cadmo
center for algorithms, discrete mathematics & optimization

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Evolutionary Algorithms in Deep Learning and Reinforcement Learning

Johannes Lengler (ETH Zürich)

DMV-ÖMG Jahrestagung, September 2021

Should we use EA or classical RL?

Evolution Strategies as a Scalable Alternative to Reinforcement Learning

Tim Salimans **Jonathan Ho** **Xi Chen** **Szymon Sidor** **Ilya Sutskever**
OpenAI

Abstract

We explore the use of Evolution Strategies (ES), a class of black box optimization algorithms, as an alternative to popular MDP-based RL techniques such as Q-learning and Policy Gradients. Experiments on MuJoCo and Atari show that ES is a viable solution strategy that scales extremely well with the number of CPUs available: By using a novel communication strategy based on common random numbers, our ES implementation only needs to communicate scalars, making it possible to scale to over a thousand parallel workers. This allows us to solve 3D humanoid walking in 10 minutes and obtain competitive results on most Atari games after one hour of training. In addition, we highlight several advantages of ES as a black box optimization technique: it is invariant to action frequency and delayed rewards, tolerant of extremely long horizons, and does not need temporal discounting or value function approximation.

Advantages of EA

- simple, easy to apply
- black-box
- hillclimbing + diversity
- highly parallelizable
- very flexible framework (e.g. self-adaptation, genetic programming)
- crossover
- populations(!!)

Drawbacks of EA

- no gradients (slow)
- prefer small to medium search spaces
- black-box
- susceptible to representation

Evolving Algorithms

Neural Architecture Evolution in Deep Reinforcement Learning for Continuous Control

Jörg K.H. Franke[*,1], Gregor Koehler[*,2], Noor Awad[1], Frank Hutter[1,3]
[1]University of Freiburg
[2]German Cancer Research Center (DKFZ)
[3]Bosch Center for Artificial Intelligence

Evolved Policy Gradients

Rein Houthooft[1] Richard Y. Chen[1] Phillip Isola[1 2 3] Bradly C. Stadie[2] Filip Wolski[1]
Jonathan Ho[1 2] Pieter Abbeel[1 2]

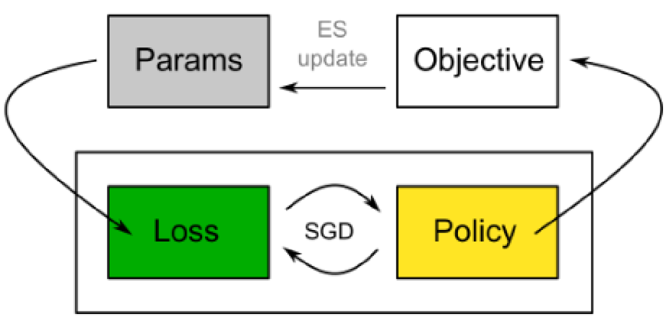

Evolutionary Architecture Search For Deep Multitask Networks

Jason Liang, Elliot Meyerson, and Risto Miikkulainen
Sentient Technologies, Inc. and
The University of Texas at Austin

Fig. 2: Sample evolved topologies of modules for the Omniglot multitask learning benchmark.

