

AGENTFORGE: A Flexible Low-Code Platform for Reinforcement Learning Agent Design

Francisco E. Fernandes Jr. ^a and Antti Oulasvirta ^b

Department of Information and Communications Engineering, Aalto University, Espoo, Finland

Keywords: Reinforcement Learning, Agents, Bayesian Optimization, Particle Swarm Optimization.

Abstract: Developing a reinforcement learning (RL) agent often involves identifying values for numerous parameters, covering the policy, reward function, environment, and agent-internal architecture. Since these parameters are interrelated in complex ways, optimizing them is a black-box problem that proves especially challenging for nonexperts. Although existing optimization-as-a-service platforms (e.g., Vizier and Optuna) can handle such problems, they are impractical for RL systems, since the need for manual user mapping of each parameter to distinct components makes the effort cumbersome. It also requires understanding of the optimization process, limiting the systems' application beyond the machine learning field and restricting access in areas such as cognitive science, which models human decision-making. To tackle these challenges, the paper presents AGENTFORGE, a flexible low-code platform to optimize any parameter set across an RL system. Available at <https://github.com/feferna/AgentForge>, it allows an optimization problem to be defined in a few lines of code and handed to any of the interfaced optimizers. With AGENTFORGE, the user can optimize the parameters either *individually* or *jointly*. The paper presents an evaluation of its performance for a challenging vision-based RL problem.


1 INTRODUCTION


Developing reinforcement learning (RL) agents is important not only for advancements in machine learning (ML) but also for fields such as the cognitive sciences, where RL increasingly serves as a computational framework to model human decision-making mechanisms (Eppe et al., 2022). For example, RL can help researchers understand how humans behave under cognitive constraints (Chandramouli et al., 2024). One of the challenging aspects of this development is optimizing a broad array of parameters that influence agent behavior and performance: this process represents a black-box optimization problem. Embodied RL agents, modeled as partially observable Markov decision processes (POMDPs) (Eppe et al., 2022), highlight this complexity in tasks such as robotic navigation (Shahria et al., 2022) and human–computer interaction (Jiang et al., 2024). In contrast against ML domains wherein the tuning focuses on a smaller set of training-related parameters, RL requires optimizing explainable ones too (e.g., reward weights), to as-

sure of trustworthiness in artificial intelligence (AI) systems.

We set out to give domain experts in cognitive sciences, alongside other fields, the ability to optimize all RL system parameters, either *jointly* or *individually*, as conditions dictate. This is essential, since even small changes in an RL algorithm's implementation, such as its reward clipping, can significantly affect performance (Engstrom et al., 2020). In some cases, careful parameter selection can improve performance more than the choice of RL algorithm itself (Andrychowicz et al., 2020). Embodied RL agents in particular employ deep neural networks with numerous parameters highly sensitive to optimization (Fetterman et al., 2023). These parameters can be roughly categorized into three classes:

1. **Agent Design:** controls internal functions such as the discount factor, entropy coefficient, and observation window size.
2. **Environment Settings:** define the task, including the field of view and the reward structure.
3. **Policy Parameters:** encompass neural-network architecture, learning rate, and activation functions.

^a  <https://orcid.org/0000-0003-2301-8820>

^b  <https://orcid.org/0000-0002-2498-7837>

Against this backdrop, we present the ongoing work on AGENTFORGE, a flexible low-code platform designed for experts across disciplines to develop RL systems without having to possess expertise in optimization or ML. This domain-agnostic ability aligns with the needs of cognitive scientists, who often require tailored experiment setups for testing hypotheses about human cognition (Eppe et al., 2022; Nobandegani et al., 2022). The development of low-code platforms is relevant to many such fields. Also, AGENTFORGE provides for rapid iteration by letting users prototype RL designs, optimize them, and analyze the results. It supports a wide range of RL systems, from simple models to complex embodied agents. Users need only define a custom environment, an objective function, and parameters in two files, where the objective function, which guides the optimization process, can employ any criterion (e.g., average reward per episode). The platform then converts these inputs into an optimization problem, automating training and evaluation. AGENTFORGE’s current version includes random search, Bayesian optimization (BayesOpt), and particle swarm optimization (PSO), with flexibility for adding more techniques.

