


Deep RL Meets Structured Prediction

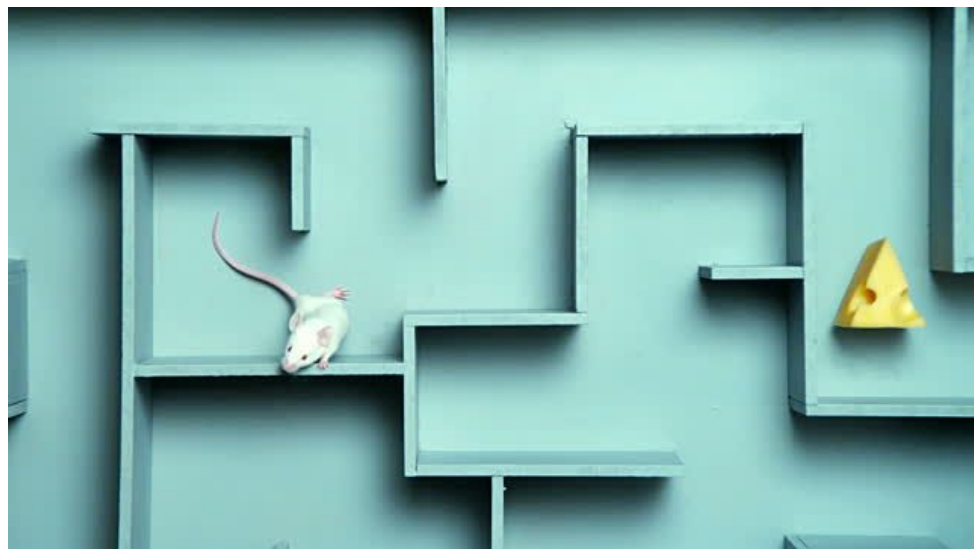
ICLR 2019 Workshop
May 6th 2019, New Orleans (USA)



Structured Prediction Often Needs to Do Search

Which fits well with RL's ability to form good search policy

- dialogue
- semantic parsing
- program synthesis
- architecture search
- machine translation
- summarization
- image caption
- knowledge graph reasoning
- information extraction
- ...



RL is attractive:

Directly Optimizing The Expected Reward
Which can be very useful for structured predictions

- **ML** optimizes the log likelihood of target sequences

$$J^{ML}(\theta) = \sum_q \log P(a_{0:T}^{best}(q)|q, \theta)$$

- **RL** optimizes the expected reward under a stochastic policy

$$J^{RL}(\theta) = \sum_q \mathbb{E}_{P(a_{0:T}|q, \theta)} [R(q, a_{0:T})]$$



[Williams 1992]


[Sutton & Barto 1998]

[Liang+ 2017]

RL has challenges:

Which we need to be aware of

- **Large search space (sparse rewards)**
 - Supervised pretraining (MLE)
 - Systematic exploration [Houthoof+ 2017]
 - Curiosity [Schmidhuber 1991][Pathak2017]
- **Credit assignment (delayed reward)**
 - Bootstrapping
 - E.g., AlphaGo uses a value function to estimate the future reward
 - Rollout n-steps
- **Train speed & stability (optimization)**
 - Trust region approaches (e.g., PPO)
 - Experience replay



[Sutton & Barto 1998]
[Abbeel & Schulman 2016]

Organizers



Audrey Durand
@Université Laval
prev. McGill U.



Wang Ling
@Deepmind
prev. Carnegie Mellon
U.



Yuandong Tian
Facebook AI Res.
prev. Carnegie
Mellon U.



Ni Lao
@Mosaix.ai
prev. Google Inc.



Zita Marinho
@Priberam Labs & U.
of Lisbon
prev. Carnegie Mellon
& U. of Lisbon




Lu Wang
@Northeastern U.
prev. Cornell U.



Chen Liang
@Google Brain
prev. Northwestern U.



André Martins
@Unbabel & U. of
Lisbon
prev. Carnegie Mellon
U. & U. of Lisbon



Jason Williams
@Siri, Apple
prev. Microsoft Res.

Schedule

8:30 AM	Poster setup (posters will be up in the room all day.)
9:45 AM	Opening remarks
9:50 AM	Jessica B. Hamrick: Structured Computation and Representation in Deep Reinforcement Learning
10:25 AM	Advertising contributed talk & Best paper award (Zhiting Hu, 5min)
10:30 AM	Coffee & Poster (sync w/ ICLR conference)
11:00 AM	Anima Anandkumar: Infusing Structure into Machine Learning Algorithms
11:35 AM	Graham Neubig: What can Statistical Machine Translation teach Neural Machine Translation about Structured Prediction?
12:10 PM	Mohammad Norouzi: Beyond Off-the-shelf Reinforcement Learning for Structured Prediction
12:45 PM	Advertising contributed talk (Wouter Kool, Zafarali Ahmed, Osbert Bastani, 3x5min)
13:00 PM	Lunch & Poster
14:20 PM	Workshop on Deep Generative Models for Highly Structured Data