6. Body-Brain Coevolution

- So far, weights and structure of neural networks are evolved for a task.
- Optimize the policy for a given physical structure.
- Evolutionary algorithms can optimize both body and brain.
- Physical body could be a substrate for open-ended evolution.

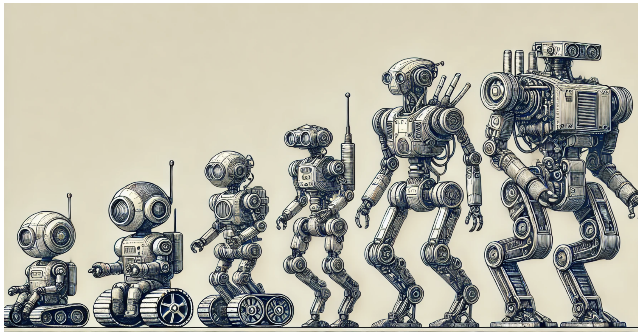

Joint Evolution of Policy and Structure

- The agent's environment, including body structure, is parameterized.
- Both the agent's policy (weights) and the body structure (parameters) evolved.
- Coevolution allows discovering optimal body design for the task.
- The policy is optimized for the body, making it believable.

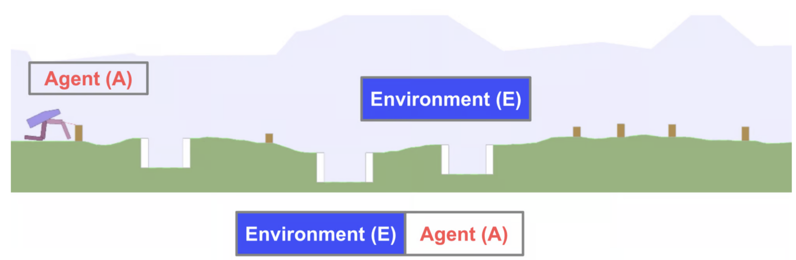

Rewarding Difficult Design Choices

- Agents are rewarded for trying more difficult designs.
- Example: Carrying the same payload using smaller legs might be rewarded more.
- Example: Balanced walk with longer legs may be rewarded more.
- Such designs can serve as stepping stones for better agents.

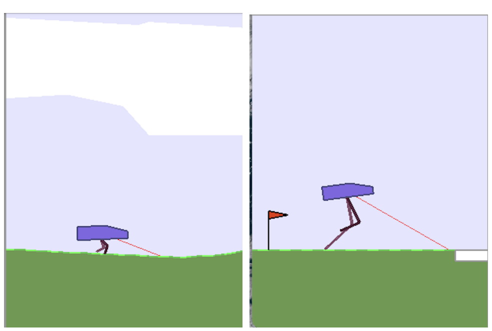

Optimizing Body for a Task

- The agent evolves a body structure better suited for specific tasks.
  - A longer and heavier rear leg helps maintain balance and get over obstacles and gaps.
- The policy coevolves with the body, resulting in natural behavior for the body.

Demo:

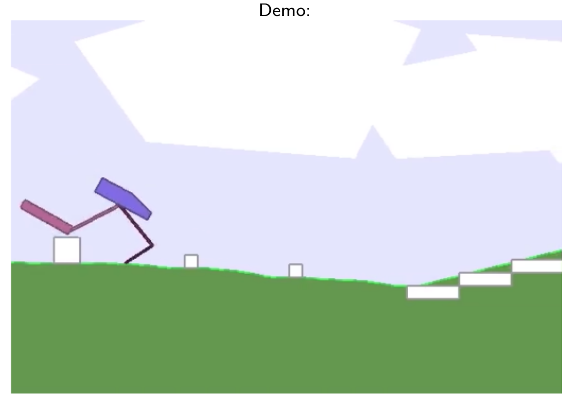

Optimizing Under a Constraint

- What if evolved to optimize material use?
  - Smallest legs that still allow running fast.

