
Concept Learning For Robot Intelligence

Justus Piater, University of Innsbruck

Abstract

Current robot learning, and machine learning in general, requires carefully-engineered setups (environments, objective functions, training data, etc.) for learning to succeed. Perception and action spaces are specially crafted to meet the requirements of the learning objective, which is specified in advance. How can we construct robot learning systems that can learn in an open-ended fashion, acquire skills not foreseen by its designers, and scale up to virtually unlimited levels of complexity? I argue that a key to achieving this lies in the robot's ability to learn abstract concepts that can be reused as a basis for future learning, both in autonomous exploration and for teaching by humans.

Table of Contents

1. Preface

1.1. Classical Machine Learning

[Image[1] from New Scientist[2]]

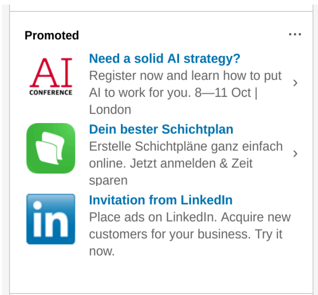

[From my LinkedIn page]

[Source[3]]

[Image[4] from The Next Web[5]]

---

[1] https://d1o50x50snmhul.cloudfront.net/wp-content/uploads/2017/05/23120156/rexfeatures_8828108ac1.jpg
[2] https://www.newscientist.com/article/2132086-deepminds-ai-beats-worlds-best-go-player-in-latest-face-off/
[3] http://5.imimg.com/data5/YN/EU/MY-54329049/face-detection-and-recognition-500x500.jpg
[4] https://cdn0.tnwcdn.com/wp-content/blogs.dir/1/files/2016/09/SwiftKey-neural-networks-hed-796x398.jpg
[5] https://thenextweb.com/apps/2016/09/16/swiftkey-improves-its-android-keyboard-predictions-with-neural-networks/

1.2. Autonomous Airshow

Apprenticeship Learning Via Inverse Reinforcement Learning [Abbeel et al. 2010]

1.3. How to learn such skills?

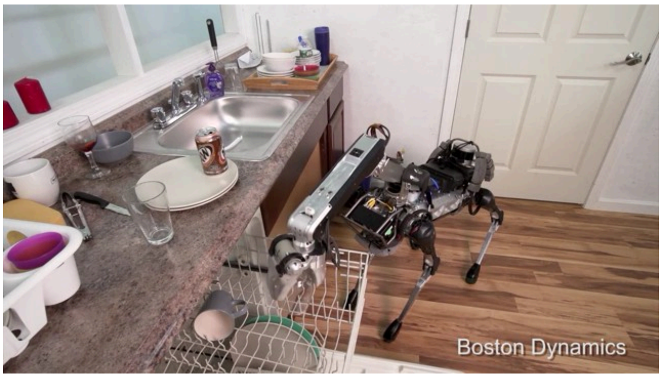

[Excerpt from Youtube[6]]

1.4. Machine Learning Paradigms

- **Supervised Learning:** Given a training set D containing N training examples (x_i, y_i), predict $y = f(x;\theta)$ for $x \notin D$.
- **Unsupervised Learning:** density estimation, clustering, data mining, dimensionality reduction
- **Reinforcement Learning:** Learn a policy π such that, at each time t, taking action $a_t = \pi(s_i)$ maximizes the expected sum of future rewards.
- **Evolutionary Learning:** optimization by stochastic alteration and recombination of parameter vector segments, guided by heuristics
- **Explanation-based Learning:** Given a domain theory (e.g., logical assertions), derive new rules, guided by training examples
- …

[6] https://www.youtube.com/watch?v=tf7IEVTDjng

1.5. Machine Learning vs. Machine Intelligence

In all cases, machine learning requires *specific, externally-provided* problems, defined in terms of objective functions *fixed a priori*.

- Can we express all we want our robot to do in terms of a single objective function?
- Can we express human behavior in terms of a single objective function?

  In evolutionary terms: survival of the species?

  Is there a difference?
- Even if we can, is it useful?

