# Machine Learning in Robotics

## 64-450 Integrated Seminar Intelligent Robotics

Oke Martensen

University of Hamburg
Faculty of Mathematics, Informatics and Natural Sciences
Department of Informatics
**Technical Aspects of Multimodal Systems**

9. June 2016

# Outline

# Why Machine Learning?

Optimization is concerned with **mathematical problems** which are mathematically well-defined thus have verifiable solutions.

Machine Learning is

- concerned with **engineering problems** (often not well-defined)
- about building mathematical models

ML algorithms can be seen as being composed of:

1. representation
2. evaluation
3. optimization

# Machine Learning

## Types

### Supervised Learning

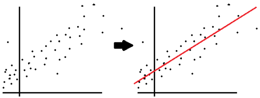

learn IO mapping

regression, SVMs, ..

### Unsupervised Learn.

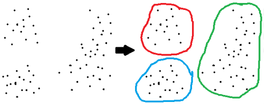

clustering

nearest neighbour, ..

### Reinforcement Learn.

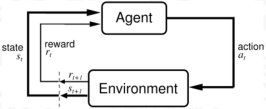


Sutton and Barto (1998)

trial and error

Q-learning, SARSA, ..

# Machine Learning

## Important Notions

Generalizability: overfitting vs. underfitting

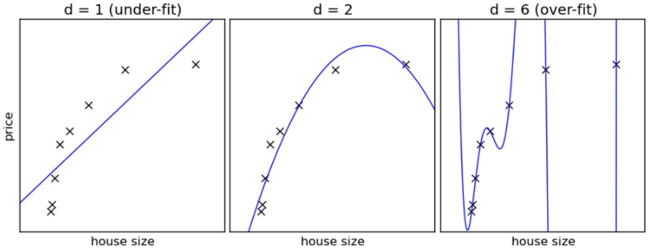


www.medium.com/@nomadic_mind/new-to-machine-learning-avoid-these-three-mistakes-73258b3848a4

Inductive bias: prior assumptions about the task at hand

No-free-lunch theorem: there is no algorithm superior for all tasks

# Robot Learning
## Problem

Again, **in contrast to optimization** (e.g. inverse kinematics).


Levine (2015)

- solve subtasks in advance e.g. via computer vision $\Rightarrow$ states
- explore, **learn** a policy through experience (RL)

# Robot Learning

## Challenges

- high-dimensional spaces
- scarce real-world data
- high variability / noise
- high-level targets
- many distinct tasks
- ...


https://www.youtube.com/watch?v=g0TaYhjpOfo

# Machine Learning in Robotics

## More Relevant ML Types

Supervised Learning

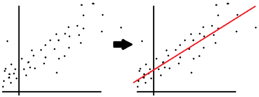

**Combinations**:

Behavior Cloning

Apprenticeship Learn.

...

Reinforcement Learn.

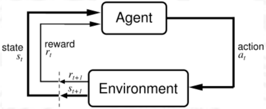


Sutton and Barto (1998)

**Problems with RL**: limited data, dimensionality, few parameters, ...

Ideas:

- initialize learning process with data of a successful execution
- interim policy evaluation by user

# Deep Learning in a Nutshell

## Nothing much new actually..

- mostly deep neural networks ($> 1$ hidden layers)
- **plethora of** old and newer **methods for tweaking**: dropout, batch normalization, data augmentation, ..
- large improvements in various domains: computer vision, speech recognition, game-playing, ..

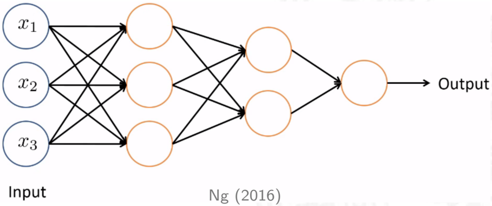


Ng (2016)

# Deep Learning in a Nutshell

## Characteristics

**learned hierarchy of features**

performance scales with **data**

heavy computations

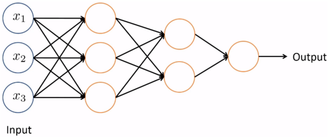


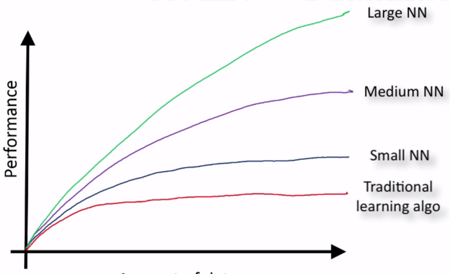


Ng (2016)

