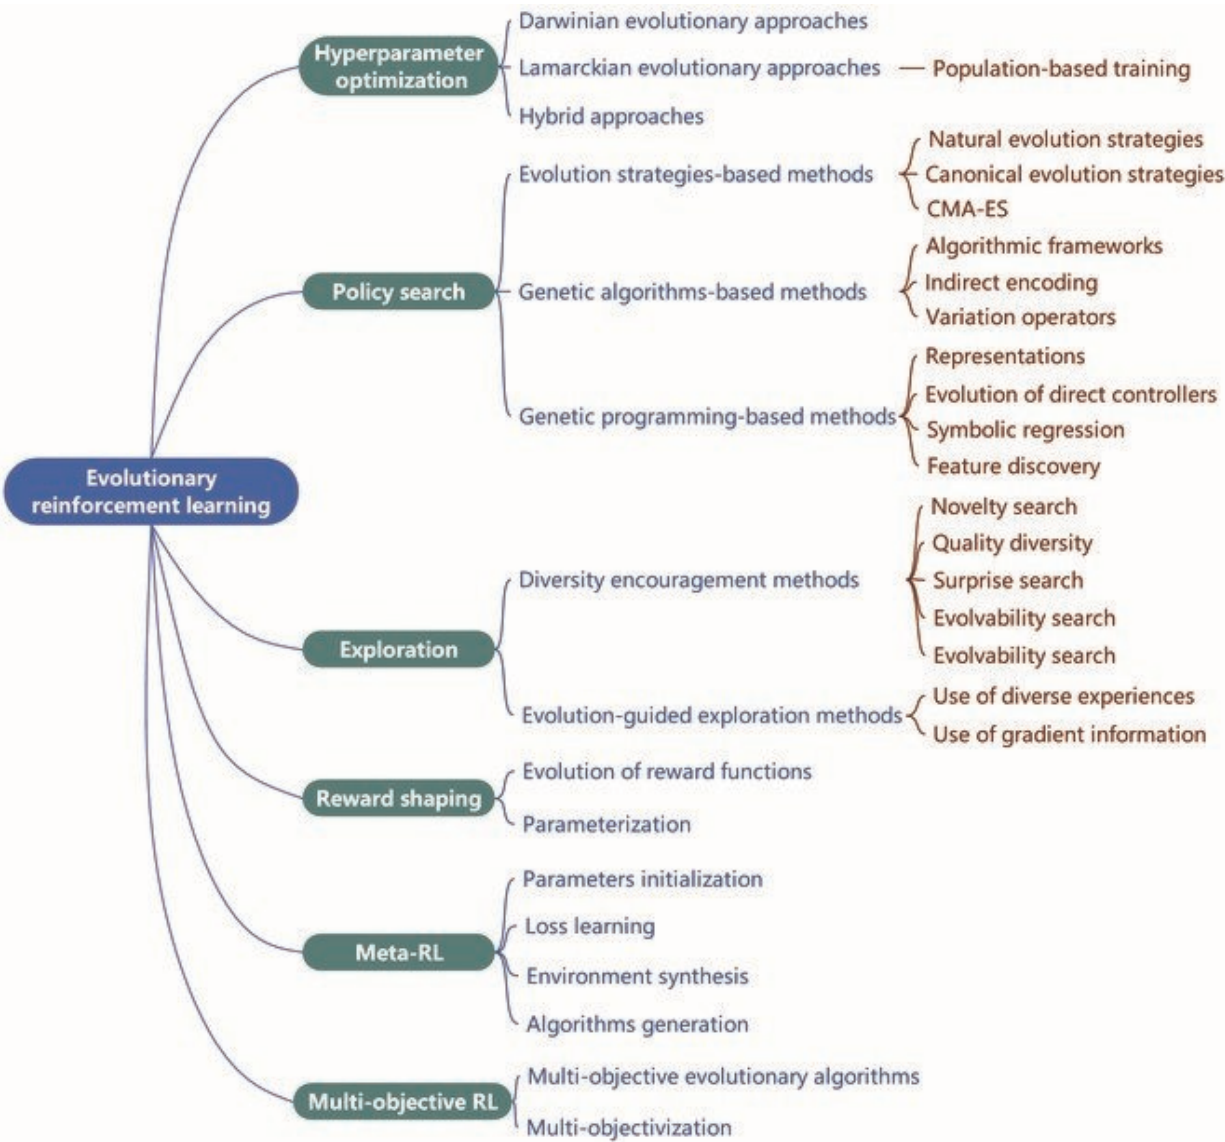


Evolutionary reinforcement learning promises further advances in machine learning

May 19 2023



Key research areas in evolutionary reinforcement learning. Credit: Hui Bai et al.

Evolutionary reinforcement learning is an exciting frontier in machine learning, combining the strengths of two distinct approaches: reinforcement learning and evolutionary computation. In evolutionary reinforcement learning, an intelligent agent learns optimal strategies by actively exploring different approaches and receiving rewards for successful performance.

