Semantically Interpretable Predictive State Representation

Johannes A. Stork, Carl Henrik Ek, and Danica Kragic
KTH, Stockholm, Sweden

ICRA 2015 Workshop on Sensorimotor Learning

Robots that are able to learn models of themselves and their environments. • Most real-world sensorimotor control problems are situated in continuous or high-dimensional environments and require real-time interaction, which can be problematic for classical learning techniques. • Is it possible to learn sensorimotor dynamics of robots or animals directly from the raw data? If not, what prior knowledge is necessary? • How can one balance the representation accuracy and the speed of inference? How much data is needed? • How can successful supervised or unsupervised learning techniques be used in sensorimotor control problems? • How can prior knowledge, including expert knowledge, user demonstrations, or distributional assumptions be incorporated into the learning/ planning framework?

http://sensorimotor-learning.mit.edu/

ICRA 2015 Workshop on Sensorimotor Learning

Robots that are able to learn models of themselves and their environments • Most real-world sensorimotor control problems are situated in continuous or high-dimensional environments and require real-time interaction, which can be problematic for classical learning techniques. • Is it possible to learn sensorimotor dynamics of robots or animals directly from the raw data If not, what prior knowledge is necessary? • How can one balance the representation accuracy and the speed of inference How much data is needed? • How can successful supervised or unsupervised learning techniques be used in sensorimotor control problems? • How can prior knowledge, including expert knowledge, user demonstration, or distributional assumptions be incorporated into the learning/ planning framework?

http://sensorimotor-learning.mit.edu/

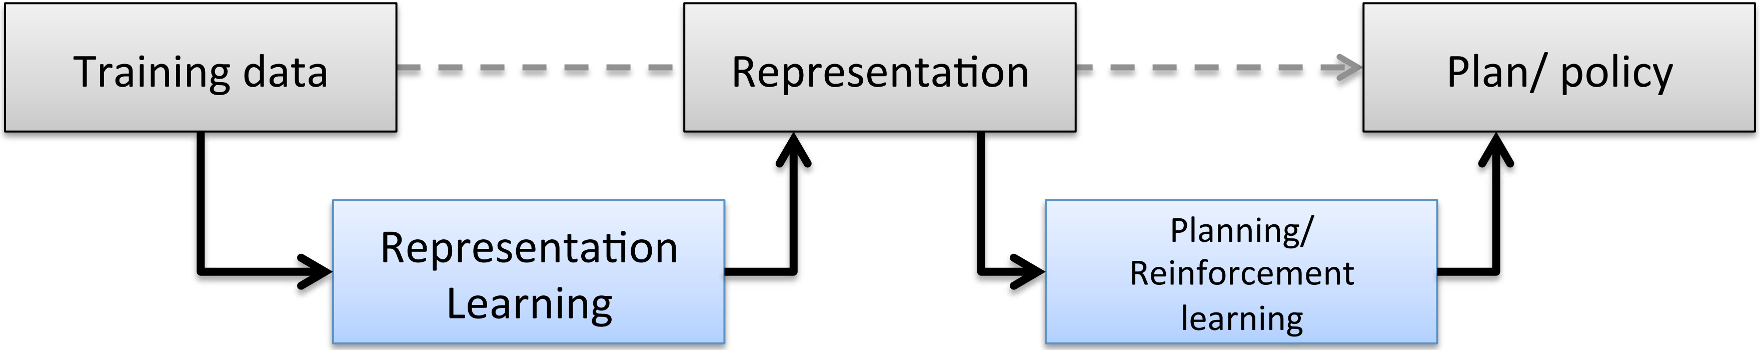
Training data
Representation
Plan/ policy
Representation Learning
Planning/ Reinforcement learning

Representation Learning for Robotics

- Sensor data
  - Little → Hidden information
  - Much → Meaning, semantic?
- Dynamics, interaction
- Processing time
- Stream of data
- Sparse training data, expensive
  - demonstrations, exploration