Demo:

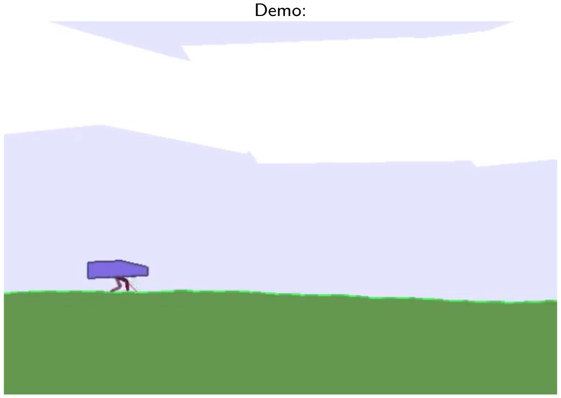

Combining Constraints

- Lightest legs while getting over obstacles.
- Still longer rear leg, but as short as feasible.

Demo:

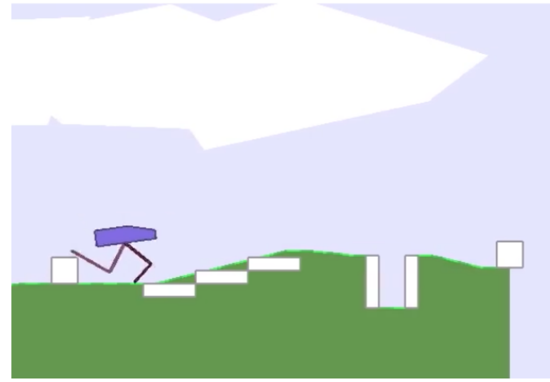

Removing Design Constraints

- What happens if design constraints are removed?
- Extreme designs may evolve that utilize loopholes.
- Example: A very tall agent falls over and lands near the exit.
- Design constraints and performance goals need to be balanced.

Demo:

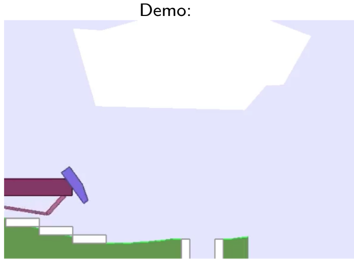

Demo:

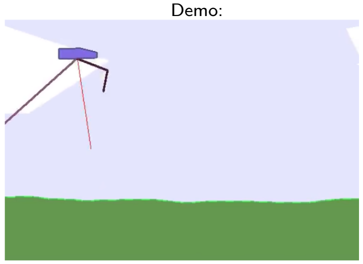

7. Co-Evolving Agents and Static Environments

- In addition to the body, the environment can also change
- E.g. the track for the bipedal walker.
- Co-evolution of problems and solutions.
- Open-ended evolution may result.

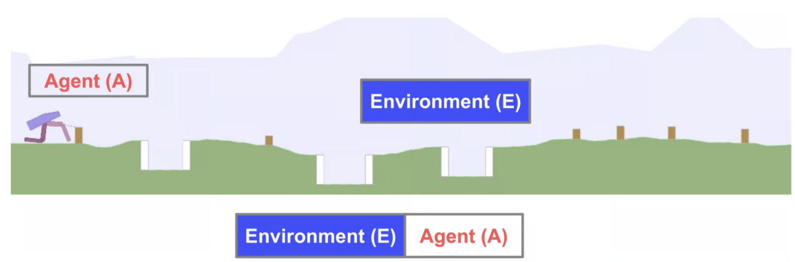

POET Algorithm Overview

- POET (Paired Open-Ended Trailblazer) is designed for co-evolving agents and environments.
- It generates new environments and agents, optimizing their performance over time.
- Three key tasks in each iteration:
  - Environment generation.
  - Agent optimization.
  - Agent transfer between environments.

