



Jiayi Weng

CST Department, Tsinghua University

February 29, 2020



- 1 Bio
- 2 Research Experience (RL)
- 3 Future Direction

## Jiayi Weng

- B.Eng. in Computer Science, Tsinghua University, 2016 - 2020.  
GPA 3.75/4.00 (Rank 17/158)
- Research Assistant to Prof. **Yoshua Bengio**, Mila  
Jul. 2019 - now, topic: NLP, RL, Neural-Symbolic Reasoning
- Research Assistant to Prof. **Jun Zhu**, Tsinghua University  
Mar. 2018 - now, topic: RL
- Research Assistant to Dr. **Hongwei Qin**, SenseTime Inc.  
Mar. 2018 - Jun. 2018, topic: CV (HDR+, denoising)
- Research Assistant to Prof. **Hailong Yao**, Tsinghua University  
May 2017 - Oct. 2017, topic: VLSI (escape routing)
- **GitHub** 1000+ personal stars, 600+ followers

① Bio

② Research Experience (RL)

VizDoom AI Competition 2018

Relational Markov Decision Process

Rule-Transformer

③ Future Direction

① Bio


② Research Experience (RL)

VizDoom AI Competition 2018

Relational Markov Decision Process

Rule-Transformer

③ Future Direction



- birth place
- navigation road
- exit door



(a) Navigation



(b) Resources: weapon and health bonus



(c) Shoot enemies



(d) Final goal: open exit door

Figure 1: Example of single-player scenarios in VizDoom.

# Challenges & Solutions

- 1 Visual understanding from raw images
  - ⇒ Set up an object detection system (YOLO-v3<sup>1</sup>) aiming to efficiently identify the obstacles and enemies.
- 2 Visual navigation in the 3D environment
  - ⇒ Incorporated the depth signal into the navigation network, which facilitated the agent to plot its paths. (+SNAIL<sup>2</sup>)
- 3 Sparse and delayed reward for the agent
  - ⇒ Formulated the problem within a hierarchical reinforcement learning framework by dividing the challenge into subtasks: navigation, attack, tool, and resources.

---

<sup>1</sup>Joseph Redmon et al. "YOLOv3: An Incremental Improvement". In: *CoRR* abs/1804.02767 (2018). arXiv: 1804.02767.

<sup>2</sup>Nikhil Mishra et al. "A Simple Neural Attentive Meta-Learner". In: *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. 2018.

VizDoom Competition Result<sup>3</sup>

- All teams submitted an agent to finish the game on ten unseen and more difficult maps.
- Our method achieved first place in both public-rank and private-rank leaderboards.

---

Team	Ours	DoomNet	VIPLAB	ddangelo
Total Time (min)	<b>25.34</b>	29.86	31.54	37.33
# of Best Record	<b>8</b>	4	3	4

---

Table 1: VDAIC 2018 Track(1) Competition Result (top 4 team).

---

<sup>3</sup>Detailed result can be found at

<http://vizdoom.cs.put.edu.pl/competitions/vdaic-2018-cig/results>.



## ① Bio

## ② Research Experience (RL)

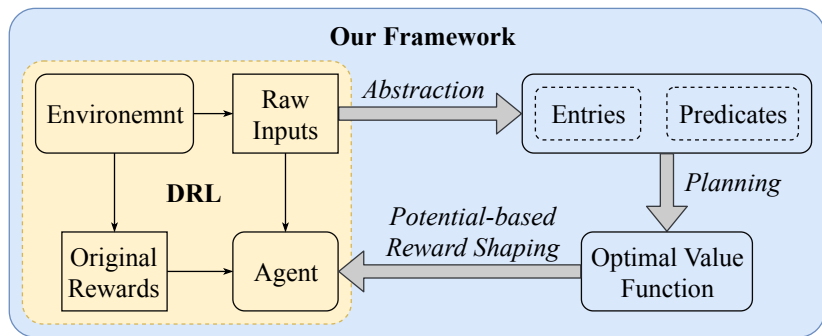
VizDoom AI Competition 2018

Relational Markov Decision Process

Rule-Transformer


## ③ Future Direction

We aim to accelerate the deep reinforcement learning procedure by state abstraction (using first-order logic) and reward shaping.




**Figure 2:** The working procedure of RMDP. Generate the credits for state-action pairs of the deep reinforcement learning in three stages: abstraction, planning, and potential-based reward shaping.

- Our method can accelerate the convergence of value-based RL method.
- RMDP is more robust under harder games.



(a) Easy Map



(b) Hard Map

**Figure 3:** Test RMDP on Health Gathering scenarios. Vertical axis indicates the accumulated reward per episode through training phase.

## ① Bio

## ② Research Experience (RL)

VizDoom AI Competition 2018

Relational Markov Decision Process

Rule-Transformer

## ③ Future Direction

- We aimed to develop a system with Consciousness Prior<sup>4</sup>, which integrated the advantage of both neural and symbolic models.
- To make Transformers more like reasoning with rules, instead of a fixed rule per layer, allow to also pick a rule, at each level, and share that set across levels.


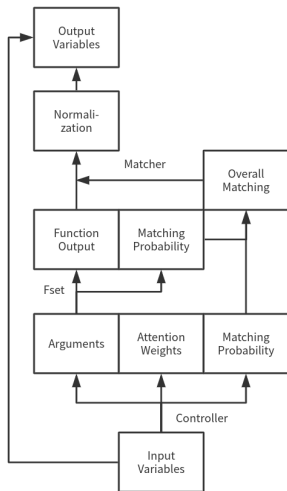


Figure 4: The Consciousness Prior, credit<sup>5</sup>.

<sup>4</sup>Yoshua Bengio. "The Consciousness Prior". In: *CoRR* abs/1709.08568 (2017). arXiv: 1709.08568.

<sup>5</sup>Yoshua Bengio. *From System 1 Deep Learning to System 2 Deep Learning*. 2019.



- Allow picking a rule at each level
- Sharing knowledge across the level
- Update different variables with different functions
- The rule could potentially be dynamically defined

Figure 5: Architecture of Rule-Transformer.

# What we have tried and are trying

- × Machine Translation
- × Language Modeling
- = GLUE Benchmark
- = Multitask Imitation Learning
- × Mathematical Dataset
- √ LeetCode-style algorithmic problems
- ? Brainfuck programming language explainer (ongoing)

- ① Bio
- ② Research Experience (RL)
- ③ Future Direction



- Generally, I would like to continue my research in the field of reinforcement learning.
- I am interested in addressing the challenges of sample complexity (incorporating human-prior knowledge for better disentanglement) / robustness and generalization (combining with meta-learning, game-theory).
- I am also keen to work on the real-world application of reinforcement learning: learning-based control system, self-driving cars, etc.

*Thank you!*