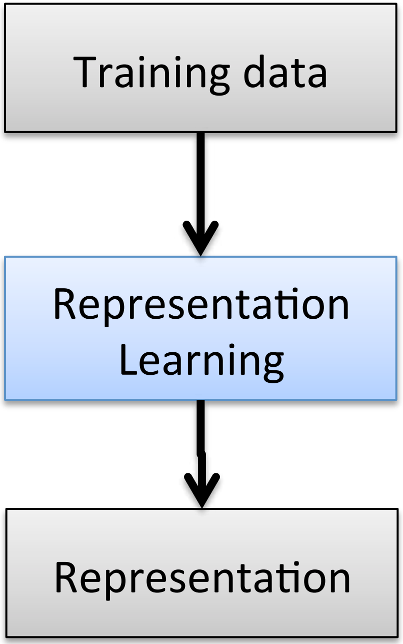

Robot Navigation

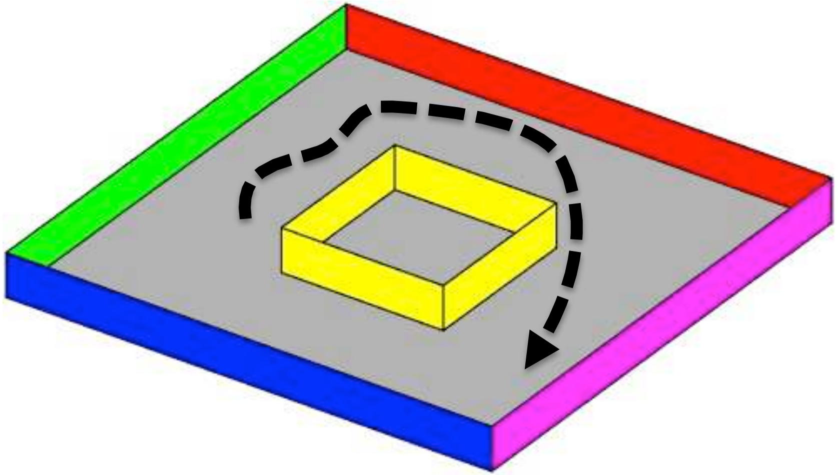

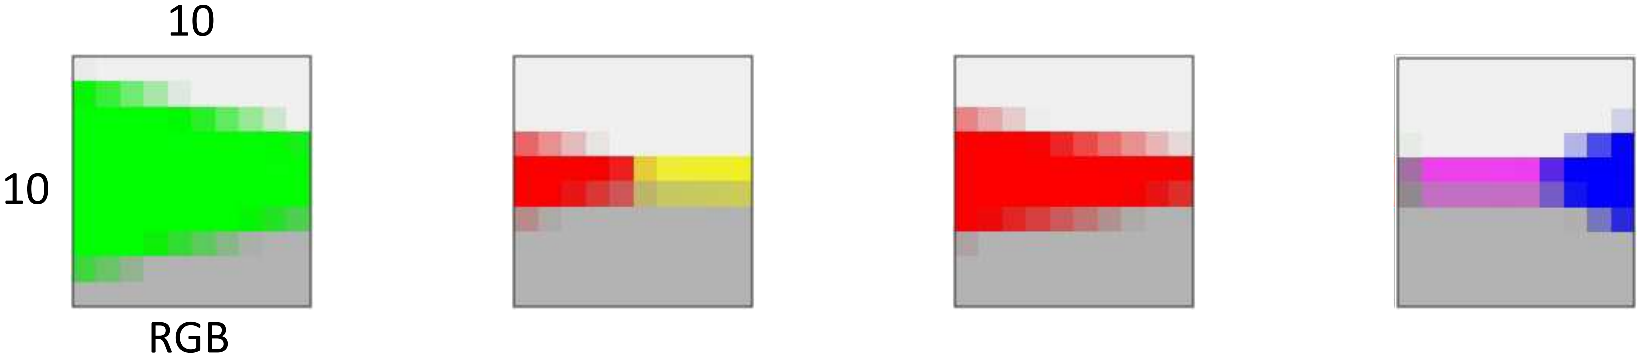

Predictive State Representation

- Representation Learning
- Recent
  (e.g. Boots2011)
- No Latent Space
- Observable
- Features
- For planning

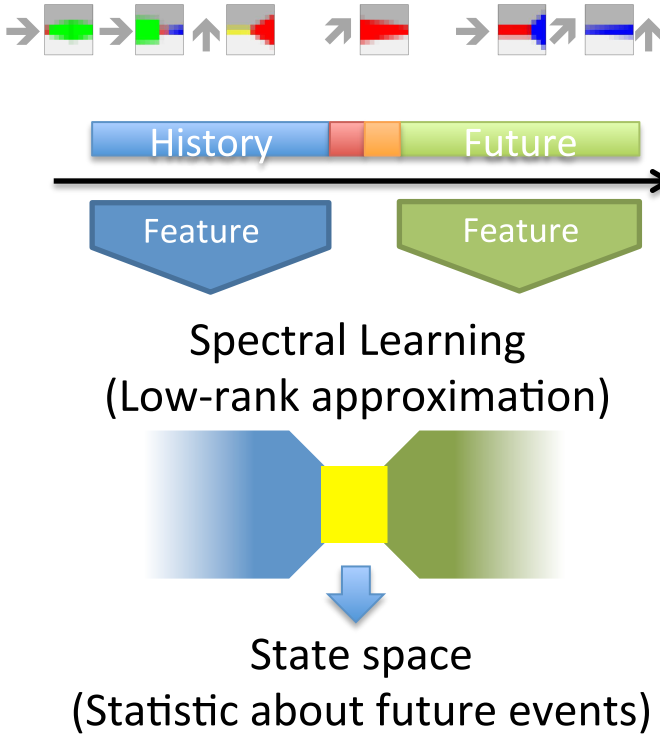

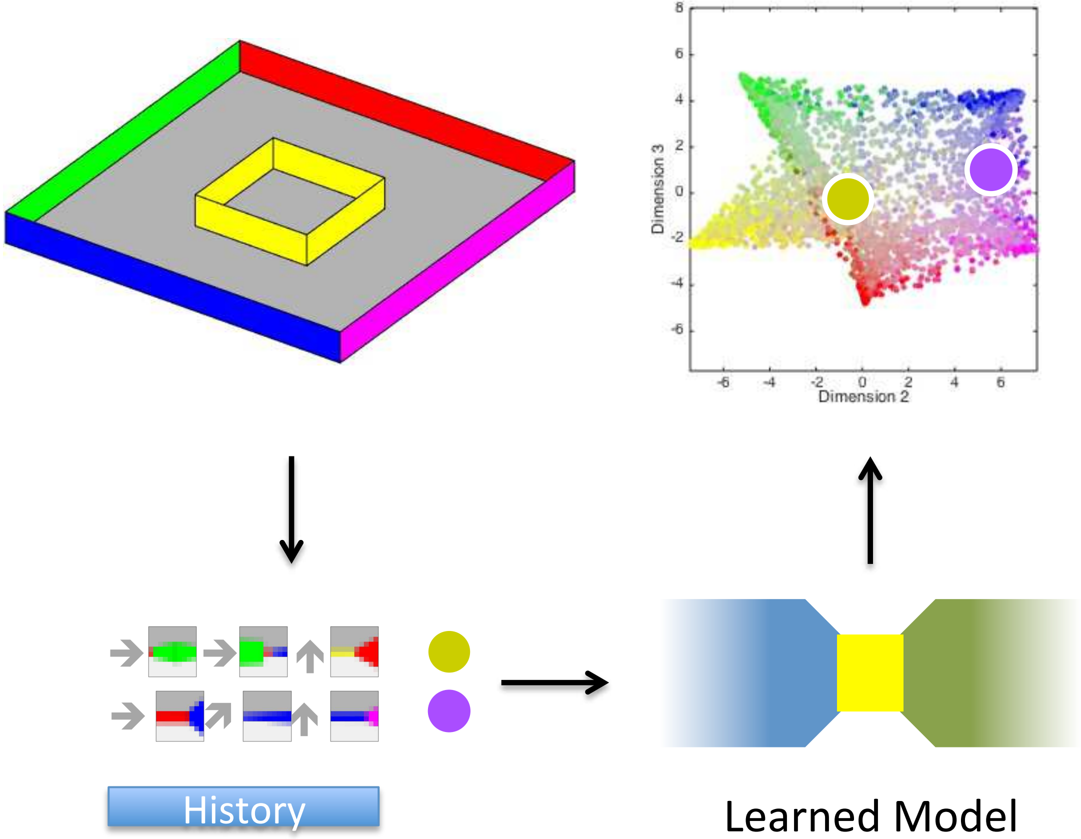

Dimension 3
Dimension 2
8
6
4
2
0
-2
-4
-6
-6
-4
-2
0
2
4
6
History
Learned Model