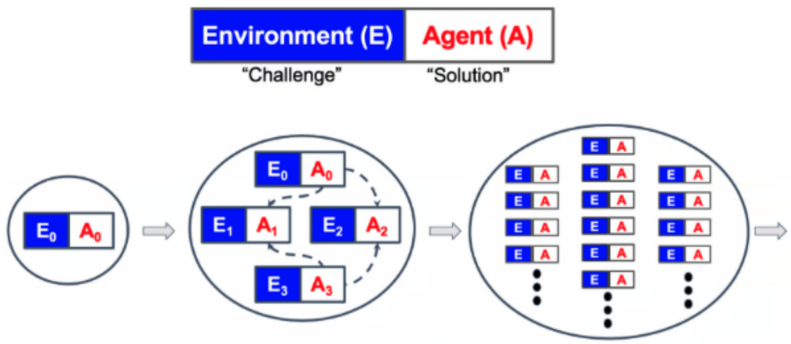

Environment Generation in POET

- Environments are generated by mutating parameters of existing ones.
- Parameters include stump height, gap width, stair height, number of stairs, and surface roughness.
- Only environments that provide suitable challenges and novelty are added.

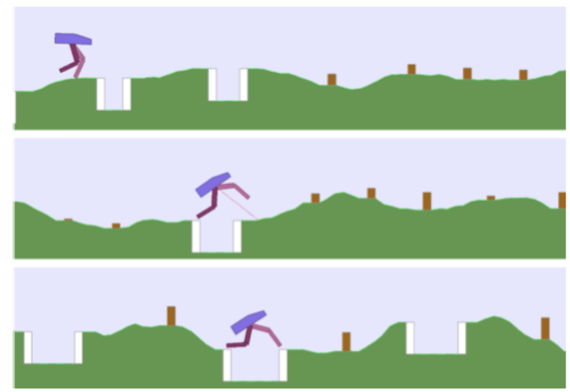

Environment Generation in Enhanced POET

- Use a CPPN to generate y for each x.
- More varied and natural environments.

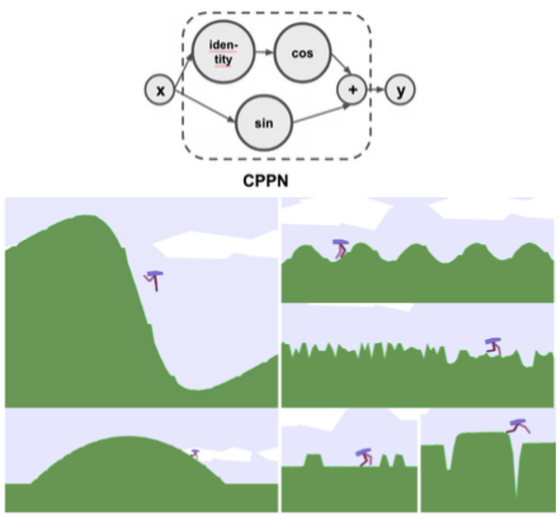

Agent Optimization in POET

- Agents optimized through neuroevolution (ES).
- The goal is to maximize the agent's performance in traversing the environment.
- Agent optimization is independent, which facilitates parallel processing.

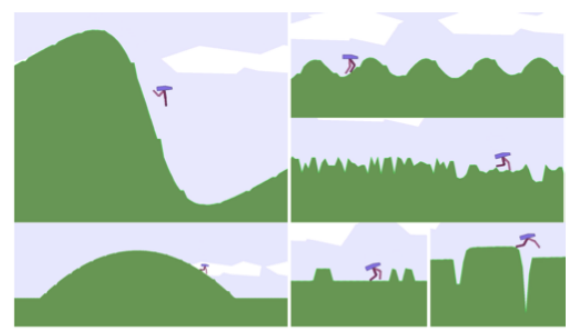

Agent Transfer in POET

- Agents are transferred between environments to foster adaptation.
- Successful strategies from one environment may help in another.
- Transfer helps agents escape local optima.
- Diversity important: Different agents may perform well in different environments.

