

Hierarchical Reinforcement Learning

Introduction

Cliff Li
17.03.2020

Cake

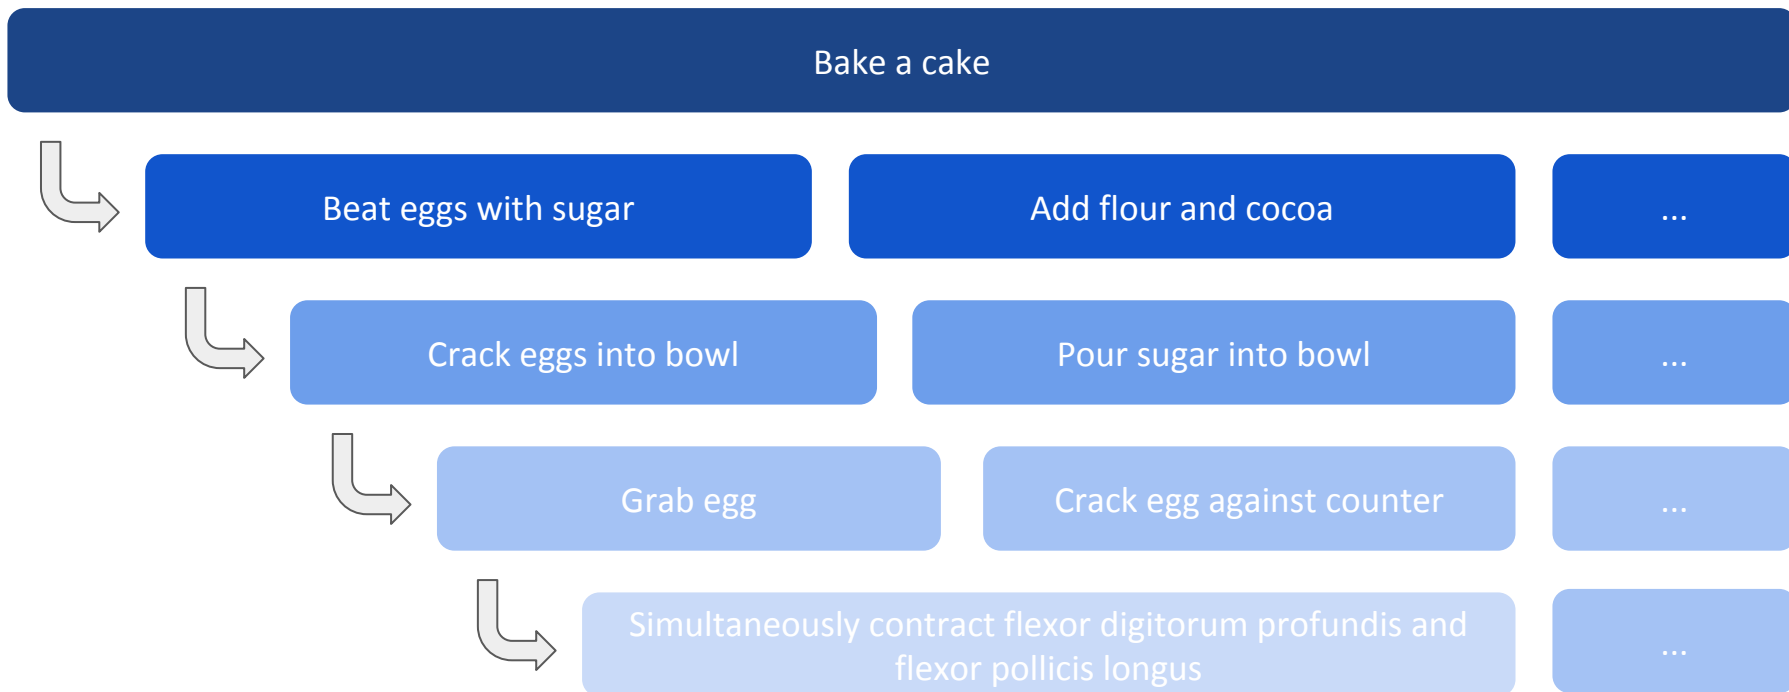


Source: [reddit.com/u/GreekElisium](https://www.reddit.com/u/GreekElisium)