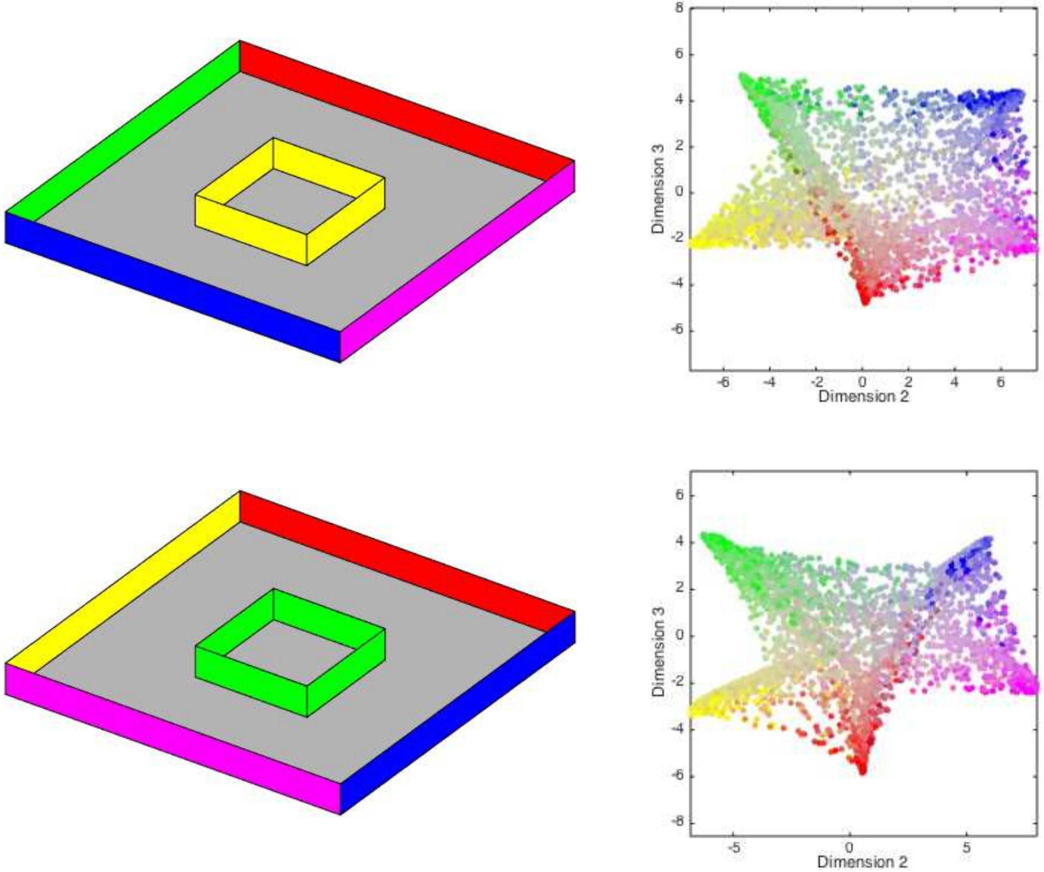

Dimension 3
Dimension 2
8
6
4
2
0
-2
-4
-6
-6
-4
-2
0
2
4
6
Dimension 3
Dimension 2
6
4
2
0
-2
-4
-6
-8
-5
0
5

Inspiration

- Engineers must […] build in architectural constraints and fundamental truths […]
- Prior + Experience = Learned competence
- Agents must learn niche-specific competences […] sensory-motor loops, world model […]

Leslie P. Kaelbling,
Keynote Lecture, AAAI 2010

- Feature engineering […] most common approach […] mapping from observation to state
- Mapping […] designed by hand, using human intuition.
- Using prior knowledge […] about world, robots can learn representations […] consistent with physics.
- Identify five robotic priors […]

R. Jonschkowski and O. Brock. "State Representation Learning in Robotics: Using Prior Knowledge about Physical Interaction". In: RSS. Berkeley, USA, 2014

Inspiration

- **Engineers** must […] build in **architectural constraints** and **fundamental truths** […]

- **Feature engineering** […] most common approach […] mapping from **observation to state**
- Mapping […] **designed by hand, using human intuition**.

- Prior + Experience = Learned competence
- Agents must learn niche-specific competences […] sensory-motor loops, world model […]

Leslie P. Kaelbling,
Keynote Lecture, AAAI 2010

- Using prior knowledge […] about world, robots can learn representations […] consistent with physics.
- Identify five robotic priors […]

R. Jonschkowski and O. Brock.
"State Representation Learning in Robotics: Using Prior Knowledge about Physical Interaction".
In: RSS. Berkeley, USA, 2014

Inspiration

- Engineers must […] build in architectural constraints and fundamental truths […]

- **Prior + Experience = Learned competence**
- Agents must **learn** niche-specific competences […] **sensory-motor loops, world model** […]

Leslie P. Kaelbling,
Keynote Lecture, AAAI 2010

- Feature engineering […] most common approach […] mapping from observation to state
- Mapping […] designed by hand, using human intuition.

- Using **prior knowledge** […] about **world**, robots can learn representations […] **consistent with physics.**
- Identify **five robotic priors** […]

R. Jonschkowski and O. Brock.
"State Representation Learning in Robotics: Using Prior Knowledge about Physical Interaction".
In: RSS. Berkeley, USA, 2014

Linear PSR Theory

- Test prediction $P(\tau^{\mathcal{O}}|h||\tau^{\mathcal{A}})$
- State
- Core set

$$s(h) = \begin{bmatrix} P(\tau_1^{\mathcal{O}}|h||\tau_1^{\mathcal{A}}) \\ \vdots \\ P(\tau_k^{\mathcal{O}}|h||\tau_k^{\mathcal{A}}) \end{bmatrix}$$