  The learning problem is *massive*.

  We don't have millions of years to evolve and train capable robots.

1.6. Artificial Intelligence

Formerly: The quest for artificial systems with human-level cognitive capabilities.

Today: The quest for solutions to problems that cannot be easily solved by hand.

Tongue in cheek:

In Practice: An unsuccessful meta-science that spawns successful scientific disciplines (symbolic planning, automated theorem proving, computer vision, machine learning, data mining, ...)

- Once we understand how to solve a problem, it is no longer considered to require intelligence.
- Thus, AI never gets credit for its achievements.

Important

Open-Endedness

1.7. Systems of Abstract Concepts

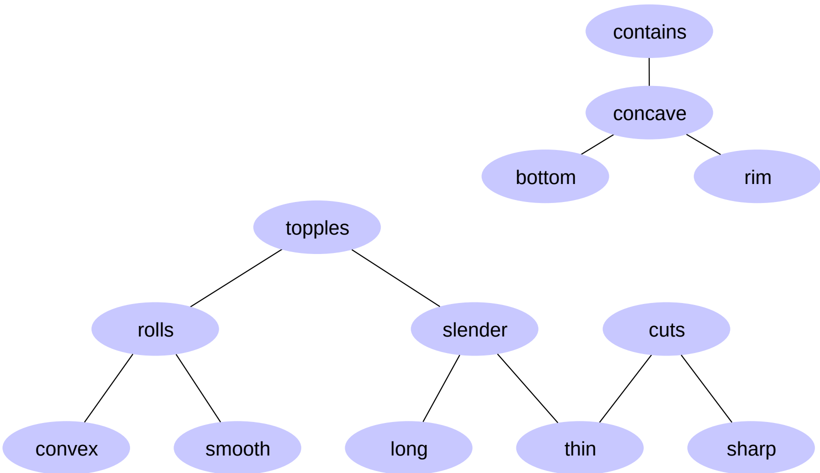

2. Skill Learning Using Stacked Classifiers

Emre Ugur

2.1. Learning About Objects

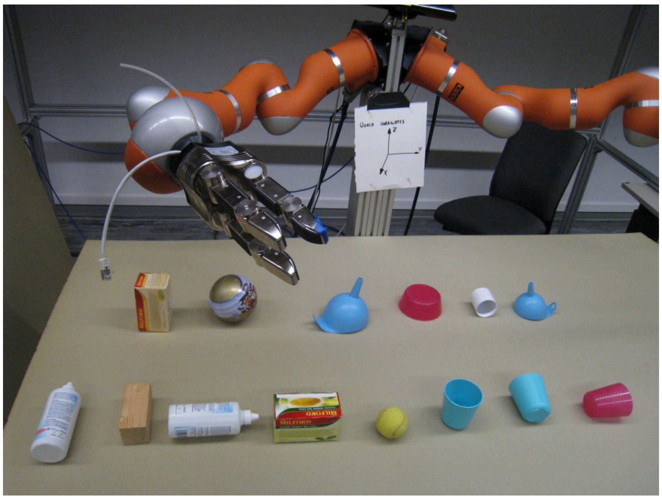

[Ugur et al. 2014]

2.2. Sensorimotor Exploration: Poking

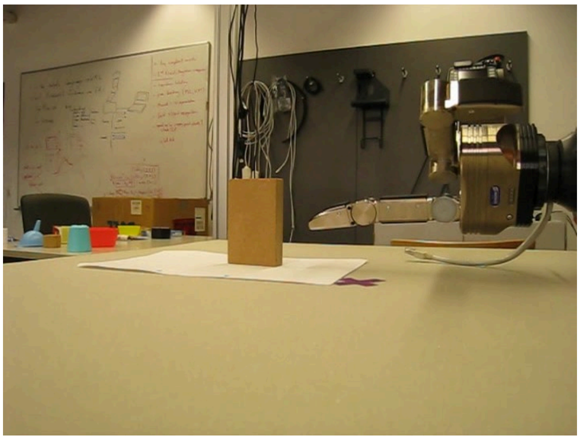

Simple Konzept: How does *one object behave* under a manipulation?

