

Hyperparameter Optimization for Multi-Objective Reinforcement Learning

Florian Felten^{a,*}, Daniel Gareev^b, El-Ghazali Talbi^{a,c} and Grégoire Danoy^{a,b}

^aSnT, University of Luxembourg

^bFSTM/DCS, University of Luxembourg

^cCNRS/CRIStAL, University of Lille

ORCID ID: Florian Felten <https://orcid.org/0000-0002-2874-3645>,

El-Ghazali Talbi <https://orcid.org/0000-0003-4549-1010>, Grégoire Danoy <https://orcid.org/0000-0001-9419-4210>

Abstract. Reinforcement learning (RL) has emerged as a powerful approach for tackling complex problems. The recent introduction of multi-objective reinforcement learning (MORL) has further expanded the scope of RL by enabling agents to make trade-offs among multiple objectives. This advancement not only has broadened the range of problems that can be tackled but also created numerous opportunities for exploration and advancement. Yet, the effectiveness of RL agents heavily relies on appropriately setting their hyperparameters. In practice, this task often proves to be challenging, leading to unsuccessful deployments of these techniques in various instances. Hence, prior research has explored hyperparameter optimization in RL to address this concern.

This paper presents an initial investigation into the challenge of hyperparameter optimization specifically for MORL. We formalize the problem, highlight its distinctive challenges, and propose a systematic methodology to address it. The proposed methodology is applied to a well-known environment using a state-of-the-art MORL algorithm, and preliminary results are reported. Our findings indicate that the proposed methodology can effectively provide hyperparameter configurations that significantly enhance the performance of MORL agents. Furthermore, this study identifies various future research opportunities to further advance the field of hyperparameter optimization for MORL.

1 Introduction

RL has become a popular approach to solving challenging problems in recent years. This is due in part to its success in a range of domains, such as playing Atari games [19] and Go [24]. Within RL, the domain of multi-objective RL (MORL) has recently emerged, where agents are tasked with making compromises between multiple rewards.

The field of MORL has seen the development of several algorithms and important theoretical findings. These can be found in recent surveys such as [23, 12]. Additionally, software tools have been released to facilitate reproducibility, adoption, and further advancement of MORL methods [3, 11].

Even though open-source implementations and tools are readily accessible, promoting the use of these techniques even among non domain experts, it is widely recognized that the effectiveness of RL (and MORL) agents is highly dependent on the values of their

hyperparameters [13, 4, 22]. This implies that in a new environment, an existing algorithm may exhibit poor performance when run with its default hyperparameters. Furthermore, baseline comparisons frequently involve contrasting new algorithms with state-of-the-art techniques. In this regard, [22] pinpoint the problem of unfairness when untuned baseline algorithms are pitted against finely tuned new algorithms.

While manually tuning these hyperparameters may seem appealing, the large number of parameters and the time required to obtain training results often make manual optimization impractical [21]. Consequently, having an automated approach to tuning hyperparameters for MORL, given a predefined level of effort, not only facilitates fairer comparisons within the field but also streamlines the implementation of these techniques for real-world applications. Yet, to the best of our knowledge, although recent endeavors have focused on automating such decisions in the context of RL [21, 9], no similar effort has been applied to MORL. Moreover, existing solutions for single-objective RL do not trivially apply to MORL.

The objective of this work is to address this gap by formally recognizing the challenge of hyperparameter optimization for MORL and introducing a systematic procedure to undertake such an endeavor. Section 2 introduces the necessary background knowledge and notation related to MORL and hyperparameter optimization (HPO). Section 3 formally defines the problem of hyperparameter optimization for multi-objective RL and proposes a methodology to solve it. Section 4 showcases preliminary results of the proposed methodology applied to a state-of-the-art method. Lastly, Section 5 provides a conclusion to our research and discusses potential leads for future work.

2 Background

This section introduces notations and background information for multi-objective reinforcement learning and hyperparameter optimization.

2.1 Multi-objective reinforcement learning

MORL utilizes the multi-objective Markov decision process (MOMDP) framework to capture an agent’s interaction with its environment [23, 12]. This framework differs from Markov decision processes (MDPs) [25] in that it expands the reward function to generate

* Corresponding Author. Email: florian.felten@uni.lu.