- State update $s(hao) \propto M_{ao} s(h)$

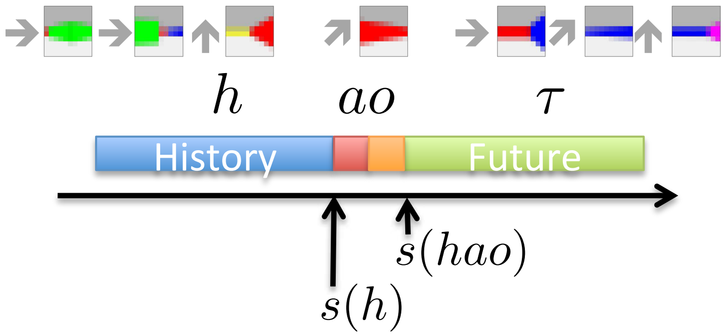

$o \in \mathcal{O}$ Observations and
$a \in \mathcal{A}$ Actions

Flow of Information and Observables

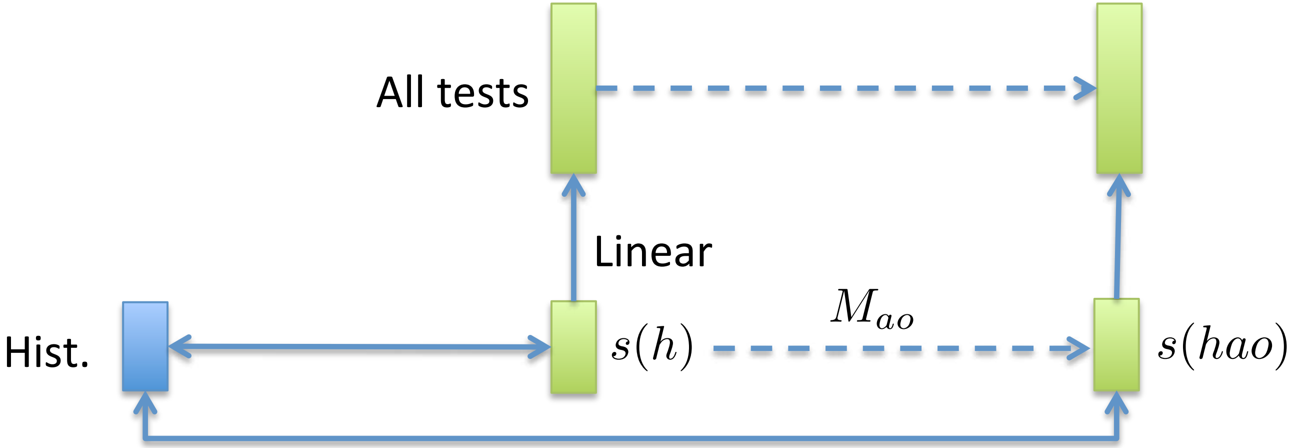

- Discovery of sufficient set
- Parameters update parameters

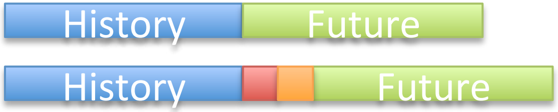

Flow of Information and Observables

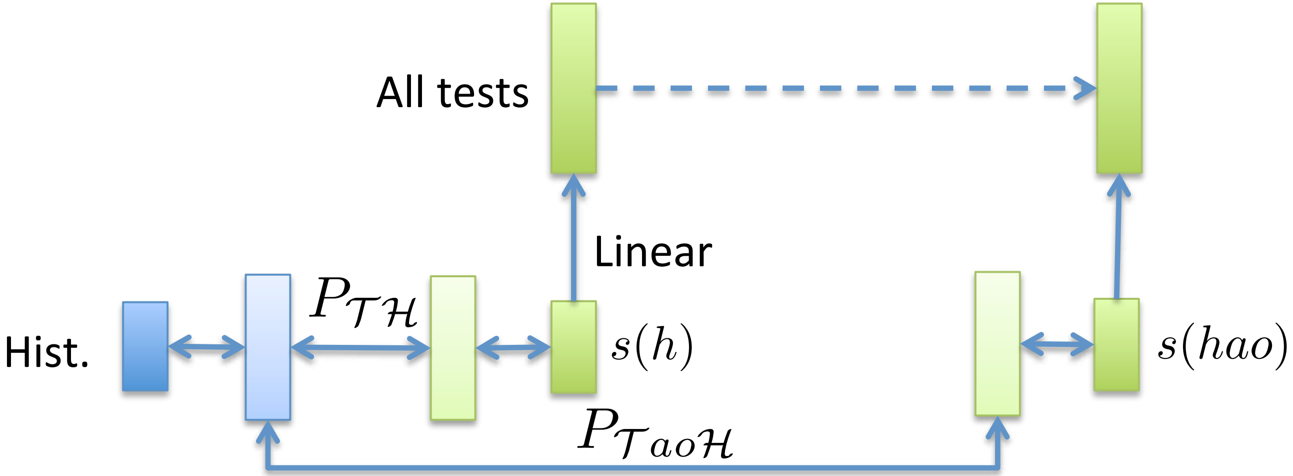

$$Zs(hao) \propto ZM_{ao}Z^{-1}s(h)$$
