

Dropout Q-Functions for Doubly Efficient Reinforcement Learning

Takuya Hiraoka^{1,2}, Takahisa Imagawa², Taisei Hashimoto^{2,3},
Takashi Onishi^{1,2}, and Yoshimasa Tsuruoka^{2,3}

¹ NEC Corporation, ² AIST, ³ The University of Tokyo

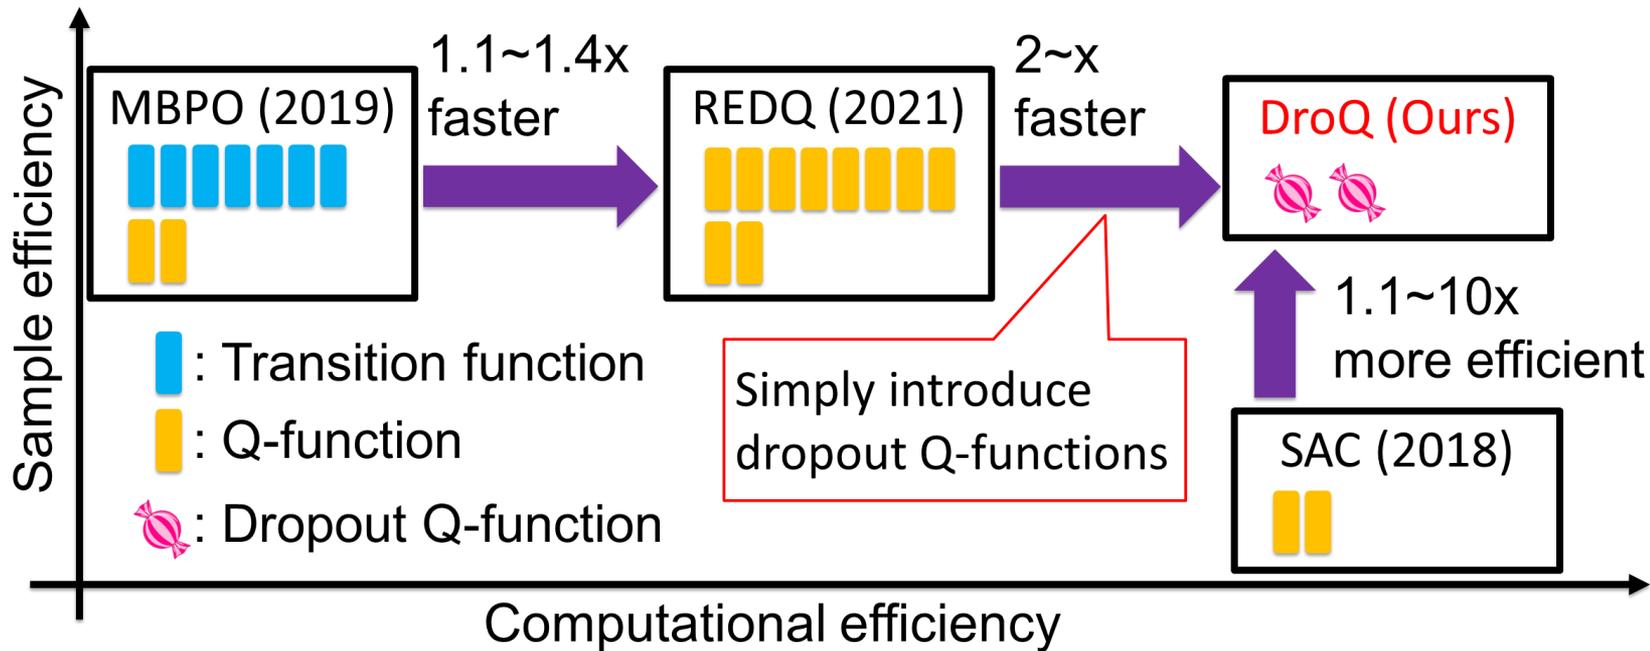
In general, **doubly efficient** RL methods that are not only **sample efficient** but also **computational efficient** are preferable.



EVERYONE SAYS THIS RL ALGORITHM IS SAMPLE EFFICIENT,
BUT IT'S TOO SLOW AND TOO HEAVY TO RUN ON MY LAPTOP.
NOT SURE WHEN MY HYPER-PARAMETER TUNING ENDS...