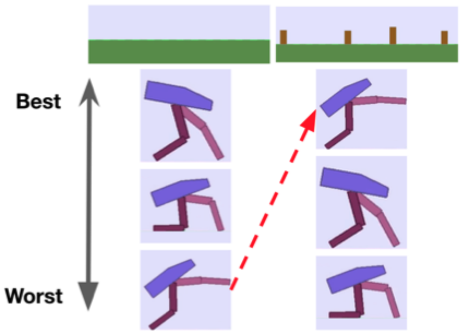

POET vs. ES

- POET generates novel solutions that ES alone cannot achieve.

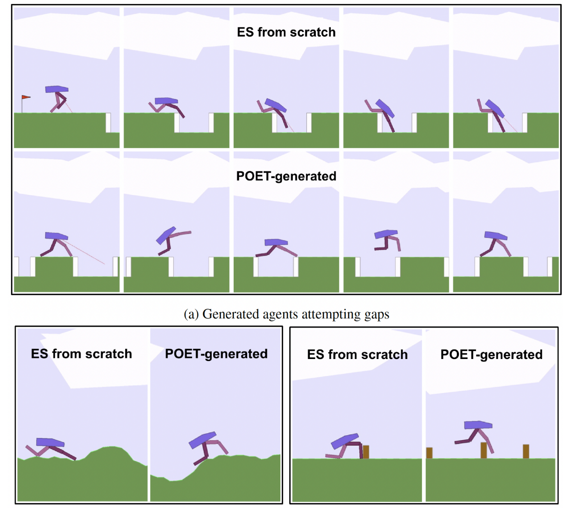

(a) Generated agents attempting gaps

(b) Generated agents on rough surfaces

(c) Generated agents attempting stumps

POET vs. Curricular Learning

- Curricular learning set up with gradually more challenging environments.
- POET finds solutions to much more complex environments.
  - Transfers from other environments form stepping stones.
  - POET utilizes stepping stones; curricular cannot.

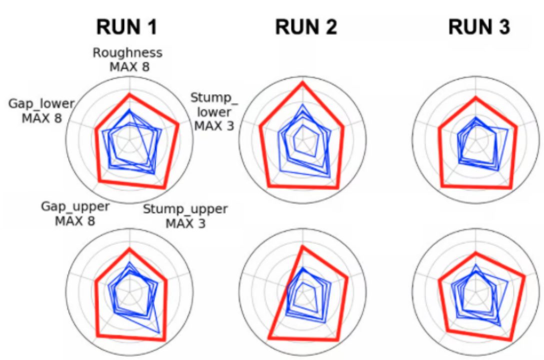

POET Results

- Agents successfully navigate complex terrain.

Demos:

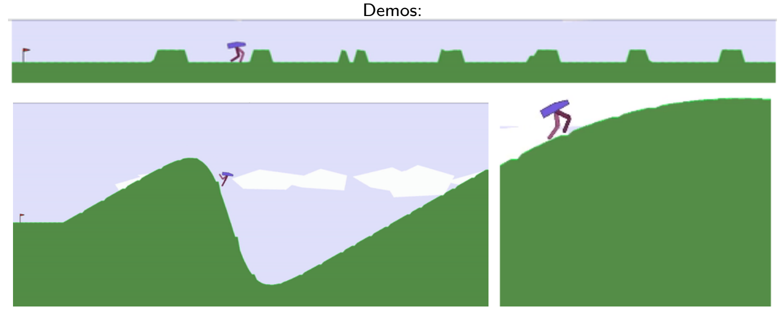

POET Results

- Agents successfully navigate complex terrain.

