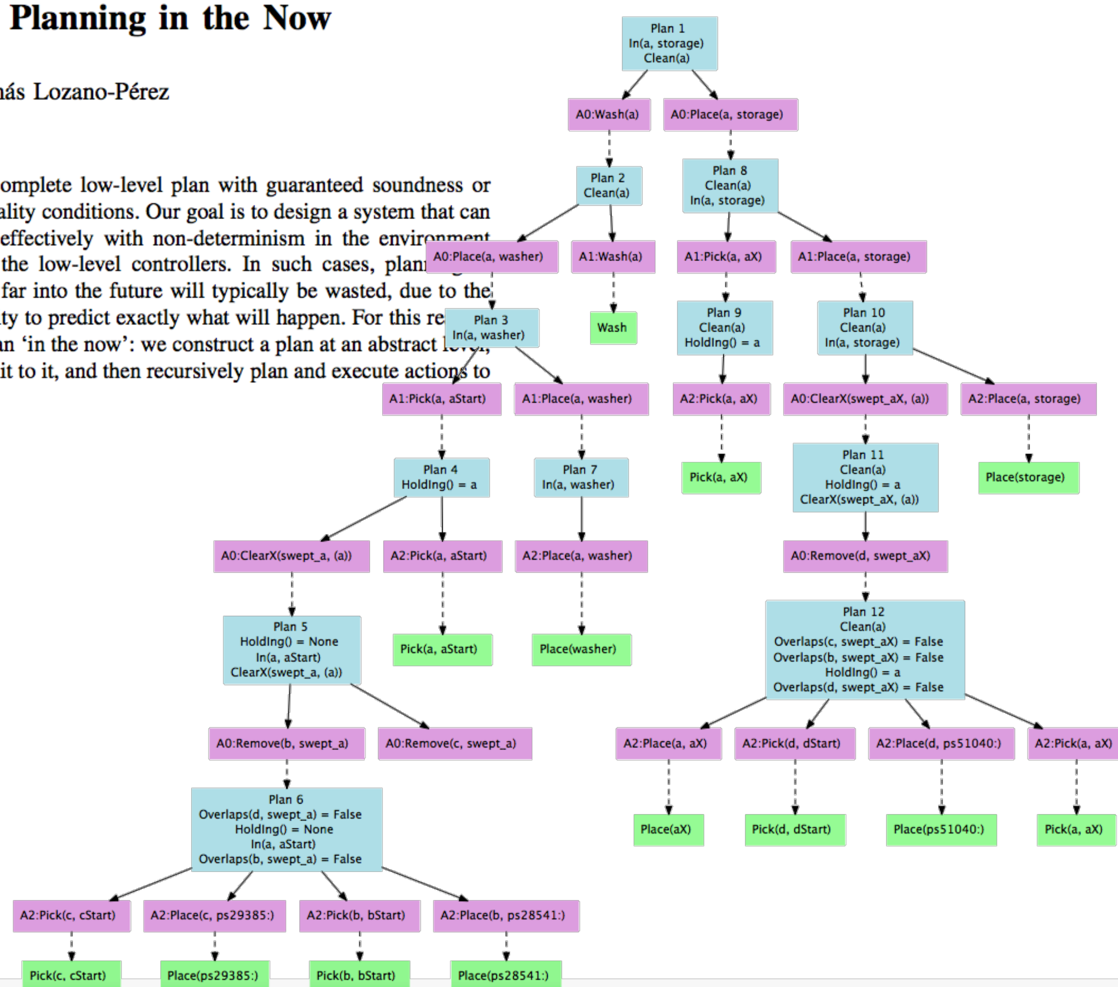

Many different approaches to applying hierarchy or structure to achieve higher level autonomy.