2.3. Sensorimotor Exploration: Stacking

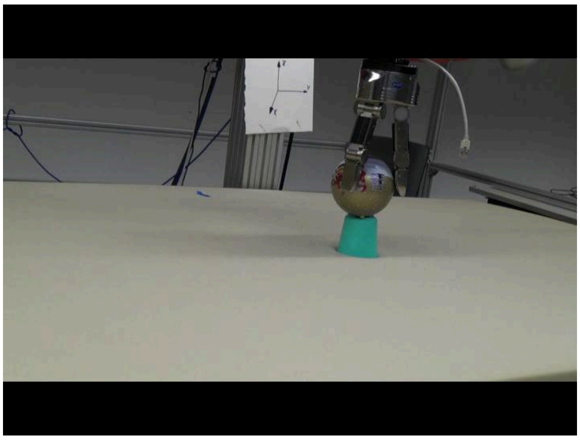

More Complex Concept: How do *two objects interact* under a manipulation?

2.4. Self-Organized Learning: 0 Objects

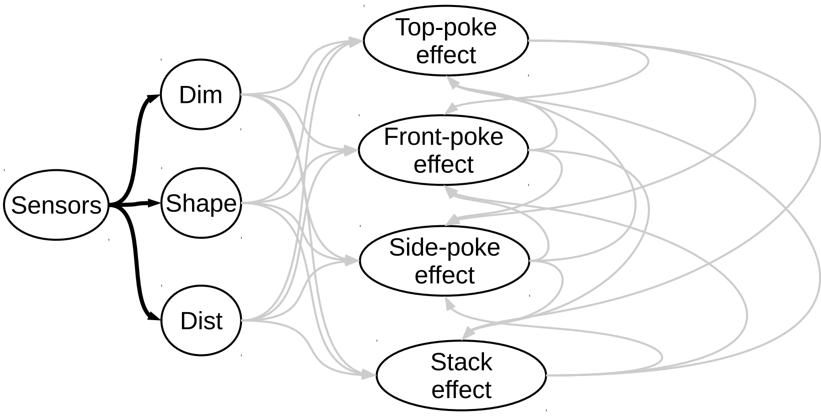

[Ugur and Piater 2014]

2.5. Self-Organized Learning: Actions

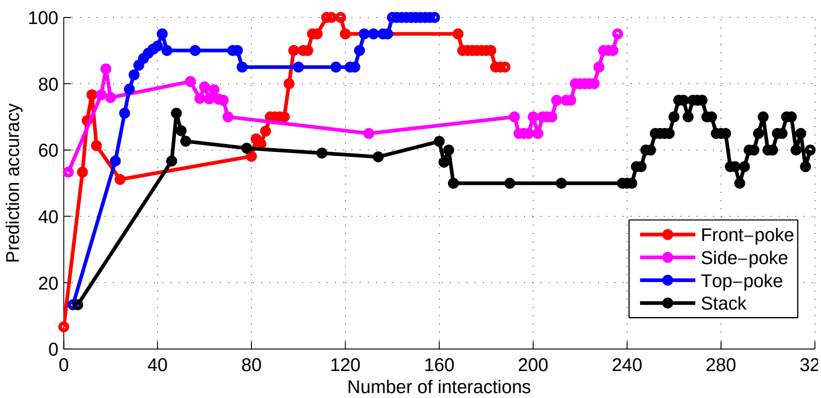

First, the robot chooses simple actions whose results are easy to predict, before focusing on stacking actions.

