

Reinforcement Learning by Guided Safe Exploration

Yang, Qisong; Simão, Thiago D.; Jansen, Nils; Tindemans, Simon H.; Spaan, Matthijs T.J.

DOI

[10.3233/FAIA230598](https://doi.org/10.3233/FAIA230598)

Publication date

2023

Document Version

Final published version

Published in

ECAI 2023 - 26th European Conference on Artificial Intelligence, including 12th Conference on Prestigious Applications of Intelligent Systems, PAIS 2023 - Proceedings

Citation (APA)

Yang, Q., Simão, T. D., Jansen, N., Tindemans, S. H., & Spaan, M. T. J. (2023). Reinforcement Learning by Guided Safe Exploration. In K. Gal, K. Gal, A. Nowe, G. J. Nalepa, R. Fairstein, & R. Radulescu (Eds.), *ECAI 2023 - 26th European Conference on Artificial Intelligence, including 12th Conference on Prestigious Applications of Intelligent Systems, PAIS 2023 - Proceedings* (pp. 2858 - 2865). (Frontiers in Artificial Intelligence and Applications; Vol. 372). <https://doi.org/10.3233/FAIA230598>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Reinforcement Learning by Guided Safe Exploration

Qisong Yang^{a,*}, Thiago D. Simão^{b,*}, Nils Jansen^b, Simon H. Tindemans^a and Matthijs T. J. Spaan^a

^aDelft University of Technology – The Netherlands

^bRadboud University, Nijmegen – The Netherlands

Abstract. Safety is critical to broadening the application of reinforcement learning (RL). Often, we train RL agents in a controlled environment, such as a laboratory, before deploying them in the real world. However, the real-world target task might be unknown prior to deployment. Reward-free RL trains an agent without the reward to adapt quickly once the reward is revealed. We consider the *constrained* reward-free setting, where an agent (the guide) learns to explore safely without the reward signal. This agent is trained in a controlled environment, which allows unsafe interactions and still provides the safety signal. After the target task is revealed, safety violations are not allowed anymore. Thus, the guide is leveraged to compose a safe behaviour policy. Drawing from transfer learning, we also regularize a target policy (the student) towards the guide while the student is unreliable and gradually eliminate the influence of the guide as training progresses. The empirical analysis shows that this method can achieve safe transfer learning and helps the student solve the target task faster.

1 Introduction

Despite the numerous achievements of reinforcement learning (RL) [45, 35], safety concerns still prevent the wide adoption of RL [11]. The lack of knowledge about the environment forces standard agents to rely on trial-and-error strategies. However, this approach is incompatible with safety-critical scenarios [15]. For instance, recommender systems should not suggest extremist content [10]. Constrained Markov decision processes (CMDP) [4] express such safety constraints with a cost signal indicating unsafe interactions. Such costs are decoupled from the rewards to facilitate the learning of safe behaviours.

Developments in safe RL have allowed us to learn safe policies in CMDPs. For instance, SAC-Lagrangian [19] combines the Soft Actor-Critic (SAC) [21, 22] algorithm with Lagrangian methods to learn a safe policy in an off-policy way. This algorithm solves high-dimensional problems with a sample complexity lower than on-policy algorithms. Unfortunately, it only finds a safe policy at the end of the training process and may be unsafe while learning. In terms of safety, we consider episode-wise constraints instead of step-wise constraints, so a few unsafe actions are allowed in an episode.

Some knowledge about the safety dynamics can ensure safety during learning. One can pre-compute unsafe behaviour and mask unsafe actions using a so-called shield [3, 26, 8], or start from an initially safe baseline policy and gradually improve its performance while remaining safe [2, 49, 57]. However, these approaches may necessitate numerous interactions with the environment before they

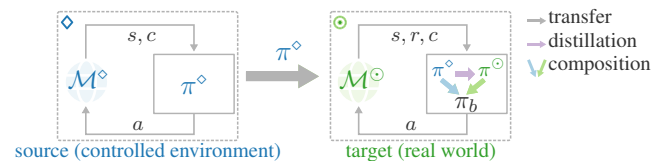


Figure 1. Transferring the Safe Guide (SaGui) policy π^\diamond from the source task (\diamond) to the target task (\odot) with three steps.

find an adequate policy [60]. Moreover, reusing a pre-trained policy can have a detrimental effect, since the agent encounters a new trajectory distribution as the policy changes [24]. Therefore, we investigate *how to efficiently solve a task without violating the safety constraints*.

We make two key observations. First, RL agents often learn in a controlled environment, such as a laboratory or a simulator, before being deployed in the real world [15]. Second, an agent can often benefit from expert guidance instead of solely relying on trial and error [36]. For instance, in autonomous driving, the driver agent can quickly learn by mimicking an expert's behaviour to handle dangerous situations. Such a process is referred to as *policy distillation*. Furthermore, under expert guidance, the agent can safely explore before taking dangerous actions.

Transfer learning [48] investigates how to improve the learning of a target task with some knowledge from a source task. In these settings, the source task may provide only partial knowledge of the target task. We adopt a transfer learning framework and refer to (i) the controlled environment as the *source task* (\diamond) and (ii) the real world as the *target task* (\odot). In our setting, the controlled environment provides only the cost signals related to safety but not the reward signals of the target task in the real world. The central problem is then to avoid safety violations after the target task has been revealed.

Our approach. We show how to transfer knowledge encoded by a policy to enhance safety. Here, we refer to the policy that has been learned in the source task as the *safe guide* (SaGui, Figure 1). The intuition is that, in the real world, the agent is guided to accomplish the target task in a safe manner. We propose to transfer SaGui from the source task to the target task. Our approach has three central steps: *i)* train the SaGui policy and *transfer* it to the target task; *ii)* *distill* the guide's policy into a *student policy* which is dedicated to the target task, and *iii)* *compose* a behaviour policy that balances safe exploration (using the guide) and exploitation (using the student).

As we train the guide in a reward-free constrained RL setting [34], the agent only observes the costs related to safety, and it does not access reward signals. This task-agnostic approach allows us to train a guide independently of the reward of the target task, so this guide can

* Equal contribution.

be useful for different reward functions. Furthermore, we assume the source task preserves the dynamics related to safety, which allows us to train a guide that can act safely when transferred to the target task. Inspired by advances in robotics where an agent is trained under strict supervision, we assume the source task is a simulated/controlled environment [40, 53]. Therefore, safety is not required while training the SaGui policy. Once the target task is revealed, SaGui safely collects the initial trajectories in the target environment and the student starts learning based on these trajectories. To ensure that the new policy quickly learns how to act safely, we also employ a policy distillation method, encouraging the student to imitate the guide.

Contributions. Our main contributions are: we *i)* formalize transfer learning for RL from a safety perspective; *ii)* propose to guide learning using a task-agnostic agent with exploration benefits; *iii)* show how to adaptively regularize the student policy to the guide policy based on the student’s safety; *iv)* investigate when to sample from the student or from the guide to ensure safe behaviour in the target environment and fast convergence of the student policy; and *v)* demonstrate empirically that, compared to learning from scratch and adapting a pre-trained policy, our method can solve the target task faster without violating the safety constraints in the target task.