Introduction (Contd.)

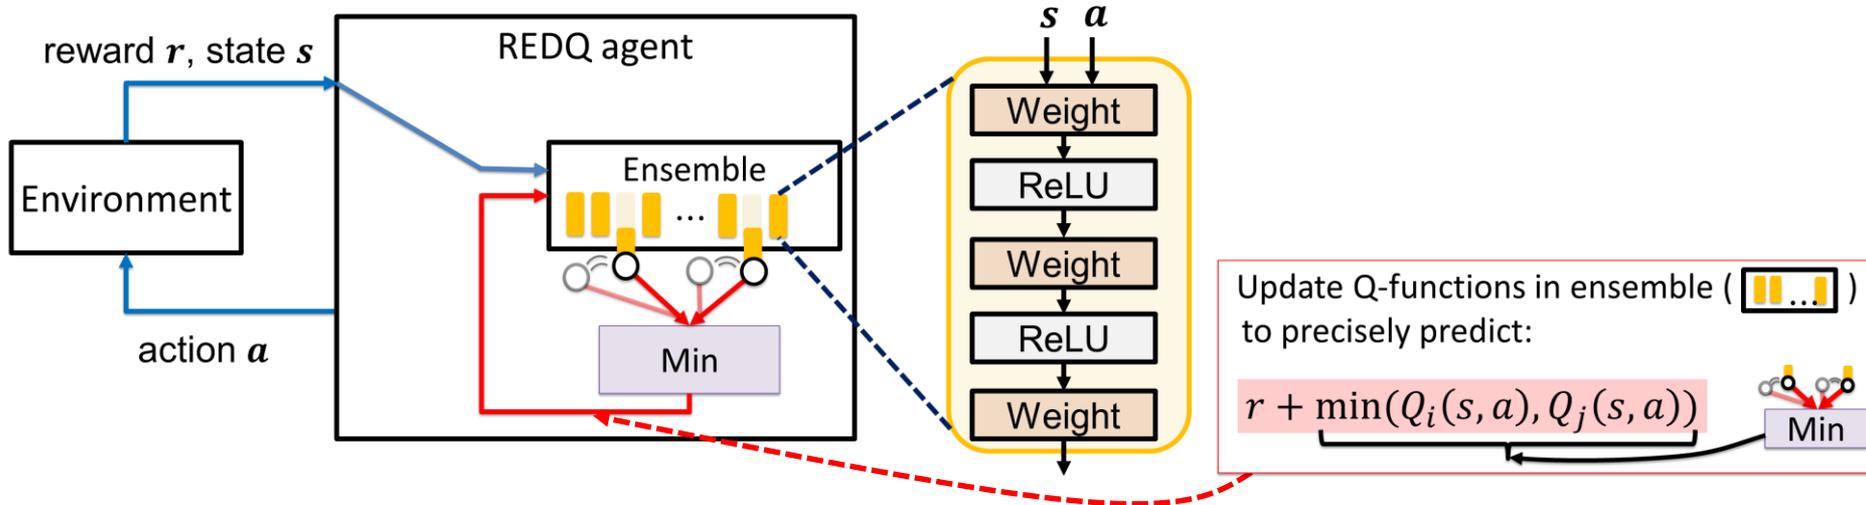
We propose **DroQ**, a simple but doubly efficient RL method, by introducing dropout Q-functions (🍬) to REDQ.