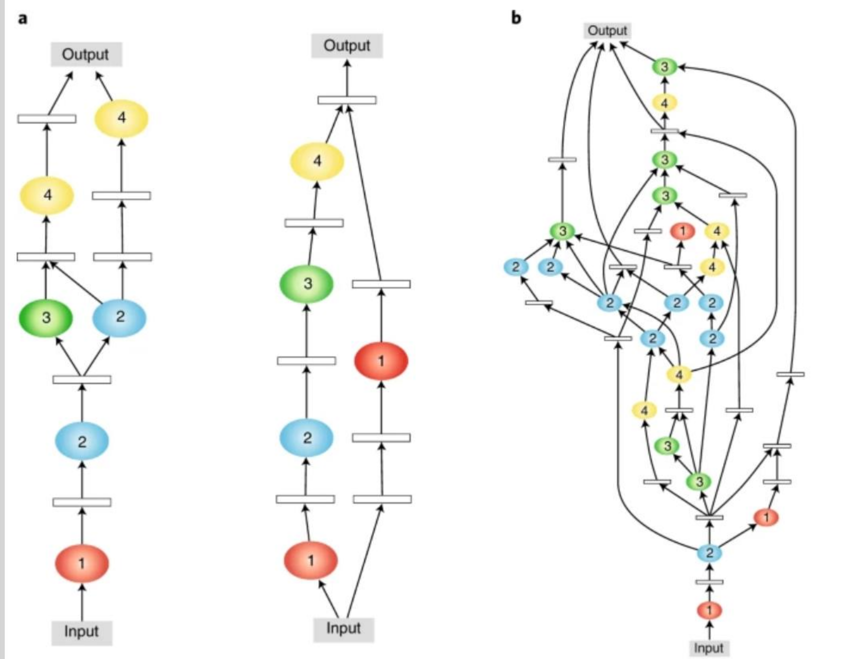

Topology for Gurmukhi and Mujarati

Topology for Angelic, an invented alphabet

Evolving Algorithms

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

AutoML-Zero: Evolving Machine Learning Algorithms From Scratch

Esteban Real [*1] Chen Liang [*1] David R. So [1] Quoc V. Le [1]