2 Related Work

Safe RL has multiple facets [15], ranging from alternative optimization criteria [55, 9] to safe exploration based on some prior knowledge [2, 3, 26, 44, 58, 42]. We review methods to train the guide and solve new tasks using a pre-trained policy.

Multiple algorithms have been proposed for generalizing policies from reward-free RL for better performance in target tasks [61, 17, 43]. However, only a few works have considered reward-free RL with constraints [34, 39]. They focus on tabular and linear settings while we consider general function approximation algorithms.

Work in transfer learning has leveraged meta-RL [14] for safe adaptation [18, 32, 30]. Our work is also related to curriculum learning [5, 51, 33]. We first train an agent to be safe and later solve a target task. However, our approach focuses on safe exploration and is able to transfer to tasks with different reward functions, so the guide’s training is ignored.

Our work resembles certain safe transfer-RL frameworks [27, 58], which also leverage prior knowledge to aid learning in a target task. However, the SaGui framework differs from them in terms of safety definition, knowledge acquisition in the source task, or knowledge usage in the target task. Our prior knowledge is more effective for various downstream tasks, and SaGui is the only framework that is safe while learning in the target task.

3 Background

We formalize the safe RL problem and describe typical approaches.

3.1 Constrained Markov Decision Processes

We consider tasks formulated by constrained Markov decision processes (CMDPs) [4, 7]. A CMDP is defined as a tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{P}, r, c, d, \gamma \rangle$: a state space \mathcal{S} , an action space \mathcal{A} , a probabilistic transition function $\mathcal{P}: \mathcal{S} \times \mathcal{A} \mapsto \text{Dist}(\mathcal{S})$, a reward function $r: \mathcal{S} \times \mathcal{A} \mapsto [r_{min}, r_{max}]$, a cost function $c: \mathcal{S} \times \mathcal{A} \mapsto [c_{min}, c_{max}]$, a safety threshold $d \in \mathbb{R}^+$, and a discount factor $\gamma \in [0, 1)$. We also consider an initial state distribution $\iota: \mathcal{S} \mapsto [0, 1]$. In a *constrained RL* problem, an agent interacts with a CMDP without

knowledge about the transition, reward, and cost functions, generating a trajectory $\tau = \langle (s_0, a_0, r_0, c_0, s'_0), (s_1, a_1, r_1, c_1, s'_1), \dots \rangle$. A trajectory starts from $s_0 \sim \iota(\cdot)$. Then, at each timestep t the agent is in a state $s_t \in \mathcal{S}$, and takes an action $a_t \in \mathcal{A}$. It subsequently gets a reward $r_t = r(s_t, a_t)$, a cost $c_t = c(s_t, a_t)$, and steps into a new state $s'_t \sim \mathcal{P}(\cdot | s_t, a_t)$. This process repeats starting from $s_{t+1} = s'_t$ until a terminal condition is met and a new trajectory starts. The goal is to learn a policy π that maximizes the expected discounted return such that the expected discounted cost-return remains below d :

$$\max_{\pi} \mathbb{E}_{\rho_{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t r_t \right] \quad \text{s.t.} \quad \mathbb{E}_{\rho_{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t c_t \right] \leq d, \quad (1)$$

where ρ_{π} indicates the trajectory distribution induced by $s_0 \sim \iota(\cdot)$, $a_t \sim \pi(\cdot | s_t)$, and $s_{t+1} \sim \mathcal{P}(\cdot | s_t, a_t)$. We define the discounted *return* starting from s, a and following π as $Q_{\pi}^r(s, a) = \mathbb{E}_{\rho_{\pi}} [\sum_{t=0}^{\infty} \gamma^t r_t | s_0 = s, a_0 = a]$, and, similarly, the discounted *cost-return* $Q_{\pi}^c(s, a)$.

From the safe RL perspective, if a policy has an expected cost-return lower than the safety-threshold d , then this policy is considered safe. Therefore, the objective of a safe RL agent is to find a policy, among the safe policies, that has the highest expected return.

3.2 Maximum Entropy Reinforcement Learning

A common strategy to improve the exploration and robustness of RL agents is to favour policies that induce diverse behaviours [63, 13]. We can incorporate it in the safe RL main objective by augmenting the problem with a term that aims to maximize the policy entropy [20]:

$$\max_{\pi} \mathbb{E}_{\rho_{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t (r_t + \alpha \mathcal{H}(\pi(s_t))) \right] \quad \text{s.t.} \quad \mathbb{E}_{\rho_{\pi}} \left[\sum_{t=0}^{\infty} \gamma^t c_t \right] \leq d, \quad (2)$$

where $\mathcal{H}(\cdot)$ is the entropy of a distribution over a random variable, and α is the entropy weight. In general, this objective encourages the agent to use maximally stochastic policies. Alternatively, we can encourage the policy to have at least a minimum entropy $\bar{\mathcal{H}}$ [22] by adding the following constraint to (1):

$$\mathbb{E}_{\rho_{\pi}} [-\log(\pi(a_t | s_t))] \geq \bar{\mathcal{H}}, \quad \forall t, \quad (3)$$

where $\bar{\mathcal{H}}$ is the given entropy threshold to ensure a minimum degree of randomness. This approach allows the policy to converge to a more deterministic behaviour than (2). Besides, it only requires the system’s designer to define $\bar{\mathcal{H}}$ and it lets the RL agent automatically find a trade-off between the policy’s entropy and rewards. Therefore, α becomes an intrinsic parameter of the RL algorithm.

The maximum entropy RL with safety constraint (2) can be solved by the SAC-Lagrangian (SAC- λ) [19] method. SAC- λ is a SAC-based method that has two critics and uses an adaptive entropy weight α (parameterized by θ_{α}) and an adaptive safety weight β (parameterized by θ_{β}) to manage a trade-off among exploration, reward, and safety. The reward critic estimates the expected return Q^r (parameterized by θ_R), possibly with an entropy bonus to promote exploration, while the safety critic estimates the cost-return Q^c (parameterized by θ_C) to encourage safety. The policy π is parameterized by θ_{π} . Appendix A [54] provides a detailed description of how to learn each component, including the losses. Throughout the paper, we represent learning rates with η , replay buffers with \mathcal{D} , and losses with J . We only update α when a desirable $\bar{\mathcal{H}}$ is given, so α is fixed whenever we use the formulation (2).