REDQ: Randomized Ensembled Double Q

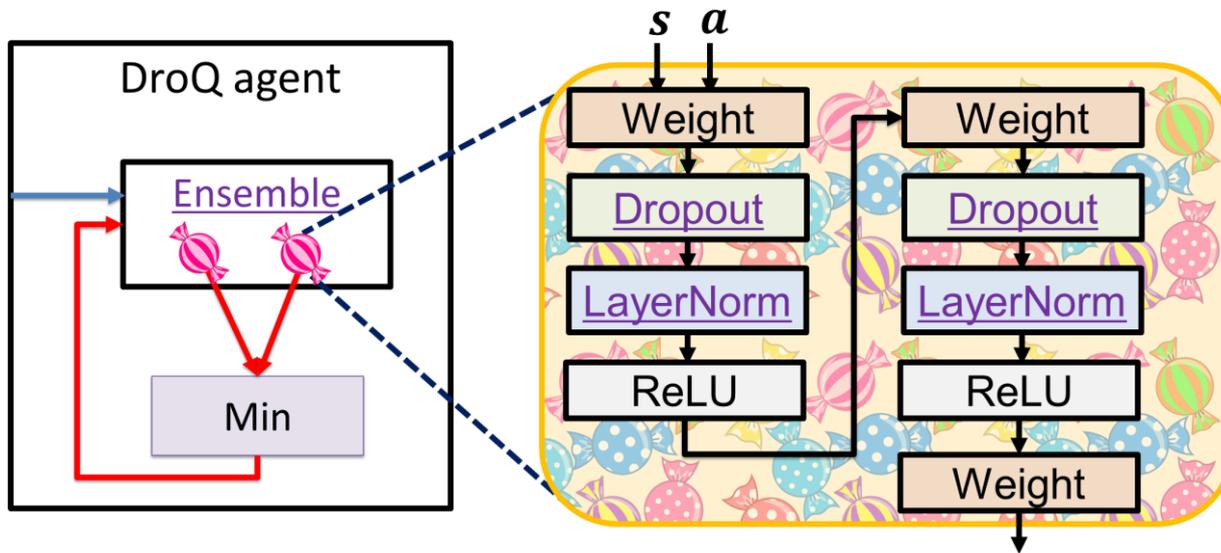
REDQ (Chen, 2021) is a sample-efficient RL method equipped with

- **High update-to-data (UTD) ratio:** number of Q updates (\rightarrow) per environment interaction (\rightarrow) is high (e.g., 20 updates per interaction).
- **Randomized ensemble:** a randomly selected subset () of ensemble () is used at the target (**Min**) in the Q update (\rightarrow).



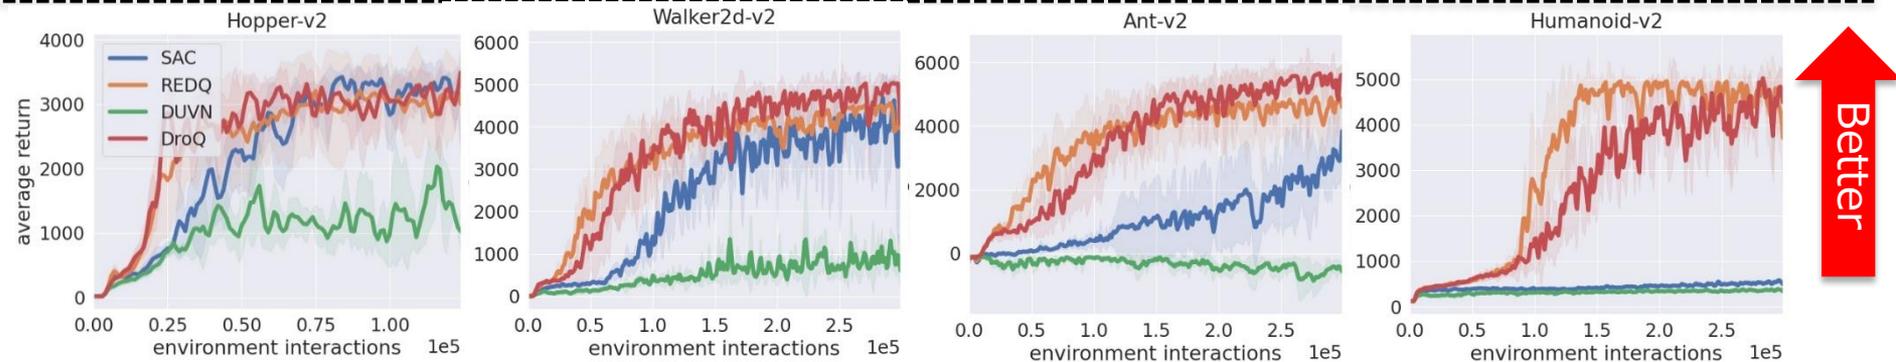
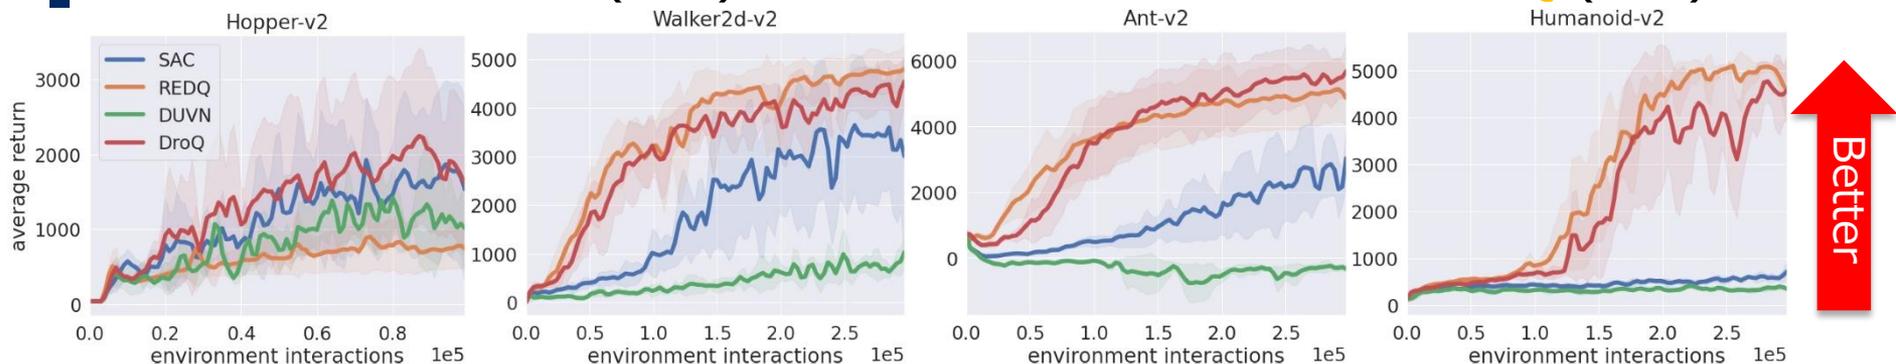
DroQ, the proposed method