Deficiencies of classical Reinforcement Learning

- Huge state and action spaces
- Credit assignment
- Transfer learning
- Overfitting (Overspecialization)
- Knowledge representation


Hierarchical Reinforcement Learning in a nutshell



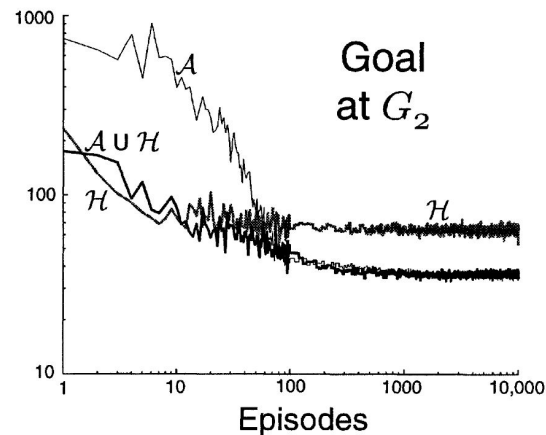
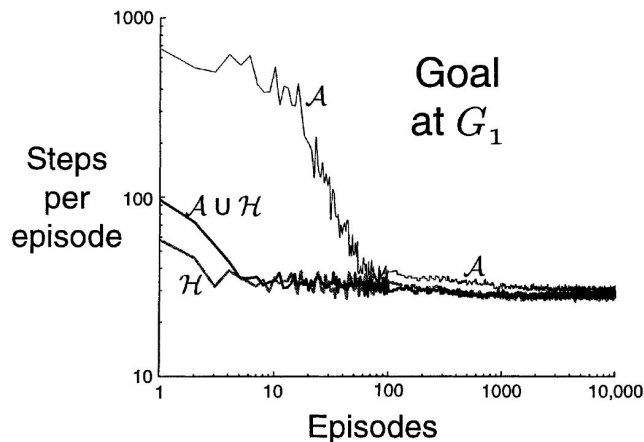

Benefits of Hierarchical Reinforcement Learning

- Structured exploration in state and action spaces
- Easier propagation of rewards
- Enables transfer learning
- Generalization through abstraction
- Better knowledge representation

Semi-Markov Decision Processes (SMDP) and Options




Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning (Sutton et al., 1999)