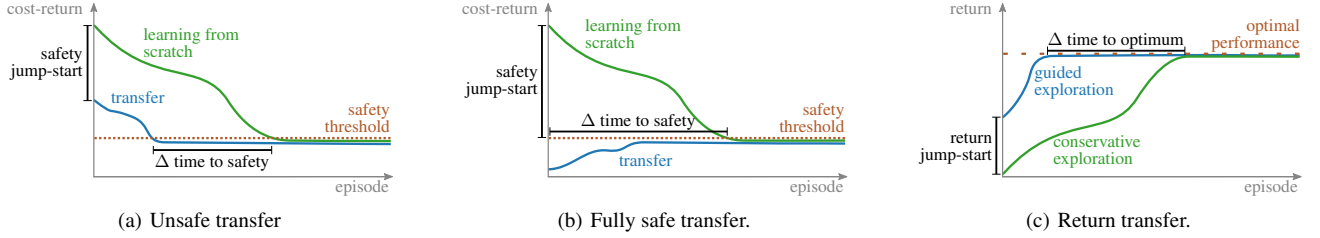


Figure 2. Transfer metrics for safe reinforcement learning. Usually, we consider *safety jump-start* and Δ *time to safety*. If we can develop agents that learn without violating the safety requirements, we can also consider *return jump-start* and Δ *time to optimum*.

4 Safe and Efficient Exploration

Naturally, to train RL agents without violating the safety constraints, some prior knowledge is required [44]. Often, a safe initial policy collects the initial trajectories [2, 49, 57]. However, these approaches largely neglect how this policy is computed or what makes it effective. Therefore, we consider the problem of how to obtain an initial policy that can safely expedite learning in the target task. Next, we formalize the problem and provide an overview of our approach.

4.1 Problem Setting

We formalize our problem setting using the transfer learning (TL) framework. In general, TL allows RL agents to use expertise from *source* tasks to speed up the learning process on a *target* task [48, 62]. The source tasks $\{\mathcal{M}^\diamond\}$ should provide some knowledge \mathcal{K}^\diamond to an agent learning in the target task \mathcal{M}° , such that, by leveraging \mathcal{K}^\diamond , the agent learns the target task \mathcal{M}° faster.

As we are particularly interested in the safety properties of the transfer, we consider a reward-free source task, which only provides information about the safety dynamics. Moreover, we use a policy to encode the knowledge transferred. Formally, given a source task $\mathcal{M}^\diamond = \langle \mathcal{S}^\diamond, \mathcal{A}^\diamond, \mathcal{P}^\diamond, \emptyset, c^\diamond, d^\diamond, \iota^\diamond, \gamma \rangle$, we compute the policy π^\diamond in the absence of a reward signal. This provides knowledge $\mathcal{K}^\diamond = \{\pi^\diamond\}$ to help solving the target task $\mathcal{M}^\circ = \langle \mathcal{S}^\circ, \mathcal{A}^\circ, \mathcal{P}^\circ, r^\circ, c^\circ, d^\circ, \iota^\circ, \gamma \rangle$.

To apply the source policy π^\diamond in the target task \mathcal{S}° , we have a mapping from the source state space to the target state space $\Xi : \mathcal{S}^\diamond \rightarrow \mathcal{S}^\circ$. Then, we can define a target policy $\pi^{\diamond \rightarrow \circ}$ as follows: $\pi^{\diamond \rightarrow \circ}(s) = \pi^\diamond(\Xi(s))$. We assume the source task \mathcal{M}^\diamond and target task \mathcal{M}° share the same action space. Appendix B.1 [54] describes how to obtain the source task based on Ξ and the target task.

Assumption 1. $\mathcal{A}^\diamond = \mathcal{A}^\circ = \mathcal{A}$.

To enable the knowledge transferable between tasks, having the same action spaces ensures that the policy learned in the source task is directly applicable to the target task.

4.2 Transfer Metrics

To evaluate a safe transfer RL algorithm, Figure 2(a) presents a schematic of metrics related to safety (inspired by transfer in RL [48]): *safety jump-start* indicates how much closer to the safety threshold the expected cost-return of an agent learning with the source knowledge is compared to the expected cost-return of an agent learning from scratch in the first episodes, and Δ *time to safety* is the difference in the number of interactions required to become safe.

Notice that a trained agent might start with an expected cost-return lower than the safety threshold, for instance, when the safety threshold in the source task is lower than in the target task (Figure 2(b)).

In this case, *safety jump-start* would be the difference between the safety threshold and the cost-return of an agent learning from scratch. Similarly, the Δ *time to safety* would be the number of interactions an agent learning from scratch needs to become safe.

In the case of two methods that can solve the target task without violating the safety constraints, we can also consider the usual metrics of transfer learning with respect to the reward [48]. For instance, Figure 2(c) shows the initial improvement in terms of performance which we call *return jump-start*, and the time necessary to reach an optimum performance, which we call the Δ *time to optimum*.

Problem statement. We aim to maximize the *safety jump-start* (potentially preventing safety violations in the target task) and to reduce the *time to optimum* (improving exploration) when transferring a policy π^\diamond from a source task \mathcal{M}^\diamond to a target task \mathcal{M}° .

4.3 Method Overview

Recall that for our transfer setting, we consider a single source task that only provides the safety signals, which we use to train the guide. Without the reward signal, the guide aims to explore the world safely and efficiently. We are interested in using the guide’s safe exploration capabilities to train the student on the target task without violating the safety constraints. Notably, *i*) the guide and the student are trained separately; *ii*) the guide is only trained once and can support the training of different students; and *iii*) the guide only has access to safety information and no knowledge about the student’s task.

To ensure the source policy is safe when deployed in the target task, we assume that the source task has a safety threshold lower than or equal to the target task, and Ξ is a state abstraction that preserves the safety dynamics, as formalized next.

Assumption 2. *The safety threshold of the target task upper bounds the safety threshold of the source task: $d^\diamond \leq d^\circ$.*

Assumption 3. Ξ is a $Q_{\pi^\diamond}^c$ -irrelevance abstraction [31], therefore

$$\Xi(s) = \Xi(s') \Rightarrow Q_{\pi^\diamond}^c(s, a) = Q_{\pi^\diamond}^c(s', a), \forall s, s' \in \mathcal{S}^\diamond, a \in \mathcal{A}, \pi^\diamond.$$

Now, we can connect the expected cost-return of a policy on the source task to the expected cost-return on the target task.

Lemma 1. *Given Assumption 1 and Assumption 3, we have*

$$Q_{\pi^\diamond}^{c,\circ}(\Xi(s), a) = Q_{\pi^{\diamond \rightarrow \circ}}^{c,\circ}(s, a) \quad \forall s \in \mathcal{S}^\diamond, a \in \mathcal{A}, \pi^\diamond.$$

That is, the expected cost of a source policy is the same in the source task and in the target task.

