

Autoverse: Evolving Symbolic Neural Cellular Automata Environments to Train Player Agents



Sam Earle, Julian Togelius
New York University



Motivation


We want to explore Open-Ended Learning (OEL)—and Unsupervised Environment Design (UED)—in game environments with **diverse mechanics**

For this, we need a **fast simulator** that can express a wide variety of games and lets us easily **mutate mechanics**.

Methods


We define a Domain-Specific Language (DSL) of local interactions over grid-worlds, encompassing **mazes, dungeons, sokoban-like**s, and more.

(In “Power Puzzle”, for example, the player builds electrical circuits.)



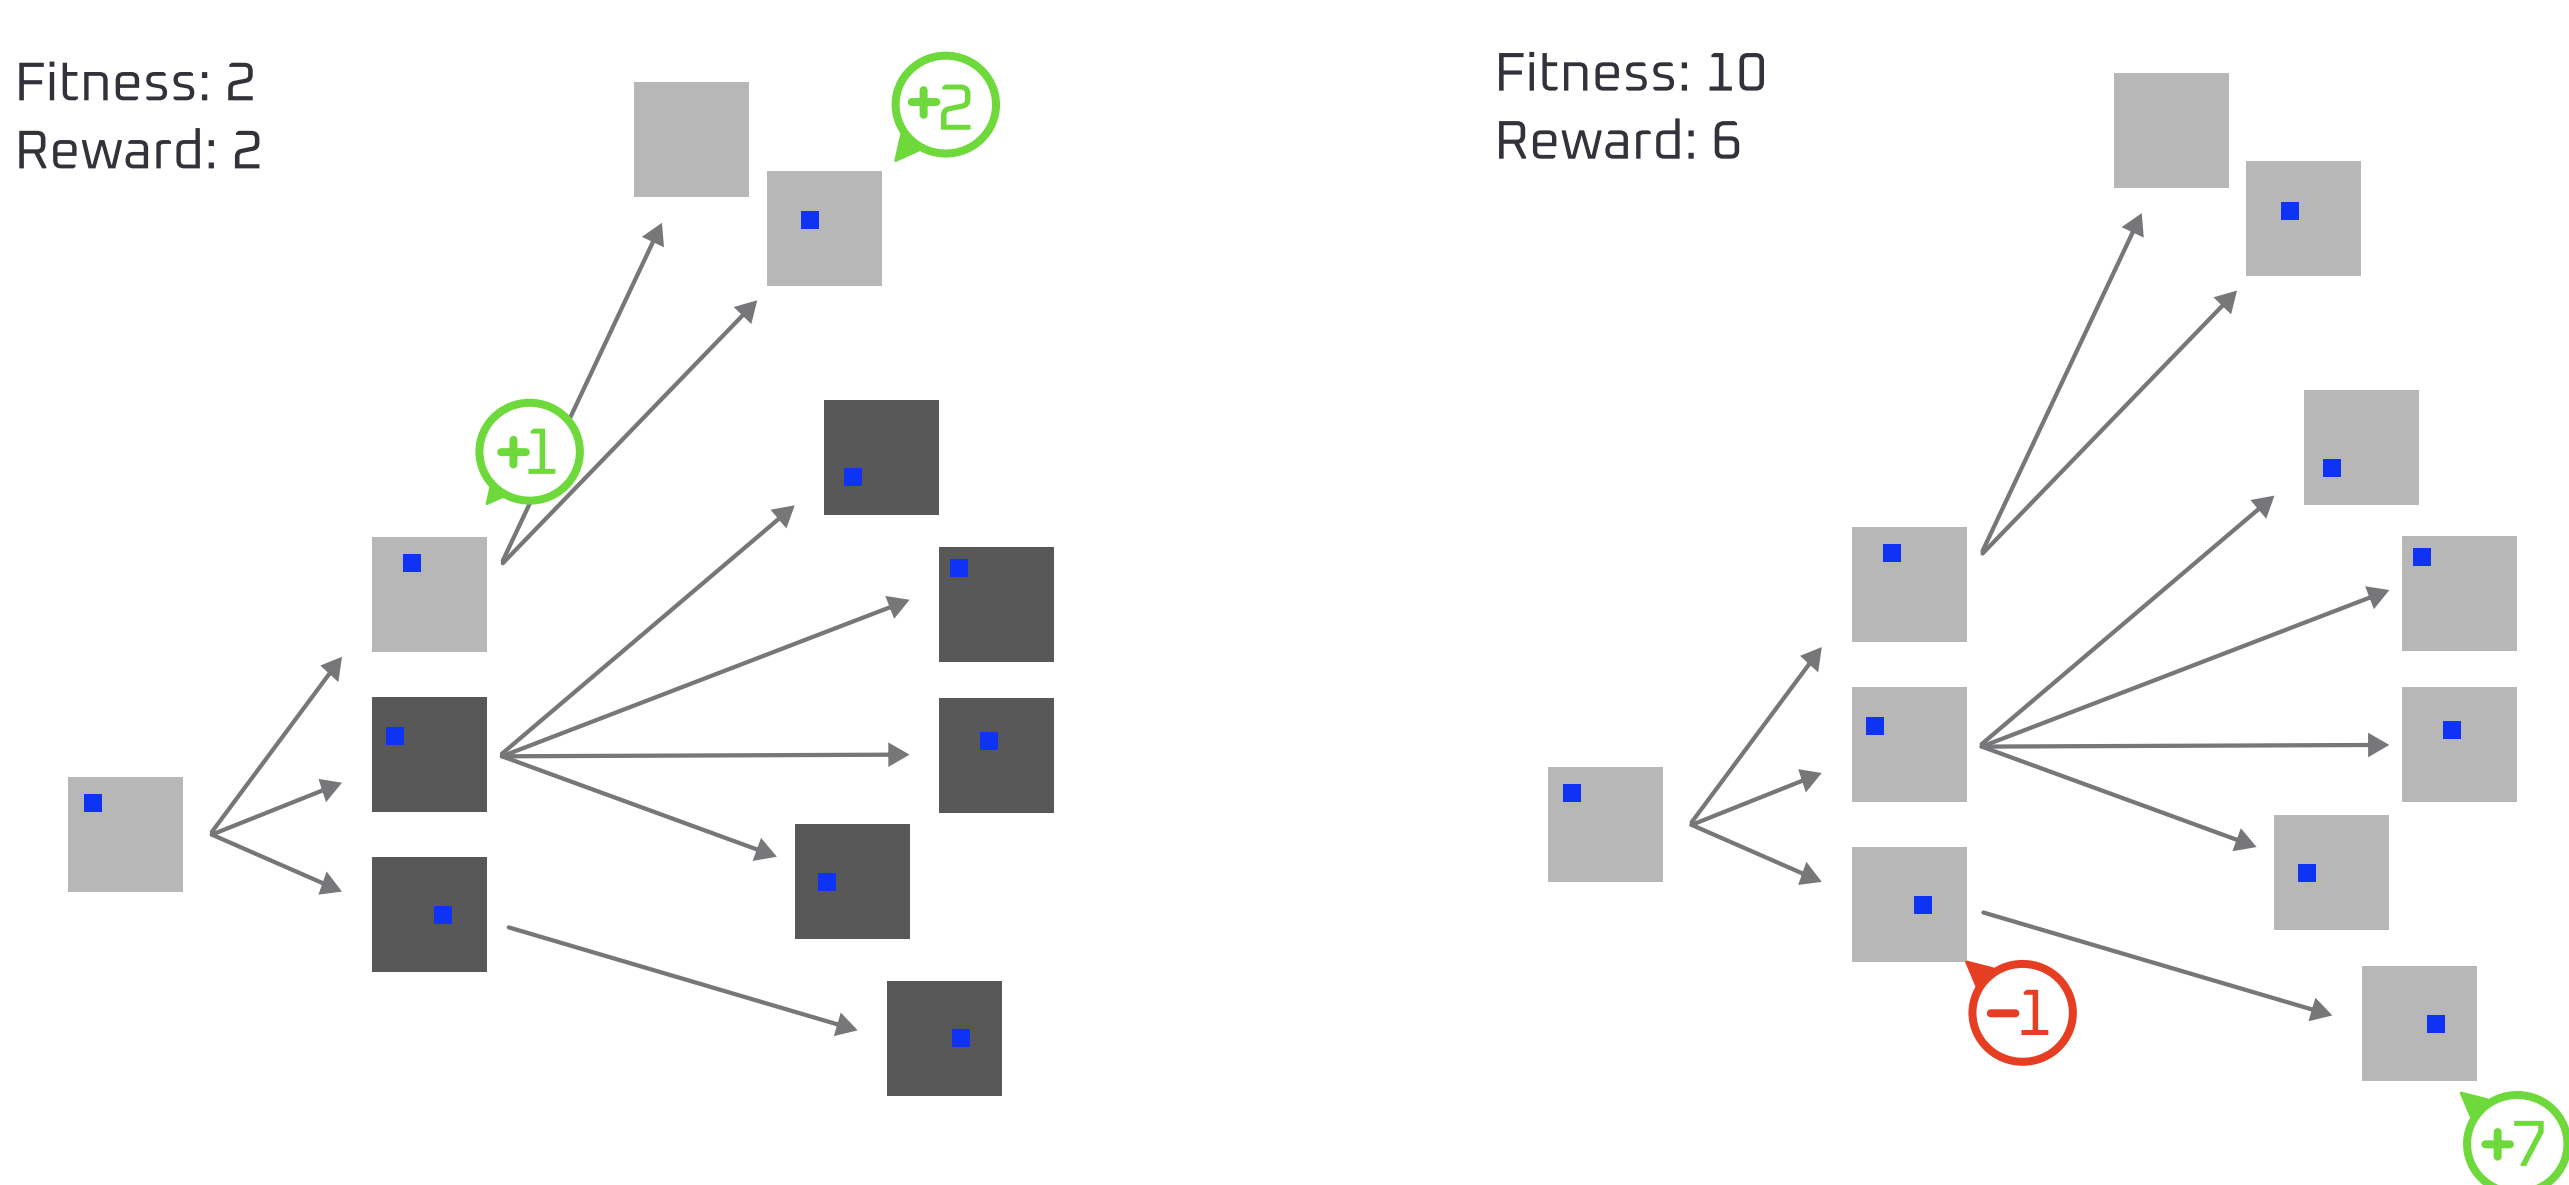
Local rules, global dynamics.

