

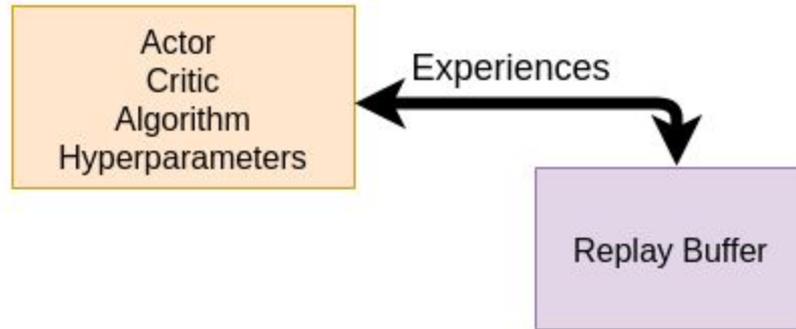


Collaborative Evolutionary Reinforcement Learning

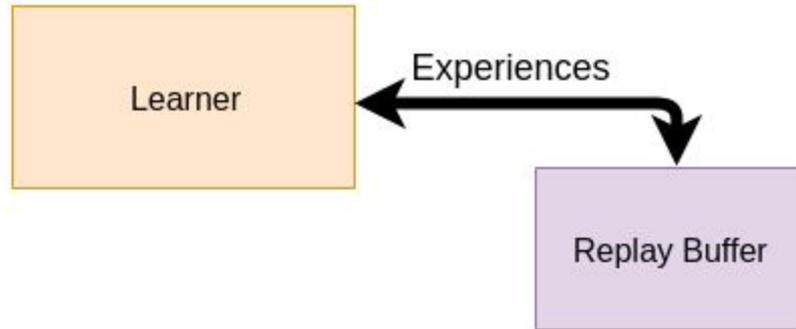
Shauharda Khadka, Somdeb Majumdar, Tarek Nassar, Zach Dwiell, Evren Tumer, Santiago Miret, Yinyin Liu, Kagan Tumer*

Artificial Intelligence Products Group, Intel Corporation
Oregon State University*

A simple actor-critic policy gradient setup



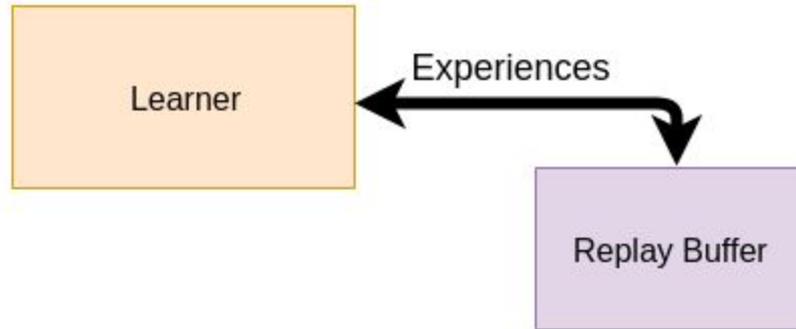
Learner



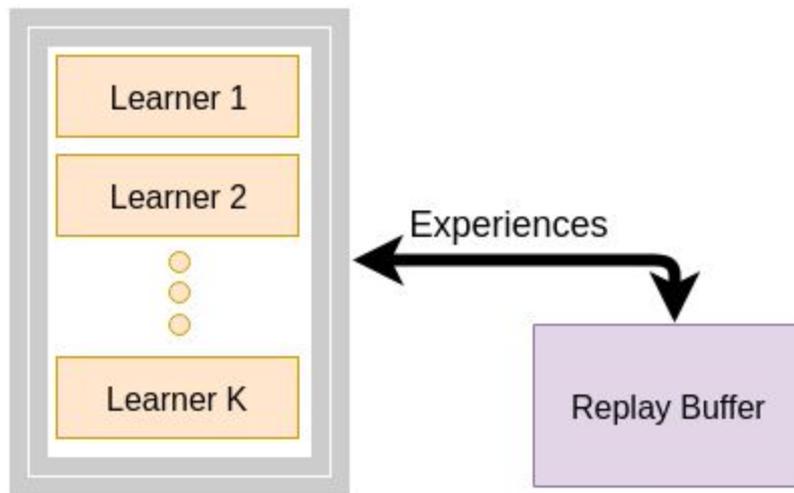
What do we optimize exactly?

$$R_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$$

Learner



Portfolio of Learners (varying discount rates)