[Ugur and Piater 2014]

2.6. Self-Organized Learning: 10 Objects

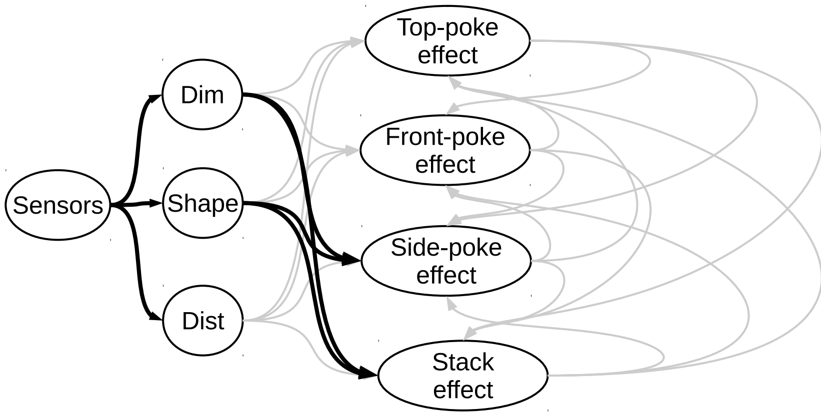

[Ugur and Piater 2014]

2.7. Self-Organized Learning: 20 Objects

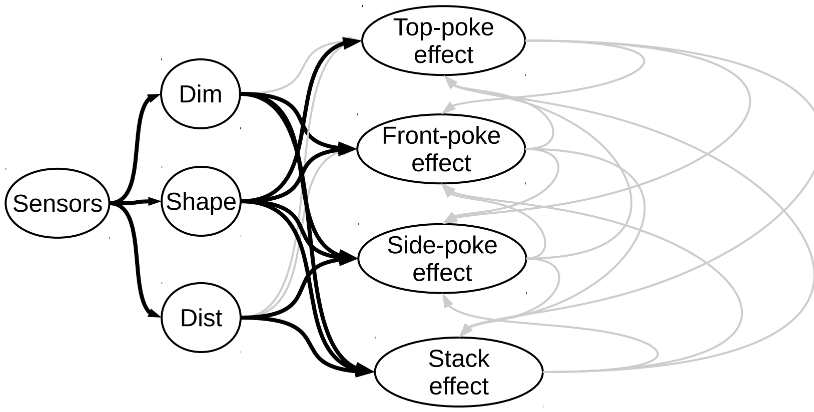

[Ugur and Piater 2014]

2.8. Self-Organized Learning: 80 Objects

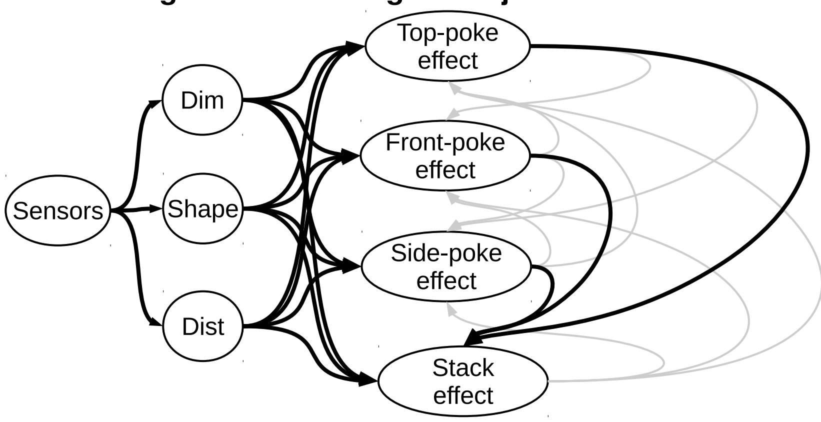

[Ugur and Piater 2014]