Obviously, parameter optimization is a perennial challenge in ML. Frameworks such as BoTorch (Balandat et al., 2020), Optuna (Akiba et al., 2019), and Google Vizier (Song et al., 2023) represent attempts to address it. However, these tools are targeted at users with advanced ML knowledge and are tricky to apply to RL agents on account of the complexity of defining objective functions for parameters across system components. This shortcoming highlights the need for *low-code solutions*, such as ours, that allow users to design effective agents without having mastered the optimization techniques crucial for handling high-dimensional inputs and dynamic RL environments.

We tested AGENTFORGE by jointly optimizing the parameters of a pixel-based Lunar Lander agent from the Gymnasium library (Towers et al., 2024), a POMDP in which the agent infers the state from raw pixel values. Since this agent includes parameters from all three categories – agent, environment, and policy – it permitted testing AGENTFORGE’s ability to handle complex parameter sets. Also, the paper demonstrates how easily an optimization problem can be defined, thus highlighting our platform’s simplicity and flexibility.

2 RELATED WORKS

Among the many fields showing extensive use of RL are robotics, cognitive science, and human–computer interaction in general. Robotics focuses on optimizing real-world sensorimotor interactions (Shang and Ryoo, 2023), while cognitive scientists model human decision-making (Chandramouli et al., 2024). Table 1 highlights the breadth of tasks and needs in RL development. Across the board, these share the challenge of optimizing large parameter spaces and adapting to dynamic environments, a goal that drives demand for flexible, low-code platforms that simplify RL development across disciplines.

The work in response forms part of the field known as AutoRL, aiming to automate RL algorithm and parameter optimization. While already showing clear progress, AutoRL is in its infancy (Parker-Holder et al., 2022; Afshar et al., 2022). Platforms such as AutoRL-Sim (Souza and Ottoni, 2024) and methods that use genetic algorithms (Cardenoso Fernandez and Caarls, 2018) support simple RL tasks but lack support for more complex real-world problems with diverse parameters across agents, environments, and policies. Among the optimization techniques commonly employed in AutoRL are random search, BayesOpt, and PSO, with the choice of technique depending on task-specific requirements. Furthermore, while such state-of-the-art RL solvers as Stable-Baselines3 (Raffin et al., 2021) and CleanRL (Huang et al., 2022) offer powerful tools for agent training, they demand deep knowledge of RL’s internals, plus extensive coding, and are not suitable for parameter optimization.

Platforms of AGENTFORGE’s ilk can address these challenges by simplifying RL agent design and providing automated optimization for diverse needs. Its provision for random search, BayesOpt, and PSO furnishes a low-code interface for joint parameter optimization, whereby users become able to compare and apply strategies for various RL agents with ease.

3 THE PROPOSED FRAMEWORK

3.1 General Characterization

For enabling flexible low-code specification of an optimization problem in RL, the following engineering objective guided the design of AGENTFORGE:

- **A Low-Code Setting.** Users provide simple config files that state optimization parameters and

Table 1: Example RL agents and corresponding parameterized modules – highlighting the diversity of parameter types, this non-exhaustive list emphasizes the complexity of parameter optimization in RL systems.