# Deep Learning Illustration

## Convolutional Neural Network (CNN/ConvNet)

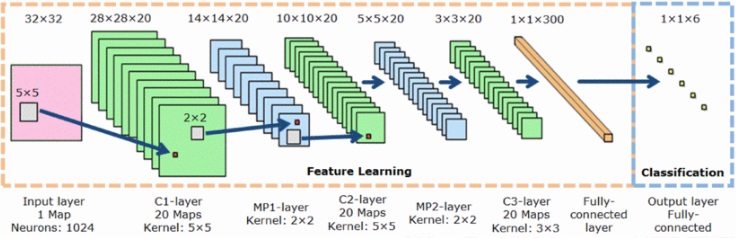


Lee et al. (2009)

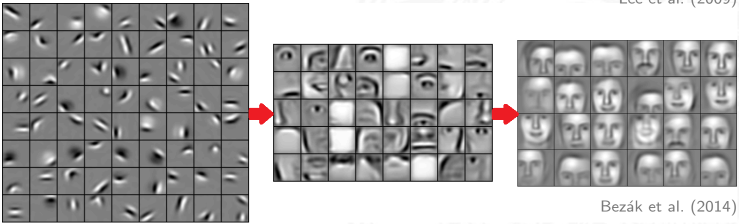

Bezák et al. (2014)

# Robot Learning

## Standard Robotic Learning vs. Deep Learning Approach

Subtasks solved with **domain-specific** approaches:

standard robotic learning

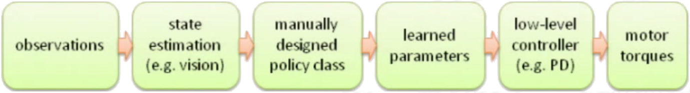


**vs.**

Largely **domain-agnostic** Deep Learning pipeline:

deep sensorimotor learning

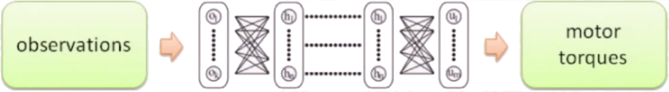


Levine (2015)

# End-to-End Training of Deep Visuomotor Policies

Levine et al. (2015)

- map image pixels & joint angles to motor torques
- guided policy search:
  - transforms policy search into SL
  - alternating between trajectory and policy optimization
- full torque control of 7-DoF robotic arms

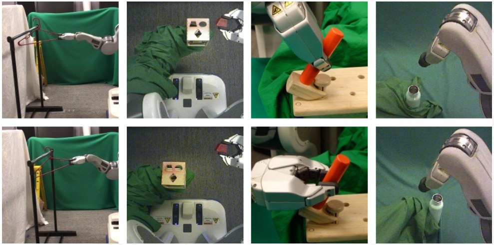

Levine et al. (2016)

# End-to-End Training of Deep Visuomotor Policies

## Visuomotor Policy Architecture

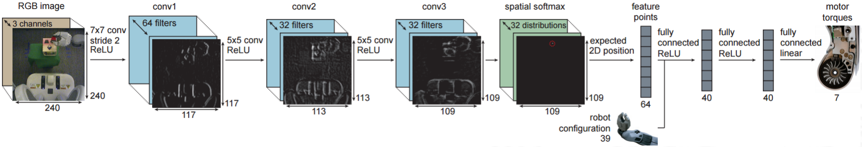


Levine et al. (2015)

- ∼92k parameters, 7 layers

# End-to-End Training of Deep Visuomotor Policies

## Results

**covered vision** leads to estimated manipulation and subsequent correction attempts → reliance on visual feedback

smaller **changes** will be adjusted for; bigger ones cause problems

**recovery** attempts after perturbations

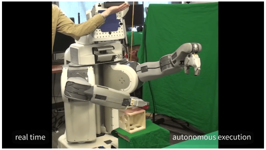

Finn (2015)

# Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection

Levine et al. (2016)

**use** grasp success prediction network
**with** continuous servoing mechanism
**for** continuous manipulator control

trained on >800k grasp attempts
from 14 distinct robots

Levine et al. (2016)