DroQ is a REDQ variant using a small ensemble of dropout Q-functions () in which dropout (**Dropout**) and layer normalization (**LayerNorm**) are used.



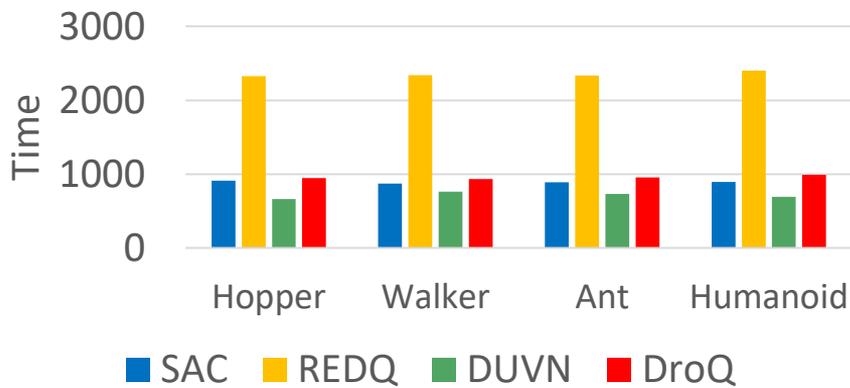
Q. How sample-efficient is **DroQ** (—)?