Agent/env.	Modules	Use case	Typical user
Gymnasium toy environments (Towers et al., 2024)	- Environment dynamics - Reward structure - Policy network - Value network	Achieving the highest reward in a well-defined environment	- Machine learning researcher - Robotics researcher
ActiveVision-RL (Shang and Ryoo, 2023)	- Observation window size - Persistence of vision - Sensorimotor reward - Q-Network	Modeling visual attention in complex environments	- Machine learning researcher - Robotics researcher
CRTypist (Shi et al., 2024)	- Foveal vision - Peripheral vision - Finger movement - Working memory - Reward structure - Neural networks	Simulating touchscreen typing behavior	- Designer - Accessibility researcher - Cognitive scientist - Behavioral scientist
Autonomous aircraft landing (Ladosz et al., 2024)	- Descent velocity - Frame stack - Image resolution - Q-Network - Reward structure	Automating aircraft landings with vision-based reinforcement learning	- Machine learning researcher - Robotics researcher
EyeFormer (Jiang et al., 2024)	- Inhibition of return (IoR) - Vision encoder - Fixation decoder - Optimization constraints - Neural networks - Reward structure	Modeling visual attention and eye-movement behavior for UI optimization	- Designer - Accessibility researcher - Cognitive scientist - Behavioral scientist

performance-evaluation methods. This approach lets users exploit the platform without deep technical expertise.

- **Integration.** Support for custom RL environments compatible with the Gymnasium library enables use across diverse application scenarios.
- **Flexibility.** Thanks to its multiple optimization methods, users can easily swap and set up methods through config files without changing the core code. This simplifies adaptation to various problem requirements.
- **Parallelization.** Allowing concurrent evaluation of multiple solutions reduces optimization time. By using Optuna’s parallelization and our custom methods, the platform can explore many configurations simultaneously, which is especially beneficial for large-scale RL problems.

The overall optimization process comprises two loops, illustrated in Figure 1. The outer loop optimizes parameters via the chosen method, while the inner one trains the RL agent through environment interactions and performance evaluations. With the next few sections, we give more details on using the framework and on its optimization algorithms.

3.2 Specifying the Optimization Task

To use AGENTFORGE effectively, the user must furnish three key components: configuration settings, a customized environment, and an evaluation method.

Users start by providing a YAML configuration file to define the optimization setup. This file specifies the parameters to be optimized, each categorized as agent, environment, or policy, along with the type for each (integer or float) and value ranges. It also includes settings for agent training hyperparameters, such as learning rates and discount factors, and the choice of optimization algorithm: in the current implementation random search, BayesOpt, or PSO. Additionally, AGENTFORGE supports optimizing neural network architectures for policy and value networks, thereby letting users set the range for the number of layers and of neurons per layer. Thus it enables a thorough search for optimal structures.

Next, via a Python script, the user implements a custom RL environment compatible with the Stable-Baselines3 library. It must adhere to the standard interface, defining action and observation spaces, and include methods for resetting and stepping through the environment. For this, the user writes a function that can create a training environment for the agent. While AGENTFORGE currently supports only the popular Python library Stable-Baselines3, future versions may extend compatibility to other libraries.

Finally, users must implement an evaluation method, again in a Python script, to assess the trained agent’s performance. The function should calculate and report metrics for factors such as reward, success rate, etc., to make sure that AGENTFORGE can collect and use these values during optimization. Optionally, they may also provide a function that records a video of the agent’s performance over a specified number of

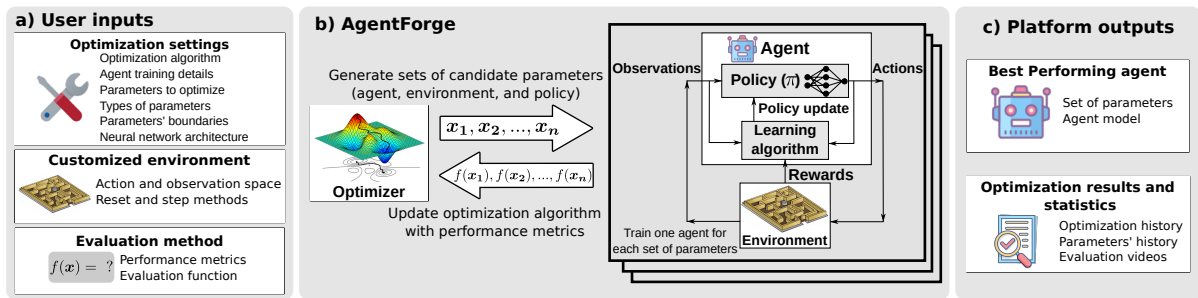


Figure 1: The AGENTFORGE platform in overview. a) The inputs consist of three elements: optimization settings (including search parameters), a customized environment, and an evaluation method. b) AGENTFORGE executes joint optimization of RL parameters through a two-loop process. c) Then, AGENTFORGE provides the user with the best-performing agent along with the optimization results and performance statistics.

episodes in the environment.