This innovative paradigm combines [reinforcement](#) learning's trial-and-error learning with evolutionary algorithms' ability to mimic [natural selection](#), resulting in a powerful methodology for artificial intelligence development that promises breakthroughs in various domains.

A review article on evolutionary reinforcement learning was published in *Intelligent Computing*. It sheds light on the latest advancements in the integration of evolutionary computation with reinforcement learning and presents a comprehensive survey of state-of-the-art methods.

Reinforcement learning, a subfield of [machine learning](#), focuses on developing algorithms that learn to make decisions based on feedback from the environment. Remarkable examples of successful reinforcement learning include AlphaGo and, more recently, [Google DeepMind robots that play soccer](#).

However, reinforcement learning still faces several challenges, including the exploration and exploitation trade-off, reward design, generalization and credit assignment.

Evolutionary computation, which emulates the process of natural evolution to solve problems, offers a potential solution to the problems

of reinforcement learning. By combining these two approaches, researchers created the field of evolutionary reinforcement learning.

Evolutionary reinforcement learning encompasses six key research areas:

- **Hyperparameter optimization:** Evolutionary computing methods can be used for hyperparameter optimization. That is, they can automatically determine the best settings for reinforcement learning systems. Discovering the best settings manually can be challenging due to the multitude of factors involved, such as the learning speed of the algorithm and its inclination towards future rewards. Furthermore, the performance of reinforcement learning relies heavily on the architecture of the neural network employed, including factors like the number and size of its layers.
- **Policy search:** Policy search entails finding the best approach to a task by experimenting with different strategies, aided by [neural networks](#). These networks, akin to powerful calculators, approximate task execution and make use of advancements in deep learning. Since there are numerous task execution possibilities, the search process resembles navigating a vast maze. Stochastic gradient descent is a common method for training neural networks and navigating this maze. Evolutionary computing offers alternative "neuroevolution" methods based on evolution strategies, genetic algorithms and genetic programming. These methods can determine the best weights and other properties of neural networks for reinforcement learning.
- **Exploration:** Reinforcement learning agents improve by interacting with their environment. Too little exploration can lead to poor decisions, while too much exploration is costly. Thus there is a trade-off between an agent's exploration to discover good behaviors and an agent's exploitation of the discovered good behaviors. Agents explore by adding randomness to their

actions. Efficient exploration faces challenges: a large number of possible actions, rare and delayed rewards, unpredictable environments and complex multi-agent scenarios. Evolutionary computation methods address these challenges by promoting competition, cooperation and parallelization. They encourage exploration through diversity and guided evolution.

- **Reward shaping:** Rewards are important in reinforcement learning, but they are often rare and hard for agents to learn from. Reward shaping adds extra fine-grained rewards to help agents learn better. However, these rewards can alter agents' behavior in undesired ways, and figuring out exactly what these extra rewards should be, how to balance them and how to assign credit among multiple agents typically requires specific knowledge of the task at hand. To tackle the challenge of reward design, researchers have used evolutionary computation to adjust the extra rewards and their settings in both single-agent and multi-agent reinforcement learning.
- **Meta-reinforcement learning:** Meta-reinforcement learning aims to develop a general learning algorithm that adapts to different tasks using knowledge from previous ones. This approach addresses the issue of requiring a large number of samples to learn each task from scratch in traditional reinforcement learning. However, the number and complexity of tasks that can be solved using meta-reinforcement learning are still limited, and the computational cost associated with it is high. Therefore, exploiting the model-agnostic and highly parallel properties of evolutionary computation is a promising direction to unlock the full potential of meta-reinforcement learning, enabling it to learn, generalize and be more computationally efficient in real-world scenarios.
- **Multi-objective reinforcement learning:** In some real-world problems, there are multiple goals that conflict with each other. A multi-objective evolutionary algorithm can balance these goals

and suggest a compromise when no solution seems better than the others. Multi-objective reinforcement learning methods can be grouped into two types: those that combine multiple goals into one to find a single best solution and those that find a range of good solutions. Conversely, some single-goal problems can be usefully broken down into multiple goals to make problem-solving easier.

Evolutionary reinforcement learning can solve complex reinforcement learning tasks, even in scenarios with rare or misleading rewards. However, it requires significant computational resources, making it computationally expensive. There is a growing need for more efficient methods, including improvements in encoding, sampling, search operators, algorithmic frameworks and evaluation.

While evolutionary reinforcement learning has shown promising results in addressing challenging reinforcement learning problems, further advancements are still possible. By enhancing its computational efficiency and exploring new benchmarks, platforms and applications, researchers in the field of evolutionary reinforcement learning can make evolutionary methods even more effective and useful for solving complex [reinforcement learning](#) tasks.

More information: Hui Bai et al, Evolutionary Reinforcement Learning: A Survey, *Intelligent Computing* (2023). [DOI: 10.34133/icomputing.0025](#)

Provided by Intelligent Computing

Citation: Evolutionary reinforcement learning promises further advances in machine learning (2023, May 19) retrieved 28 April 2024 from <https://techxplore.com/news/2023-05-evolutionary->

[advances-machine.html](#)

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.