$$= B_{ao}\beta(h)$$

State Update Unrolled

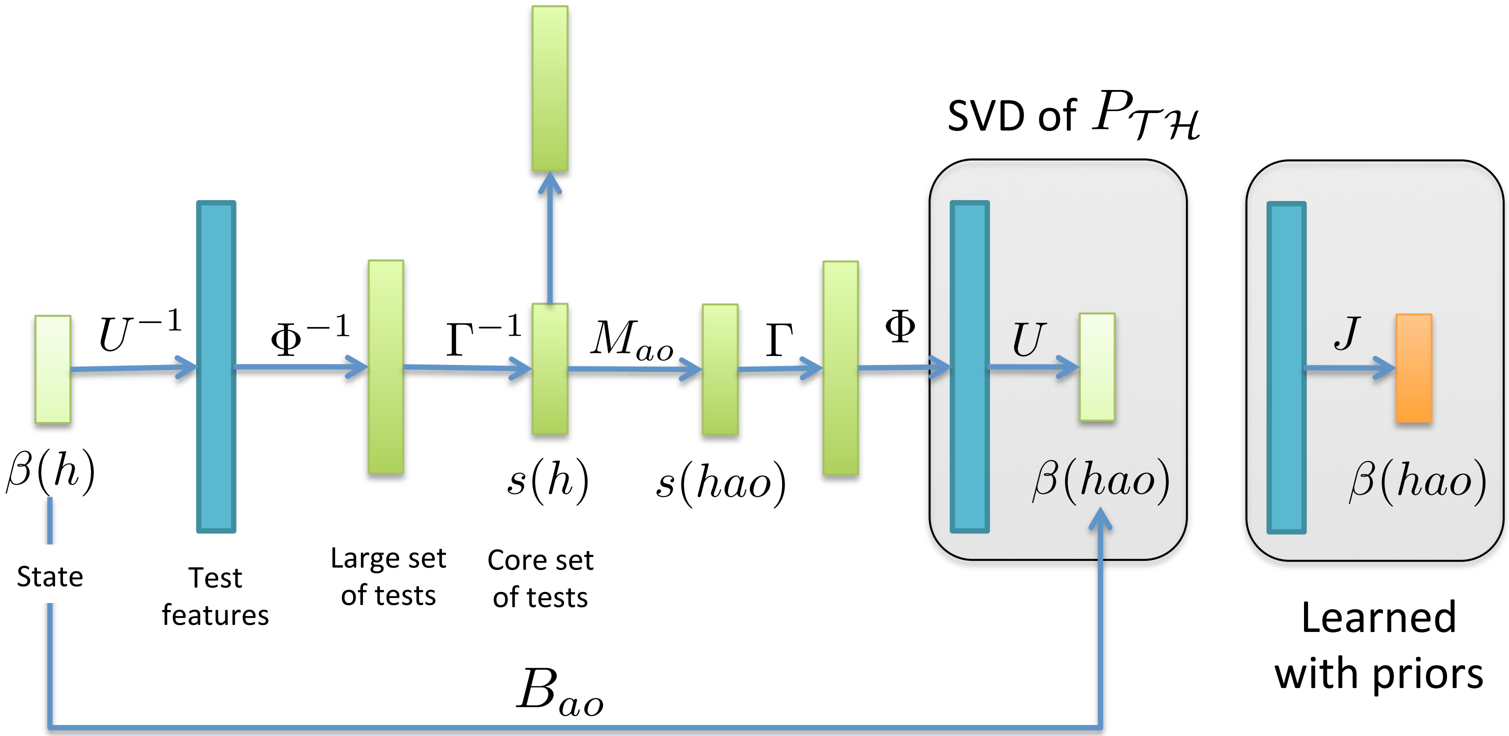

Priors, where do they come from?

- Semantic for planning
  - Simplicity
  - Temporal coherence
  - Label consistency

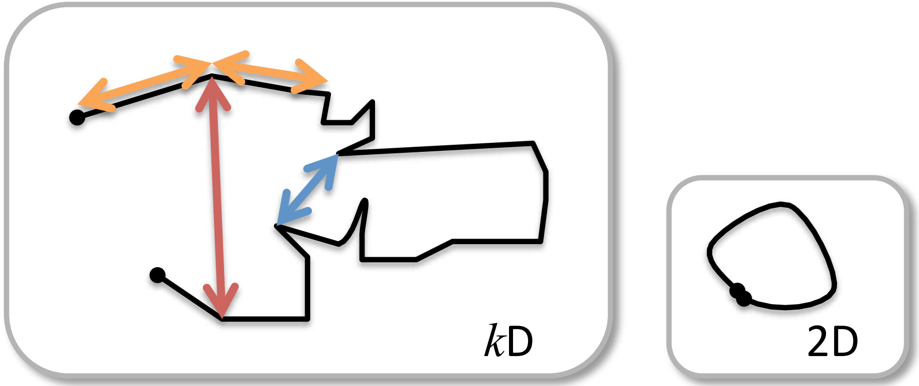

- Loss-functions
- Optimization problem
- Parameter $J = AU^{\intercal}$

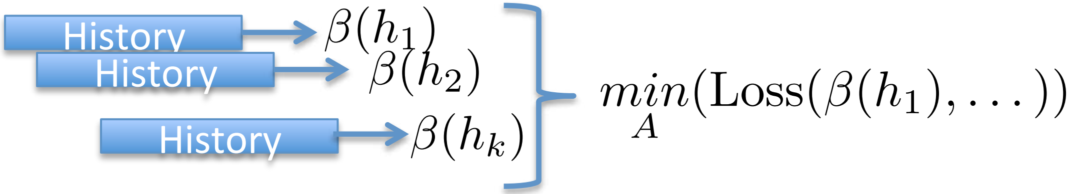

Simplicity

- Occam's razor
- Generalization
- Orthogonal concepts?