A. Better than **SAC** (—) and almost the same as **REDQ** (—).

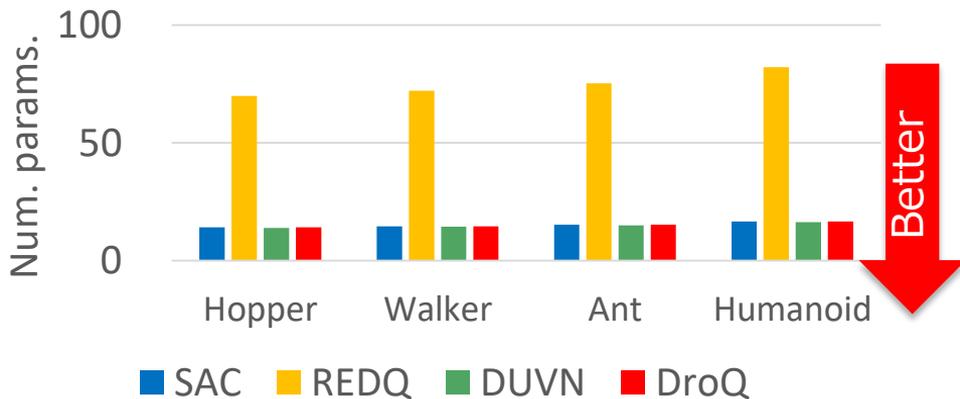


Q. How computationally efficient is **DroQ**?

A. Much better than **REDQ** and almost the same as **SAC**.



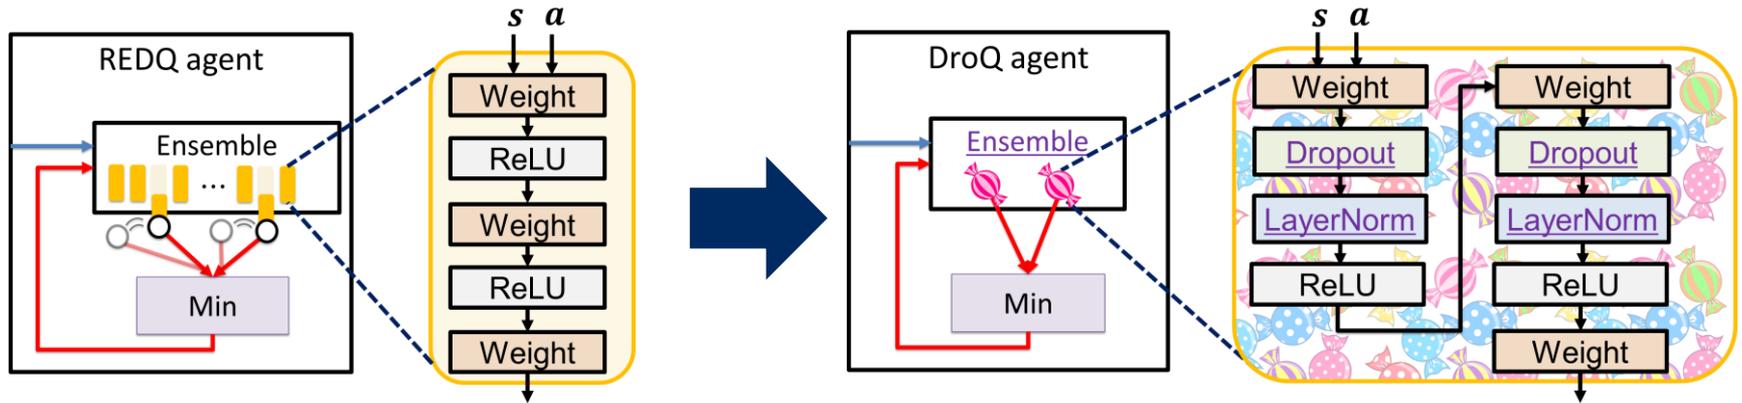
**Times per 20 updates
+ 1 interaction (in msec)**



Number of parameters (/1e4)

Conclusion

DroQ (REDQ + ) is simple but doubly efficient.