vectors instead of scalar values. Formally, a MOMDP is defined as a tuple $(\mathcal{S}, \mathcal{A}, p, \mathbf{r}, \mu, \gamma)$, where \mathcal{S} represents a state space, \mathcal{A} an action space, $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$ is the distribution over next states given state s and action a , $\mathbf{r} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto \mathbb{R}^m$ is a multi-objective reward function containing m objectives, μ is an initial state distribution, and $\gamma \in [0, 1)$ is a discounting factor. A *policy* $\pi : \mathcal{S} \mapsto \mathcal{A}$ is a function mapping states to actions, which is typically what the agent aims at learning. Let S_t and A_t denote the variables corresponding to state and action at time step t . The vector value function of a policy π in state s is defined as

$$\mathbf{v}^\pi(s) \equiv \mathbb{E}_\pi \left[\sum_{t=0}^{\infty} \gamma^t \mathbf{r}(S_t, A_t, S_{t+1}) \mid S_t = s \right],$$

where $\mathbb{E}_\pi[\cdot]$ denotes expectation with respect to trajectories induced by the policy π . We denote $\mathbf{v}^\pi \equiv \mathbb{E}_{s_0 \sim \mu}[\mathbf{v}^\pi(s_0)]$ as the *value vector* of policy π .

Solution concepts. Unlike single-objective RL, where the value of a policy is a scalar value, comparing the values of two different policies in MORL is not a straightforward task. Indeed, in MORL, it is possible for a policy to excel in achieving higher expected return for one objective while simultaneously exhibiting lower performance for other objectives. Hence, it is impossible to determine a total ordering on the value vectors, e.g. how to know if $\mathbf{v}^\pi = (1, 0)$ is better than $\mathbf{v}^{\pi'} = (0, 1)$ without knowing the user preferences with regards to the objectives at hand?

In such settings, algorithms usually rely on the concept of Pareto dominance as a partial ordering of the solutions. A policy π is said to *Pareto dominate* another policy π' if its associated value vector is strictly greater for at least one objective without being worse for any other objective:

$$\mathbf{v}^\pi \succ \mathbf{v}^{\pi'} \iff (\forall i : v_i^\pi \geq v_i^{\pi'}) \wedge (\exists i : v_i^\pi > v_i^{\pi'}).$$

Since this ordering is partial, some solutions may not dominate or be dominated by others. A *Pareto frontier* (PF) defines the set of all value vectors which are nondominated, representing all achievable trade offs between the objectives found by the MORL algorithm. Formally:

$$\mathcal{F} \equiv \{\mathbf{v}^\pi \mid \nexists \pi' \text{ s.t. } \mathbf{v}^{\pi'} \succ \mathbf{v}^\pi\}.$$

This set is usually presented to the users after the learning process to let them choose which compromise to make before execution of the corresponding policy. Hence, **the solution returned by MORL algorithms consists in a set of policies** (that we denote as Π^*) linked to these nondominated value vectors, called *Pareto set* (PS).

In contrast, note that when the user preferences are known before the learning process (*a priori*), MORL algorithms return only one policy since the user preferences provide a total ordering of the value vectors. Nevertheless, this study focuses on the setting where the user preferences are unknown.

Performance metrics In the unknown preferences setting, it is often challenging to evaluate the superiority of one algorithm over another because their respective PFs may (partially) overlap or intersect. In the field of multi-objective optimization and MORL, the use of *performance metrics* is common practice. These metrics serve the purpose of transforming the PFs into scalar values, thereby facilitating comparison between algorithms. Two primary aspects are typically taken into consideration when assessing the quality of a PF: *convergence* and *diversity*. Convergence measures how closely

the discovered PF aligns with an optimal PF, indicating the quality of the policies. Diversity, on the other hand, quantifies the variety of compromises presented, which allows for a broad range of trade-offs to be presented to the end user [26]. There exists many metrics that quantify either of those aspects, or even hybrid metrics quantifying both at the same time.

Inverted generational distance (IGD) [8] is a convergence metric that characterizes the distance between the PF found by the algorithm and a known optimal PF. Let \mathcal{Z} be such a PF for the given domain. Then, $\text{IGD}(\mathcal{F}, \mathcal{Z}) = \frac{1}{|\mathcal{Z}|} \sqrt{\sum_{\mathbf{v}^* \in \mathcal{Z}} \min_{\mathbf{v}^\pi \in \mathcal{F}} \|\mathbf{v}^* - \mathbf{v}^\pi\|^2}$.

Sparsity [29] measures the diversity based on the distance between the policies found by the MORL algorithm. It can be computed using the following formula: $S(\mathcal{F}) = \frac{1}{|\mathcal{F}|-1} \sum_{j=1}^m \sum_{i=1}^{|\mathcal{F}|-1} (\mathcal{L}_j(i) - \mathcal{L}_j(i+1))^2$, where \mathcal{L}_j is the sorted list of the values of the j -th objective considering all policies in \mathcal{F} , and $\mathcal{L}_j(i)$ is the i -th value in \mathcal{L}_j .