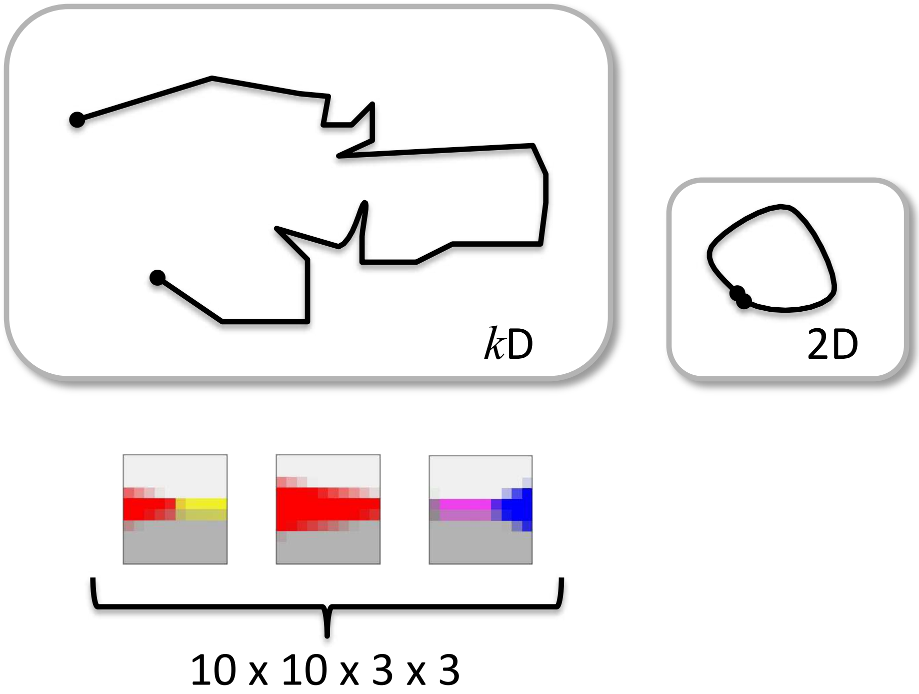

Temporal Coherence

- Slowness
- Smoothness
- Gradual change
- Inertia

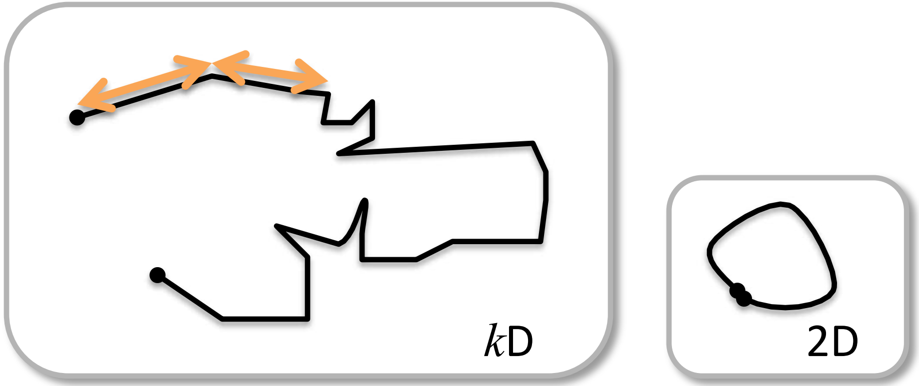

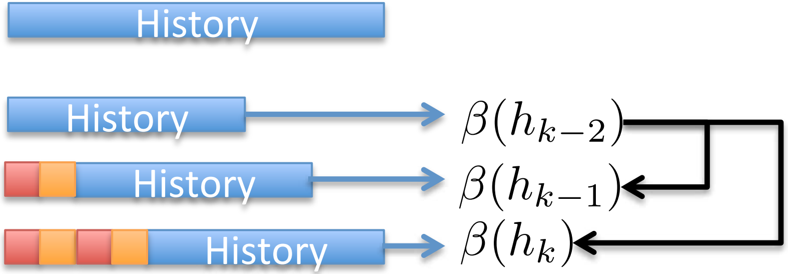

Label Consistency

- Semantic for planning
- Structure
- Labels
- Metric/ distance measure

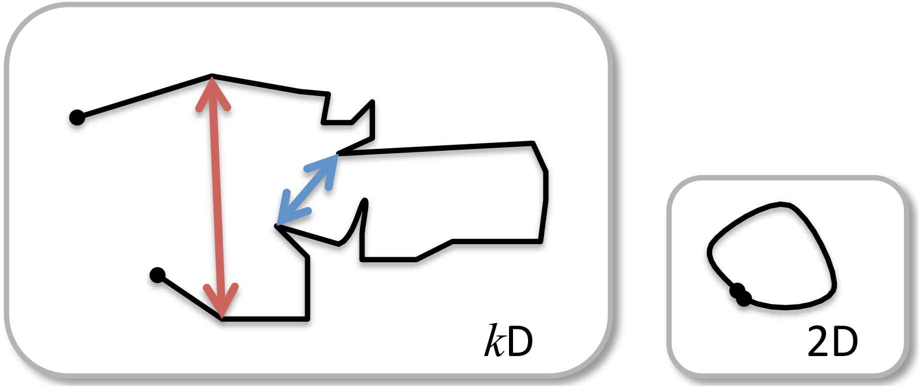

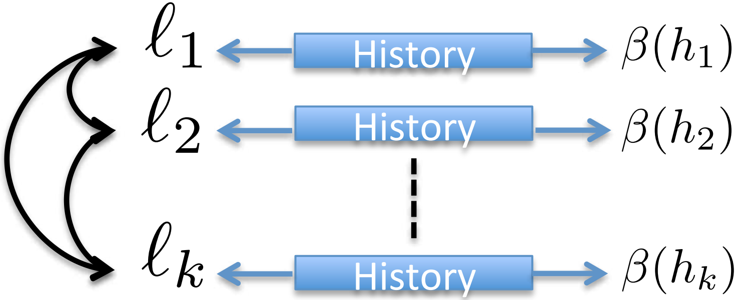

Results: Embedding