2.9. From Sensorimotor Interaction to Symbolic Planning

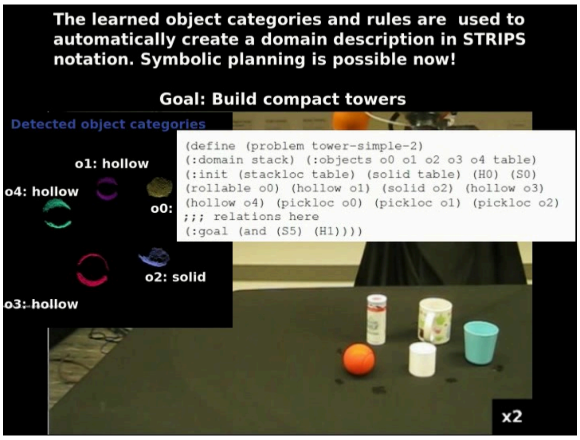

2.10. Learning about the top of the stack

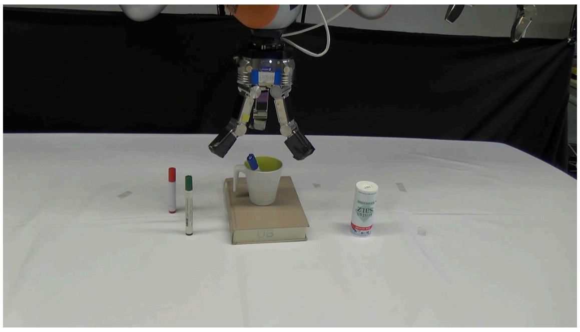

2.11. Playing With Building Blocks!

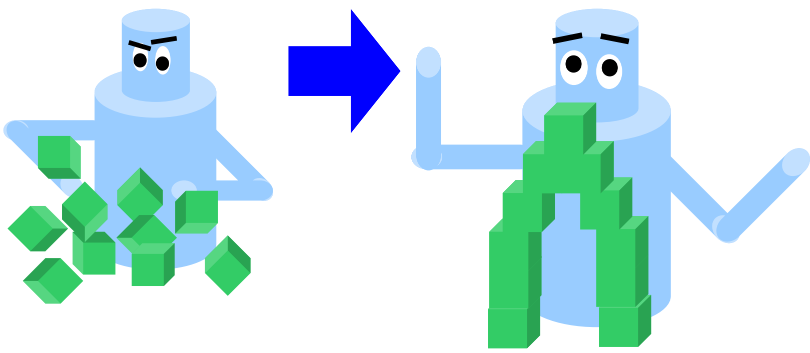

- Symbol formation by sensorimotor interaction

  Emre's prior work among the pioneers

- Hierarchical concept learning

3. Skill Learning Using Projective Simulation

Simon Hangl

3.1. Projective Simulation

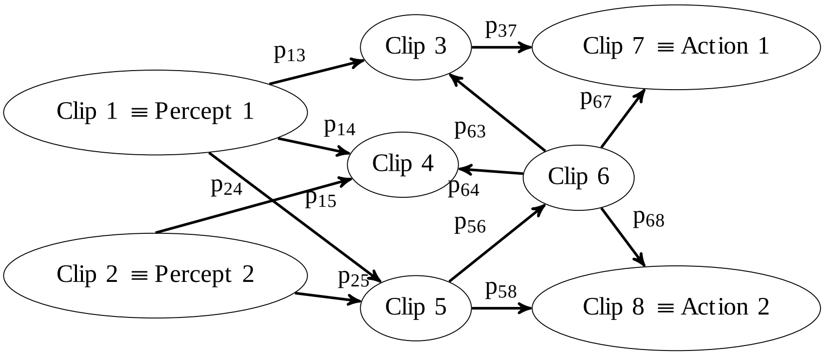

- *Episodic Compositional Memory (ECM)*: Markov network plus machinery for learning transition probabilities from experience
- *Clip*: elementary piece of experience
- *Learning*: random walk with ECM parameters updated according to rewards
- *Execution*: random walk

3.2. Picking and Placing Books: ECM

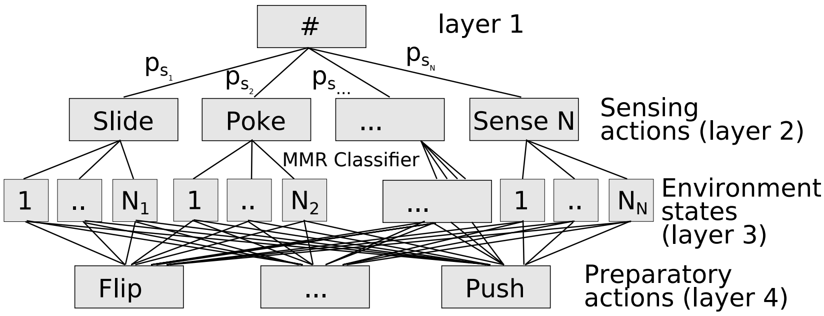

3.3. Picking and Placing Books: Resulting Behavior

[Hangl et al. 2016]