Proof. Appendix B.2 provides the proof [54]. \square

Theorem 1. *If Ξ is a Q_π^c -irrelevant state abstraction, then any policy that is safe on the source task \mathcal{M}^\diamond is also safe when deployed on the target task \mathcal{M}° .*

Proof.

$$Q_{\pi^\circ \rightarrow \circ}^{c,\circ}(s, a) \stackrel{\text{Lemma 1}}{=} Q_{\pi^\diamond}^{c,\circ}(\Xi(s), a) \stackrel{\text{Premise}}{\leq} d^\diamond \stackrel{\text{Assumption 2}}{\leq} d^\circ. \quad \square$$

It is important to note, however, that the reward function r° in the target task may be unrelated to the state space of the source task \mathcal{S}^\diamond . Therefore, although a policy that is safe on the source task is also safe on the target task, the behaviour required to accomplish the target task may not be defined on the source task. Consider, for instance, an agent with access to its position and the position of a threat. In each target task, the agent might need to visit a different goal position, which is not defined in the source task. Then, a safe policy may be conditioned only on the positions of the agent and the threat, but to achieve the target, the agent must consider the goal position. This highlights the need to compute a policy dedicated to the target task.

5 Guided Safe Exploration

In this section, we consider how to train the *safe guide* (SaGui) policy. Then, we describe how the student learns to imitate the SaGui policy after the task is revealed while learning to complete the target task. Finally, we investigate how to prevent safety violations while the student has not yet learned how to act safely.

5.1 Training the Safe Guide

Since the source task does not provide information regarding the reward of the target task, we adopt a reward-free exploration approach to train the guide. To efficiently explore the world, we first consider maximizing the policy entropy under safety constraints. Then, we can solve the problem defined in Equation 2 with $r(s, a) = 0 : \forall s \in \mathcal{S}, a \in \mathcal{A}$ to get a guide MAXENT. However, although MAXENT tends to have diverse behaviours, that does not imply efficient exploration of the environment. Especially for continuous state and action spaces, it is possible that a policy provides limited exploration even if it has high entropy.

To enhance the exploration of the guide, we adopt an auxiliary reward that motivates the agent to visit novel states. To measure the novelty, we first define the metric space $(\mathcal{S}^\ddagger, \delta)$, where \mathcal{S}^\ddagger is an abstracted state space and $\delta : \mathcal{S}^\ddagger \times \mathcal{S}^\ddagger \rightarrow [0, \infty)$ is a distance function:

$$\begin{aligned} \delta(s, s') &= 0 \Leftrightarrow s = s', \\ \delta(s, s') &= \delta(s', s), \text{ and} \\ \delta(s', s'') &\leq \delta(s, s') + \delta(s, s''), \quad \forall s, s', s'' \in \mathcal{S}. \end{aligned}$$

Note that \mathcal{S}^\ddagger may not be the original state space \mathcal{S} . Especially when \mathcal{S} is high-dimensional, \mathcal{S}^\ddagger can be some selected dimensions from \mathcal{S} , or a latent space from representation learning. Next, we define the auxiliary reward as the expected distance between the current state and the successor state:

$$r_t^\delta(s_t, a_t) = \mathbb{E}_{s_{t+1} \sim \mathcal{P}(\cdot | s_t, a_t)} \left[\delta(f^\ddagger(s_t), f^\ddagger(s_{t+1})) \right], \quad (4)$$

where we may apply a potential abstraction $f^\ddagger : \mathcal{S} \rightarrow \mathcal{S}^\ddagger$. So, we train the *guide* agent by solving the constraint optimization problem (2) based on the auxiliary reward r^δ . Then, we can use SAC- λ directly to solve (2), as Algorithm 2 shows (Appendix A [54]). In

Algorithm 1 Guided Safe Exploration

Input: $\mathcal{M}^\circ, \pi^\diamond, \overline{\mathcal{H}}, d$

Initialize: $\mathcal{D} \leftarrow \emptyset, \theta_\chi^\circ$ for $\chi \in \{\pi, R, C, \alpha, \beta\}$

Output: Optimized parameters θ_π° for π°

```

1: for each iteration do
2:   for each environment step do
3:     if linear-decay then
4:        $b \leftarrow f_{\text{id}}(\diamond, \circ)$   $\triangleright$  linearly eliminate the effect of  $\pi^\diamond$ 
5:     else if control-switch then
6:        $b \leftarrow f_{\text{cs}}(\diamond, \circ)$   $\triangleright \pi^\diamond$  takes control if unsafe
7:     end if
8:      $a_t \sim \pi_b(\cdot | s_t)$   $\triangleright$  Composite sampling (6)
9:      $\mathcal{I}_t \leftarrow \mathcal{I}(s_t, a_t)$   $\triangleright$  IS ratio (7)
10:     $r_t^\circ \leftarrow r^\circ(s_t, a_t)$ 
11:     $r_t^\diamond \leftarrow \log \pi^\diamond(a_t | \Xi(s_t))$ 
12:     $c_t^\circ \leftarrow c^\circ(s_t, a_t)$ 
13:     $s_{t+1} \sim \mathcal{P}^\circ(\cdot | s_t, a_t)$ 
14:     $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t, r_t^\circ, r_t^\diamond, c_t^\circ, \mathcal{I}_t, s_{t+1})\}$ 
15:  end for
16:  for each gradient step do
17:    Sample experience from  $\mathcal{D}$ 
18:    for  $\chi \in \{\pi, R, C, \alpha, \beta\}$  do
19:       $\theta_\chi^\circ \leftarrow \theta_\chi^\circ - \eta_\chi \hat{\nabla}_{\theta_\chi^\circ} \mathcal{L} J_\chi(\theta_\chi^\circ)$   $\triangleright$  Updating  $\theta_\chi^\circ$ 
20:    end for
21:  end for
22: end for

```

future research, we will also investigate different distance functions to understand their effects on exploration.

This auxiliary reward does not explicitly promote exploration, but we find that increasing the step size and policy entropy significantly improves exploration in practice. Overall, our experiment with the auxiliary reward aimed to evaluate the impact of the exploration of the guide on how safely and quickly the student learns.

We could also consider more sophisticated reward-free exploration strategies such as maximizing the entropy of the state occupancy distribution [41, 47, 23]. We leave this as future work and focus on using the guide to improve how the student learns.

5.2 Policy Distillation From the Safe Guide

When the agent is trained for a certain task, it is difficult to generalize when faced with a new task [24]. Similarly, it is not trivial to adjust the guide's policy that was trained to explore the environment to perform the target task. Therefore, we train a new policy, referred as the student, dedicated to the target task.

