

The Edward S. Rogers Sr. Department of Electrical & Computer Engineering UNIVERSITY OF TORONTO

Salma Emara, Daniel Liu, Fei Wang, Baochun Li **Department of Electrical and Computer Engineering** University of Toronto

> **INFOCOM 2024 FOGML** Presenter: Fei Wang

Cascade: Enhancing Reinforcement Learning with Curriculum Federated Learning and Interference Avoidance — A Case Study in Adaptive Bitrate Selection

Reinforcement Learning









In Curriculum Learning ...



Benefits

Avoids interference

Xia et al., "Automatic curriculum learning for learning adaptation in networking," ACM SIGCOMM, 2022.











Difficulty Metric

Absolute value of loss

Rate of change in loss

Second derivative of loss

Higher cumulative average reward Faster convergence

Inference Avoidance (IA)

Introduces a regularization term into loss

to penalize changes to important parameters

Change to weights $l_{i,\text{model}} = L_i(\theta) + \lambda \Omega_{i,t,model} (\theta_i - \theta_t)^2,$

Important weight matrix



A Case Study: Adaptive Bitrate Selection

Adaptive Bitrate (ABR) Selection

time, low variation in video quality

NETFLIX



The objective of each client: high video quality, low rebuffering





Mao et al., "Neural adaptive video streaming with Pensieve," ACM SIGCOMM, 2017.





Each client observes environments simulated using traces on a





11

Selected Experimental Results

Evaluation Metrics

- (environments) the model was trained on
- **2 Model Skewness**

Avg. reward over task *i* $\sum_{i=1}^{L} (G_i - G_{avg})^2$ Number of tasks learned Avg. reward over Z tasks 3 Average test reward: average reward on traces not used during training to test generalization

1 Asymptotic reward: summation of average reward in all tasks





1	Δ	

Generalization of Cascade FedAvg



RedABR

Cascade

-45 10 65 120 -100 Average episodic reward over 30 test traces

Avaraga anisodic reward over 20



$\cdot \mathbf{\nabla}$	1	5
-------------------------	---	---

Ablation Study



Lower model skewness Higher test reward





Concluding Remarks

Cascade: a new federated reinforcement learning framework

Curriculum learning + Interference avoidance

Outperforms FL algorithms by up to

20% in asymptotic performance

21% in model skewness

At a cost of 11% more communication rounds