Hypervolume [32] is a hybrid metric representing the volume of the area created between each point in the approximated Pareto front and a reference point in the objective space. This reference point, \mathbf{v}_{ref} , is (carefully) chosen to be a lower bound in each objective. Formally, $\text{HV}(\mathcal{F}, \mathbf{v}_{\text{ref}}) = \bigcup_{\mathbf{v}^\pi \in \mathcal{F}} \text{volume}(\mathbf{v}_{\text{ref}}, \mathbf{v}^\pi)$.

Finally, some metrics aim at quantifying the utility of the final user by modelling it using a function. In that vein, *expected utility* (EU) [31] reflects the user utility across various potential preferences, under the assumption that it can be captured by a weighted sum utility function. Let \mathcal{W} be an m -dimensional simplex representing a set of weight vectors characterizing potential user preferences, then $\text{EU}(\mathcal{F}) = \mathbb{E}_{\mathbf{w} \sim \mathcal{W}} [\max_{\mathbf{v}^\pi \in \mathcal{F}} \mathbf{w} \cdot \mathbf{v}^\pi]$.

2.2 Hyperparameter optimization for RL

Hyperparameter optimization for RL is the process of finding a set of hyperparameter values for the RL algorithm which leads to improved learning performances. The work of [21, 9] provide formal definitions of the objectives in hyperparameter optimization for single-objective RL. In general, for a given RL algorithm denoted A , this problem can be defined as:

$$\max_{\zeta} f(\zeta, \pi^*, v^{\pi^*}) \quad (1)$$

$$\text{s.t. } \pi^*, v^{\pi^*} = A(\zeta) \quad (2)$$

$$\zeta \in Z \wedge \text{valid}(\zeta), \quad (3)$$

where ζ denotes the hyperparameters values, Z denotes the space of hyperparameter values, π^* and v^{π^*} denote the (potentially optimal) policy found by the RL algorithm and its corresponding value¹. f denotes the objective function for the result of the learning process using the given hyperparameters.

Aside from finding hyperparameter values leading to good performance, it is interesting to remark that the data collected during the optimization process may be useful to understand the behaviour of the studied RL algorithm. Indeed, the search memory can be seen as a dataset where features consist of hyperparameters and labels are the resulting objective values. From this, one can study the relative importance of hyperparameters on the final performances using sensitivity analysis tools from supervised learning [7].

While the mathematical definition of the problem seems straightforward, solving such an optimization problem is in practice very challenging due to the following aspects:

¹ Notice that v^π is just the counterpart of \mathbf{v}^π in single-objective RL.

Stochastic. The effectiveness of RL algorithms can vary based on the random seeds employed during training [13, 9]. As a result, aggregated metrics are often utilized to report these performances across multiple runs for statistical purposes. To achieve this in the HPO context, the RL algorithm needs to be executed on various seeds for each possible set of hyperparameter values and aggregated values of the resulting performance can be used as an objective function. Formally,

$$f(\zeta) = \text{Agg}_{\sigma \in \Sigma_{\text{search}}}(\text{Eval}(A, \zeta)), \quad (4)$$

where $\text{Agg}_{\sigma \in \Sigma_{\text{search}}}$ represents an aggregation operator over a set of random seeds² used in the search phase (typically the mean), and Eval gives the objective value of training one policy given the hyperparameter values. In general, this evaluation is linked to the value of the returned policy, *e.g.* $\text{Eval}(A, \zeta) = v^{\pi^*}$, but more complex schemes exist and are discussed below.

Moreover, it is also crucial to ensure that the seeds utilized during the search phase are distinct from those employed for the final validation of the best hyperparameters discovered [9], *i.e.* $\Sigma_{\text{search}} \cap \Sigma_{\text{validation}} = \emptyset$. This precautionary measure prevents the optimization process from overfitting the search seeds.

Multi-objective. As mentioned above, the HPO objective function f can be different depending on the context; it is in general linked to the agent performance on one environment, but it could also be the total training time required, or the performance of a policy on different environments if the user wants to optimize for policies able to generalize, *e.g.* $\text{Eval}(A, \zeta) = \mathbb{E}_{i \sim \mathcal{I}} [v^{\pi^*}(i)]$, with \mathcal{I} representing a distribution of instances. Although most works choose to optimize only one objective for simplicity, some have already identified HPO for machine learning and RL to be multi-objective problems [21, 16, 6, 14, 7].