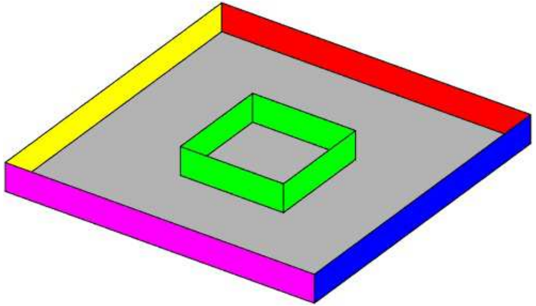

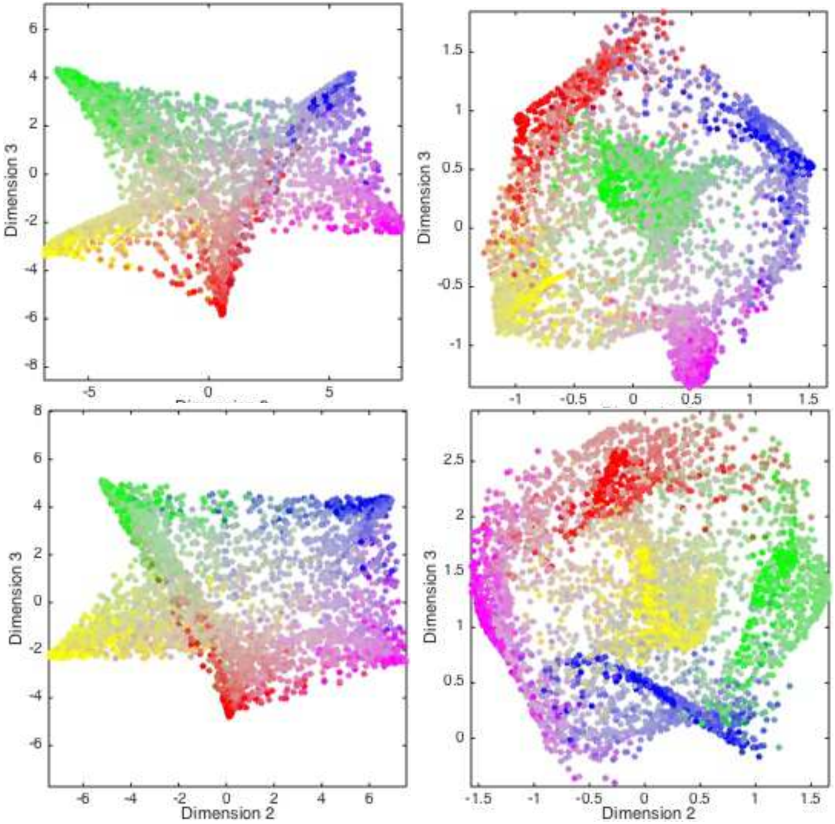

Results: Embedding

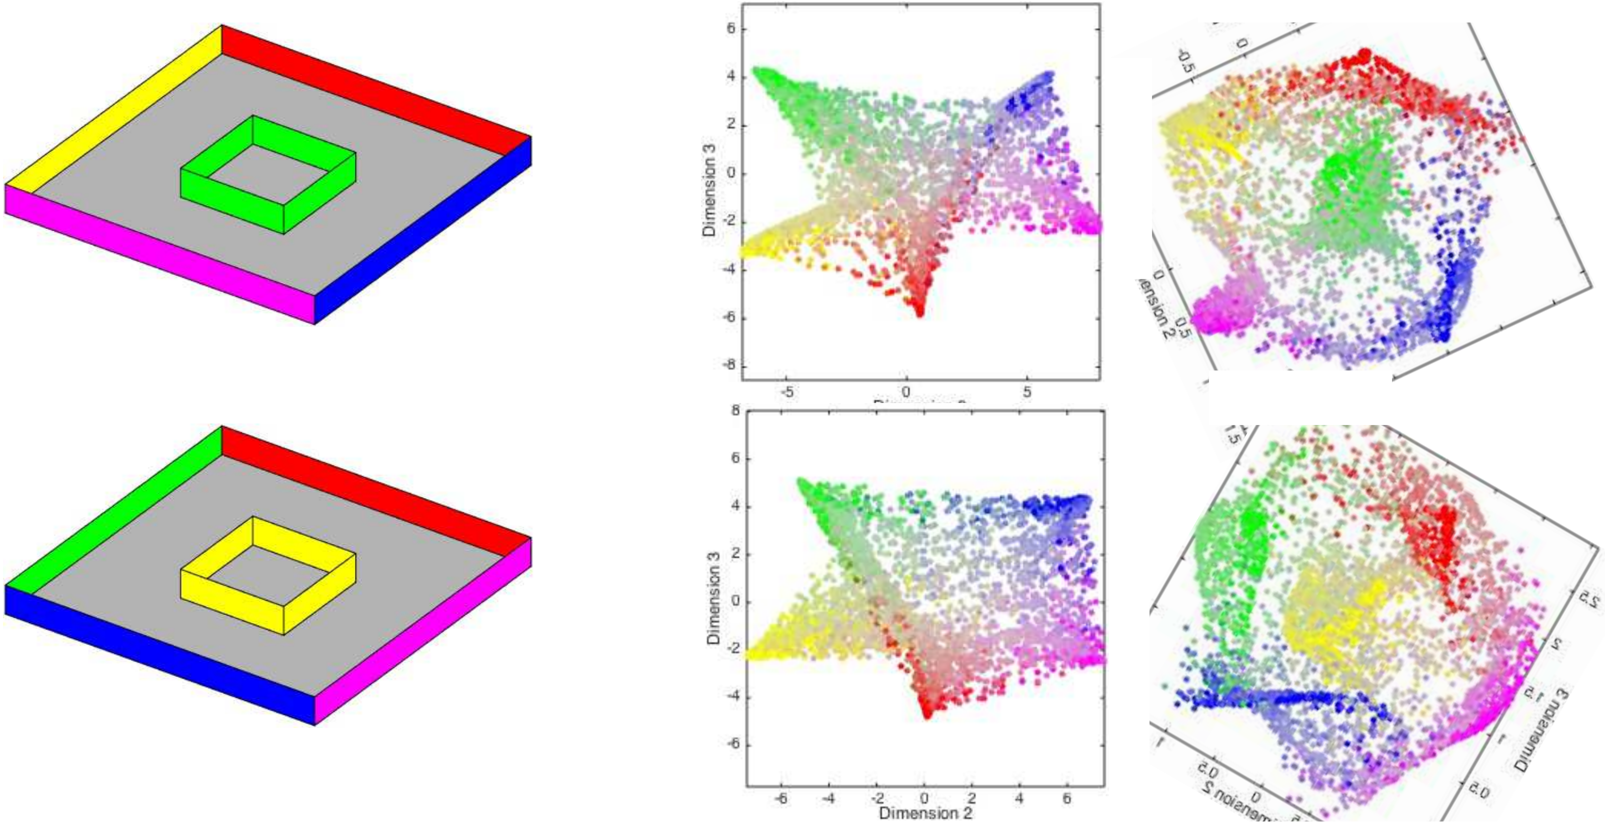

Results: Actions

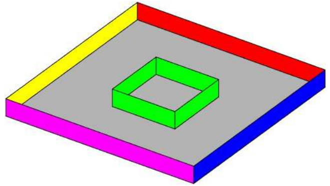

Spectral

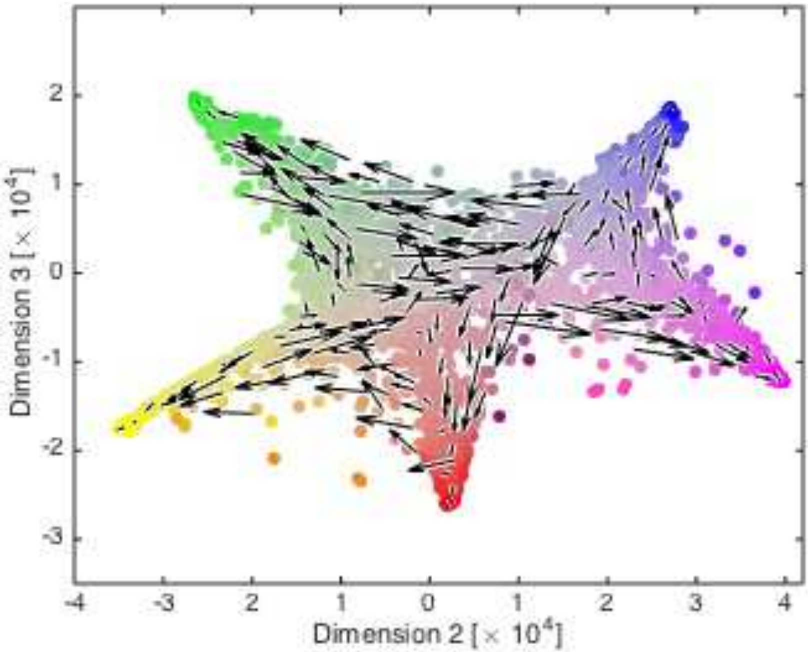

Learned with priors

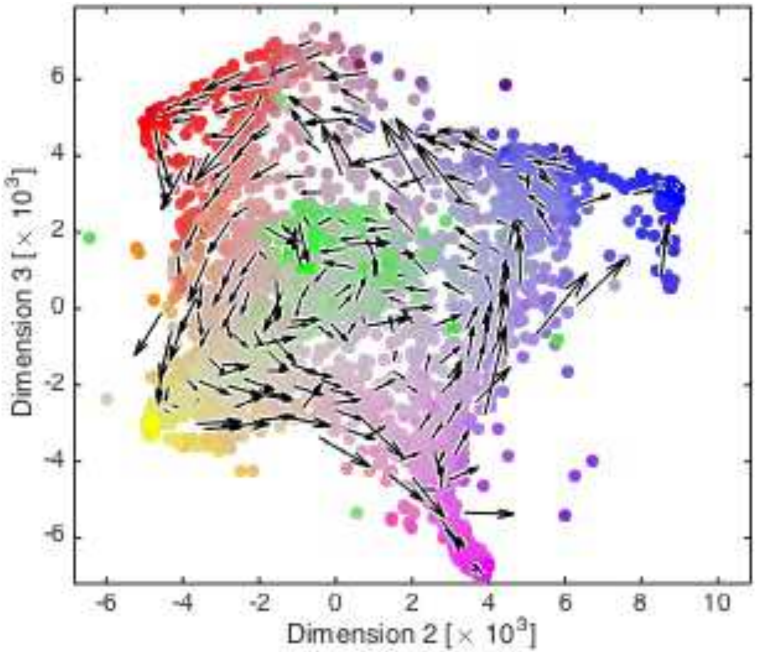

Action: Forward + Turn Left

Results: Path