3.4. Picking and Placing Books: Learned ECM

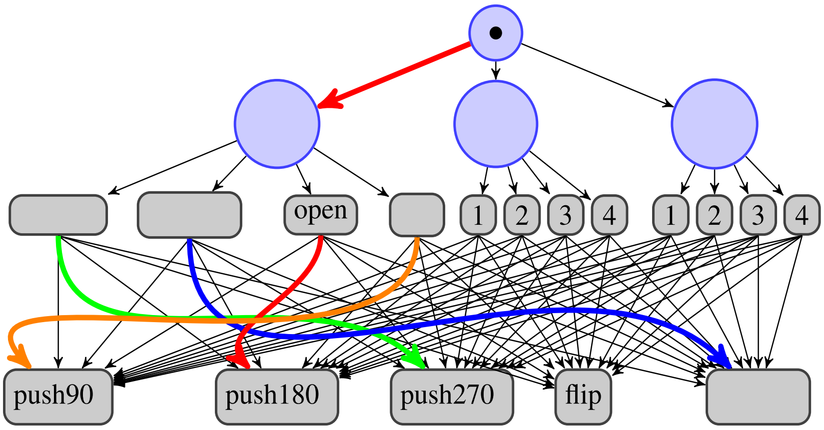

3.5. Enhancement: Environment Model

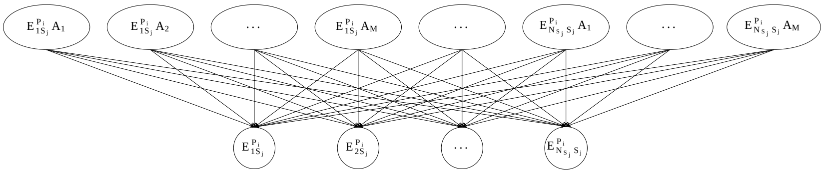

3.6. Active Learning and Skill Creation

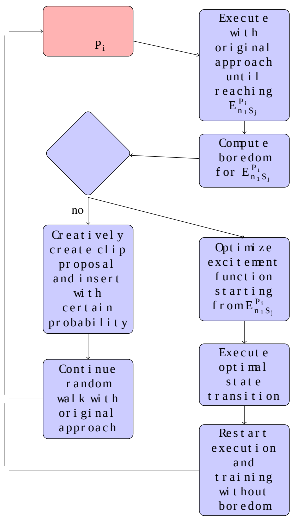

Active Learning

- Entropy $H(A \mid E_k)$ is low, i.e., the agent knows what to do in state E_k.
- Find a higher-entropy state, and determine, using the Environment Model, an action to transition to it.

Creative Skill Creation

- Using the Environment Model, synthesize a new compound preparatory skill clip, and add it to the ECM.
- Akin to compiling cognitively-controlled (cortical) complex skills into automated (cerebellar) routines.

3.7. Simulated Experiments

Without Active Learning | With Active Learning

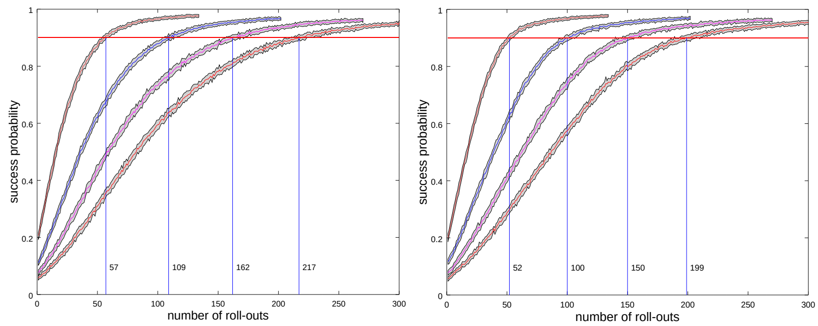

- **Skills:** void, flip, rotate 90°, rotate 180°, rotate 270°

  From left to right, each curve adds 5 distractor skills.
- Skill success rate of 95%
- Means and standard deviations of 1000 simulated robots