# Learning Hand-Eye Coordination for Robotic Grasping

## Architecture of CNN Grasp Predictor

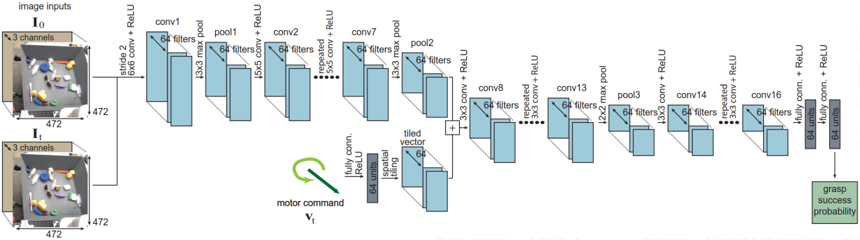


Levine et al. (2016)

# Learning Hand-Eye Coordination for Robotic Grasping

## Results

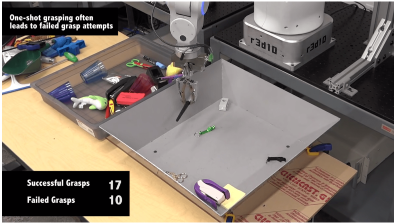


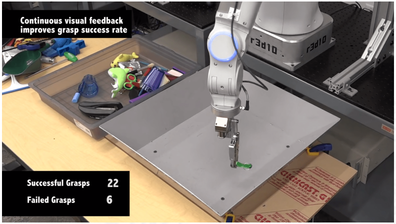


Pastor (2016c,a,b)

failure reduction from 34% to 18%

corrections after mistakes

recovery after perturbations/changes

# Conclusion

Problems for DL/ML in robotics:

- **data sparsity**
- high dimensionality/variability
- some generalizability but limited (related to the previous points)

Benefits of end-to-end DL approaches:

- more **natural movements**
- learned from scratch
- discovery of unconventional/ **non-obvious behaviour** (e.g. grasping of soft vs. hard objects)

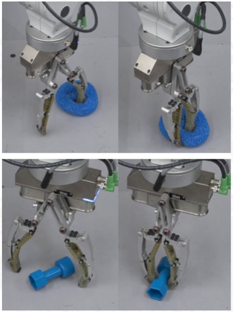

Levine et al. (2016)

# Thanks for your attention!

Questions?

# References

Bezák, P., Nikitin, Y. R., and Božek, P. (2014). Robotic grasping system using convolutional neural networks. *American Journal of Mechanical Engineering*, 2(7):216–218.

Finn, C. (2015). End-to-end training of deep visuomotor policies [video file]. Retrieved from `https://www.youtube.com/watch?v=Q4bMcUk6pcw#t=172`.

Lee, H., Grosse, R., Ranganath, R., and Ng, A. Y. (2009). Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 609–616. ACM.

Levine, S. (2015). Deep learning for decision making and control [video file]. Talk at CSE 519 Colloquium. Retrieved from `https://www.cs.washington.edu/events/colloquia/search/details?id=2686`.

Levine, S., Finn, C., Darrell, T., and Abbeel, P. (2015). End-to-end training of deep visuomotor policies. *arXiv preprint arXiv:1504.00702v5*.

# References (cont.)

Levine, S., Pastor, P., Krizhevsky, A., and Quillen, D. (2016). Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *arXiv preprint arXiv:1603.02199*.

Ng, A. (2016). How scale is enabling deep learning [video file]. Retrieved from https://www.youtube.com/watch?v=LcfLo7YP8O4.

Pastor, P. (2016a). Continuous visual feedback improves grasp success rate [video file]. Retrieved from https://www.youtube.com/watch?v=H4V6NZLNu-c.

Pastor, P. (2016b). Learning hand-eye coordination for robotic grasping [video file]. Retrieved from https://www.youtube.com/watch?v=l8zKZLqkfII#t=13.

Pastor, P. (2016c). One-shot grasping often leads to failed grasp attempts [video file]. Retrieved from https://www.youtube.com/watch?v=Q9tDHuidzak.

Sutton, R. S. and Barto, A. G. (1998). *Reinforcement learning: An introduction*, volume 1. MIT press Cambridge.