



EvoTorch

Next-generation Evolutionary Search, Learning & Planning

EvoTorch™ is an advanced optimization library created at NNAISENSE to solve industrial optimization use cases at scale by harnessing the power of evolutionary algorithms (EAs) and deep neural networks. It provides:

- Fast, high-quality implementations of various state-of-the-art distribution-based and population-based evolutionary algorithms.
- Built-in support for large-scale parallelization from single machines to clusters of CPUs and/or GPUs of any size.
- Support for modern machine learning logging tools.
- A modular design that readily allows client-specific custom solutions.
- Enterprise support from experts with decades of experience with EAs.

Why Evolutionary Algorithms?

Generality. EAs solve optimization problems by repeatedly recombining the best (most "fit") solutions in a population of candidates to generate increasingly fit solutions based loosely on the principles of natural selection. Because EAs rely solely on the cost function of the optimization problem (the fitness), they require no information about the particular structure of the problem. However, this black-box approach typically comes at a cost: to perform well, EAs require a large number of candidate solutions to be evaluated. EvoTorch solves this problem, and has already enabled NNAISENSE engineers to use unprecedented population sizes of hundreds of thousands of solutions, all evaluated efficiently in parallel.

Optional Customization. To further boost efficiency, EAs can be enhanced with problem-specific operators that leverage crucial information about the search space to generate new solutions (i.e. "crossover" and "mutation"). EvoTorch is designed to easily implement such customizations, and NNAISENSE works closely with client experts to incorporate this valuable domain knowledge into the algorithms.

Deployment

Deploying EvoTorch is simple since it's built on Python and PyTorch, two technologies designed with portability front and center. EvoTorch can be deployed both in the cloud to exploit large computational resources, or on-prem, without an internet connection. It can be used through any language that can interface either directly to Python or to its command-line interface.

How we use EvoTorch

NNAISENSE has been using EvoTorch internally for many years, applying it to industrial control, optimization and planning problems according to client requirements. Examples include:

- Planning for complex chemical processes with safety constraints while handling uncertainty.
- Robot control with transfer to a real robot with a significant sim-to-real gap.
- Autonomous multi-agent vehicle control, balancing multiple conflicting objectives and constraints.

What we offer

The team behind EvoTorch are experts in evolutionary computation with decades of experience and citations in this field. Over the last five years, NNAISENSE has devoted significant resources to devising the most efficient approach to implementing EAs that scale to real-world industrial problems. Our Research Scientists, Engineers and Software Infrastructure team can provide you with:

- Algorithms customized for the constraints and requirements specific to your problem. This includes a variety of features not available in the open-source EvoTorch release, such as domain-specific solution structures, constraint handling and the combination of EAs with other methods.
- Custom problem implementations that squeeze as much performance as possible out of target hardware platforms to solve problems faster.
- Integration with NNAISENSE's Deep Digital Twin technology that learns data-driven surrogate models of your industrial process, which can dramatically reduce iteration times for solution development.
- Deployment support for both cloud-based and on-site deployments.
- Integration support, helping your engineers interface EvoTorch into your existing software ecosystem, regardless of the implementation language.