Demos:

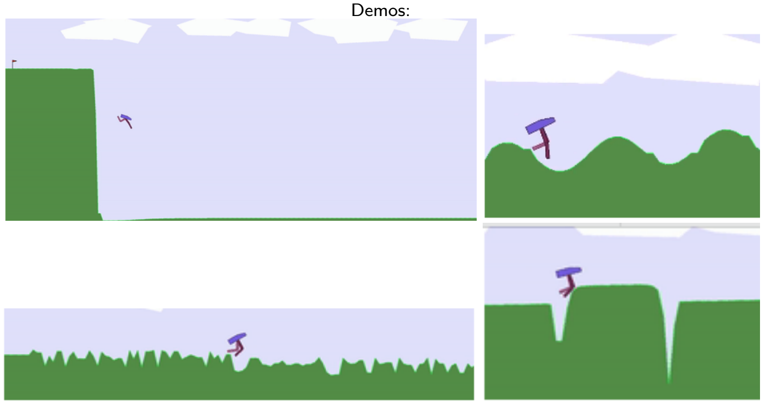

POET Performance

- Quantitatively, POET can solve harder problems.
- Enhanced POET can keep discovery going longer.

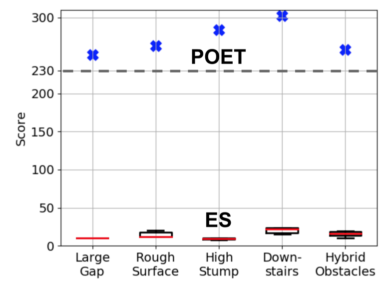

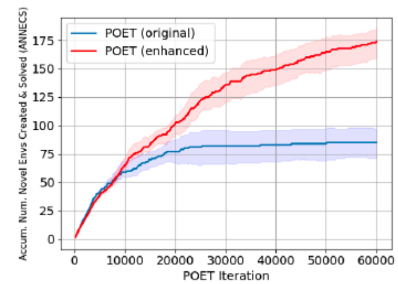

Coevolution of Agent and Dynamic Environments

- A quadruped chaser robot is coevolved with an escapee.
- As the chaser evolves to catch, the escapee evolves to evade.
- This continuous feedback loop enhances the capabilities of both agents.
- A competitive coevolution process.

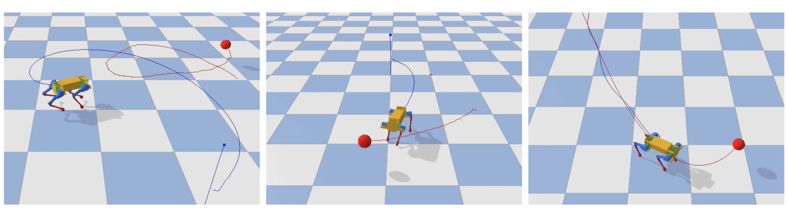

Encouraging Behavioral Diversity

- Diverse strategies are encouraged through different $d_{\min}$ (i.e. catch) thresholds for escapee robots.
- Smaller thresholds encourage quick dodges, while larger ones promote broader evasive movements.
- This diversity allows for the development of robust chaser strategies.

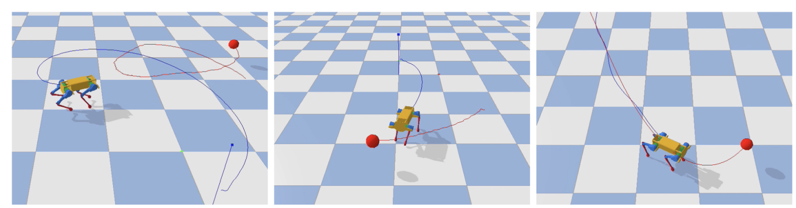

Baseline Comparisons

- Three baseline methods emulate escapee movement:
  - Cone: chasing forward
  - Circular: chasing random moves
  - Zigzag: chasing a zigzagging escapee
- Evolve against each one and compare to coevolution.