These three components are summarized in Figure 1 and in Listing 1.

3.3 Optimization Algorithms

AGENTFORGE supports several popular optimization algorithms. While random search, BayesOpt, and PSO form the set of default options, the framework is flexible; it can be extended to include other methods, for various RL scenarios.

- **Random search:** is a simple baseline method that randomly samples the parameter space. Useful for initial comparisons in testing of more sophisticated algorithms, it guarantees rich exploration

```

1 ...
2 # Training parameters
3 gae_lambda:
4     type: default
5     searchable: true
6     integer: false
7     user_preference: 0.9
8     start: 0.9
9     stop: 0.95
10 ...
11 def user_train_environment(seed, n_train_envs,
12                             **env_kwargs):
13     """
14     Set up RL environment for training.
15     Args:
16     - seed (int): Random seed.
17     - n_train_envs (int): Number of training envs.
18     - **env_kwargs: Additional args for env
19       creation.
20     Returns:
21     - train_env (VecEnv): Configured training
22       environment.
23     """
24 def user_evaluate_policy(seed, model, **env_kwargs):
25     """
26     Evaluate the policy with a separate evaluation
27     environment.
28     Args:
29     - model (BaseAlgorithm): Model to eval.
30     Returns:
31     - mean_objective_value (float): Mean objective
32       value obtained during evaluation.
33     - extra_info (dict): Additional metrics defined
34       by the user.
35     """

```

Listing 1: Sample code snippet illustrating parameter definition, creation of the environment, and evaluation setup.

but is less efficient at finding optimal regions.

- **BayesOpt:** is a structured approach using probabilistic models (e.g., Gaussian processes) to focus on promising areas of the parameter space. Our platform uses a tree-structured Parzen estimator (TPE), well suited to high-dimensional spaces (Bergstra et al., 2011). It efficiently balances exploration and exploitation for fast convergence.
- **PSO:** is a population-based approach suitable for continuous optimization. Its stochastic nature aids in escaping local optima. We use inertia (w), cognitive (c_1), and social (c_2) weights of 0.9694, 0.099381, and 0.099381, respectively, for better exploration of high-dimensional spaces, as recommended by (Oldewage et al., 2020).

In AGENTFORGE, Random search and BayesOpt are implemented with Optuna, while PSO is custom-designed. The mechanism of letting users select and configure optimization strategies through a user-defined configuration file (see the previous subsection) renders the framework adaptable to various applications.

3.4 Parallel Evaluation of Trials

Taking advantage of parallel evaluation improves the efficiency and speeds up optimization. By running each trial as an independent Python process, we avoid issues with Python’s Global Interpreter Lock. For all optimization methods, parameter sets are distributed across nodes for evaluation, with multiple processes accessing a centralized database to get a new set for each trial. Parallelization is crucial since each trial involves training agents for millions of timesteps, which often takes hours. By utilizing parallel computing, the platform significantly reduces the time required for a full optimization loop, thus facilitating extensive RL experimentation.

Table 2: General optimization configuration, applied across all experiments.

Training setup for each trial	
Number of training timesteps	5,000,000
Evaluation frequency	Every 10 batches
Maximum steps per episode	1024
Vectorized environments for training	10
Training algorithm	PPO
Policy architecture	CNN
Random search	
Number of trials	400
Pruning of trials	After 50 trials
Pruning strategy	MedianPruner
Wait before trial is pruned	At least 48 evaluations
BayesOpt	
Number of trials	400
Number of initial random trials	50
Number of improvement candidates expected	80
Multivariate TPE	Enabled
Pruning of trials	After 50 trials
Pruning strategy	MedianPruner
Wait before trial is pruned	At least 48 evaluations
PSO	
Number of generations	20
Population size	20
Inertia weight (w)	0.9694
Cognitive component weight (c_1)	0.099381
Social component weight (c_2)	0.099381

4 EVALUATION

We evaluated AGENTFORGE on the basis of *agent performance*, judged in terms of the maximum mean reward achieved by a pixel-based POMDP agent.