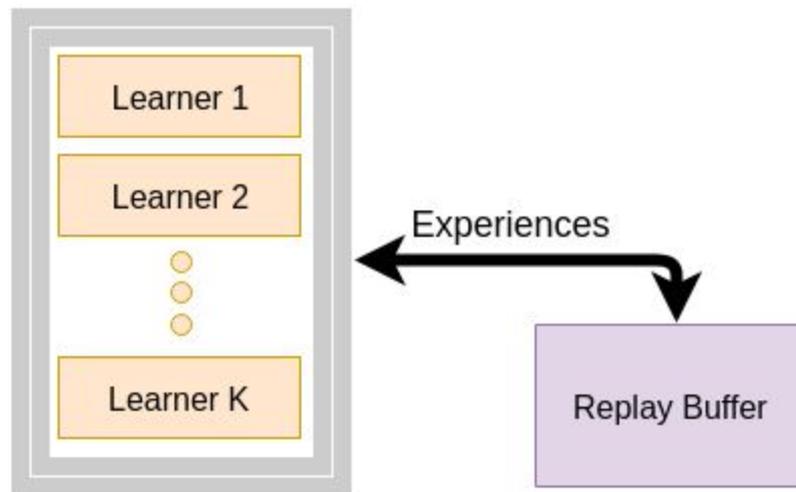
Why varying discount rates?



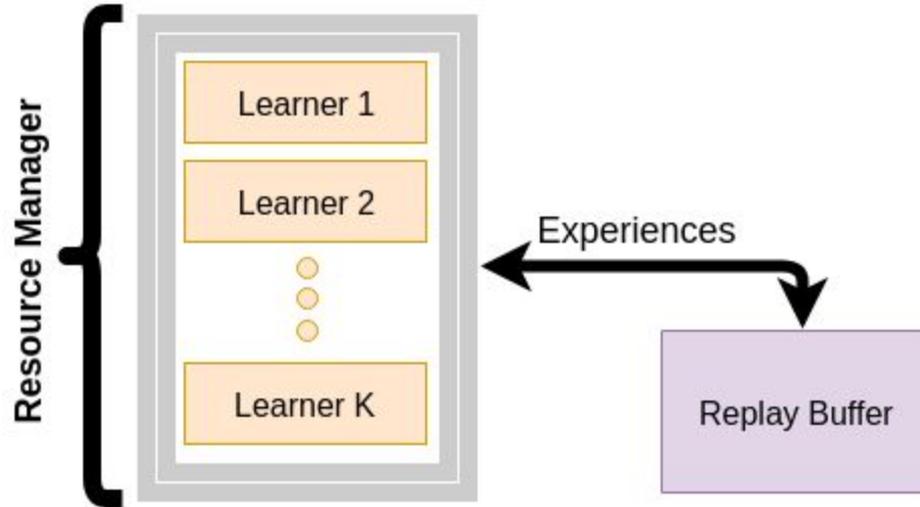
Why varying discount rates?