3.8. Simulated Experiments

With Active Learning and Creativity

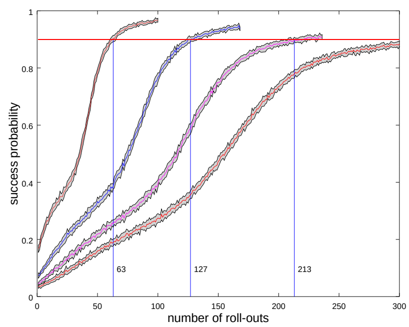

- **Initial Skills:** void, flip, rotate 90°

  From left to right, each curve adds 5 distractor skills.
- Skill success rate of 95%
- Means and standard deviations of 1000 simulated robots

3.9. Learning Complex Manipulation Sequences!

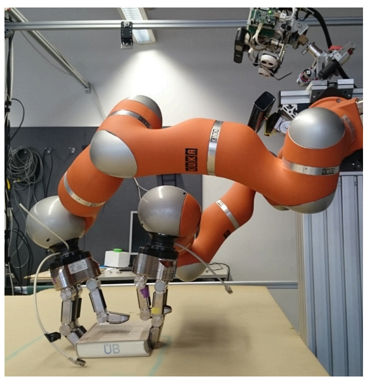

Complex skill learning by sequencing actions

- with unified learning and execution,
- guided by reinforcement signals,
- adaptively creating compound skills.

4. Search Space Management

How can we scale this up?

4.1. Open-Ended Learning: The Ambition

The robot can learn things its designer did not specifically foresee.

Example: Teach my dishwasher-un/loading robot to unscrew jam jar caps.

The robot can learn a quasi-infinity of things.

Gianluca Baldassarre's *difficult goal space*

> ***Note***
>
> **Q:** Why not just *program* or *train* the robot to do its new job?
>
> **A:** Too complicated and time consuming; teaching should be quick and effortless.
>
> Olivier Sigaud nicely motivated why the problem cannot easily be solved by neural networks.

Pragmatic standpoint driven by utility to humans.

4.2. Open-Ended Learning: The Challenge

In open-ended learning in the real world, both perception and action spaces are necessarily very large.

Any perceptual detail; any motor pattern may become important in a novel task!

> ***Note***
>
> To be solvable, a learning problem requires
>
> - either a strong bias,
> - or a large number of training instances.

Corollary of the *No Free Lunch* theorems [Wolpert 1996]

The only way out: *Structure difficult learning problmes into a (partially-ordered) set of **simple** learning problems!*

Search Space Management

4.3. How to Build an Open-Ended Learner

Scientific Agenda:

1. Conventional Machine Learner
2. Learn Reusable Concepts
3. Shape the Hypothesis Space
4. Choose Learning Goals

Close this loop → autonomous learning strategy

4.4. Conventional Machine Learner

- Can learn one task at a time.

- Generalization capabilities critically depend on the ***representations*** of its inputs and outputs.

4.5. Learn Reusable Concepts

Learn perceptual *(features)* or motor *(options)* ***concepts*** useful for future learning problems

Examples:

- *Representation Learning*

  representational redescription

  (possibly during sleep), as noted by Stéphane Doncieux and Kevin O'Regan [Karmiloff-Smith, 1992]

- *Parametrized Skills*

- Hierarchical Learning

  Curiosity-Driven Discovery of Tool Use [Forestier & Oudeyer, arXiv 2017]: forward-chaining of concepts

- *Transfer Learning* relies on concepts shared by the source and target tasks

  Jochen Triesch

  > ***Note***
  >
  > during the IMOL 2017 panel discussion

  Example: *keep salad in bowl → keep water in glass* relies on

  - *containment* (matter kept in place by surrounding compartment)

  - *upright* (normal opposite to gravity vector)

  These (reusable) concepts should also generally be learned.