Constrained. There is a potential for erratic behaviors and exceptions to occur during the RL learning process due to various combinations of hyperparameter values. In order to represent this aspect of the problem, Equation 3 in the above optimization problem formulation establishes a space of configurations Z that is then checked for validity. In practice, these configurations can be verified either within the optimizer or within the RL algorithm itself, for example by returning a penalized objective value if an invalid configuration is detected.

Computationally expensive. It is widely acknowledged that training an RL algorithm on a specific environment demands a significant amount of computing time. Additionally, it is important to note that during the HPO process, the RL algorithm needs to be executed with multiple random seeds for each potential combination of hyperparameter values. This practical requirement renders the problem impractical for exhaustive search methods.

In practice, such expensive problems are usually solved by relying on Bayesian Optimization (BO) [20, 17], where a *surrogate* model of the objective value is learned by the optimizer based on the collected samples during the search process. This guides the optimizer to choose the most promising areas of the hyperparameter space. Moreover, early-stopping mechanisms such as hyperband [10, 18] help allocating more resources to well performing hyperparameters to better use the budget given to the optimization process.

² To avoid confusion with the RL states, the seeds are denoted using sigma.

3 Hyperparameter optimization for MORL

The previous sections presented background information on the problems of multi-objective RL and hyperparameter optimization for RL. Building on these, this section presents the problem of hyperparameter optimization applied to the field of MORL.

3.1 Formal definition

The main difference with HPO for RL lies in the fact that MORL algorithms return a Pareto set and its corresponding Pareto front. Hence, Equations 1–3 can be adapted as follows:

$$\max_{\zeta} f(\zeta, \Pi^*, \mathcal{F}) \quad (5)$$

$$\text{s.t. } \Pi^*, \mathcal{F} = A(\zeta) \quad (6)$$

$$\zeta \in Z \wedge \text{valid}(\zeta), \quad (7)$$

where Π^* and \mathcal{F} represent the Pareto set of policies being learned by the MORL algorithm and its corresponding Pareto front.

Dealing with sets as output entails an additional step in the HPO process: transforming the PFs into scalars for comparison. As discussed in Section 2.1, there are two aspects (or objectives) to consider in such a case: convergence and diversity. This means that the problem of HPO for MORL is also multi-objective. An initial simple solution to such an issue is to rely on performance metrics. In particular, hybrid performance metrics have the ability to quantify both aspects at the same time. Yet these metrics also suffer from their ability to summarize the information; for example, it is possible to have a large hypervolume score with only one well-placed point in the PF, while a PF composed by many points may have a small hypervolume score.

3.2 Proposed methodology

Building on all these aspects and extending the work of [9], we argue that the problem of HPO for MORL can be solved in two phases: a first phase which searches for hyperparameter values on a range of searching seeds, and a second phase which validates the final performance of the found hyperparameters on a range of validation seeds.

Algorithm 1 illustrates our proposed methodology. The algorithm starts by instantiating an optimizer, which is in charge of proposing hyperparameter values and keeping track of the associated final performances (line 1). Such an optimizer can be based on Bayesian Optimization such as [10], but also could rely on simpler mechanisms such as grid search or random search. Next, the algorithm enters its search phase for a given amount of time specified by a stopping criterion. For each trial, the optimizer generates a new hyperparameter values candidate (line 3). These hyperparameters are then used to train the MORL algorithm on various search seeds for a given budget b_{search} (lines 4–6). As outlined in Section 3.1, a key distinction from HPO for RL emerges. In the MORL scenario, the outcomes of diverse runs manifest as PFs (sets of vectors), necessitating conversion into comparable entities. Consequently, each of these PFs are initially converted into scalar values through the utilization of a performance metric akin to those introduced in Section 2.1. Next, to deal with the stochastic aspect of the optimization problem, these result metrics are aggregated as discussed in Section 2.2 and reported to the optimizer (lines 7–8). Then, the hyperparameters leading to the best evaluation are returned by the optimizer to start the validation phase (line 11). In this phase, the MORL algorithm is trained over

various validation seeds for a given budget $b_{\text{validation}}$ (lines 12–14). Finally, the validation performances and best-performing hyperparameters are returned to the user.

Algorithm 1 Hyperparameter optimization for MORL.

Input: Stopping criterion $stop$, MORL algorithm A , hyperparameter space Z , environment env , searching budget b_{search} , validation budget $b_{\text{validation}}$, optimizer Opt , set of search seeds Σ_{search} , set of validation seeds $\Sigma_{\text{validation}}$.

