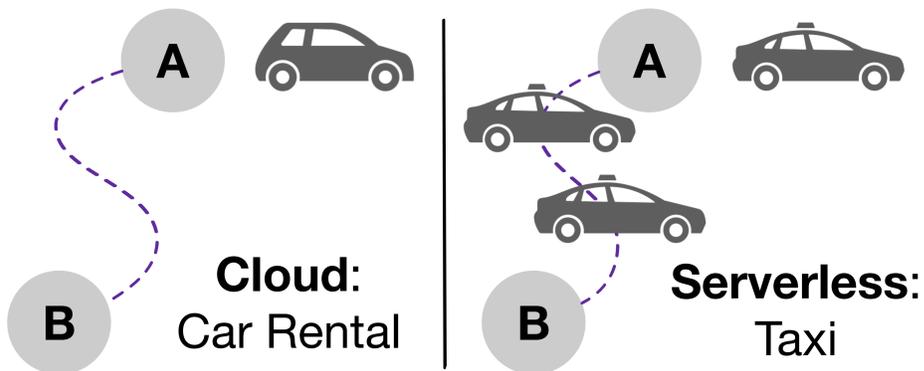


# Nitro: Boosting Distributed Reinforcement Learning with Serverless Computing

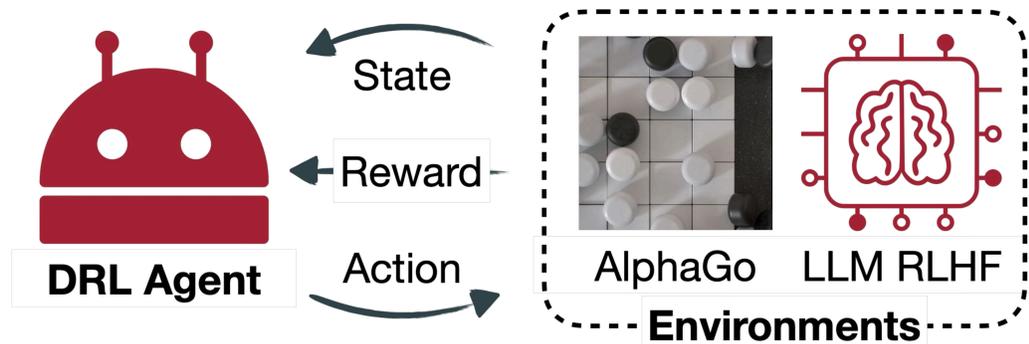
Hanfei Yu<sup>1</sup>, Jacob Carter<sup>2</sup>, Hao Wang<sup>1</sup>, Devesh Tiwari<sup>3</sup>, Jian Li<sup>4</sup>, Seung-Jong Park<sup>5</sup>

Stevens Institute of Technology<sup>1</sup>, Louisiana State University<sup>2</sup>, Northeastern University<sup>3</sup>,  
Stony Brook University<sup>4</sup>, Missouri University of Science & Technology<sup>5</sup>