  What about folding laundry? Requires completely different features/ actions/concepts.

Objective:

- Add the new concepts to the perception/action repertoire
- Learn constraints and biases (*scaffolding*)

to *reduce the difficulty* of subsequent learning problems.

4.6. Shape the Hypothesis Space

Begin with a reduced perception-action space or augmented learning bias; relax these constraints as needed.

Examples:

- Maturation

  - Gradual reduction of learning bias
    - body height and weight for walking
    - brain plasticity
    - proximodistal exploration of body motor control

- Existing structure in the physical world
  - Saliency
  - Spatial coherence; physical contact

    Kevin O'Regan's toys attached to rakes

- Teaching
  - Explicitly point out informative features or effective actions

4.7. Choose Learning Goals

Choose new learning problems expected to yield reusable concepts.

Examples:

- autonomously (*Artificial Curiosity*)
  - Pierre-Yves Oudeyer's *Intrinsically-Motivated Goal Exploration Processes*
  - highly powerful: *goal babbling* (Matthias Rolf) and *empowerment* (Daniel Polani)

- parenting

4.8. And More...

- Meta-Learning – learn the above four skills

- Form associations from passive observation
  - powerful way of adding structure to the perceptual space
  - can be done without attention, activity, supervision, goals, etc.
  - totally underused in robotics

- Knowledge Mining

5. Conclusions

5.1. Conclusions

- All practical Machine Learning systems are designed and/or trained for *given tasks*; higher levels of open-ended learning are not yet within reach.

- Open-Ended Learning requires learning of ***reusable concepts***.

- This will allow, and success will require, a combination of multiple paradigms, including *autonomous, exploratory learning* (hierarchical/ transfer learning, artificial curiosity, …) and *teaching*.

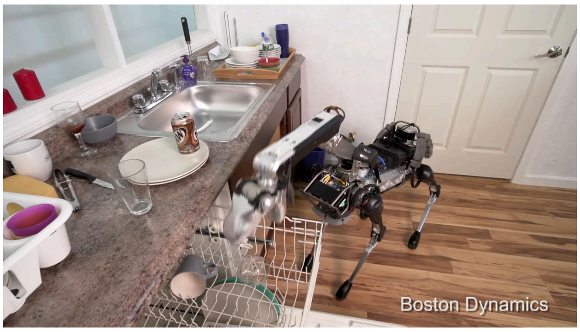

5.2. A Learning Robot

[From the 1986 movie *Short Circuit*]

5.3. References

P. Abbeel, A. Coates, A. Ng, "Autonomous Helicopter Aerobatics through Apprenticeship Learning[1]". *International Journal of Robotics Research* 29(13), pp. 1608–1639, 2010.

S. Hangl, E. Ugur, S. Szedmak, J. Piater, "Robotic Playing for Hierarchical Complex Skill Learning[2]". *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2799–2804, 2016.

E. Ugur, S. Szedmak, J. Piater, "Bootstrapping paired-object affordance learning with learned single-affordance features[3]". *The Fourth Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*, pp. 476–481, 2014.

E. Ugur, J. Piater, "Emergent Structuring of Interdependent Affordance Learning Tasks[4]". *The Fourth Joint IEEE International Conference on Development and Learning and on Epigenetic Robotics*, pp. 489–494, 2014.

D. Wolpert, "The Lack of A Priori Distinctions Between Learning Algorithms[5]". *Neural Computation* 8, pp. 1341–1390, 1996.

[1] http://dx.doi.org/10.1177/0278364910371999
[2] https://iis.uibk.ac.at/public/papers/Hangl-2016-IROS.pdf
[3] https://iis.uibk.ac.at/public/papers/Ugur-2014-ICDLEPIROB-119.pdf
[4] https://iis.uibk.ac.at/public/papers/Ugur-2014-ICDLEPIROB-97.pdf
[5] https://web.archive.org/web/20161220125415if_/http://www.zabaras.com/Courses/BayesianComputing/Papers/lack_of_a_priori_distinctions_wolpert.pdf