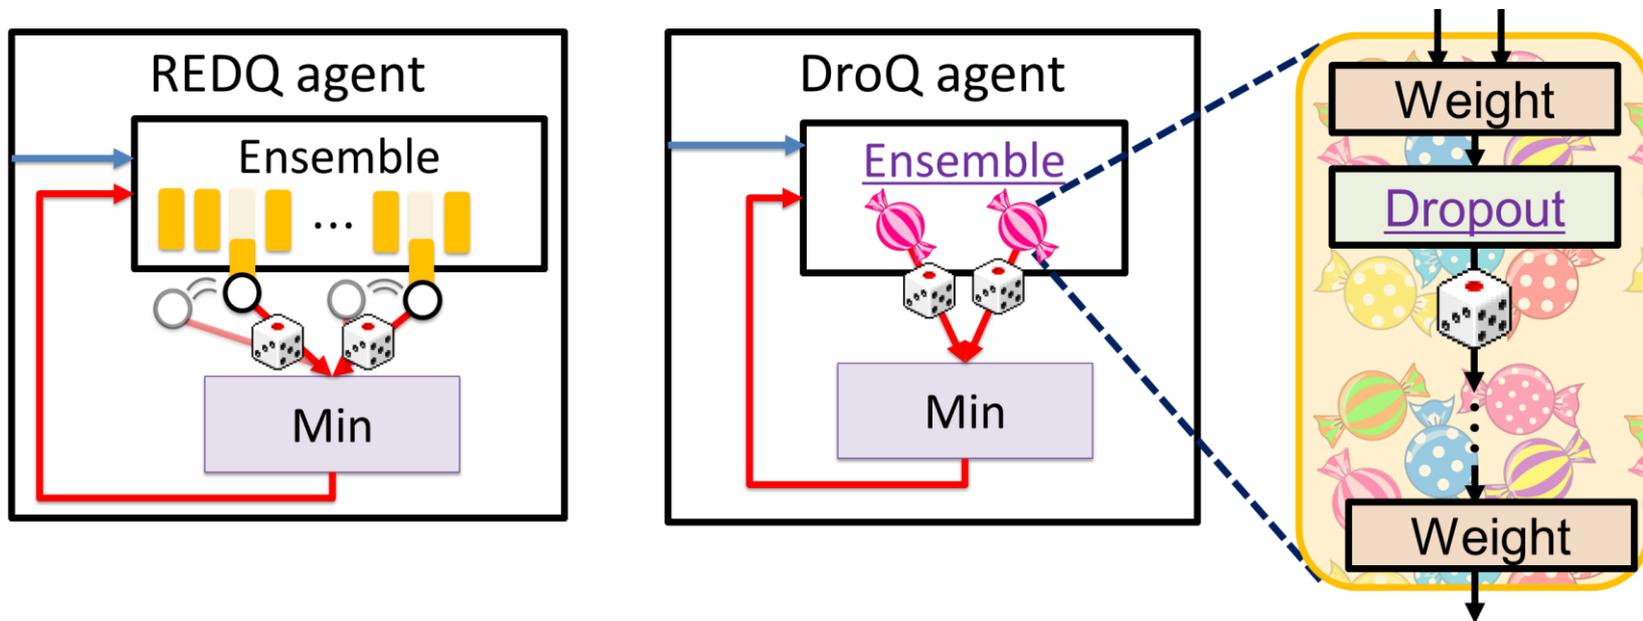


Are you interested in our work?
Or feel that all we did was just randomly changing modules
of the existing RL method?

INTRODUCTION
TO
OUR POSTER

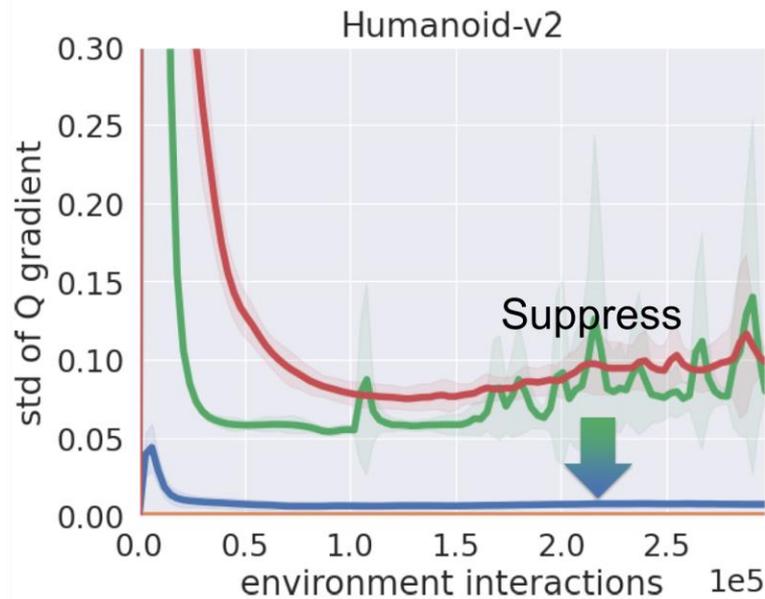
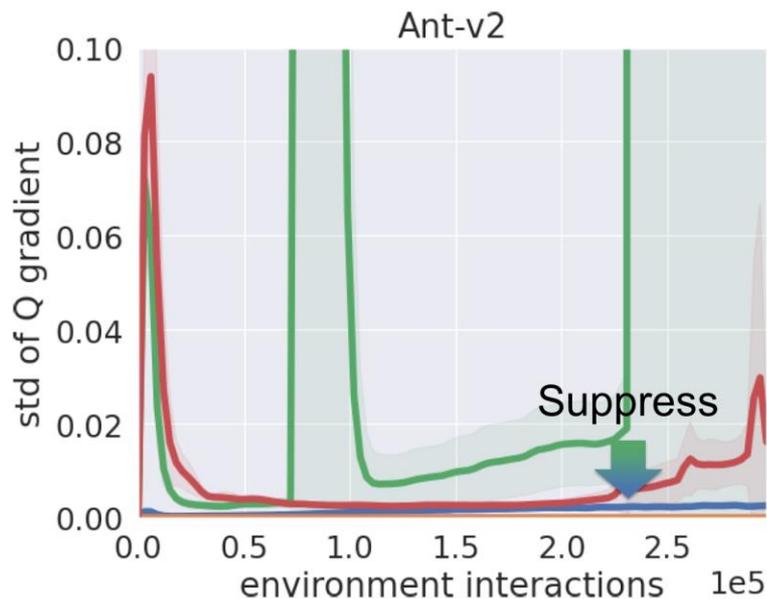
Q. Why is dropout (Dropout) needed?