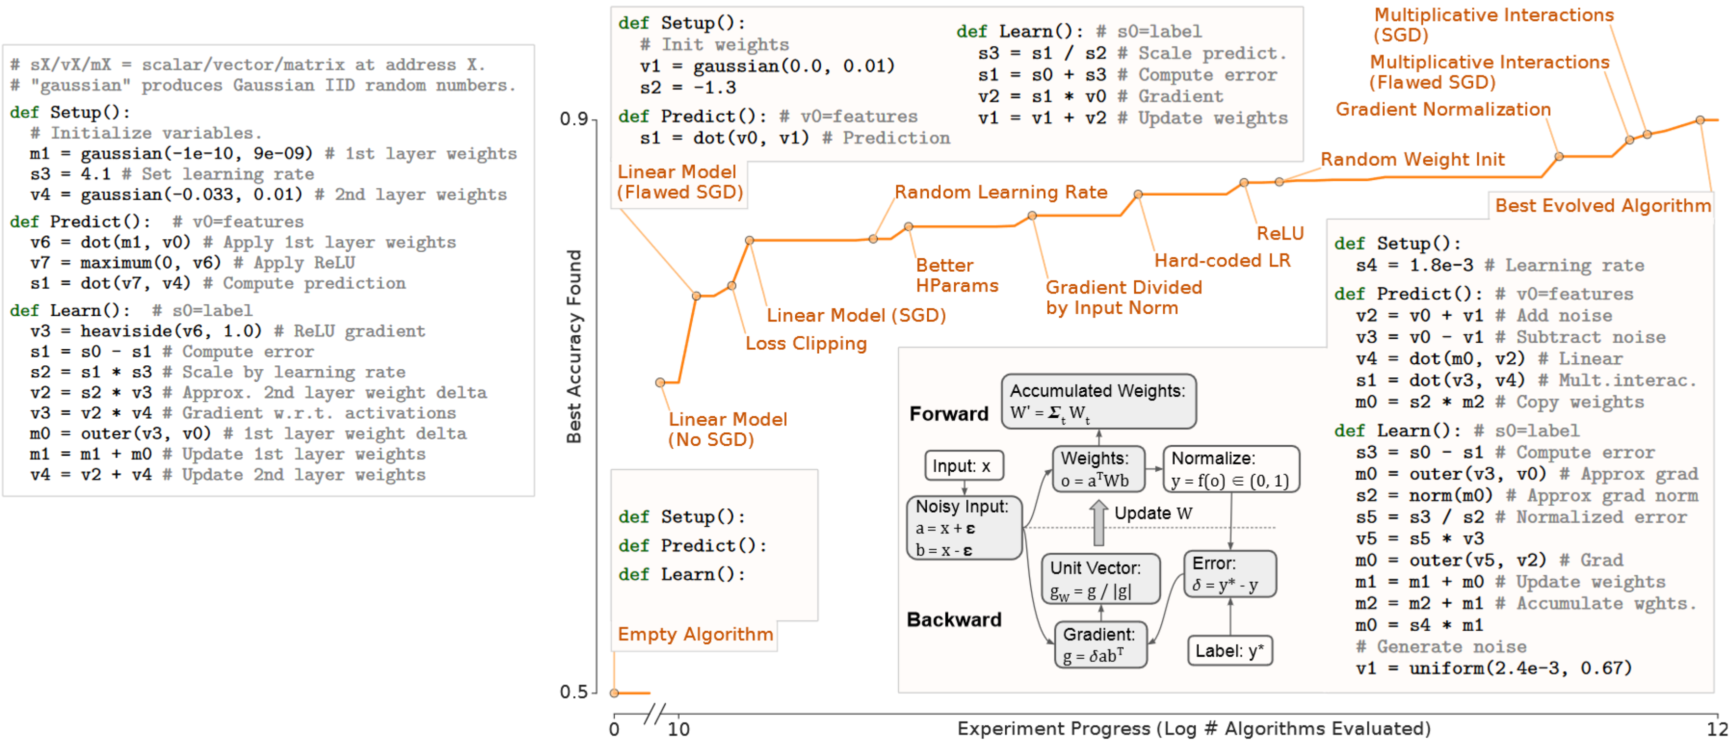

Search in Latent Space

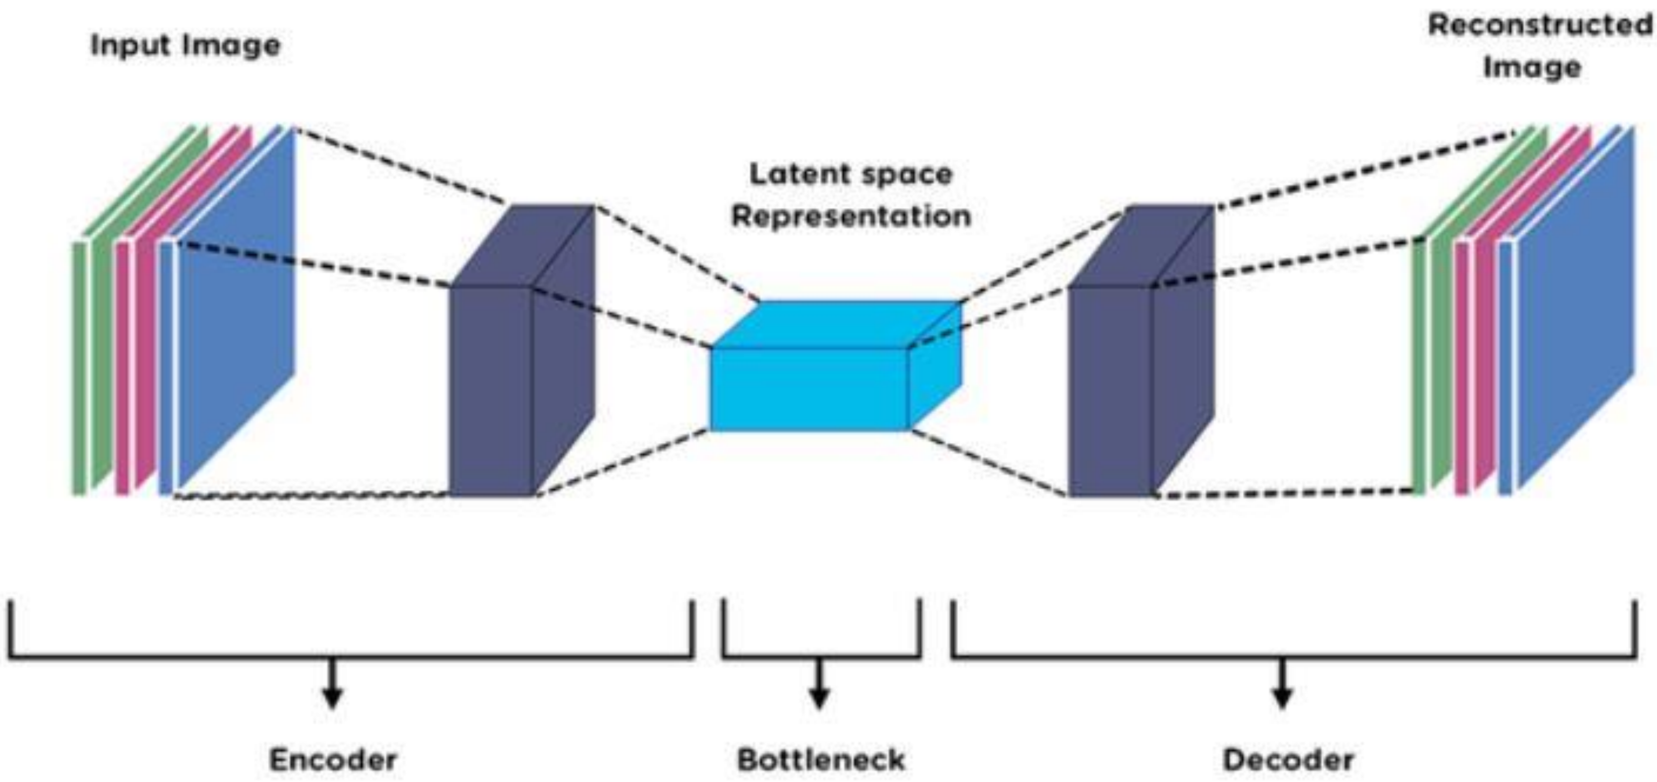

Search in Latent Space

Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

V. Volz, J. Schrum, J. Liu, S. M. Lucas, A. Smith, and S. Risi

GECCO '18, July 15–19, 2018, Kyoto, Japan

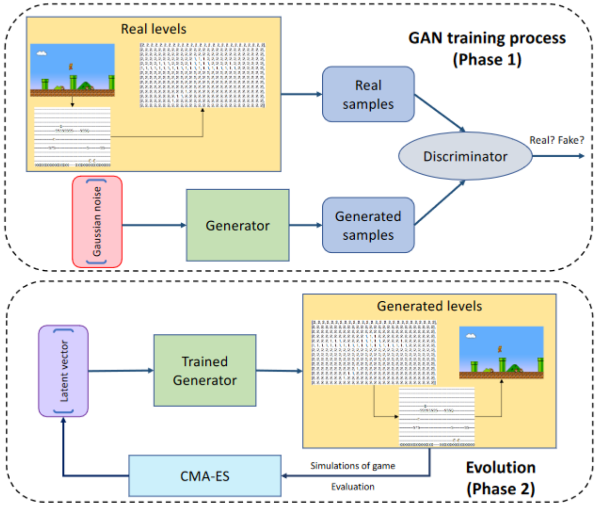

Figure 1: Overview of the GAN training process and the evolution of latent vectors. The approach is divided into two distinct phases. In Phase 1 a GAN is trained in an unsupervised way to generate Mario levels. In the second phase, we search for latent vectors that produce levels with specific properties.