Output: The best hyperparameter values for algorithm A on env , ζ^* and corresponding Pareto fronts.

```

1:  $opt = new\ Opt(Z)$  ▷ Searching phase
2: while  $\neg stop$  do
3:    $\zeta = opt.next()$ 
4:   for  $\sigma \in \Sigma_{\text{search}}$  do
5:      $\mathcal{F}_\sigma, \Pi_\sigma^* = A(\zeta, \sigma, b_{\text{search}}, env)$ 
6:   end for
7:    $f = Agg_{\sigma \in \Sigma_{\text{search}}}(Eval(\mathcal{F}_\sigma))$  ▷ See Equation 4
8:    $opt.report(\zeta, f)$ 
9: end while
10:
11:  $\zeta^* = opt.best()$  ▷ Validation phase
12: for  $\sigma \in \Sigma_{\text{validation}}$  do
13:    $\mathcal{F}_\sigma, \Pi_\sigma^* = A(\zeta^*, \sigma, b_{\text{validation}}, env)$ 
14: end for
15: return  $\mathcal{F}, \zeta^*$ 

```

4 Experiments

This section presents some early results of the proposed methodology on a state-of-the-art MORL algorithm. Our main goals are: (1) to validate that HPO provides an advantage compared to relying on default hyperparameter values, (2) to showcase the potential benefits of studying hyperparameters through the lense of sensitivity analysis on the search memory.

In practice, we optimized 15 hyperparameters of Envelope Q-Learning [30] according to the hyperparameter space described in Table 1. Having such a large number of hyperparameters to consider makes it virtually impossible for a human to tackle the problem. We study this algorithm on the *minecart-v0* environment [1, 3]. In this

Table 1. Hyperparameters search space (Z).

Hyperparameter	Value range
learning_rate	[0.001, 0.0001]
learning_starts	[1..1, 000]
buffer_size	[1, 000..2, 000, 000]
max_grad_norm	[0.1, 10.0]
gradient_updates	[1..10]
target_net_update_freq	[1..10, 000]
tau	[0.0, 1.0]
num_sample_w	[2..10]
per_alpha	[0.1, 0.9]
initial_homotopy_lambda	[0.0, 1.0]
final_homotopy_lambda	[0.0, 1.0]
homotopy_decay_steps	[1..100, 000]
initial_epsilon	[0.01, 1.0]
final_epsilon	[0.01, 1.0]
epsilon_decay_steps	[1..100, 000]

Table 2. HPO parameters for our experiments.

Parameter	Value
Z	See Table 1
$stop$	48 hours (37 trials)
Σ_{search}	[10..12]
b_{search}	100,000 steps
$Eval$	HV(\mathcal{F}_σ), $\mathbf{v}_{\text{ref}} = (-1, -1, -200)$
Agg	Mean
$\Sigma_{\text{validation}}$	[0..9]
$b_{\text{validation}}$	400,000 steps

environment, the agent controls a cart and has three objectives: it must collect two types of ore from different mines as well as minimize fuel consumption. Finding a diverse PF in such a domain is particularly challenging due to the delayed rewards associated with both ores. Indeed, the agent must first go to a mine to gather ore, then return to its base to obtain the related rewards. This necessitates a significant amount of exploration.

Our experiments were conducted using the implementations in MORL-Baselines [11], MO-Gymnasium [3], and we used the Bayesian optimizer of Weights and Biases (W&B) [5]. Table 2 presents a global overview of the parameters linked to the optimization process³. To tackle the stochastic aspect of the problem, we chose to run each trial on 3 search seeds (in parallel) in order to balance between speed and accuracy of the trials’ estimates. Hyper-volume has been chosen as performance metric because of its ability to characterize both diversity and convergence. The total optimization budget was 2 days, which was equivalent to running 37×3 trainings (each one for 100,000 steps). For reference, each validation training (400,000 steps) required approximately 3.5 hours. All our experiments have been conducted using the high-performance computer of the University of Luxembourg [28].

4.1 Sensitivity analysis

The integration with W&B provides dashboards allowing analysis of the search phase. These dashboards utilize the search memory as a dataset, where hyperparameter values serve as features and resulting performances as labels. Then, a random forest is trained to calculate parameter importance, while linear correlations between hyperparameters and final results are also presented. In the case under study, we extracted the four most significant hyperparameters from the dashboard and listed these in Table 3.

Table 3. Hyperparameters importance and correlations.

Parameter	Importance	Correlation
final_homotopy_lambda	0.288	0.678
tau	0.222	-0.617
homotopy_steps	0.101	0.558
epsilon_decay_steps	0.064	0.457

The table indicates that the homotopy optimization approach used by Envelope has a strong influence on the final results. This involves optimizing a linear combination of two

³ The W&B hyperparameter optimization dashboard can be found at: <https://api.wandb.ai/links/openrlbenchmark/9ffv3e5c>. Our code can be found at <https://github.com/LucasAlegre/morl-baselines/tree/paper/hpo-morl-modem>.