4.1 A Pixel-Based POMDP Agent

Our evaluation of AGENTFORGE used a pixel-based version of the Lunar Lander environment from the Gymnasium library (Towers et al., 2024), a POMDP wherein the agent relies on raw-pixel observations instead of state-based inputs. The agent processes four stacked frames as visual input and uses continuous actions to control the main engine and lateral boosters. Reward structure remains unchanged from the original environment’s, with the task considered solved when the agent reaches a reward of at least 200 in a single episode.

This experiment demonstrated AGENTFORGE’s flexibility in optimizing parameters across all three categories (agent, environment, and policy), including neural network architectures. The breadth of optimization attests to AGENTFORGE’s capacity to handle complex RL configurations effectively. Below, we provide details of its parameter optimization.

4.2 Optimization Configuration

The optimization configuration for all experiments, covering the parameters for individual training trials and the random search, BayesOpt, and PSO methods, is detailed in Table 2. Each trial involved training for five million timesteps, with performance evaluations after every 10 batches. This setup used 10 parallel, vectorized environments with the proximal policy optimization (PPO) algorithm and a convolutional neural network (CNN) feature extractor. Agent performance was measured as the mean reward over 300 independent episodes. The table includes all the optimization settings: the number of trials, pruning strategies, and algorithm-specific parameters. All configurations can be easily customized through the user-defined configuration file introduced in Section 3.

4.3 The Parameters Optimized

We optimized the parameters listed in Table 3. The environment parameter chosen for optimization here was the size of the field of view (in pixels). This parameter determines how much the original environment window is downsampled into a square window suitable for the feature extractor. To define its dimensions, only one value is needed. Further parameters optimized were the discount factor (γ), generalized advantage estimation (λ), learning rate, number of epochs, entropy coefficient, and clipping range. These are critical for training because they influence reward discounting, learning stability, and the efficiency of policy updates. Also tuned were neural network architecture parameters, such as the activation function, number of layers, and number of neurons per layer for the policy and value networks both. The activation function was defined as a float in the range [0.0, 1.0], with Tanh used for values below 0.5 and ReLU for those of 0.5 or above.

4.4 Joint vs. Individual-Parameter Optimization

The BayesOpt results presented below demonstrate the flexibility of AGENTFORGE well. We report these for three parameter-optimization configurations:

- **Joint Optimization with NAS:** Simultaneous optimization of all parameters listed in Table 3, addressing the policy and value networks’ architecture
- **Optimization Without NAS:** Optimizing all parameters except for the network architecture (activation function, number of layers, and neuron

Table 3: The parameters optimized and the best values found for the pixel-based Lunar Lander agent – the wide range of parameters being jointly optimized highlights the flexibility of AGENTFORGE.

Parameter	Type	Category	Range	Random search	BayesOpt	PSO
Field of view's size (pixels)	Integer	Environment	[40, 92]	71	92	92
Discount factor (γ)	Float	Agent	[0.4, 0.8]	0.7934	0.7984	0.8
Generalized advantage estimation (λ)	Float	Agent	[0.9, 0.95]	0.9433	0.9299	0.95
Learning rate	Float	Agent	$[3.5 \cdot 10^{-4}, 3.5 \cdot 10^{-3}]$	0.0006	0.0034	0.0035
Number of epochs	Integer	Agent	[3, 10]	6	3	5
Entropy coefficient	Float	Policy	[0.01, 0.1]	0.0284	0.0279	0.1
Clipping range	Float	Policy	[0.01, 0.3]	0.1407	0.0178	0.3
Activation function	Float	Policy	[0.0, 1.0]	0.5478	0.9259	1.0
No. of layers in policy	Integer	Policy	[1, 4]	1	3	4
No. of neurons per layer in policy	Integer	Policy	[64, 128]	83	89	127
No. of layers in value network	Integer	Policy	[2, 4]	4	2	4
No. of neurons per layer in value	Integer	Policy	[64, 128]	100	89	128
Mean reward				84.26	172.43	119.19