We can leverage the *guide* to quickly learn how to act safely. Through the mapping function Ξ , the transferred policy can be used by most constrained RL algorithms to regularize the student policy π° towards the guide policy π^\diamond using KL divergences, as shown in Figure 3. So, with π^\diamond fixed, we have an augmented reward function $r_t' = r_t^\circ + \omega r_t^{\text{KL}} + \alpha r_t^{\mathcal{H}}$, where $r_t^{\text{KL}} = \log \frac{\pi^\circ(a_t | \Xi(s_t))}{\pi^\diamond(a_t | s_t)}$ and $r_t^{\mathcal{H}} = -\log \pi^\circ(a_t | s_t)$. The weights ω and α indicate the strengths of the KL and entropy regularization (respectively). Appendix C [54] shows that setting $r_t^\diamond = \log \pi^\diamond(a_t | \Xi(s_t))$ we obtain $\omega r^{\text{KL}} + \alpha r^{\mathcal{H}} = \omega r^\diamond + (\omega + \alpha) r^{\mathcal{H}}$. Therefore, we can define the student's objective:

$$\max_{\pi^\circ} \mathbb{E}_{\tau \sim \rho_{\pi^\circ}} \sum_{t=0}^{\infty} \gamma^t \left[r_t^\circ + \omega r_t^\diamond + (\omega + \alpha) r_t^{\mathcal{H}} \right]. \quad (5)$$

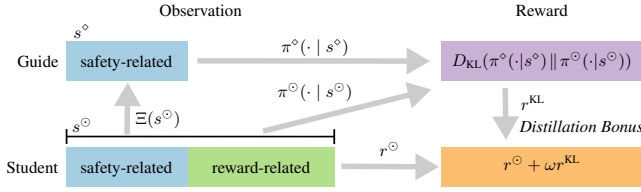


Figure 3. Overview of the policy distillation. Through the mapping function Ξ , the transferred policy can be used to regularize the student policy π° towards the guide policy π^\diamond .

To find an appropriate ω , our goal is to follow the guide more for safer exploration if the student’s policy is unsafe, but eliminate the influence from the guide and focus more on the performance if the student’s policy is safe. Therefore, we propose to set $\omega = \beta$ to determine the strength of the KL regularization since the adaptive safety weight β reflects the safety of the current policy.

In summary, we have an entropy regularized expected return with redefined (regularized) reward $r_t^{\prime\prime} = r_t^\circ + \beta r_t^\diamond$. This augmented reward encourages the student to yield actions that are more likely to be generated by the guide. Then, SAC- λ can be directly used to solve (5) with the additional entropy constraint (Algorithm 1, lines 16-19).

5.3 Composite Sampling

To enhance safety and improve the student during training (Algorithm 1, lines 2-14), we leverage a *composite sampling* strategy, which means our behaviour policy (π_b) is a mixture of the guide’s policy (π^\diamond) and the student’s policy (π°). So, at each environment step, $a_t \sim \pi_b(\cdot | s_t)$, $s_t \in \mathcal{S}^\circ$ where

$$\pi_b(\cdot | s_t) = \begin{cases} \pi^\diamond(\cdot | \Xi(s_t)), & \text{if } b = \diamond, \\ \pi^\circ(\cdot | s_t), & \text{otherwise.} \end{cases} \quad (6)$$

We investigate two strategies to define the stochastic value of b .

Linear-decay (Algorithm 3 in Appendix D [54]). This strategy, denoted as $b = f_{ld}(\diamond, \circ)$, linearly decreases the probability of using π^\diamond with a constant decay rate after each iteration of the algorithm, conversely increasing the probability of using π° . We have two modes with *linear-decay*: *step-wise*, where in each time step we may change π_b ; and *trajectory-wise*, where π_b only changes at the start of a trajectory. The mode is decided before executing an episode, and smoothly switches from the complete *step-wise* to the complete *trajectory-wise* over the training process.

Control-switch (Algorithm 4 in Appendix D [54]). To balance between the safe exploration and the sample efficiency (the samples from the target policy is relatively more valuable), the student policy keeps sampling, i.e., $\pi_b = \pi^\circ$ at the start of a trajectory; after we meet the first $c_{t-1} > 0$, we have $\pi_b = \pi^\diamond$ until the end of the trajectory. Therefore, the guide policy serves as a *rescue policy* to improve safety during sampling. We denote this strategy as $b = f_{cs}(\diamond, \circ)$.

With the *composite sampling* strategy, the function approximation may diverge, because π° and π_b are too different, especially when we collect most data following π^\diamond . This phenomenon is related to the *deadly triad* [46]. To eliminate its negative effect, we endow each sample with an *importance sampling* (IS) ratio:

$$\mathcal{I}(s, a) = \min \left(\max \left(\frac{\pi^\circ(a | s)}{\pi_b(a | s)}, \mathcal{I}_l \right), \mathcal{I}_u \right). \quad (7)$$

The clipping hyper-parameters \mathcal{I}_u and \mathcal{I}_l are introduced to reduce the variance of the off-policy TD target. Notice that if π_b is using the

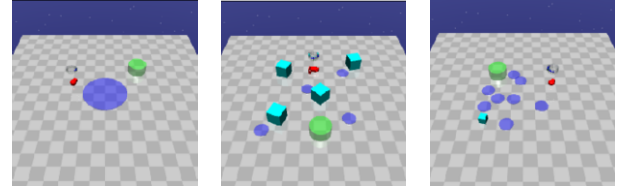


Figure 4. Navigation tasks with different complexity levels: **static** where all objects are fixed (left), **semi-dynamic** where the goal is randomly initialized before each episode (center), and **dynamic** where all objects are randomly initialized locations before each episode (right).

student π° then $\mathcal{I}(s, a) = 1$. Here, in addition to use the IS ratio \mathcal{I} for learning values (the *critics*), we also use it in the policy update, as shown in line 19 of Algorithm 1.

6 Empirical Analysis

We evaluate how well our method transfers from the reward-free setting using the SafetyGym engine [38], where a random-initialized robot navigates a 2D map to reach target positions while trying to avoid dangerous areas and obstacles (Figure 4). These tasks are particularly complex due to the observation space; instead of observing its location, the agent observes the relative location of other objects with a lidar sensor. We considered three environments with different complexity levels. A **static** environment with a point robot and a hazard. The locations of the hazard and goal are fixed in all episodes. A **semi-dynamic** environment with a car robot, four hazards, and four vases. The locations of the hazards and vases are the same in all episodes. The location of the goal is randomly-initialized in each episode. A **dynamic** environment with a point robot, eight hazards, and a vase. The locations of the goal, vase, and hazards are randomly-initialized in each episode.

The *guide* agent is trained without the goals, and its auxiliary reward is the magnitude of displacement at each time step. We provide a detailed description of the safety-mapping function in Appendix G [54]. Since our focus is on the target task and the guide is trained in a controlled environment, we do not consider the guide’s training in the evaluation. In the target tasks, we use the original reward signal from Safety Gym, i.e., the distance towards the goal plus a constant for finishing the task [38]. In all environments: $c = 1$, if an unsafe interaction happens, and $c = 0$, otherwise. We repeat each experiment 10 times with different random seeds and the plots show the mean and standard deviation of all runs.

To evaluate the performance during training, we use the following metrics: safety of the behaviour policy (Cost-Return π_b), performance of the behaviour policy (Return π_b), safety of the target policy (Cost-Return π°), and performance of the target policy (Return π°). To check the convergence of the target policy, we have a test process with 100 episodes after each epoch (in parallel to the training) to evaluate Return π° and Cost-Return π° . Appendix F reports the evaluation of π° and Appendix G the hyper-parameters used [54]. The code of the experiments is available at <https://github.com/qisong-yang/SaGui>.