Figure 1. Average performance metrics and 95% confidence intervals over 10 seeded runs of the tuned algorithm and existing baselines on *mincart-v0*.

loss functions to smoothen the objective landscape, given by $L(\theta) = (1 - \lambda_k) \cdot L^A(\theta) + \lambda_k \cdot L^B(\theta)$, where L^A is a loss based on vectorial Q-values and L^B is based on scalarized values. The parameter $\lambda_k \in [\text{initial_homotopy_lambda}, \text{final_homotopy_lambda}]$ is increased throughout the learning process based on a schedule given by the parameters listed in the table. It seems that on this environment, a large λ (and thus relying more on L^B) is correlated with good performances.

The parameter τ is linked to the soft update of the target network, as in DQN [19]. Keeping this parameter relatively low appears to be linked to better performance. Finally, $\text{epsilon_decay_steps}$ regulates the decay rate of the agent’s exploration strategy. Maintaining a high level of exploration appears to be associated with favorable final performance. This observation is consistent with the earlier discussion concerning the environment’s delayed rewards. For a more detailed explanation of these parameters, refer to the original Envelope paper by Yang et al. [30].

4.2 Validation

To validate the effectiveness of the proposed method, we compare the results of training Envelope with the tuned hyperparameters (Envelope tuned) against training with the default hyperparameters (Envelope default) defined in MORL-Baselines. Furthermore, to establish an upper baseline, we compare the results against GPI-PD [2], one of the most sample-efficient MORL algorithm available at the time of writing. For the baseline results, we reuse the ones provided by the MORL-Baselines authors in `openrlbenchmark` [15]. To compare these algorithms, we rely on the performance metrics discussed in Section 2.1. These allow assessing the diversity and convergence aspects of the resulting PFs as well as the final expected user utility.

Figure 1 illustrates the learning curves of the three different algorithms over training steps. There is a drastic difference between Envelope default and its tuned counterpart. Indeed, Envelope default seems to fail to discover new policies over the course of the learning process, as all metrics stay rather flat. On the other hand, Envelope tuned improves on all metrics, meaning it improves the learned PF. Upon comparison with GPI-PD, Envelope tuned exhibits remarkably strong performance, achieving results almost on par with this more recent technique. It is worth noting that Envelope benefits from a faster implementation in practice, resulting in even closer performance when measured against walltime instead of environment steps. Naturally, to ensure a fair comparison between Envelope and GPI-PD, the latter algorithm should also be tuned on this environment.

5 Conclusion

In this work, we discussed the benefits and challenges of hyperparameter optimization for multi-objective reinforcement learning. We presented the existing literature on MORL and HPO for RL. Then, we derived a methodology for conducting HPO for MORL based on these. As early results, our methodology has been applied on a state-of-the-art algorithm and shown to greatly improve its final performances on a classical MORL environment. We also present an analysis of the importance of some hyperparameters using the newly introduced tools. This work is opening avenues for a significant amount of future research opportunities.

First, our implementation could be further improved in various ways. For example, one could improve the optimizer by including more advanced techniques such as hyperband [18] or by distributing the search phase over various computing nodes. Moreover, the set of tools for sensitivity analysis of the optimization process to gain better understanding of the hyperparameters’ influence could be enhanced by relying on the extensive literature from the supervised learning community.

Second, it is important to evaluate the adequacy of utilizing three search seeds as a general practice. This assessment involves confirming whether the optimization consistently identifies comparable values for high-performing hyperparameters across multiple runs employing varying search seeds.

Third, an avenue for exploration could involve an investigation into the generality of the discovered hyperparameters. This exploration would entail determining whether the identified hyperparameters yield favorable performance exclusively within the studied environment, across a range of related environments (e.g., Mujoco [27]), or even universally across all environments.

Fourth, this work could also be extended to automated algorithm design, where entire algorithmic choices, such as the Bellman update components, can be decided by the optimizer.

Finally and more importantly, it would be interesting to apply the proposed methodology to a broader set of applications, e.g. other algorithms and other environments, and study its results. In particular, we believe that the process of HPO for MORL will play a pivotal role in methodically accumulating significant empirical evidence regarding MORL algorithms’ performance. This, in turn, will aid in enabling a thorough comparative analysis of various MORL algorithms.

Acknowledgements

This work was financed by the Fonds National de la Recherche Luxembourg (FNR), CORE program under the ADARS Project, ref. C20/IS/14762457.