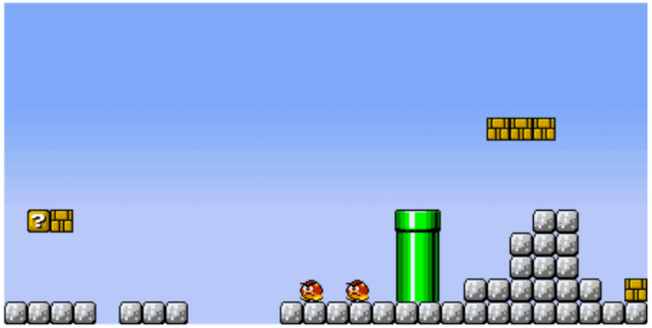

(a) Playable level maximizing jumps

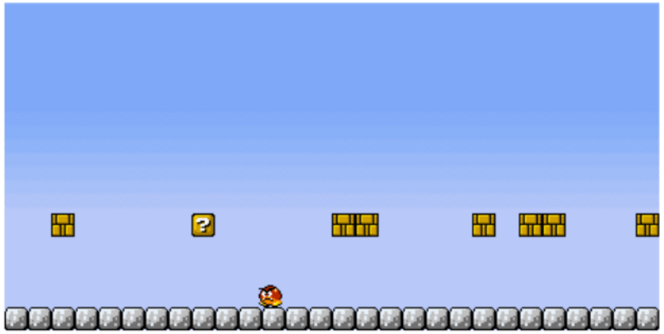

(b) Playable level minimizing jumps

Enhanced POET: Open-Ended Reinforcement Learning through Unbounded Invention of Learning Challenges and their Solutions

Rui Wang [1] Joel Lehman [1] Aditya Rawal [1] Jiale Zhi [1] Yulun Li [1] Jeff Clune [*2] Kenneth O. Stanley [*1]

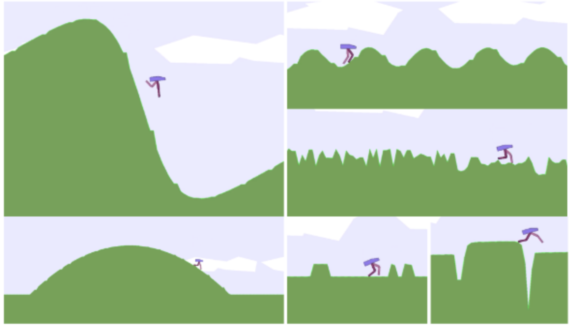

(b) Sample environments from a single run of Enhanced POET.

Populations of Agents

Article | Published: 30 October 2019

Grandmaster level in StarCraft II using multi-agent reinforcement learning

Oriol Vinyals, Igor Babuschkin, […]David Silver

Nature **575**, 350–354 (2019) | Cite this article

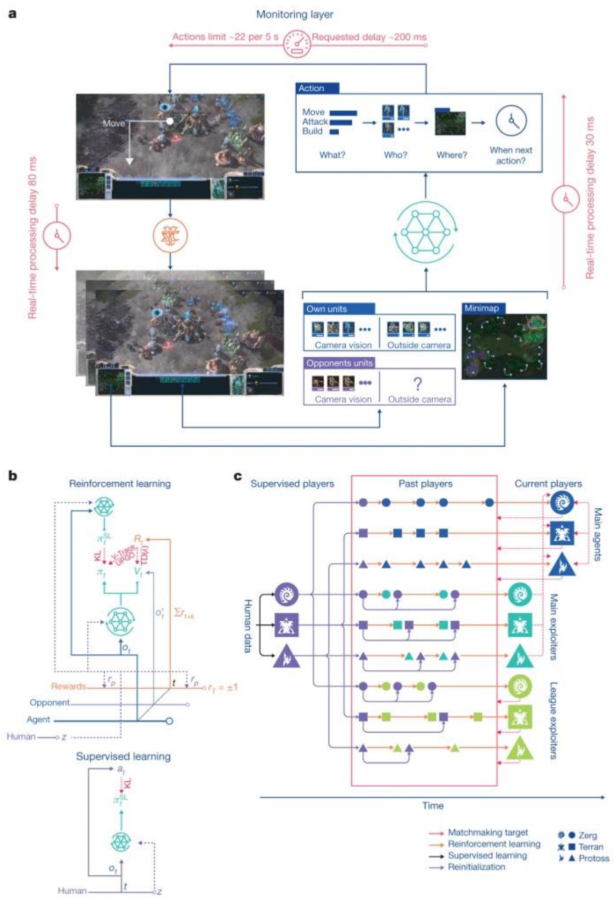

Published: 27 May 2015

Robots that can adapt like animals

Antoine Cully, Jeff Clune, Danesh Tarapore & Jean-Baptiste Mouret

Nature **521**, 503–507 (2015) | Cite this article

Figure 3: Main experiments and results.