6.1 Ablation Study

We investigate each component of the proposed SAGUI algorithm individually to answer the following questions: *i*) Does the *auxiliary reward* enlarge the exploration range? *ii*) Does a better *guide* agent result in a better student in the target task? *iii*) How does the *adaptive*

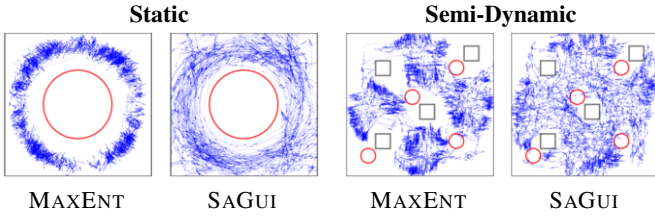


Figure 5. Exploration analysis with trajectories collected by the different guide agents in Static and Semi-Dynamic.

strength of the KL regularization affect the performance? *iv*) How does the *composite sampling* benefit the safe transfer learning?

i) Auxiliary reward leads to more diverse trajectories. We performed an ablation of our approach where no auxiliary reward is added while training the *guide* agent, called MAXENT. We refer to the agent with the auxiliary reward as SAGUI. This teases apart the role the designed auxiliary task plays in the exploration. In Figure 5, we can see that SAGUI can explore larger areas in *Static* and *Semi-Dynamic*, which have the same layout in each episode. We notice that MAXENT is safe, but the explored space is limited. That is also the case in *Dynamic*, as shown in the attached videos.

ii) An effective guide can speed up the student’s training. We compare how these guides (MAXENT and SAGUI) affect the learning in the target task. In Figure 7 (Appendix E [54]), we notice that both methods can collect samples safely, but the agent using the auxiliary reward needs fewer interactions to find high-performing policies.

iii) Safety-adaptive regularization improves the student’s convergence rate. To combine the original reward with the bonus to follow the guide (ω), we have the following choices: fix the weights of the bonus and make it to be a hyperparameter to tune (FIXREG); apply a decay rate to linearly decrease the weights during training (DECREG); and, adapt the weights of the bonus based on the safety performance (SAGUI). In Figure 7(a) (Appendix E [54]), we observe that this weight does not affect the safety of the agent, but both FIXREG and DECREG cause the student to converge slower in terms of performance (Figure 7(b) in Appendix E [54]).

iv) Composite sampling enhances safety and final performance. We modify the composite sampling approach, sampling only from the guide (GUISAM) or the student (STUSAM) instead. From the results in Figure 7(a) (Appendix E [54]), we can see that GUISAM can ensure safety, but the student does not learn a safe optimal policy (Figure 7(b) in Appendix E [54]). Compared to our method, STUSAM performs similarly converging to a safe target policy, but fails to satisfy the constraint at the early stage of training. So, *composite sampling* is necessary to avoid the dangerous actions from a naive policy and to ensure the target task is solved.

6.2 Comparison with Baselines

Finally, we compare ■ SAGUI (control-switch) and ■ SAGUI (linear-decay) with five baselines, divided into three groups.

Learning from scratch. (1) ■ SAC- λ [19] shows the performance when starting to learn from scratch, representing an off-policy algorithm. Similarly, (2) ■ CPO [2] is an on-policy algorithm that maximizes the reward in a small neighbourhood to enforce the safety constraints.

Pre-training. (3) ■ CPO-PRE and (4) ■ SAC- λ -PRE demonstrate how CPO and SAC- λ perform after being pre-trained in a task that replaces the target reward by the auxiliary reward. This way,

we also encourage exploration in the task for pre-training, which shares the same observation space with the target task.

Expert-in-the-loop. (5) As an upper bound, we also consider the Expert Guided Policy Optimization (■ EGPO) [36] algorithm, which uses knowledge from the target task in the form of an expert to train a student policy. EGPO proposes a guardian mechanism that replaces the actions of the student by the expert when the student takes actions too different from the expert. In summary, EGPO uses an expert policy as a demonstrator as well as a safety guardian (see Appendix H [54] for more details).

Notice, for CPO-PRE, SAC- λ -PRE and EGPO we adapt the source task to have the same observation space as the target task, which gives them an advantage compared to SAGUI. Furthermore, EGPO has access to a policy trained on the target task, while SAGUI only has access to the source task without the goal observations.

Safety during training. In Figure 6, we observe that SAGUI (control-switch) and EGPO are the only methods that exhibit safe behaviour during the full training process.

Learning from scratch is unsafe and may converge to sub-optimal and even unsafe policies. SAC- λ and CPO can learn safe policies in relatively simpler environments (*Static* and *Semi-Dynamic*) but they violate the safety constraints at the beginning of training, which is expected. In *Dynamic*, SAC- λ and CPO fail to attain safe performance. However, with benefits from the *guide*, SAGUI (control-switch), on the basis of SAC- λ , attains a better balance between safety and performance.

Pre-training is insufficient. With pre-training, a safe initialization cannot benefit CPO-PRE and SAC- λ -PRE in safety, and may have negative effects. We infer that it is difficult to generalize a task when faced with a new reward signal [24]. Especially for SAC- λ -PRE with an initialized Q^r , the difficulty to adapt is evident.

Fast convergence rates. Benefiting from the targeted expert policy, the behaviour policy of EGPO has a high return throughout the training in the target environment. But SAGUI (control-switch) quickly finds policies with similar performance despite lack of knowledge of the target task (Figure 6).

The distillation mechanism ensures the safety of the target policy. Figure 8 (Appendix F [54]) shows that SAGUI (control-switch) can learn a well-performing target policy in a safe way. Without the policy distillation mechanism like SAGUI, EGPO (learning only from the expert demonstrations) fails to find a safe target policy. This indicates that the target policy computed with SAGUI may eventually take full control of the target task, while the policy computed by EGPO may still require interventions from the expert.

Control-switch can be more effective than linear-decay. SAGUI (linear-decay), which lacks samples from π^\odot at the early stage of training, does not achieve similar performance as SAGUI (control-switch). Figures 6(b) and 6(c) show that *linear-decay* fails to compose the behaviour policy π_b safely.

Summary. Overall, SAGUI does not violate the safety constraints on the target environment, quickly finds high-performing policies, and can train a student able to act independently from the guide.