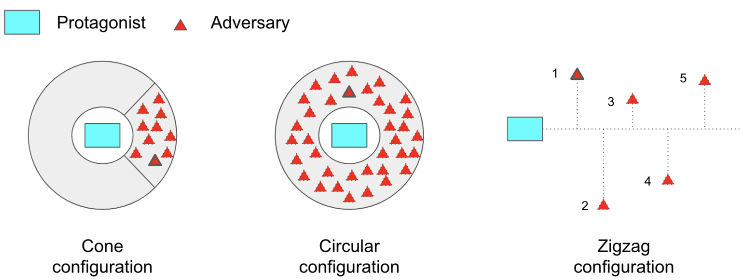

Coevolution Performance vs. Baseline

- Escapee moves with constant speed in a sine pattern (with varying amplitude and frequency).
- Coevolved agents catch it early (indicated by a dot).
- Baseline policies struggle to catch, and often the robot falls (x)

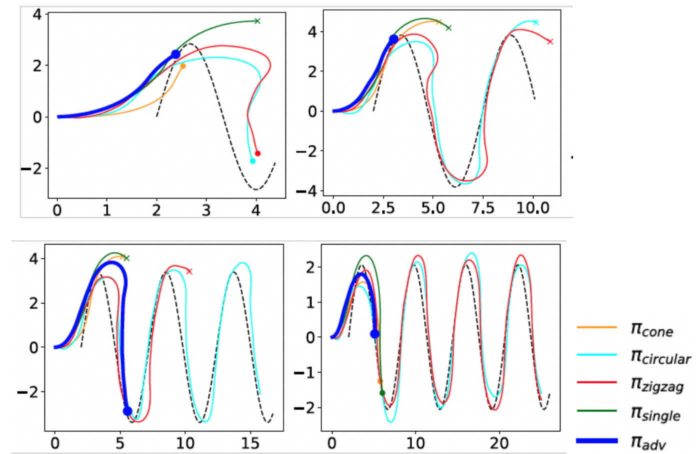

https://youtu.be/XNWYjOJZGJA

Is Neuroevolution Open-ended?

- Ingredients for open-ended neuroevolution already exist.
  - Large populations, weak selection, and neutral mutations can be scaled up with computational power.
  - Extinction events, evolvable representations can promote evolvability.
  - Expressive encodings could enable continuous exploration of more complex solutions.
  - Coevolution of agents and environments (body, task, adversaries) presents new challenges for evolution to continue.

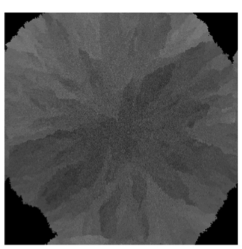

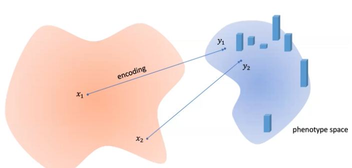

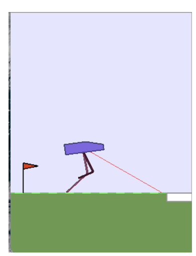

Complexity Can Be Created, but Is It Open-Ended?

- The process runs out of steam eventually.
  - E.g. not seeing major transitions.
- What is still missing?
  - E.g. do agents need to modify the environment with permanent artifacts?
- Why do even we care whether complexity is open-ended?

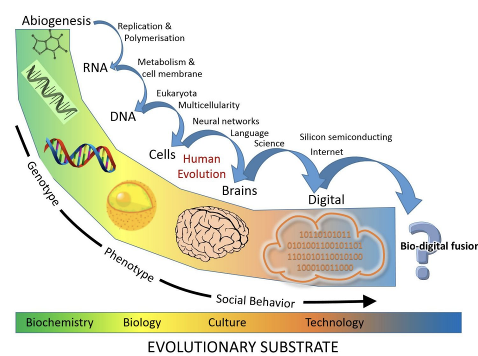

Ultimate Goal: General Intelligence

- Better understanding of biological evolution and the origins of intelligence.
- The ability to adapt indefinitely is a key feature of artificial general intelligence.
- We need open-ended neuroevolution to develop artificial systems at the same level.