Higher Level Autonomy

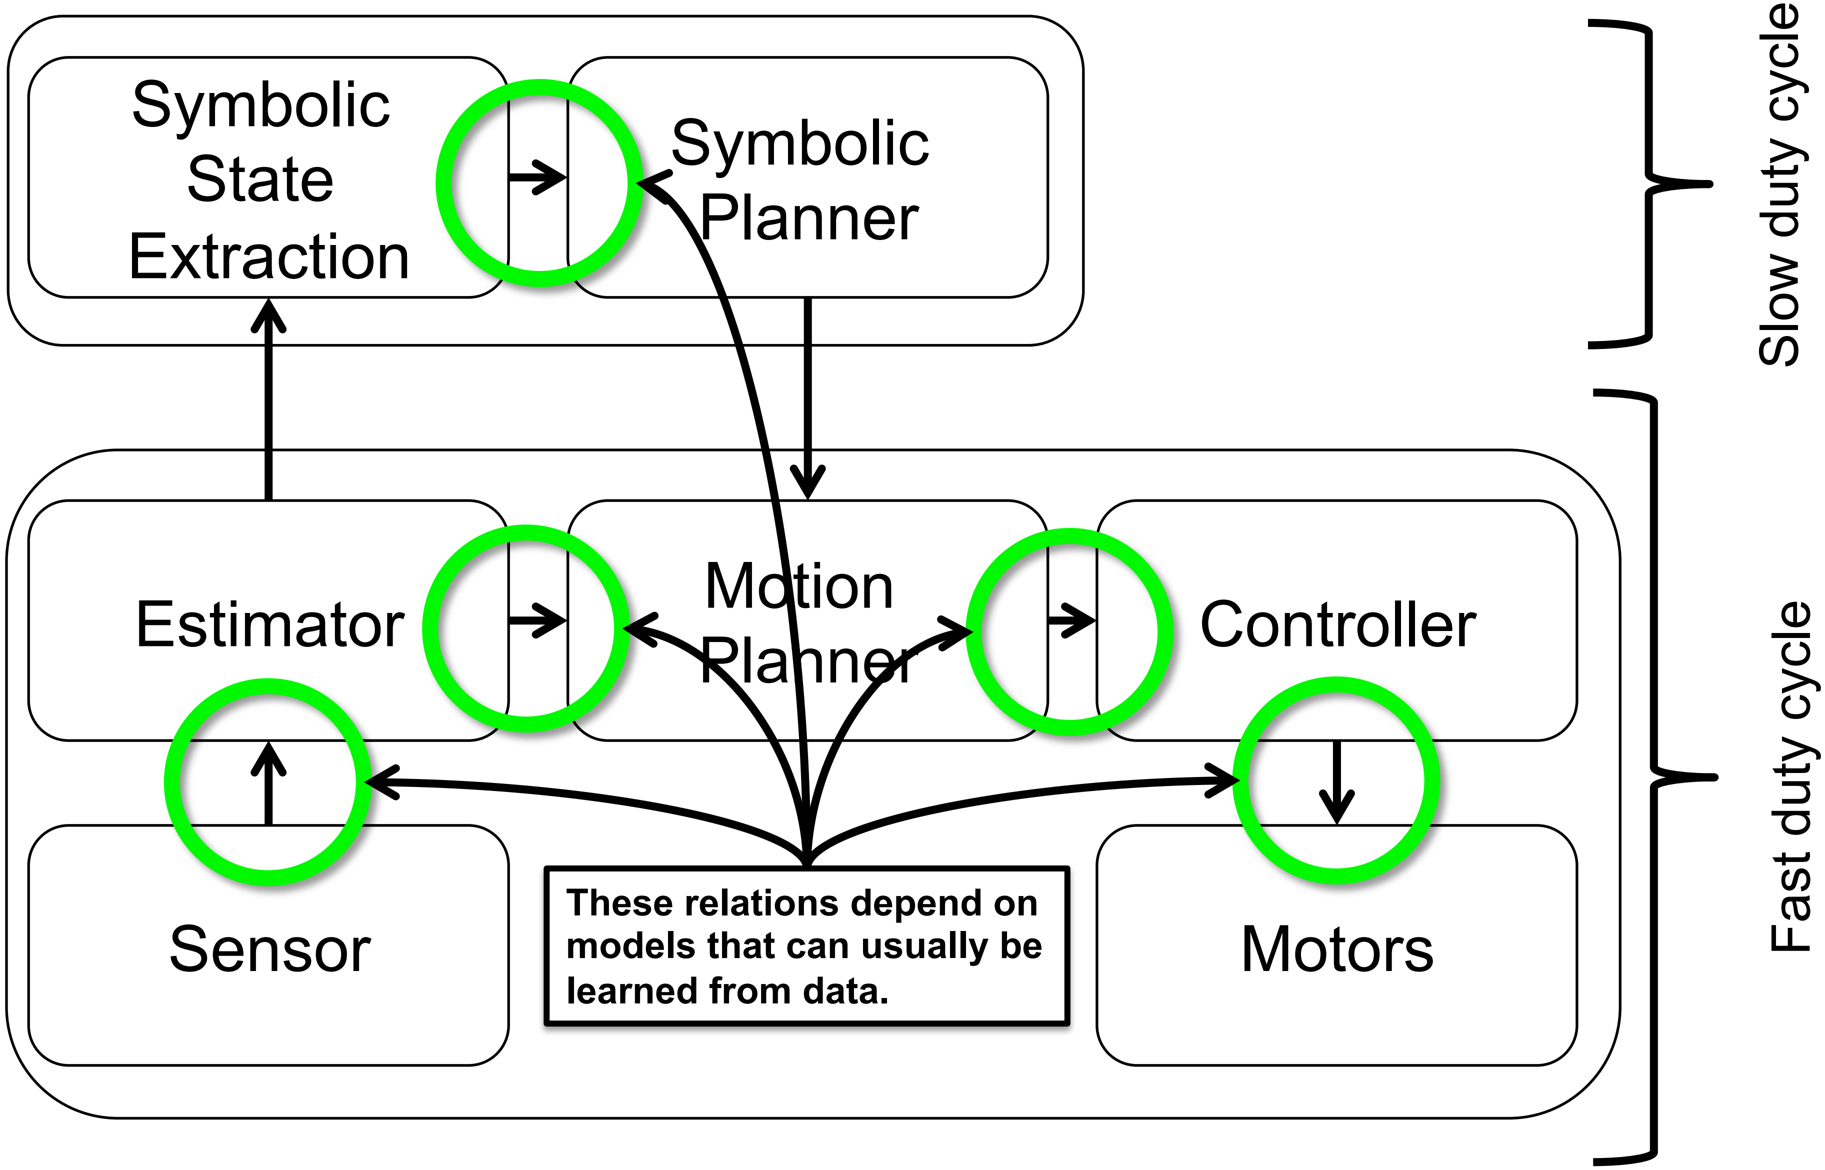

Higher Level Autonomy

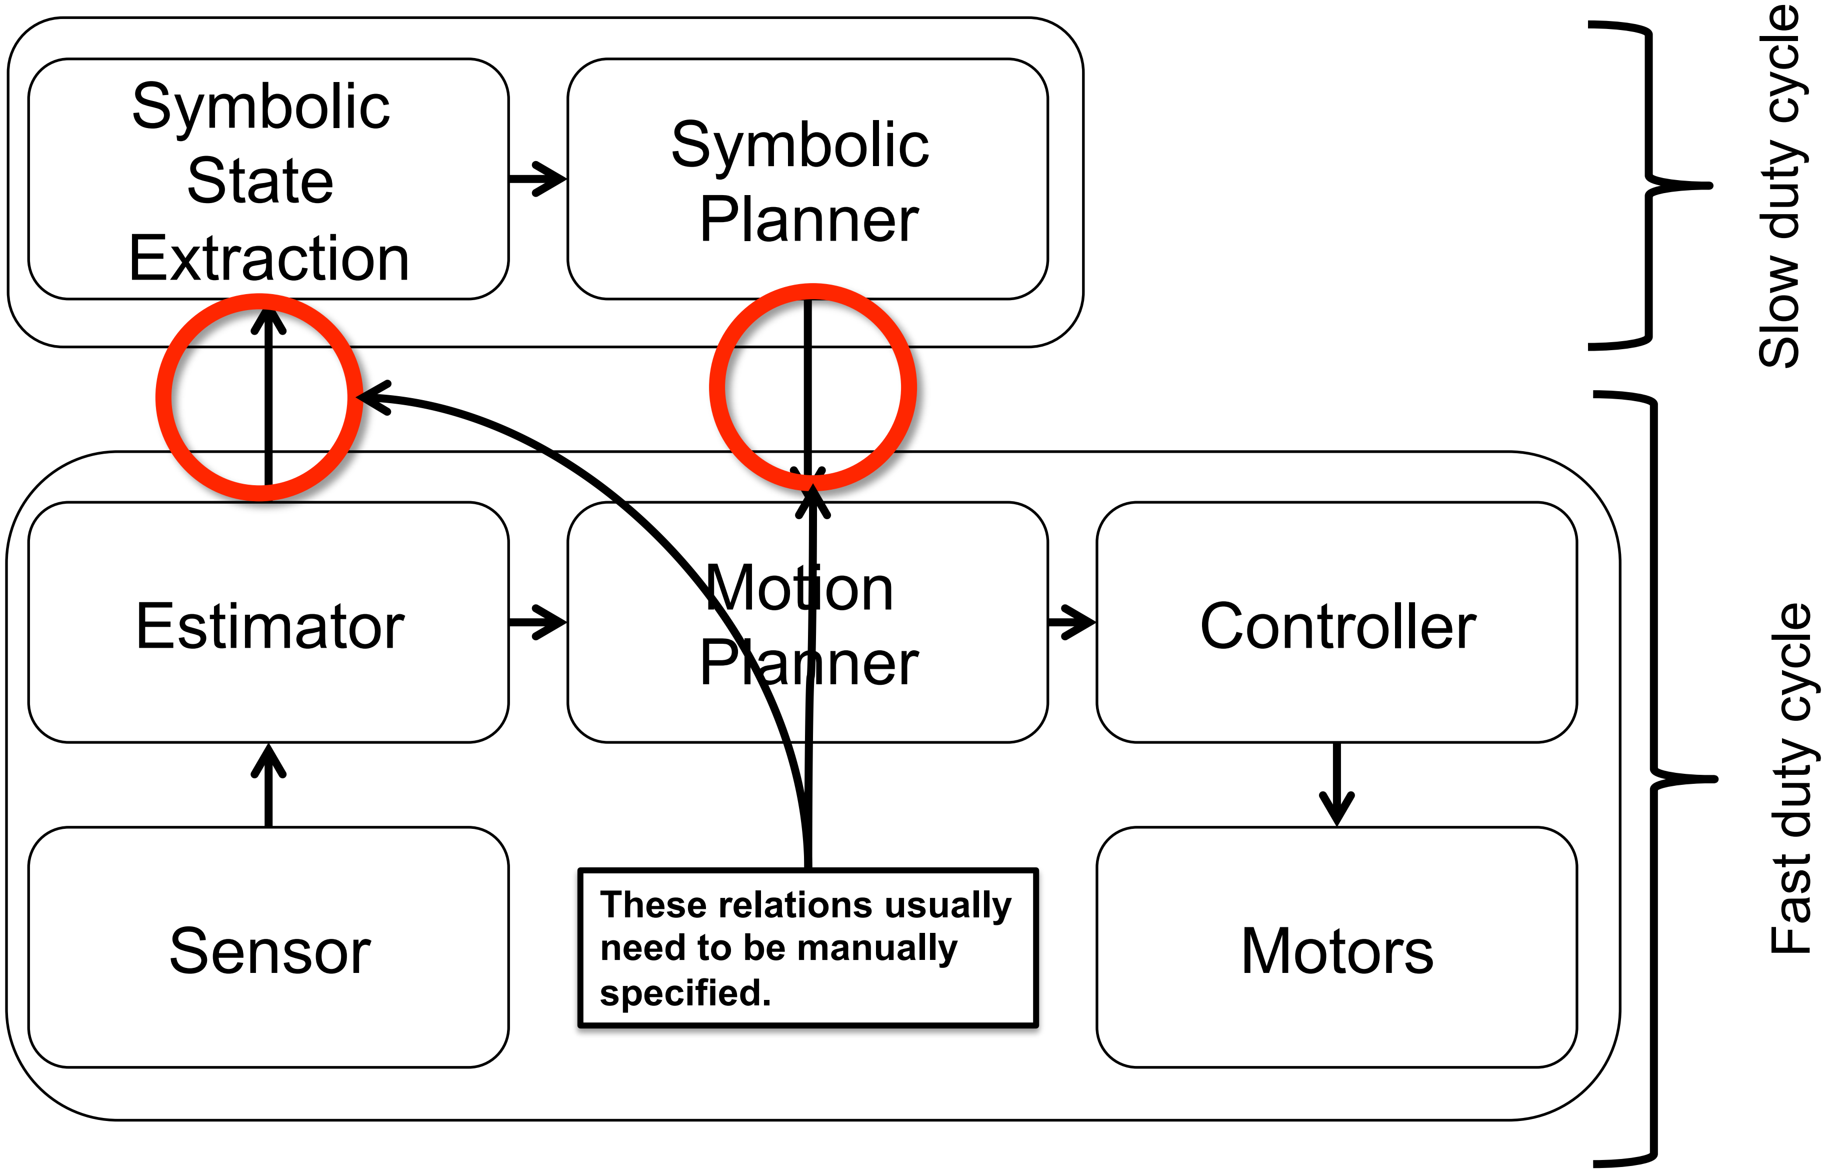

Summary

- Robust, long-term autonomy in unknown, populated environments
- Models for representing complex worlds that let us learn and plan efficiently
  - Bayesian non-parametrics
  - Leveraging domain structure for efficient learning