- A. To inject Q-function uncertainty (🎲) to the target (Min), similarly to REDQ.



Q. Why is layer normalization (**LayerNorm**) needed?

A. To suppress (↓) the learning instability caused by dropout.

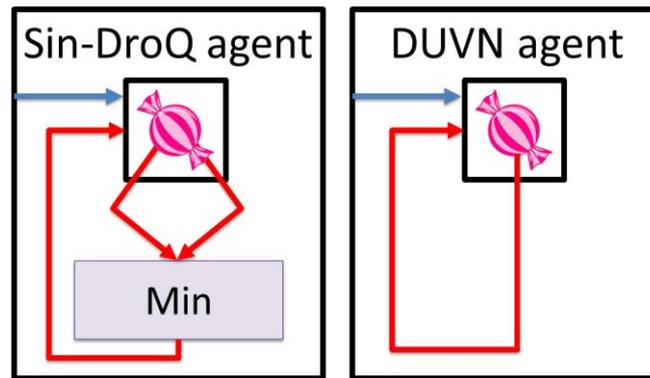
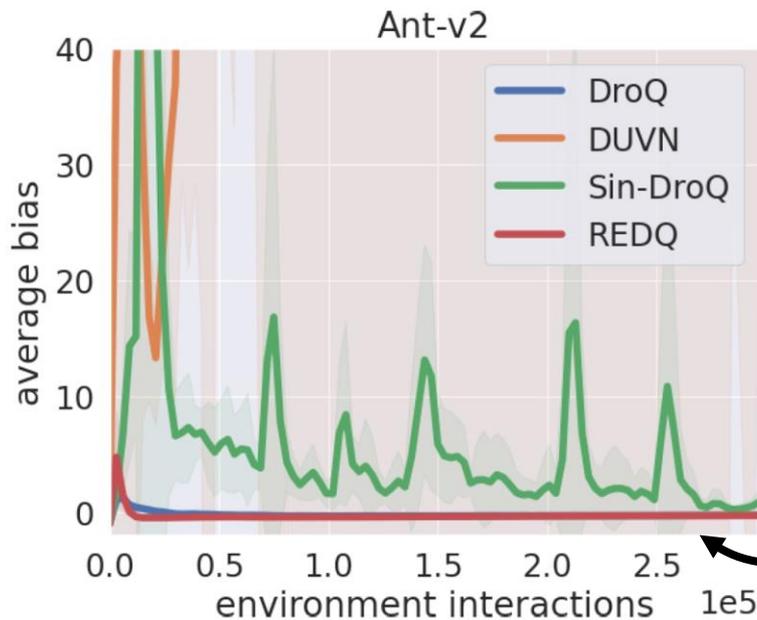


The standard deviation of the gradient of Q-loss w.r.t. parameters

— DroQ — w/o Dropout — w/o LayerNorm — w/o Dropout nor LayerNorm

Q. Why is a small ensemble () needed?

A. Using a single dropout Q-function () alone induces a large bias in Q-estimation.



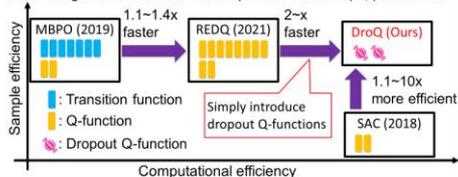
Q-estimates bias calculated as $\frac{| \text{True Q} - \text{Estimated Q} |}{\text{Normalize coefficient}}$

1. Introduction:

- In general, RL methods that are not only **sample efficient** but also **computational efficient** (i.e., **doubly efficient**) are preferable.