7 Conclusion

This work handles multiple challenges of reinforcement learning with safety constraints. It shows how we can use a safe exploration policy (the guide) during data collection and gradually switch to a policy that is dedicated to the target task (the student). It tackles the off-policy issue that arises from collecting data with a policy different from the target policy. It shows how the student can make the

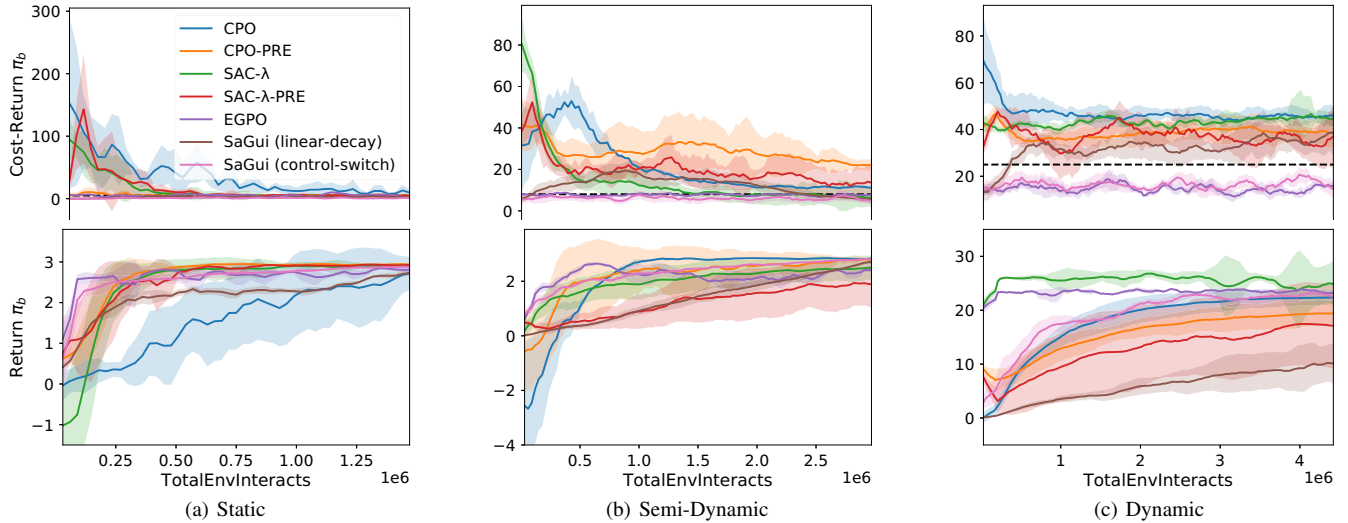


Figure 6. Evaluation of π_b for CPO, CPO-PRE, SAC- λ , SAC- λ -PRE, EGPO, and SaGui over 10 seeds. The solid lines are the average of all runs, and the shaded area is the standard deviation. The black dashed lines indicate the safety thresholds.

best use of the guide’s policy using an incentive to imitate the guide, which makes the student learn faster how to behave safely. It demonstrates that simply initializing an agent with a safe policy may not be as effective as learning a new policy dedicated to the target task through policy distillation. Finally, it proposes a method that can collect diverse trajectories, which reduces the sample complexity of the student on the target task. In summary, the framework proposed is a safe and sample-efficient way of training the agent on a target task.

Limitations. Our framework assumes that the source task provides information on the cost function, allowing the guide policy to accumulate the same cost in the target task as in the source task (Section 4.3). This assumption enables safe learning in the target task. However, if the cost function or trajectory distribution changes, the source task may not provide useful safety information for the target task. In such cases, alternative methods should be considered to ensure safe exploration. We focus on downstream tasks where pre-trained agents are utilized for safe exploration knowledge, disregarding the interactions used to train the SaGui policy. Sample efficiency in the target task is emphasized, not including samples used for source task learning. Nevertheless, the pre-trained policy can be reused for multiple target tasks, enabling us to amortize the guide’s training across them, making the number of samples required to train the guide negligible as the number of downstream tasks increases. While efficient learning of a SaGui policy is a significant challenge, we view it as a separate research direction [23].

Future work. While we consider a relatively simple strategy to achieve rich exploration, our framework allows the translation of any progress in reward-free RL into training the *guide* agent. For instance, we could adopt works with the entropy of the state density [23, 29, 41, 25, 61, 47, 52, 37, 56], or with the adaptive reward functions to explore various skills [12]. Another option to improve exploration is to find a set of diverse policies to the same problem [16, 28, 59]. Our framework could easily combine multiple guides. As to composite sampling strategies, recovery and shielding mechanisms [3, 50] could be further explored to combine with a safe guide, in particular using the control-switch mechanism that we evaluated. Nevertheless, we highlight that while a student using a recovery policy must explore alone, the safe guide can enhance the student’s exploration, accelerating the learning of the target task.

Acknowledgements

We thank the reviewers for their insightful comments. This work has been partially funded by the ERC Starting Grant 101077178 (DEUCE) and the NWO grant NWA.1160.18.238 (PrimaVera). Qisong Yang is supported by Xidian University.

References

- [1] David Abel, D. Ellis Hershkowitz, and Michael L. Littman, ‘Near optimal behavior via approximate state abstraction’, in *ICML*, pp. 2915–2923, (2016).
- [2] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel, ‘Constrained Policy Optimization’, in *ICML*, pp. 22–31, (2017).
- [3] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu, ‘Safe Reinforcement Learning via Shielding’, in *AAAI*, pp. 2669–2678, (2018).
- [4] Eitan Altman, *Constrained Markov decision processes*, volume 7, CRC Press, 1999.
- [5] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston, ‘Curriculum learning’, in *ICML*, pp. 41–48, (2009).
- [6] Dimitri P Bertsekas, *Constrained Optimization and Lagrange Multiplier Methods*, volume 1, Academic press, 1982.
- [7] Vivek S Borkar, ‘An actor-critic algorithm for constrained Markov decision processes’, *Systems & control letters*, **54**(3), 207–213, (2005).
- [8] Steven Carr, Nils Jansen, Sebastian Junges, and Ufuk Topcu, ‘Safe reinforcement learning via shielding under partial observability’, in *AAAI*, pp. 14748–14756, (2023).
- [9] Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone, ‘Risk-constrained reinforcement learning with percentile risk criteria’, *JMLR*, **18**, 6070–6120, (2017).
- [10] Tommaso Di Noia, Nava Tintarev, Panagiota Fatourou, and Markus Schedl, ‘Recommender systems under european ai regulations’, *Communications of the ACM*, **65**(4), 69–73, (2022).
- [11] Gabriel Dulac-Arnold, Nir Levine, Daniel J. Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester, ‘Challenges of real-world reinforcement learning: definitions, benchmarks and analysis’, *Mach. Learn.*, **110**(9), 2419–2468, (2021).
- [12] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine, ‘Diversity is all you need: Learning skills without a reward function’, in *ICLR*, (2019).
- [13] Benjamin Eysenbach and Sergey Levine, ‘Maximum entropy RL (provably) solves some robust RL problems’, in *ICLR*, (2022).
- [14] Chelsea Finn, Pieter Abbeel, and Sergey Levine, ‘Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks’, in *ICML*, pp. 1126–1135, (2017).