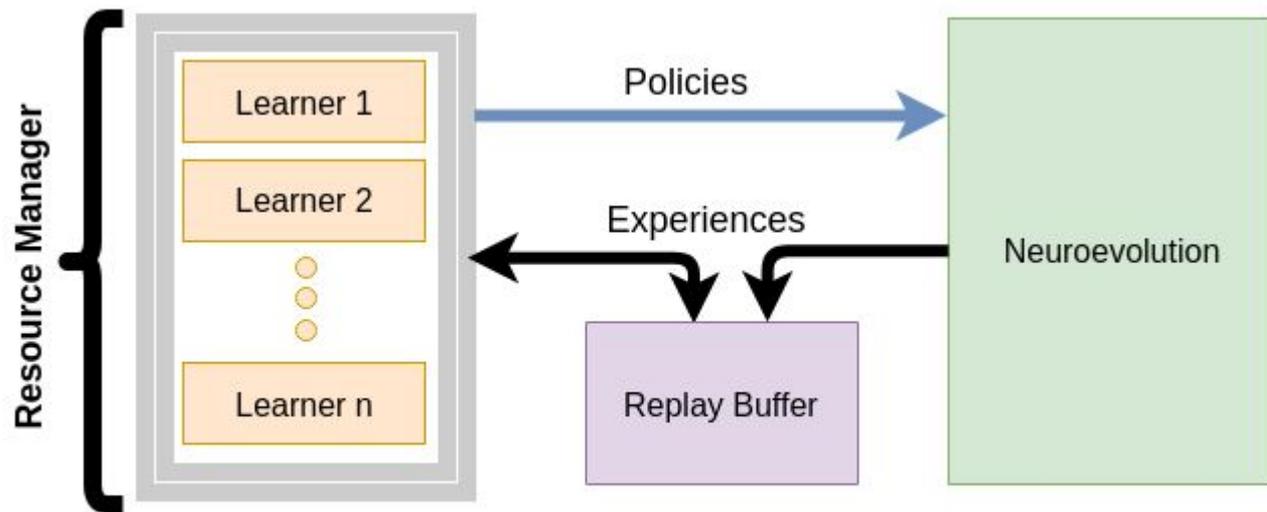
Back to Portfolio of Learners



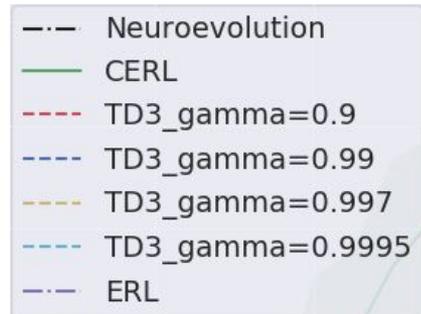
Adding a Resource Manager



Adding Neuroevolution

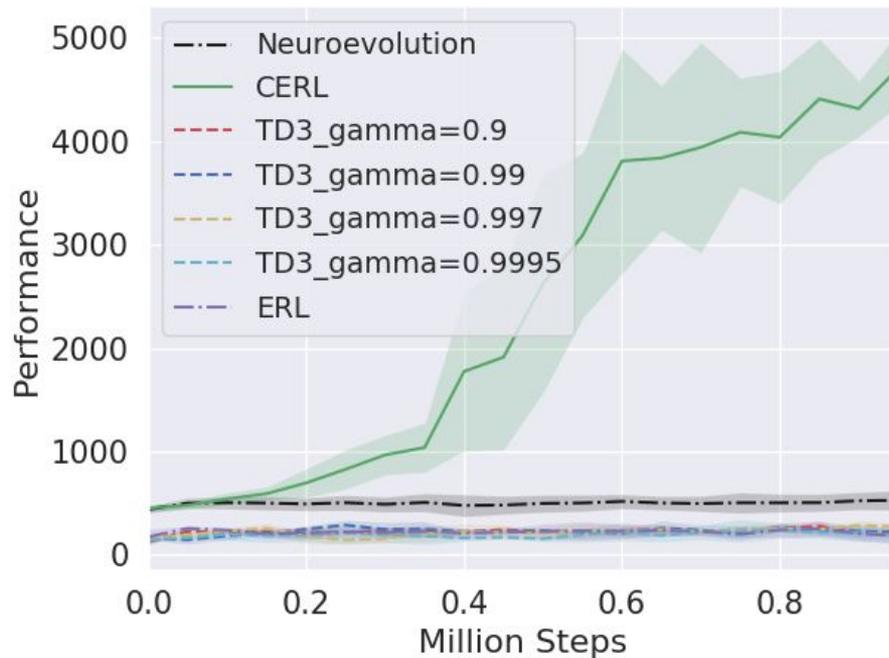


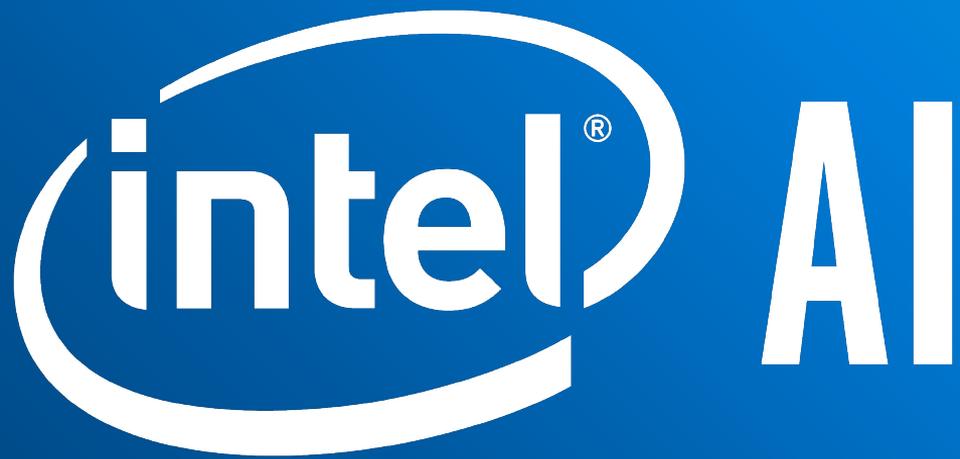
Experiment: Humanoid



Experiment: Humanoid

- Solves Humanoid under 1 million samples
- TD3 learners fail entirely
- Neuroevolution ~ 62.5 million samples



The image features the Intel AI logo in white against a blue background. The logo consists of the word "intel" in a lowercase, sans-serif font, with a registered trademark symbol (®) to its upper right. A thick white swoosh starts above the "i", loops around the top and left sides of the word, and ends below the "l". To the right of this swoosh, the letters "AI" are displayed in a large, bold, uppercase, sans-serif font. The background is a solid blue color with a faint, light blue geometric pattern of interconnected nodes and lines, resembling a network or data structure, primarily visible on the left and right sides.

intel[®] AI