count), which remain at their stable-baselines3 defaults

- **NAS only:** Optimization restricted to the network architecture while all other parameters are fixed at their Stable-Baselines3 defaults

5 RESULTS

5.1 Results from Joint Optimization

The upper pane of Figure 2 shows the evolution of the mean reward in joint optimization over these episodes for each trial under random search, BayesOpt, and PSO. One can see that BayesOpt achieved the highest cumulative rewards while PSO showed the fastest convergence. Random search lagged behind with a mean reward of 84.26, serving as a baseline for comparison. Notably, PSO reached approximately 69% of BayesOpt's mean reward, while random search represented about 49% of BayesOpt's level. These results spotlight the effectiveness of BayesOpt but also demonstrate AGENTFORGE's flexibility in integrating multiple optimization algorithms, which supports users' selection of the best approach for their specific RL agents.

BayesOpt consistently outperformed random search and PSO in terms of mean cumulative reward, achieving a mean reward of 172.43. Though PSO reached a lower final reward (mean reward: 119.19), it demonstrated the fastest convergence. Random search lagged behind with a mean reward of 84.26, serving as a baseline for comparison. Notably, PSO reached approximately 69% of BayesOpt's mean reward, while random search represented about 49% of BayesOpt's level. These results spotlight the effectiveness of BayesOpt but also demonstrate AGENTFORGE's flexibility in integrating multiple optimization algorithms, which supports users' selection of the best approach for their specific RL agents.




Figure 2: Mean reward per episode for the pixel-based Lunar Lander agent under different optimization strategies. *Top:* Joint optimization of parameters and NAS using random search, BayesOpt, and PSO. *Bottom:* A plot of joint vs. individual-parameter optimization using BayesOpt, with and without NAS, and NAS-only optimization.

5.2 Results of the Joint vs. Parameter-Wise Optimization

Figure 2's lower pane presents the results of *joint* versus *individual-parameter* optimization. Joint optimization with and without NAS achieved mean episode rewards of 172.43 and 174.19, respec-

tively, while the NAS-alone condition yielded a mean episode reward of -32.40 . These findings suggest that network architecture is not the most critical parameter set to optimize for our agent. However, users short on time can obtain satisfactory results by opting for complete joint optimization.

5.3 Platform Overhead

AGENTFORGE introduces negligible overhead relative to the time required for evaluating RL agents. A single candidate’s evaluation takes roughly three hours on a consumer-grade Nvidia RTX 4090 GPU, making any additional delays from the optimization process insignificant. However, parallel evaluation is vital; running 400 trials serially would take about 50 days. By evaluating eight trials in parallel, we can complete the same 400 trials in just seven days.

6 DISCUSSION

Our results affirm AGENTFORGE’s effectiveness as a flexible, low-code platform for RL agent design and optimization. It addresses the challenge of optimizing interrelated parameters in large numbers across the agent, environment, and policy classes by treating this as a black-box optimization problem. We found evidence that parameter optimization on its own can improve the RL agent’s performance, without reliance on changes to the training algorithms. All three optimization methods currently available with AGENTFORGE enhanced performance in our test case, with BayesOpt arriving at the highest episode mean reward observed. However, the platform’s key strength lies in its support for joint and individual-level optimization of RL agents’ parameters and neural network architectures. This crystallizes the platform’s flexibility for strongly supporting the study of critical interplay among various aspects of RL agent design.