References

- [1] Axel Abels, Diederik M. Roijers, Tom Lenaerts, Ann Nowé, and Denis Steckelmacher, ‘Dynamic Weights in Multi-Objective Deep Reinforcement Learning’, *arXiv:1809.07803 [cs, stat]*, (May 2019). arXiv: 1809.07803.
- [2] Lucas N Alegre, Ana L C Bazzan, Diederik M Roijers, and Ann Nowé, ‘Sample-Efficient Multi-Objective Learning via Generalized Policy Improvement Prioritization’, in *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems*, (2023).
- [3] Lucas N. Alegre, Florian Felten, El-Ghazali Talbi, Grégoire Danoy, Ann Nowé, Ana LC Bazzan, and Bruno C. da Silva, ‘MO-Gym: A Library of Multi-Objective Reinforcement Learning Environments’, in *Proceedings of the 34th Benelux Conference on Artificial Intelligence BNAIC/Benelearn*, (2022).
- [4] Marcin Andrychowicz, Anton Raichuk, Piotr Stańczyk, Manu Orsini, Sertan Girgin, Raphaël Marinier, Leonard Hussenot, Matthieu Geist, Olivier Pietquin, Marcin Michalski, Sylvain Gelly, and Olivier Bachem, ‘What Matters for On-Policy Deep Actor-Critic Methods? A Large-Scale Study’, (January 2021).
- [5] Lukas Biewald. Experiment Tracking with Weights and Biases, 2020.
- [6] Martin Binder, Julia Moosbauer, Janek Thomas, and Bernd Bischl, ‘Multi-objective hyperparameter tuning and feature selection using filter ensembles’, in *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*, GECCO ’20, pp. 471–479, New York, NY, USA, (June 2020). Association for Computing Machinery.
- [7] Bernd Bischl, Martin Binder, Michel Lang, Tobias Pielok, Jakob Richter, Stefan Coors, Janek Thomas, Theresa Ullmann, Marc Becker, Anne-Laure Boulesteix, Difan Deng, and Marius Lindauer, ‘Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges’, *WIREs Data Mining and Knowledge Discovery*, **13**(2), e1484, (2023). [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/widm.1484](https://onlinelibrary.wiley.com/doi/pdf/10.1002/widm.1484).
- [8] Carlos A. Coello Coello and Margarita Reyes Sierra, ‘A Study of the Parallelization of a Coevolutionary Multi-objective Evolutionary Algorithm’, in *MICAI 2004: Advances in Artificial Intelligence*, eds., Raúl Monroy, Gustavo Arroyo-Figueroa, Luis Enrique Sucar, and Humberto Sossa, Lecture Notes in Computer Science, pp. 688–697, Berlin, Heidelberg, (2004). Springer.
- [9] Theresa Eimer, Marius Lindauer, and Roberta Raileanu. Hyperparameters in Reinforcement Learning and How To Tune Them, June 2023. arXiv:2306.01324 [cs].
- [10] Stefan Falkner, Aaron Klein, and Frank Hutter, ‘BOHB: Robust and Efficient Hyperparameter Optimization at Scale’, in *Proceedings of the 35th International Conference on Machine Learning*, pp. 1437–1446. PMLR, (July 2018). ISSN: 2640-3498.
- [11] Florian Felten and Lucas N. Alegre. MORL-Baselines: Multi-Objective Reinforcement Learning algorithms implementations, 2022. Publication Title: GitHub repository.
- [12] Conor Hayes, Roxana Rădulescu, Eugenio Bargiacchi, Johan Källström, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai Irissappane, Patrick Mannion, Ann Nowe, Gabriel Ramos, Marcello Restelli, Peter Vamplew, and Diederik Roijers, ‘A practical guide to multi-objective reinforcement learning and planning’, *Autonomous Agents and Multi-Agent Systems*, **36**, (April 2022).
- [13] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger, ‘Deep reinforcement learning that matters’, in *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’18/IAAI’18/EAAI’18, pp. 3207–3214, New Orleans, Louisiana, USA, (February 2018). AAAI Press.
- [14] Daniel Horn and Bernd Bischl, ‘Multi-objective parameter configuration of machine learning algorithms using model-based optimization’, in *2016 IEEE Symposium Series on Computational Intelligence (SSCI)*, pp. 1–8, (December 2016).
- [15] Shengyi Huang, Quentin Gallouédec, Florian Felten, Antonin Raffin, Rousslan Fernand Julien Dossa, Yanxiao Zhao, Ryan Sullivan, Viktor Makoviychuk, Denys Makoviychuk, Cyril Roumégous, Jiayi Weng, Chufan Chen, Masudur Rahman, João G. M. Araújo, Guorui Quan, Daniel Tan, Timo Klein, Rujikorn Charakorn, Mark Towers, Yann Berthelot, Kinal Mehta, Dipam Chakraborty, Arjun KG, Valentin Char-