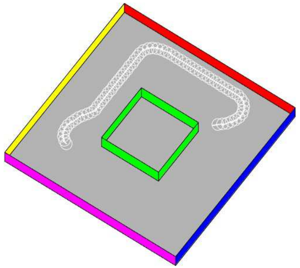

Spectral

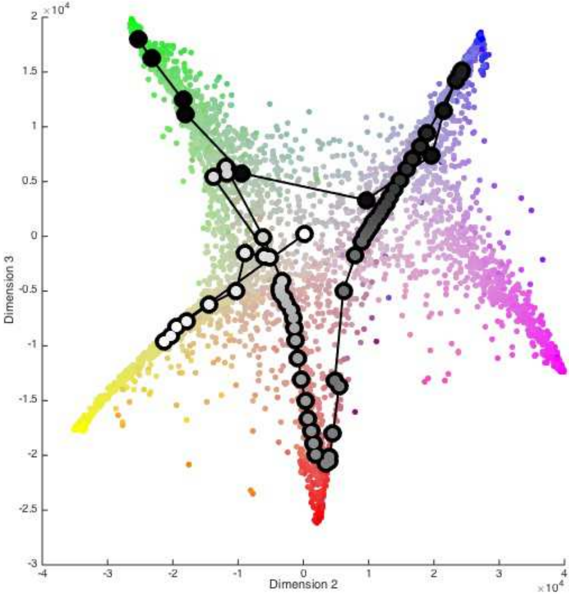

Learned with priors

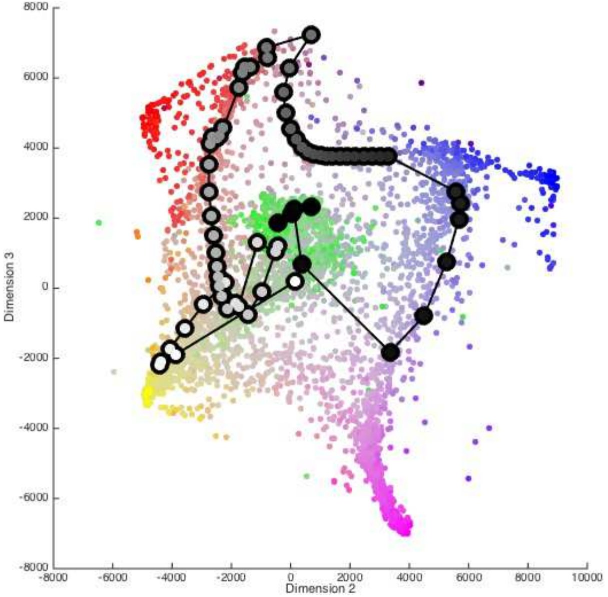

Action: Forward + Turn Left

Results: Planning

- Policy learned
- Random start
- Same goal
- Length of path

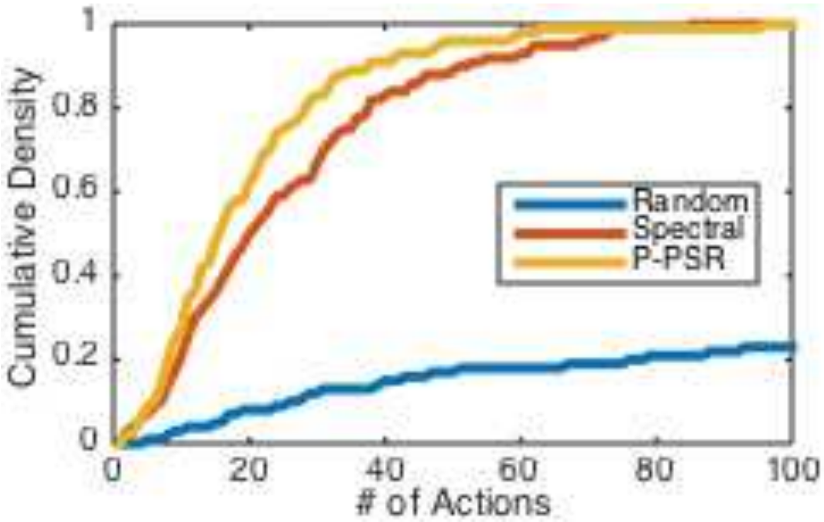

Conclusion

- Representation learning
- Predictive state representation
- Semantics
- Priors