In principle, AGENTFORGE’s integration of various optimization methods allows users to compare approaches and can streamline the prototyping and fine-tuning of RL agents. By configuring optimization settings through a simple YAML file and a Python script, they can quickly customize their agents without extensive coding. Thus helping pare back the typical trial-and-error process in RL experimentation, this frees more of practitioners’ time for refining their problem models. While our results are encouraging, further testing, with user studies and different RL systems, will be necessary to validate AGENTFORGE’s broader impact. Still, it already offers a flexible solution to effectively tackle the unique challenges of RL

systems, overcoming several issues that practitioners might face when using other platforms.

7 CONCLUSIONS AND FUTURE WORK

With this short paper, we provided a glimpse at AGENTFORGE, which simplifies joint optimization of RL parameters, accelerating agent development – especially for users with little optimization experience – through flexible optimization and easier design. Evaluation with a pixel-based POMDP agent settings proved illuminating. The setup was simple from a low-code perspective yet proved able to support finding a suitable balance (e.g., BayesOpt delivered the best optimization results, while PSO showed faster convergence but requires further tuning). The platform also gives users control over their experiment setups: it demonstrated good flexibility for both joint and individual parameters’ optimization.

This version does display several limitations, which will be addressed in a follow-up paper. For instance, we are developing a graphical interface to enable a comprehensive user study evaluating the ease of use of our solution in comparison to preexisting ones such as AutoRL-Sim. Also, the current interface is designed to assist cognitive and behavioral scientists in using RL to model human decision-making processes more effectively. We plan to test use cases additional to these, probe more agent designs, and fully address the challenge of managing RL training uncertainty during optimization.

ACKNOWLEDGEMENTS

This work was supported by the ERC AdG project Artificial User (101141916) and the Research Council of Finland (under the flagship program of the Finnish Center for Artificial Intelligence, FCAI). The calculations were performed via computer resources provided by the Aalto University School of Science project Science-IT. The authors also acknowledge Finland’s CSC – IT Center for Science Ltd. for providing generous computational resources.

REFERENCES

- Afshar, R. R., Zhang, Y., Vanschoren, J., and Kaymak, U. (2022). Automated Reinforcement Learning: An Overview. arXiv:2201.05000.

- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). Optuna: A Next-Generation Hyperparameter Optimization Framework. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Andrychowicz, M., Raichuk, A., Stańczyk, P., Orsini, M., Girgin, S., Marinier, R., Hussenot, L., Geist, M., Pietquin, O., Michalski, M., Gelly, S., and Bachem, O. (2020). What Matters in On-Policy Reinforcement Learning? A Large-Scale Empirical Study. arXiv:2006.05990.
- Balandat, M., Karrer, B., Jiang, D. R., Daulton, S., Letham, B., Wilson, A. G., and Bakshy, E. (2020). BoTorch: A Framework for Efficient Monte-Carlo Bayesian Optimization. arXiv:1910.06403.
- Bergstra, J., Bardenet, R., Bengio, Y., and Kégl, B. (2011). Algorithms for Hyper-parameter Optimization. In *Proceedings of the 24th International Conference on Neural Information Processing Systems, NIPS '11*, pages 2546–2554, Red Hook, NY, USA. Curran Associates.
- Cardenoso Fernandez, F. and Caarls, W. (2018). Parameters [sic] Tuning and Optimization for Reinforcement Learning Algorithms Using Evolutionary Computing. In *2018 International Conference on Information Systems and Computer Science (INCISCOS)*, pages 301–305, Quito, Ecuador. IEEE.
- Chandramouli, S., Shi, D., Putkonen, A., De Peuter, S., Zhang, S., Jokinen, J., Howes, A., and Oulasvirta, A. (2024). A Workflow for Building Computationally Rational Models of Human Behavior. *Computational Brain & Behavior*, 7:399–419.
- Engstrom, L., Ilyas, A., Santurkar, S., Tsipras, D., Janoos, F., Rudolph, L., and Madry, A. (2020). Implementation Matters in Deep RL: A Case Study on PPO and TRPO. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Eppe, M., Gumbsch, C., Kerzel, M., Nguyen, P. D. H., Butz, M. V., and Wermter, S. (2022). Intelligent Problem-Solving As Integrated Hierarchical Reinforcement Learning. *Nature Machine Intelligence*, 4(1):11–20.
- Fetterman, A. J., Kitanidis, E., Albrecht, J., Polizzi, Z., Fogelman, B., Knutins, M., Wróblewski, B., Simon, J. B., and Qiu, K. (2023). Tune As You Scale: Hyperparameter Optimization For Compute Efficient Training. arXiv:2306.08055.
- Huang, S., Dossa, R. F. J., Ye, C., Braga, J., Chakraborty, D., Mehta, K., and Araújo, J. G. M. (2022). CleanRL: High-Quality Single-File Implementations of Deep Reinforcement Learning Algorithms. *The Journal of Machine Learning Research*, 23(1):12585–12602.
- Jiang, Y., Guo, Z., Tavakoli, H. R., Leiva, L. A., and Oulasvirta, A. (2024). EyeFormer: Predicting Personalized Scanpaths with Transformer-Guided Reinforcement Learning. arXiv:2404.10163.
- Ladosz, P., Mammadov, M., Shin, H., Shin, W., and Oh, H. (2024). Autonomous Landing on a Moving Platform Using Vision-Based Deep Reinforcement Learning. *IEEE Robotics and Automation Letters*, 9(5):4575–4582.
- Nobandegani, A. S., Shultz, T. R., and Rish, I. (2022). Cognitive Models As Simulators: The Case of Moral Decision-Making. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.
- Oldewage, E. T., Engelbrecht, A. P., and Cleghorn, C. W. (2020). Movement Patterns of a Particle Swarm in High Dimensional Spaces. *Information Sciences*, 512:1043–1062.
- Parker-Holder, J., Rajan, R., Song, X., Biedenkapp, A., Miao, Y., Eimer, T., Zhang, B., Nguyen, V., Calandra, R., Faust, A., Hutter, F., and Lindauer, M. (2022). Automated Reinforcement Learning (AutoRL): A Survey and Open Problems. *Journal of Artificial Intelligence Research*, 74:517–568.
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., and Dormann, N. (2021). Stable-Baselines3: Reliable Reinforcement Learning Implementations. *The Journal of Machine Learning Research*, 22(1):12348–12355.
- Shahria, M. T., Sunny, M. S. H., Zarif, M. I. I., Ghommam, J., Ahamed, S. I., and Rahman, M. H. (2022). A Comprehensive Review of Vision-Based Robotic Applications: Current State, Components, Approaches, Barriers, and Potential Solutions. *Robotics*, 11(6).
- Shang, J. and Ryoo, M. S. (2023). Active Vision Reinforcement Learning under Limited Visual Observability. In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S., editors, *Advances in Neural Information Processing Systems*, volume 36, pages 10316–10338. Curran Associates.
- Shi, D., Zhu, Y., Jokinen, J. P., Acharya, A., Putkonen, A., Zhai, S., and Oulasvirta, A. (2024). CRTypist: Simulating Touchscreen Typing Behavior via Computational Rationality. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, CHI '24*. ACM.
- Song, X., Perel, S., Lee, C., Kochanski, G., and Golovin, D. (2023). Open Source Vizier: Distributed Infrastructure and API for Reliable and Flexible Blackbox Optimization. arXiv:2207.13676.
- Souza, G. K. B. and Ottoni, A. L. C. (2024). AutoRL-Sim: Automated Reinforcement Learning Simulator for Combinatorial Optimization Problems. *Modelling*, 5(3):1056–1083.
- Towers, M., Kwiatkowski, A., Terry, J., Balis, J. U., De Cola, G., Deleu, T., Goulão, M., Kallinteris, A., Krimmel, M., KG, A., Perez-Vicente, R., Pierré, A., Schulhoff, S., Tai, J. J., Tan, H., and Younis, O. G. (2024). Gymnasium: A Standard Interface for Reinforcement Learning Environments. arXiv:2407.17032.