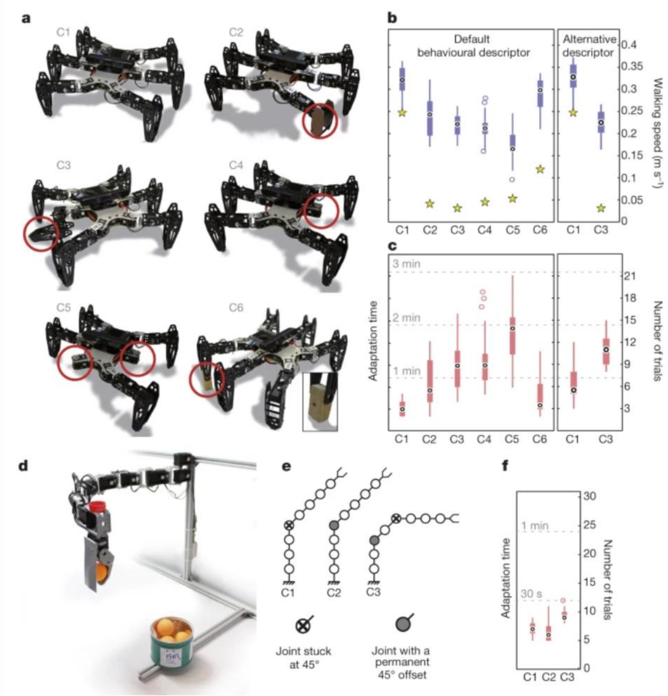

Multiobjective Optimization

A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation

Runzhe Yang
Department of Computer Science
Princeton University
runzhey@cs.princeton.edu

Xingyuan Sun
Department of Computer Science
Princeton University
xs5@cs.princeton.edu

Karthik Narasimhan
Department of Computer Science
Princeton University
karthikn@cs.princeton.edu

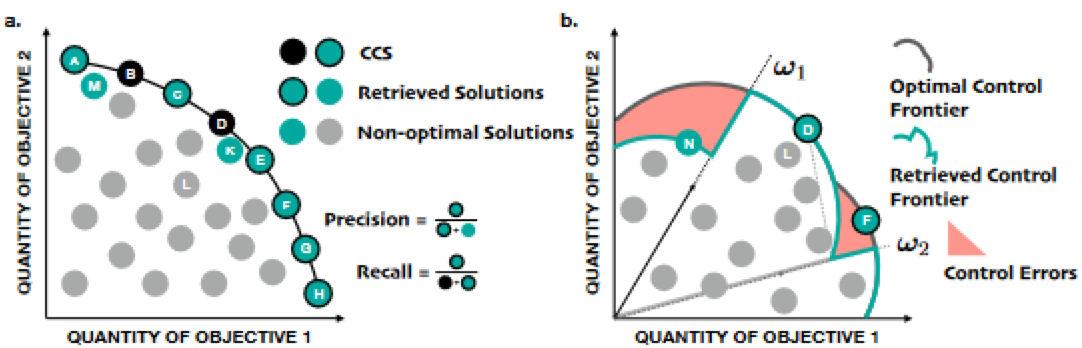

Review Article | Published: 07 January 2019

Designing neural networks through neuroevolution

Kenneth O. Stanley, Jeff Clune, Joel Lehman & Risto Miikkulainen

Nature Machine Intelligence **1**, 24–35 (2019) | Cite this article

Thank you
for your attention!