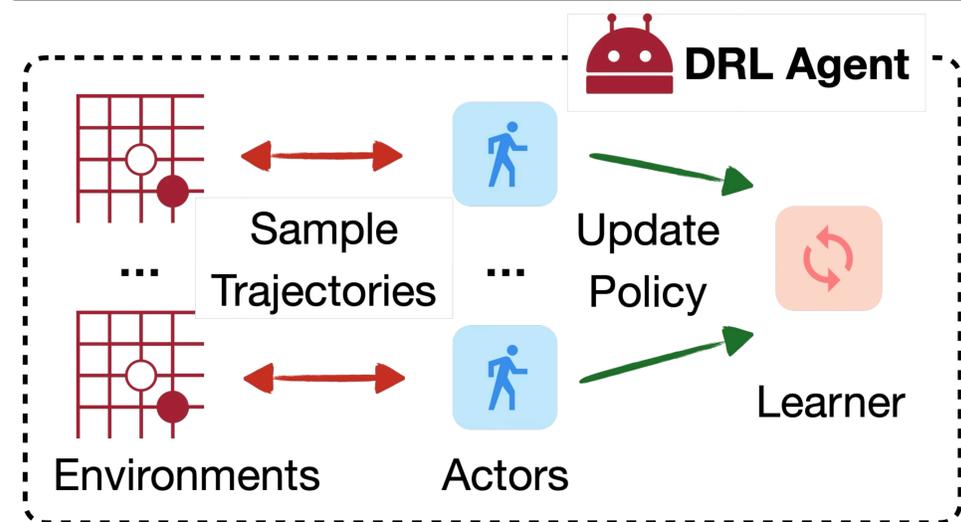
## 1. Cloud vs. Serverless Computing



## 2. Deep Reinforcement Learning (DRL)

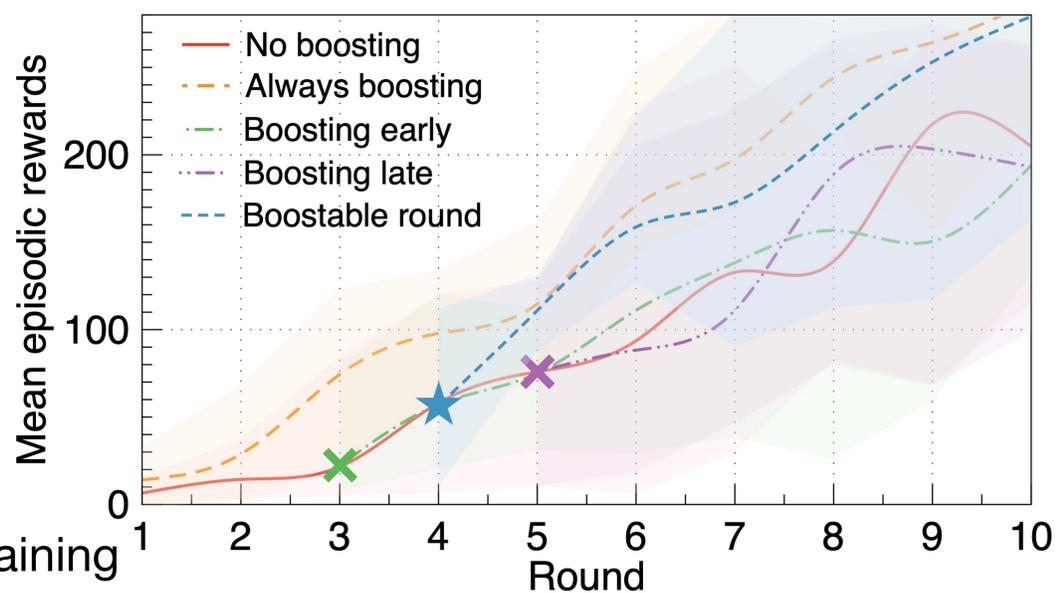


## 3. Actor-Learner Architecture

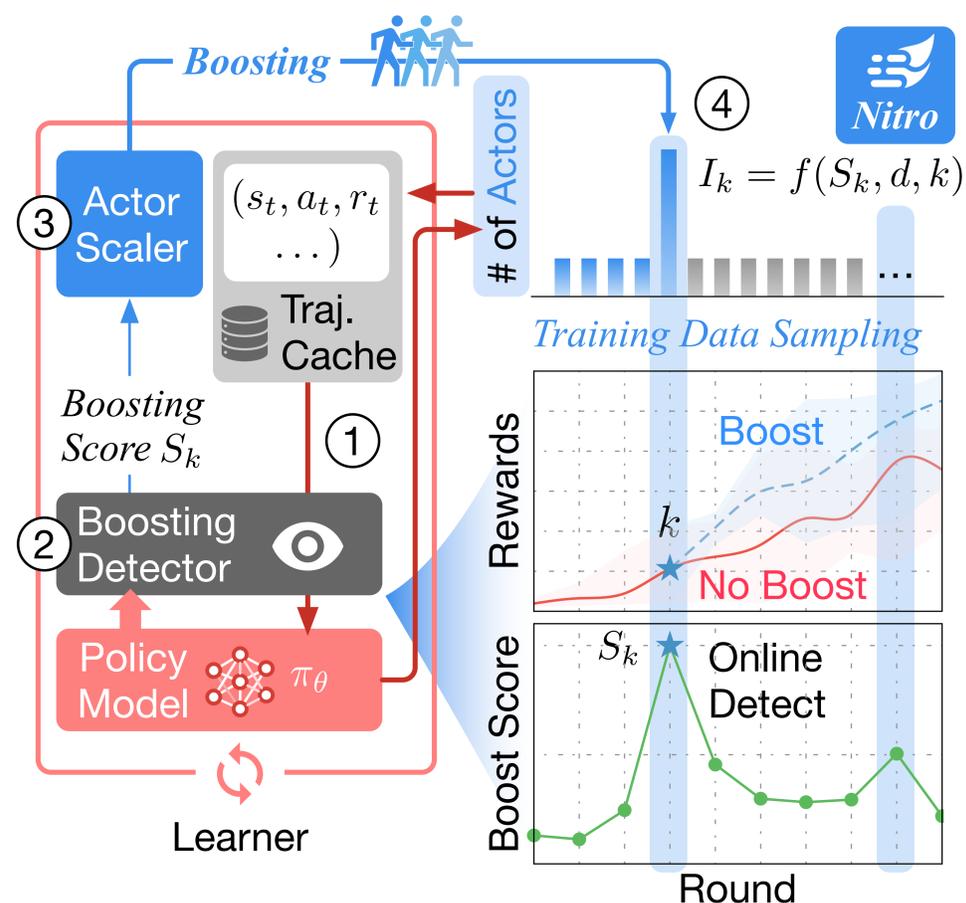


**Boosting:** increase actors to speed up DRL training

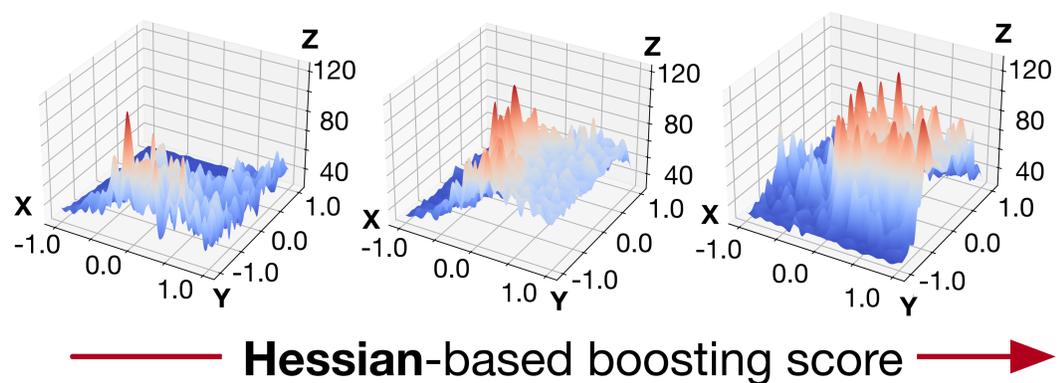
## 4. Boosting Distributed DRL Training



## 5. Nitro's Workflow and Actor Scaling



## 6. Boosting Opportunity Detection



## 7. DRL Training Improvement

