Learning to do as planned

Marc Toussaint

Intelligent Systems Lab, TU Berlin
Max Planck Institute for Intelligent Systems, Stuttgart

IROS 2020 Workshop on New Horizons for Robot Learning, Oct 29, 2020

Why it is worth to (at least try to)...

Why it is worth to (at least try to)...

understand how physical reasoning & manipulation planning works...

Why it is worth to (at least try to)...

understand how physical reasoning & manipulation planning works...

and based on that continue thinking about learning.

Why work on model-based planning?

Why work on model-based planning?

- Arguments against:

Why work on model-based planning?

- Arguments against:
  - It'll never work!

Why work on model-based planning?

- Arguments against:
  - It'll never work!
  - All these model-based assumptions are insane!

Why work on model-based planning?

- Arguments against:
  - It'll never work!
  - All these model-based assumptions are insane!
  - Perception will never "work"! Never provide precise objects, shapes, etc

Why work on model-based planning?

- Arguments against:
  - It'll never work!
  - All these model-based assumptions are insane!
  - Perception will never “work”! Never provide precise objects, shapes, etc
  - Precise tracking will never work!

Why work on model-based planning?

- Arguments against:
  - It'll never work!
  - All these model-based assumptions are insane!
  - Perception will never "work"! Never provide precise objects, shapes, etc
  - Precise tracking will never work!
  - Partial observability!

Why work on model-based planning?

- Arguments against:
  - It'll never work!
  - All these model-based assumptions are insane!
  - Perception will never "work"!   Never provide precise objects, shapes, etc
  - Precise tracking will never work!
  - Partial observability!
  - and general POMDPs will never work anyway!

Why work on model-based planning?

- But:

Why work on model-based planning?

- But: *Generalization!*
  - The generalization power of model-based reasoning is amazing!
  - Daily work with planning algorithms: Get them to work for one instance, suddenly they work for huge ensembles of related settings!
  - See work on physical reasoning & sequential manipulation planning

Toussaint, Allen, Smith, Tenenbaum: *Differentiable Physics and Stable Modes for Tool-Use and Manipulation Planning*. R:SS'18

Toussaint, Ha, Driess: *Describing Physics For Physical Reasoning: Force-based Sequential Manipulation Planning*. RAL/IROS 2020

(Also work by many other TAMP researchers: Leslie Kaelbling, Tomás Lozano-Pérez, Caelan Garrett, Neil Dantam, etc etc)

Why work on model-based planning?

- But: *Generalization!*
  - The generalization power of model-based reasoning is amazing!
  - Daily work with planning algorithms: Get them to work for one instance, suddenly they work for huge ensembles of related settings!
  - See work on physical reasoning & sequential manipulation planning

  Toussaint, Allen, Smith, Tenenbaum: *Differentiable Physics and Stable Modes for Tool-Use and Manipulation Planning*. R:SS'18

  Toussaint, Ha, Driess: *Describing Physics For Physical Reasoning: Force-based Sequential Manipulation Planning*. RAL/IROS 2020

  (Also work by many other TAMP researchers: Leslie Kaelbling, Tomás Lozano-Pérez, Caelan Garrett, Neil Dantam, etc etc)

- I believe in *that kind of generalization,* leveraging geometric and physical computation/modelling/priors/understanding

- The question is:

  How realize this kind of generalization in a real, noisy, imperfect, partially observable, hardly perceivable, unexpectedly behaving world, where "literal state estimation" might never work.

Learning to Do as Planned

- Transfer from model-based reasoning to robust, vision-based, stationary reactive policies
  - Cf. guided policy search (Levine, Koltun, ICML'13)
  - input re-mapping

- Leverage model-based planning methods to **generate data**

Examples from different projects

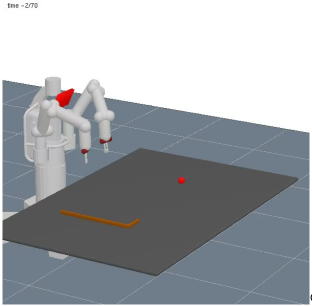

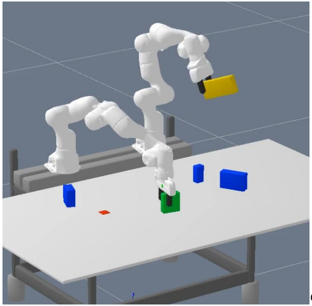

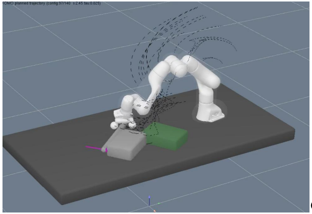

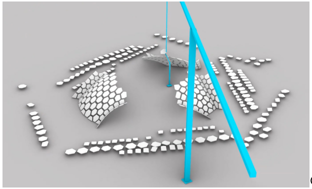

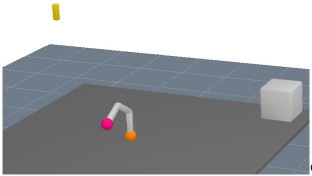

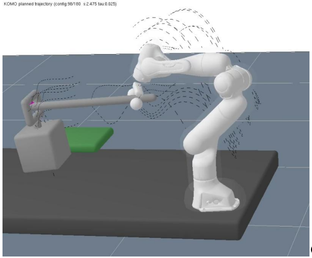

Learning to directly predict and control full solutions!