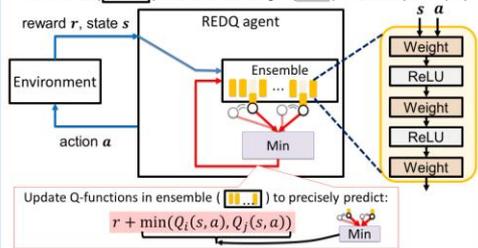


- We propose DroQ, a simple but doubly efficient RL method, by introducing a small ensemble of Dropout Q-functions (🍷) to REDQ.



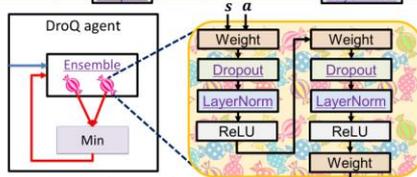
2. Randomized Ensemble Double Q-Learning (REDQ):

- REDQ (Chen, 2021) is a sample-efficient RL method equipped with **high update-to-data (UTD) ratio** and **randomized ensemble**.
- High UTD ratio:** number of Q updates (→) per environment interaction (←) is high (e.g., 20 updates per interaction).
- Randomized ensemble:** a randomly selected subset (🍷) of ensemble (🍷) is used at the target (Min) in the Q update (→).

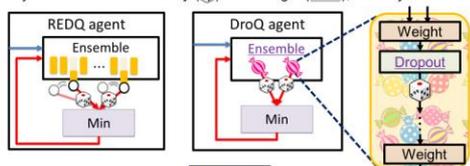


3. DroQ, the proposed method:

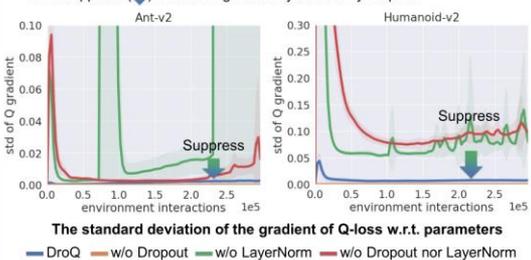
- DroQ is a REDQ variant using a **small ensemble of dropout Q-functions** (🍷) in which **dropout** (🍷) and **layer normalization** (🍷) are used.



- Q. Why is dropout (🍷) needed?**
 A. To inject Q-function uncertainty (🍷) to the target (Min), similarly to REDQ.

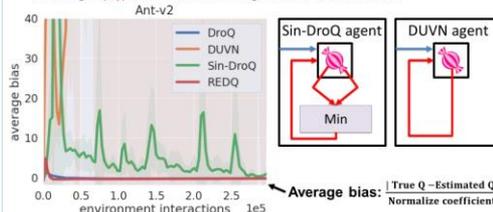


- Q. Why is layer normalization (🍷) needed?**
 A. To suppress (🍷) the learning instability caused by dropout.

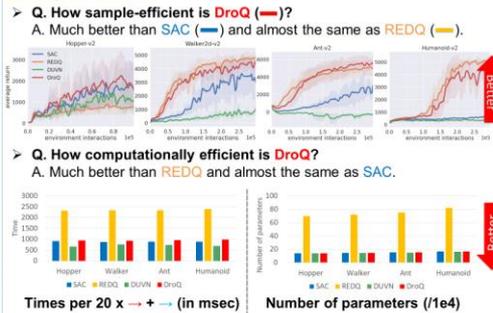


3.5. DroQ, the proposed method (Contd.):

- Q. Why is a small ensemble (🍷) needed? (Why not use a single dropout Q-function (🍷) alone?)**
 A. Using it (🍷) alone induces a large bias in Q-estimation.



4. Experiments:



5. Conclusion:

- DroQ (REDQ + 🍷) is simple but doubly efficient.

Our source code is available at



Thank you for watching this video!



NEC-AIST
AI Cooperative
Research Laboratory