DeepMind

Deep Reinforcement Learning with Plasticity Injection

Evgenii Nikishin, Junhyuk Oh, Georg Ostrovski, Clare Lyle, Razvan Pascanu, Will Dabney, André Barreto

Open Review .net

Overview

- Deep RL agents **lose plasticity** during training, but detecting it is challenging
- Plasticity Injection is a tool for careful **analysis** of the phenomenon
- Plasticity Injection **reincarnates** agents that lost plasticity without the need to re-train
- Agents with Plasticity Injection **save compute** by dynamically growing networks

Background on Plasticity Loss

On a sequence of policy evaluation tasks, a network becomes worse at predicting values

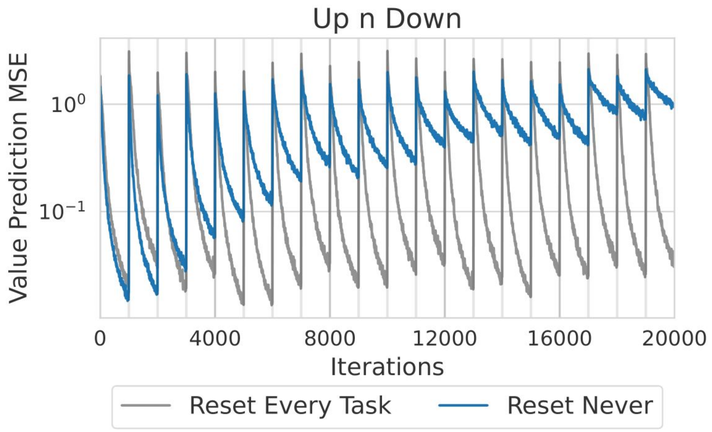

Challenges in Deep RL

- No clear measure of network plasticity
- Performance in deep RL can be bottlenecked by many factors including exploration

How to tell if an agent cannot improve due to loss of plasticity?

Plasticity Injection

At some point, freeze the network and create a new, learning a residual to predictions

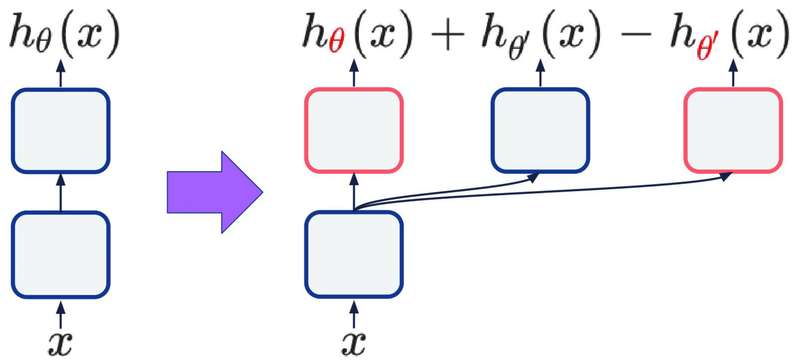

Key properties:

- Same number of trainable parameters
  → roughly **keeps representational power**
- Unaffected predictions after intervening
  → **excludes exploration** confounders

Diagnosing Loss of Plasticity

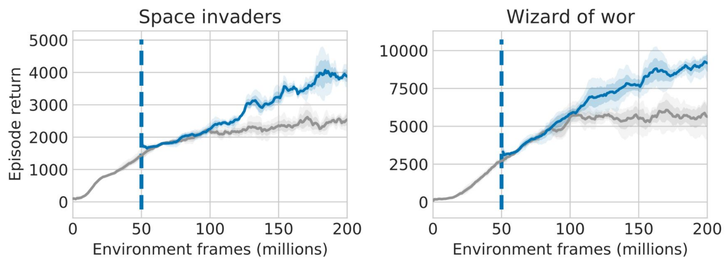

Improves performance → agent lost plasticity

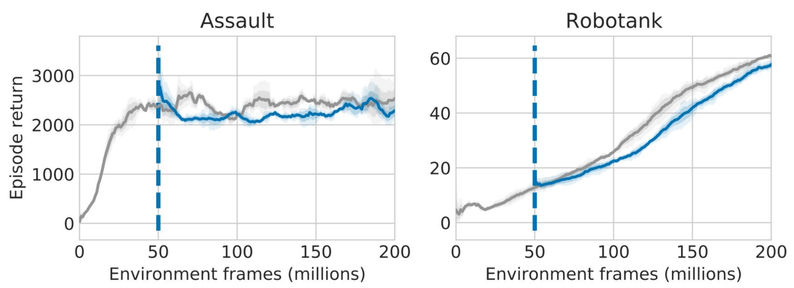

Agent is stuck due to exploration / learns well → no effect from plasticity injection

Results on 57 Atari Games

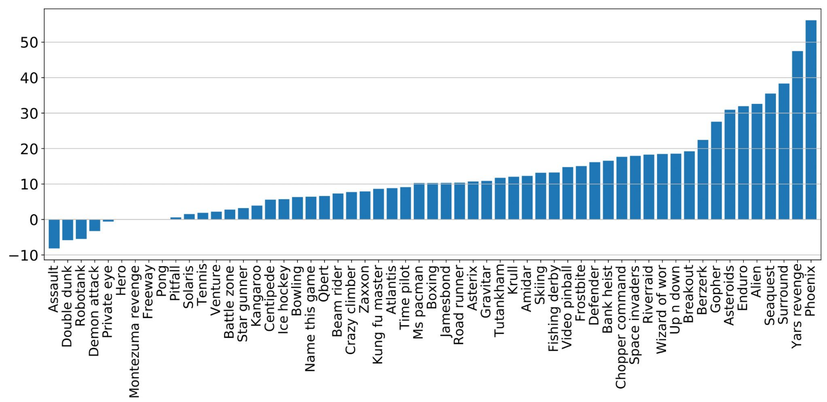

Agents lose plasticity within their lifetime

Reincarnating with Injection

Keep learning after plasticity loss efficiently

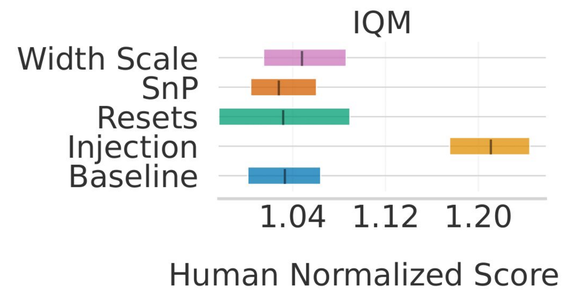

Better than re-training-free alternatives

Dynamic Network Growth

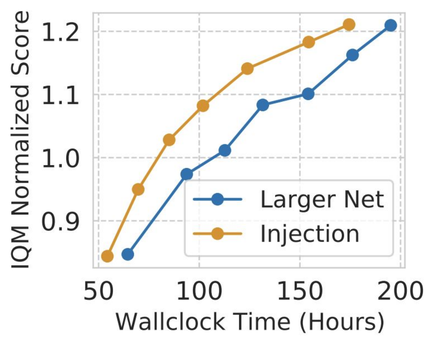

Same IQM as large nets from start but cheaper