- raut, Chang Ye, Zichen Liu, Lucas N. Alegre, Jongwook Choi, and Brent Yi. openrlbenchmark, May 2023.
- [16] *Multi-Objective Machine Learning*, eds., Yaochu Jin and Janusz Kacprzyk, volume 16 of *Studies in Computational Intelligence*, Springer, Berlin, Heidelberg, 2006.
- [17] Donald R. Jones, Matthias Schonlau, and William J. Welch, ‘Efficient Global Optimization of Expensive Black-Box Functions’, *Journal of Global Optimization*, **13**(4), 455–492, (December 1998).
- [18] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar, ‘Hyperband: A Novel Bandit-Based Approach to Hyperparameter Optimization’, *Journal of Machine Learning Research*, **18**(185), 1–52, (2018).
- [19] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fiedjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis, ‘Human-level control through deep reinforcement learning’, *Nature*, **518**(7540), 529–533, (February 2015).
- [20] J. Moćkus, ‘On bayesian methods for seeking the extremum’, in *Optimization Techniques IFIP Technical Conference Novosibirsk, July 1–7, 1974*, ed., G. I. Marchuk, Lecture Notes in Computer Science, pp. 400–404, Berlin, Heidelberg, (1975). Springer.
- [21] Jack Parker-Holder, Raghu Rajan, Xingyou Song, André Biedenkapp, Yingjie Miao, Theresa Eimer, Baohe Zhang, Vu Nguyen, Roberto Calandra, Aleksandra Faust, Frank Hutter, and Marius Lindauer, ‘Automated Reinforcement Learning (AutoRL): A Survey and Open Problems’, *Journal of Artificial Intelligence Research*, **74**, (September 2022).
- [22] Andrew Patterson, Samuel Neumann, Martha White, and Adam White. Empirical Design in Reinforcement Learning, April 2023. arXiv:2304.01315 [cs].
- [23] Diederik Marijn Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley, ‘A Survey of Multi-Objective Sequential Decision-Making’, *Journal of Artificial Intelligence Research*, **48**, 67–113, (October 2013). arXiv: 1402.0590.
- [24] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis, ‘Mastering the game of Go with deep neural networks and tree search’, *Nature*, **529**(7587), 484–489, (January 2016). Bandiera_abetest: a Cg_type: Nature Research Journals Number: 7587 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Computational science; Computer science; Reward Subject_term_id: computational-science; computer-science; reward.
- [25] Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, Adaptive Computation and Machine Learning series, A Bradford Book, Cambridge, MA, USA, 2 edn., November 2018.
- [26] El-Ghazali Talbi, *Metaheuristics: From Design to Implementation*, volume 74, June 2009. Journal Abbreviation: Metaheuristics: From Design to Implementation Publication Title: Metaheuristics: From Design to Implementation.
- [27] Emanuel Todorov, Tom Erez, and Yuval Tassa, ‘MuJoCo: A physics engine for model-based control’, in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033, (October 2012). ISSN: 2153-0866.
- [28] Sébastien Varrette, Pascal Bouvry, Hyacinthe Cartiaux, and Fotis Georgatos, ‘Management of an academic HPC cluster: The UL experience’, in *2014 International Conference on High Performance Computing & Simulation (HPCS)*, pp. 959–967. IEEE, (2014).
- [29] Jie Xu, Yunsheng Tian, Pingchuan Ma, Daniela Rus, Shinjiro Sueda, and Wojciech Matusik, ‘Prediction-Guided Multi-Objective Reinforcement Learning for Continuous Robot Control’, in *Proceedings of the 37th International Conference on Machine Learning*, pp. 10607–10616. PMLR, (November 2020). ISSN: 2640-3498.
- [30] Runzhe Yang, Xingyuan Sun, and Karthik Narasimhan, ‘A Generalized Algorithm for Multi-Objective Reinforcement Learning and Policy Adaptation_arxiv’, *arXiv:1908.08342 [cs]*, (November 2019). arXiv: 1908.08342.
- [31] Luisa M Zintgraf, Timon V Kanter, Diederik M Roijers, Frans A Oliehoek, and Philipp Beau, ‘Quality Assessment of MORL Algorithms: A Utility-Based Approach’, (2015).

- [32] E. Zitzler, 'Evolutionary algorithms for multiobjective optimization: methods and applications', in *Ph.D. Dissertation. ETH Zurich, Switzerland.*, (December 1999).