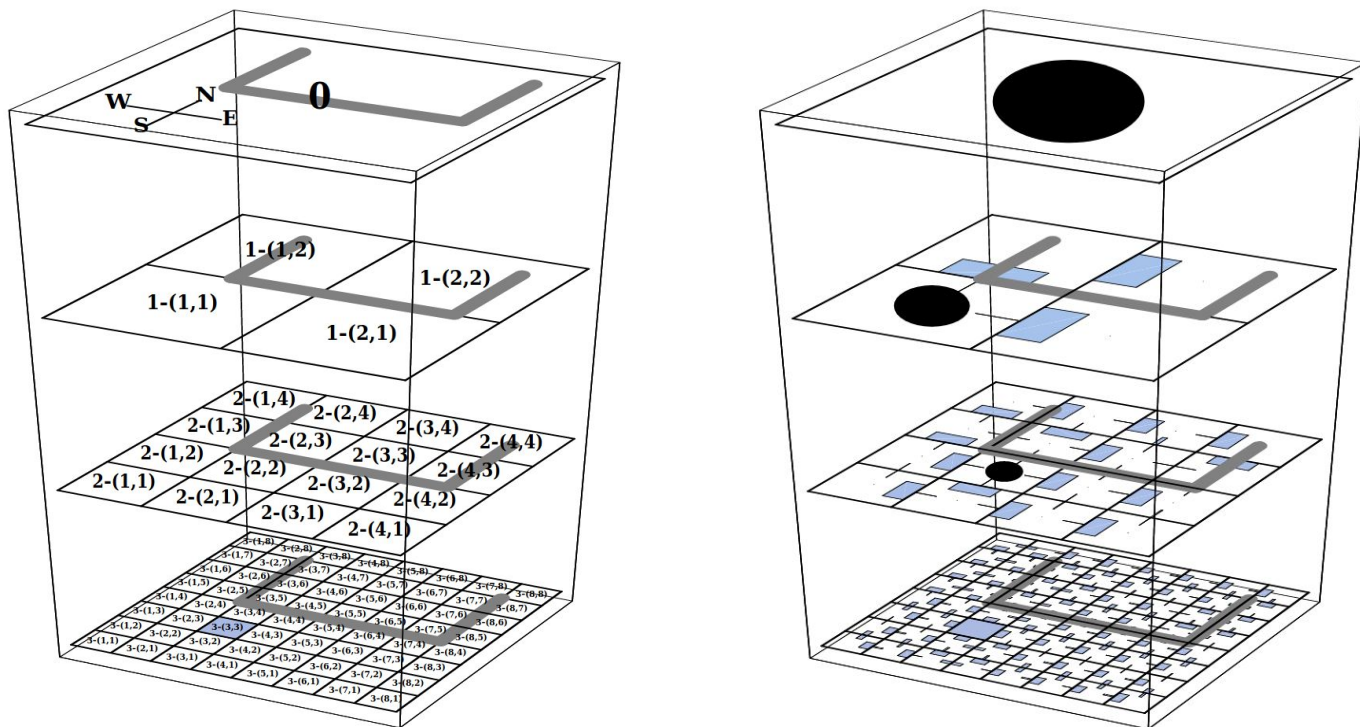


Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning (Sutton et al., 1999)

Landmarks Problem



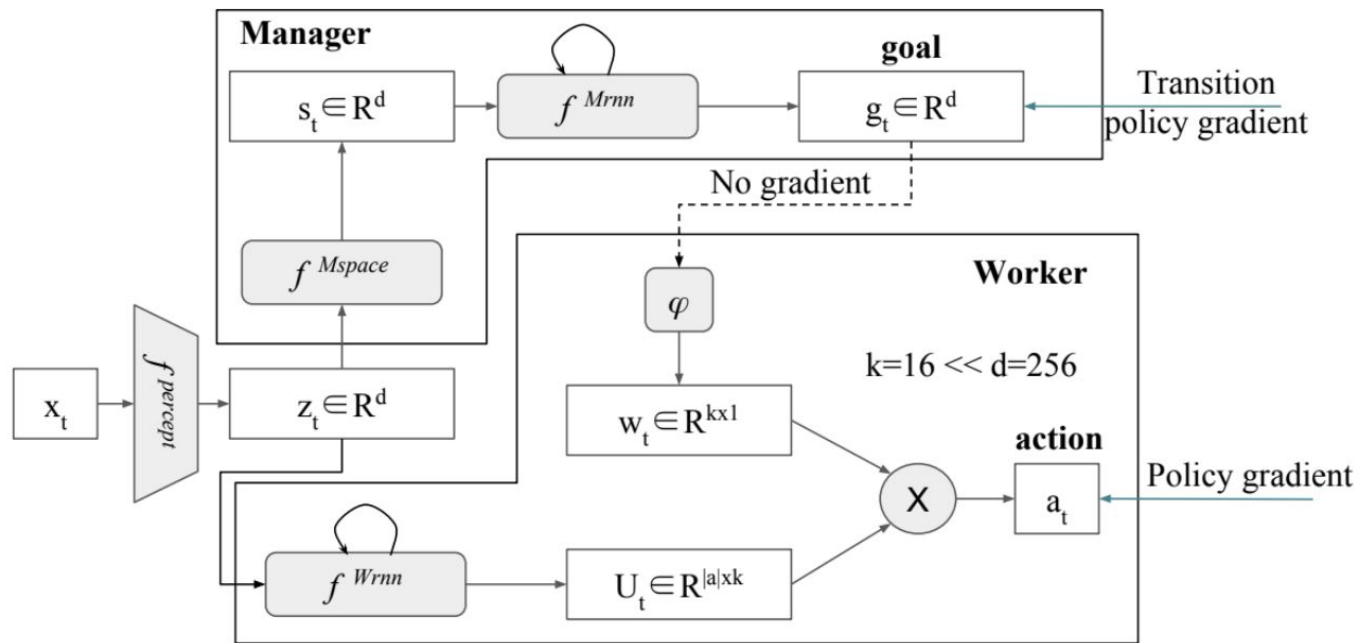
Feudal Reinforcement Learning (Dayan and Hinton, 1993)