Driess, Ha, Toussaint: *Deep Visual Reasoning: Learning to Predict Action Sequences for Task and Motion Planning from an Initial Scene Image*. RSS 2020

Driess, Ha, Toussaint: *Learning Geometric Reasoning and Control for Long-Horizon Tasks from Visual Input*. ICRA SUBMISSION 2021

Danny Driess

Jung-Su Ha

- From a **visual** scene encoding, directly predict a solution!
- ...estimate a very strong **heuristic** over discrete decisions that **generalizes** "1st order" (across objects)
- predict parameters of **vision-based reactive control** to execute

Learning to predict skeletons from scenes

- Raw data $D = \{\big(S^i, g^i, a^i_{1:K^i}, F^i\big)\}_{i=1}^n$ with scene S^i, goal g^i, skeleton $a^i_{1:K^i}$, feasibility F^i
- **Sequence training data** $\mathcal{D} = \{(S^i, g^i, a^i_{1:K^i}, f^i)\}_{i=1}^n$ with $f^i = f^i_{1:K}$:

$$f^i_j = \begin{cases} 1 & F^i = 1 \\ 1 & \exists \left(S^l, a^l_{1:K^l}, g^l, F^l\right) \in D \quad \text{s.t.} \quad F^l = 1,\ g^l = g^i,\ a^l_{1:j} = a^i_{1:j} \\ 0 & \text{else} \end{cases}$$

Learning to predict skeletons from scenes

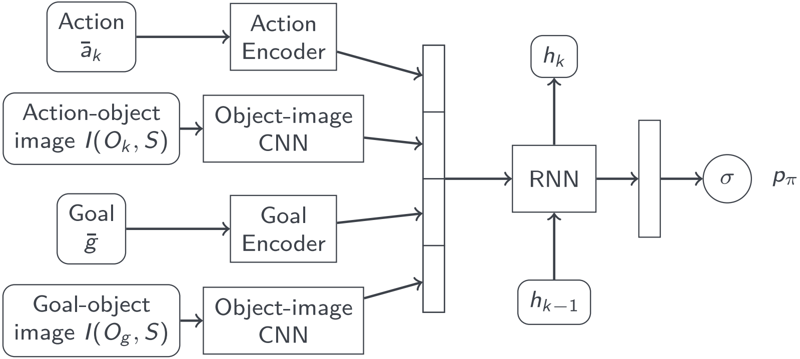

$$
\begin{aligned}
(p_\pi, h_k) &= \pi_{\mathrm{NN}}\big(\bar{a}_k, I(O_k, S), \bar{g}, I(O_g, S), h_{k-1}\big) \\
&= \pi\big(a_k, g, a_{1:k-1}, S\big)
\end{aligned}
$$

- Separate encoding of predicates $\bar{a}, \bar{g}$ and references O (as masks)

Learning to predict skeletons from scenes

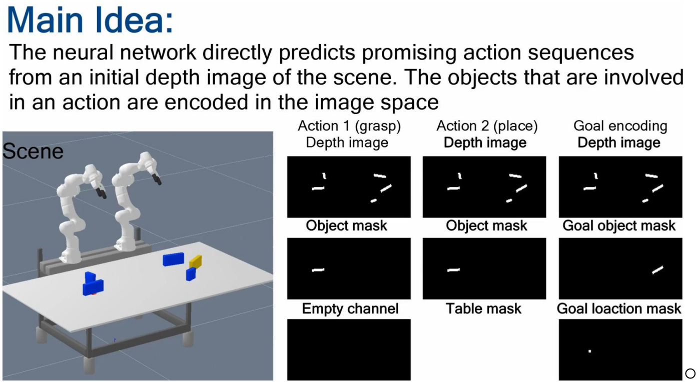

Driess, Ha, Toussaint: *Deep Visual Reasoning: Learning to Predict Action Sequences for Task and Motion Planning from an Initial Scene Image*. RSS 2020

Learning also vision-based execution control

- Use the model-based **MPC** solver to generate data for stationary reactive control of solutions
- Train a vision-based network to imitate that control behavior
  - Assume **funnel** policies where $\dot{e} = -e$; the NN defines funnel variable e.
  - Also predict model-based cost-to-go estimate

Driess, Ha, Toussaint: *Learning Geometric Reasoning and Control for Long-Horizon Tasks from Visual Input*. ICRA SUBMISSION 2021

Learning also vision-based execution control

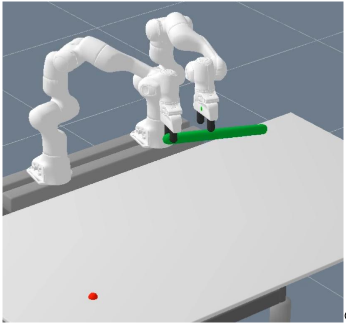

Driess, Ha, Toussaint: *Learning Geometric Reasoning and Control for Long-Horizon Tasks from Visual Input*. ICRA SUBMISSION 2021

Discussion

- “New Horizons for Robot Learning”
  - Aim for *that kind of generalization,* leveraging geometric and physical computation/modelling/priors/understanding
  - Reasoning as a core ingredient to robot learning

Thanks

- *for your attention!*

- *Deep Visual Reasoning:* (RSS '20)
- *Learning Geometric Reasoning and Control...* (ICRA'21 SUBMISSION)

Danny Driess

Jung-Su Ha

- Funding: This work was supported by the **Max Planck Fellowship** (MPI for Intelligent Systems, Stuttgart), and the **IMPRS** (MPI).