Local pattern **rewrite rules** can be implemented as **convolutions** (we use jax) allowing for parallel execution on the GPU.



Pattern rewrite rules as convolutions.

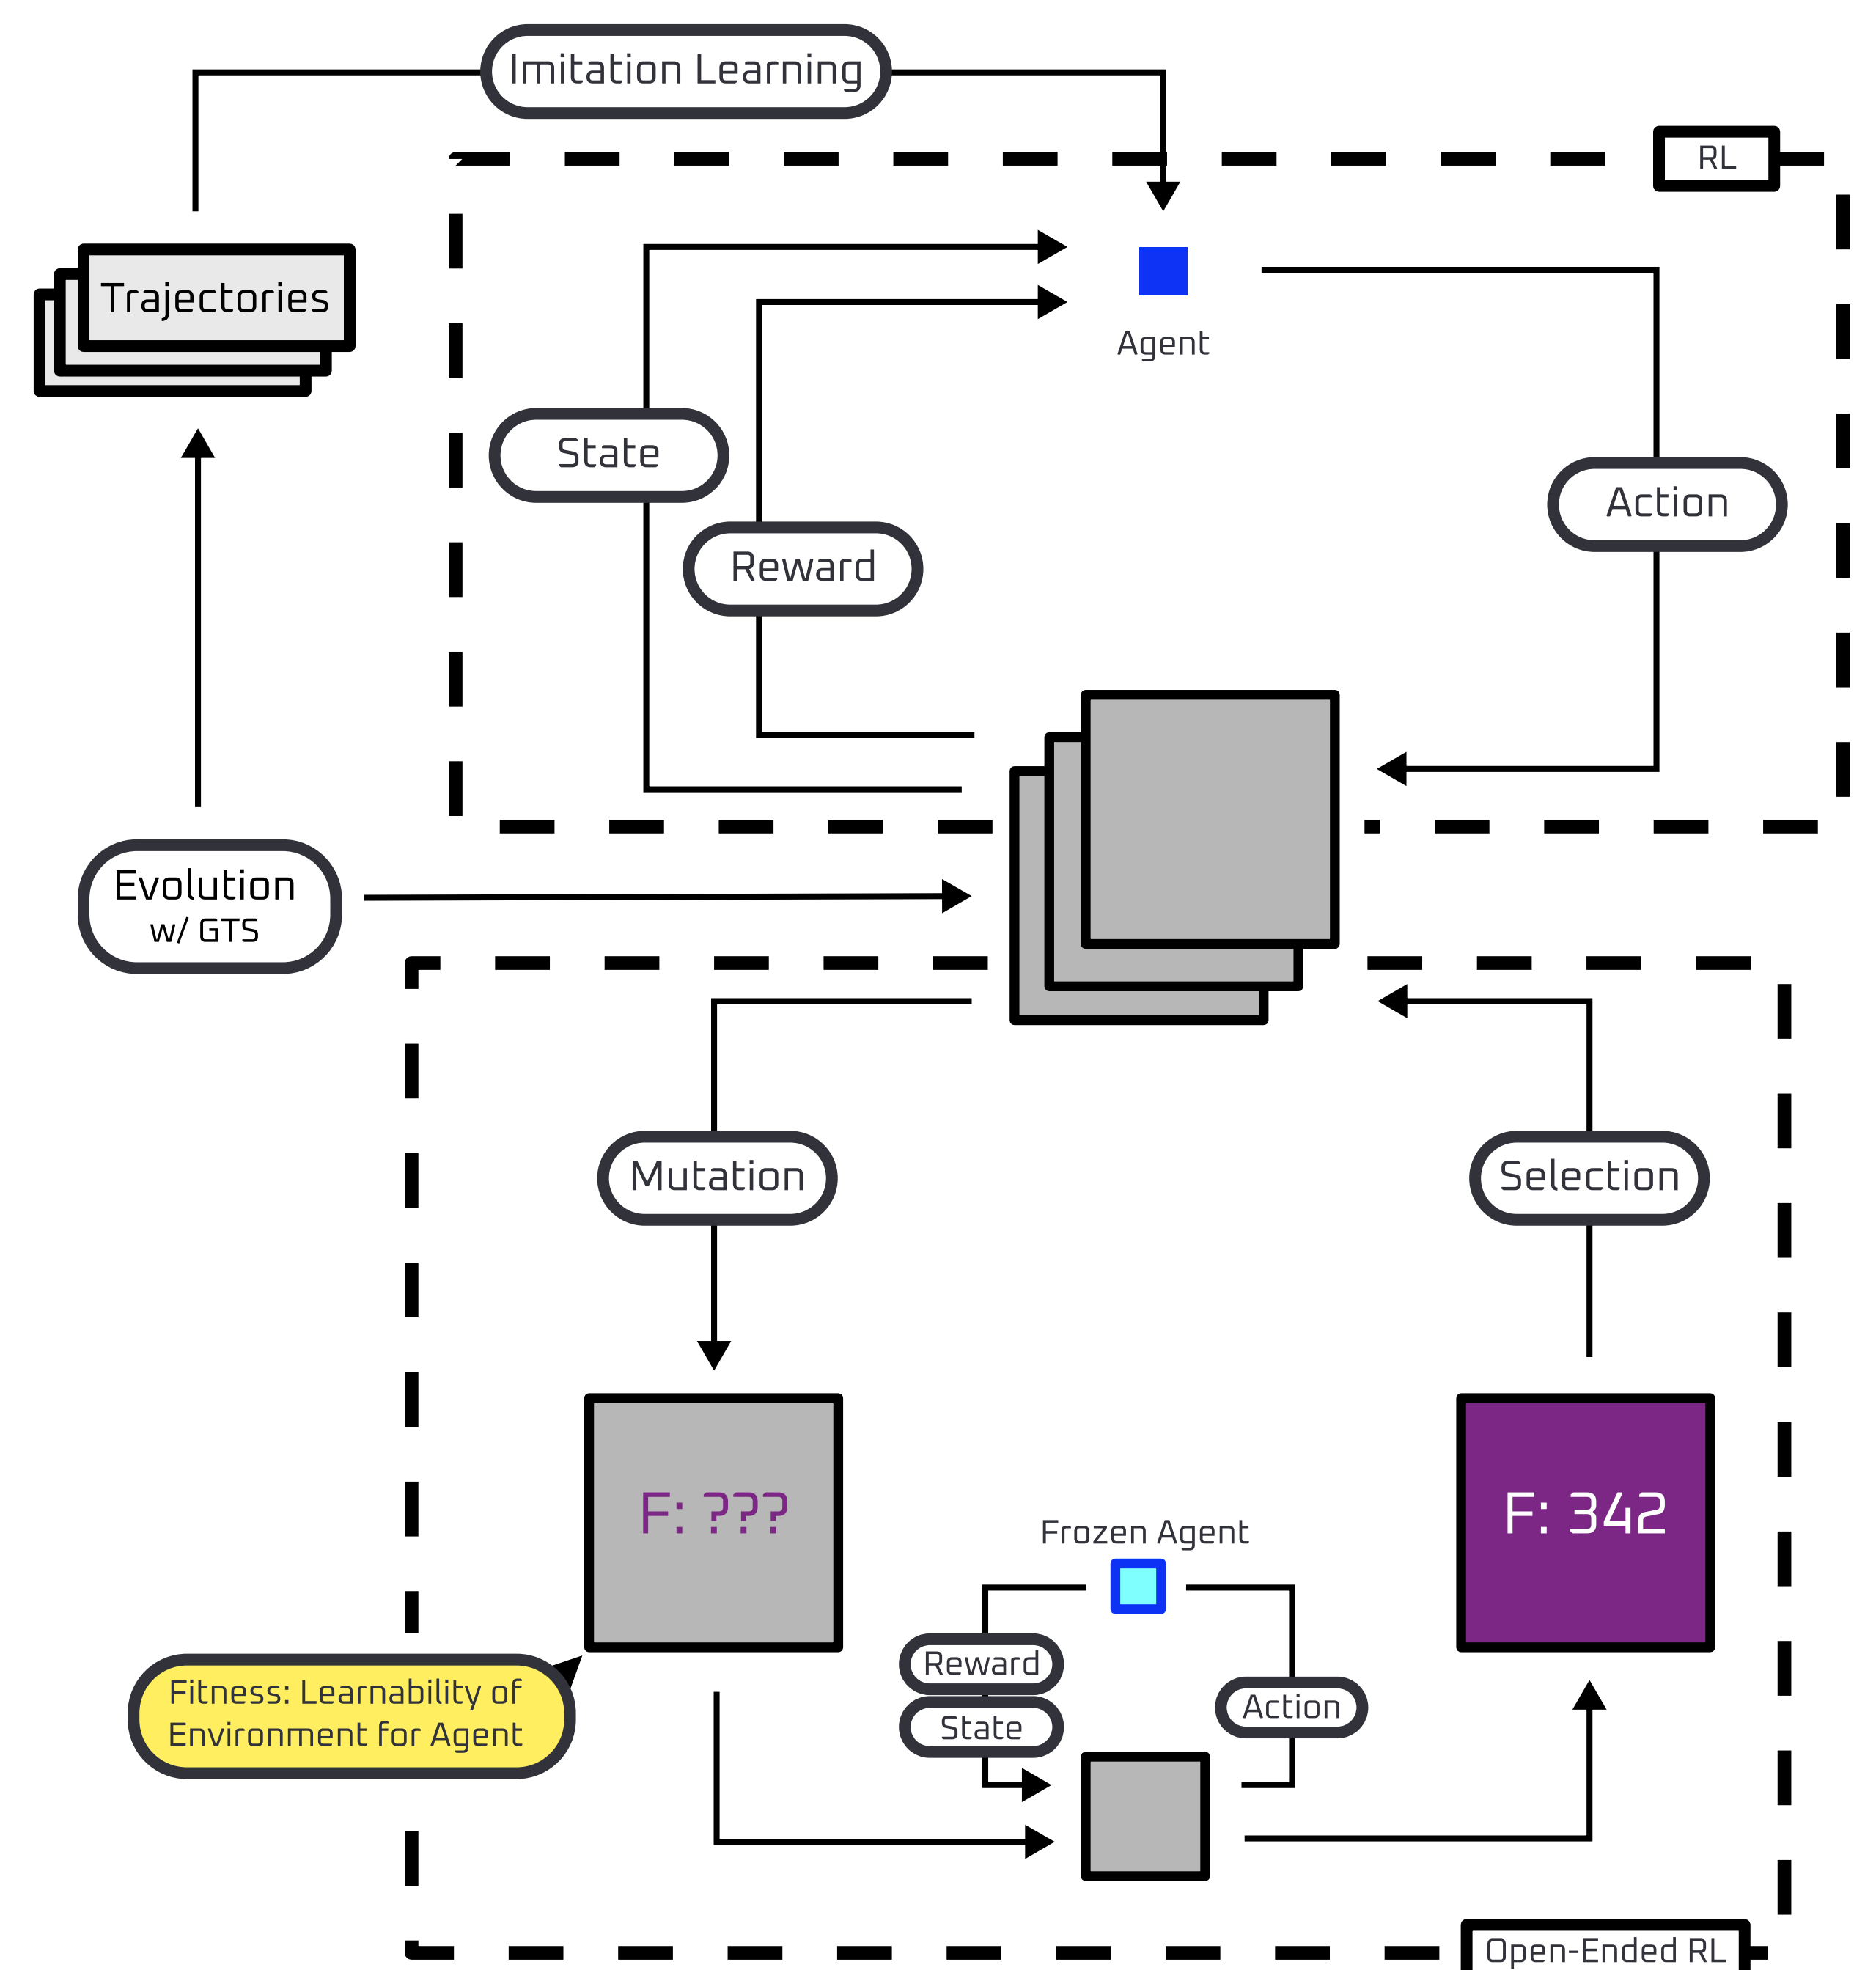
We mutate rewrite rules, selecting for environments where the best solution takes a long time to find with **greedy search**. (This favors **deceptive** environments.)



Evolve environments to maximize complexity.

We Imitation Learn on player trajectories from search, then continue to Reinforcement Learn in evolved environments.

Use search to jump-start Reinforcement Learning.



Continue evolving environments to maximize surprise.

During RL, we continue evolving environments to maximize value function error of the player agent.


Results

Running evolution (for search-based complexity) for longer, to generate a set of training environments, leads to better generalization (w.r.t. performance on environments from a separate evolutionary run).


More evolved environments, better generalization.

	train mean	val. mean	test mean
load gen			
0	130.88 ± 3.73	51.38 ± 1.24	90.79 ± 6.79
10	174.32 ± 29.53	67.03 ± 3.61	123.92 ± 6.51
20	121.71 ± 1.07	71.18 ± 0.25	133.71 ± 9.60


Many flavors (entropically speaking)...



stable



chaotic



and balanced!