Challenges of Hierarchical Reinforcement Learning

- Learning options
- Meaningful hierarchies
- Collapsing hierarchies into single policy
- Updating lower-level policies affects higher-level performance


FeUdal Networks for Hierarchical Reinforcement Learning (Vezhnevets et al., 2017)




FeUdal Networks for Hierarchical Reinforcement Learning (Vezhnevets et al., 2017)




Example frame




LSTM




Full FuN




sub-policy 1



sub-policy 2

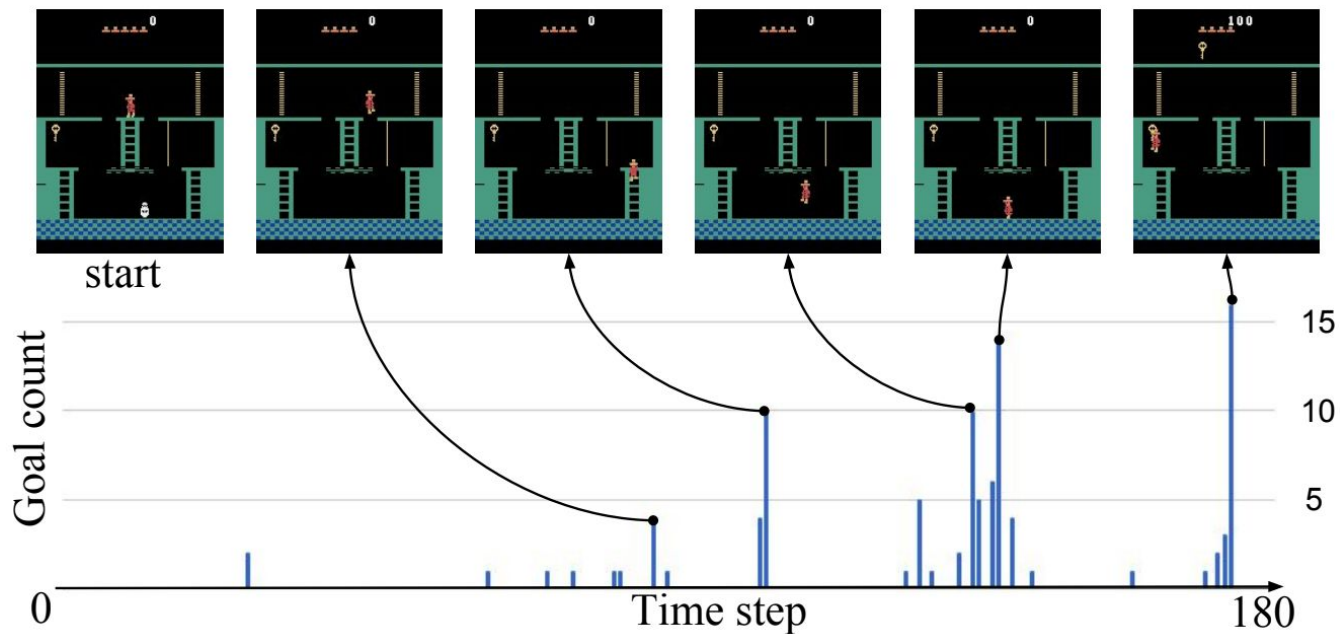


sub-policy 3



sub-policy 4

FeUdal Networks for Hierarchical Reinforcement Learning (Vezhnevets et al., 2017)



References

Richard S. Sutton, Doina Precup, Satinder Singh. *Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning*. Artificial Intelligence, 1999.

Peter Dayan, Geoffrey Hinton. *Feudal Reinforcement Learning*. Advances in Neural Information Processing Systems, 1993.

Alexander Sasha Vezhnevets, Simon Osindero, Tom Schaul, Nicolas Heess, Max Jaderberg, David Silver, Koray Kavukcuoglu. *FeUdal Networks for Hierarchical Reinforcement Learning*. 2017.