- [15] Javier García and Fernando Fernández, ‘A Comprehensive Survey on Safe Reinforcement Learning’, *JMLR*, **16**, 1437–1480, (2015).
- [16] Mahsa Ghasemi, Evan Scope Crafts, Bo Zhao, and Ufuk Topcu, ‘Multiple Plans are Better than One: Diverse Stochastic Planning’, in *ICAPS*, pp. 140–148, (2021).
- [17] Michael Gimelfarb, Andre Barreto, Scott Sanner, and Chi-Guhn Lee, ‘Risk-aware transfer in reinforcement learning using successor features’, in *NeurIPS*, pp. 17298–17310, (2021).
- [18] Djordje Grbic and Sebastian Risi, ‘Safe Reinforcement Learning through Meta-learned Instincts’, in *ALIFE*, pp. 183–291, (2020).
- [19] Sehoon Ha, Peng Xu, Zhenyu Tan, Sergey Levine, and Jie Tan, ‘Learning to Walk in the Real World with Minimal Human Effort’, in *CoRL*, pp. 1110–1120, (2020).
- [20] Tuomas Haarnoja, Haoran Tang, Pieter Abbeel, and Sergey Levine, ‘Reinforcement Learning with Deep Energy-Based Policies’, in *ICML*, pp. 1352–1361, (2017).
- [21] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine, ‘Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor’, in *ICML*, pp. 1861–1870, (2018).
- [22] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft Actor-Critic Algorithms and Applications, 2018. [arXiv:1812.05905](https://arxiv.org/abs/1812.05905).
- [23] Elad Hazan, Sham Kakade, Karan Singh, and Abby Van Soest, ‘Provably Efficient Maximum Entropy Exploration’, in *ICML*, pp. 2681–2691, (2019).
- [24] Maximilian Igl, Gregory Farquhar, Jelena Luketina, Wendelin Boehmer, and Shimon Whiteson, ‘Transient Non-stationarity and Generalisation in Deep Reinforcement Learning’, in *ICLR*, (2021).
- [25] Riashat Islam, Zafarali Ahmed, and Doina Precup. Marginalized State Distribution Entropy Regularization in Policy Optimization, 2019. [arXiv:1912.05128](https://arxiv.org/abs/1912.05128).
- [26] Nils Jansen, Bettina Könighofer, Sebastian Junges, Alex Serban, and Roderick Bloem, ‘Safe Reinforcement Learning Using Probabilistic Shields (Invited Paper)’, in *CONCUR*, pp. 1–16, (2020).
- [27] Thommen George Karimnapan, Santu Rana, Sunil Gupta, Truyen Tran, and Svetha Venkatesh, ‘Learning transferable domain priors for safe exploration in reinforcement learning’, in *IJCNN*, pp. 1–10. IEEE, (2020).
- [28] Saurabh Kumar, Aviral Kumar, Sergey Levine, and Chelsea Finn, ‘One Solution is Not All You Need: Few-Shot Extrapolation via Structured MaxEnt RL’, in *NeurIPS*, p. 8198–8210, (2020).
- [29] Lisa Lee, Benjamin Eysenbach, Emilio Parisotto, Eric Xing, Sergey Levine, and Ruslan Salakhutdinov. Efficient Exploration via State Marginal Matching, 2019. [arXiv:1906.05274](https://arxiv.org/abs/1906.05274).
- [30] Thomas Lew, Apoorva Sharma, James Harrison, Andrew Bylard, and Marco Pavone, ‘Safe Active Dynamics Learning and Control: A Sequential Exploration–Exploitation Framework’, *IEEE Transactions on Robotics*, **38**(5), 2888–2907, (2022).
- [31] Lihong Li, Thomas J Walsh, and Michael L Littman, ‘Towards a Unified Theory of State Abstraction for MDPs’, in *AI&M*, pp. 1–10, (2006).
- [32] Michael Luo, Ashwin Balakrishna, Brijen Thananjeyan, Suraj Nair, Julian Ibarz, Jie Tan, Chelsea Finn, Ion Stoica, and Ken Goldberg. MESA: Offline Meta-RL for Safe Adaptation and Fault Tolerance, 2021. [arXiv:2112.03575](https://arxiv.org/abs/2112.03575).
- [33] Luca Marzari, Davide Corsi, Enrico Marchesini, and Alessandro Farinelli, ‘Curriculum learning for safe mapless navigation’, in *SAC*, pp. 766–769, (2022).
- [34] Sobhan Miryoosefi and Chi Jin, ‘A simple reward-free approach to constrained reinforcement learning’, in *ICML*, pp. 15666–15698, (2022).
- [35] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fiedelnd, 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, (2015).
- [36] Zhenghao Peng, Quanyi Li, Chunxiao Liu, and Bolei Zhou, ‘Safe driving via expert guided policy optimization’, in *CoRL*, pp. 1554–1563, (2022).
- [37] Zengyi Qin, Yuxiao Chen, and Chuchu Fan, ‘Density Constrained Reinforcement Learning’, in *ICML*, pp. 8682–8692, (2021).
- [38] Alex Ray, Joshua Achiam, and Dario Amodei. Benchmarking Safe Exploration in Deep Reinforcement Learning, 2019. <https://cdn.openai.com/safexp-short.pdf>.
- [39] Yagiz Savas, Melkior Ornik, Murat Cubuktepe, and Ufuk Topcu, ‘Entropy Maximization for Constrained Markov Decision Processes’, in *Allerton*, pp. 911–918, (2018).
- [40] Erik Schuitema, Martijn Wisse, Thijs Ramakers, and Pieter Jonker, ‘The design of LEO: A 2D bipedal walking robot for online autonomous Reinforcement Learning’, in *IROS*, pp. 3238–3243, (2010).
- [41] Younggyo Seo, Lili Chen, Jinwoo Shin, Honglak Lee, Pieter Abbeel, and Kimin Lee, ‘State Entropy Maximization with Random Encoders for Efficient Exploration’, in *ICML*, pp. 9443–9454, (2021).
- [42] Thiago D. Simão, Nils Jansen, and Matthijs T. J. Spaan, ‘AlwaysSafe: Reinforcement Learning Without Safety Constraint Violations During Training’, in *AAMAS*, p. 1226–1235, (2021).
- [43] Krishnan Srinivasan, Benjamin Eysenbach, Sehoon Ha, Jie Tan, and Chelsea Finn. Learning to be Safe: Deep RL with a Safety Critic, 2020. [arXiv:2010.14603](https://arxiv.org/abs/2010.14603).
- [44] Yanan Sui, Alkis Gotovos, Joel Burdick, and Andreas Krause, ‘Safe Exploration for Optimization with Gaussian Processes’, in *ICML*, pp. 997–1005, (2015).
- [45] Richard S. Sutton and Andrew G. Barto, *Reinforcement Learning: An Introduction*, volume 2, MIT press, 2018.
- [46] Richard S. Sutton, A. Rupam Mahmood, and Martha White, ‘An Emphatic Approach to the Problem of Off-policy Temporal-Difference Learning’, *JMLR*, **17**, 2603–2631, (2016).
- [47] Oleg Svidchenko and Aleksei Shpilman. Maximum Entropy Model-based Reinforcement Learning, 2021. [arXiv:2112.01195](https://arxiv.org/abs/2112.01195).
- [48] Matthew E. Taylor and Peter Stone, ‘Transfer Learning for Reinforcement Learning Domains: A Survey’, *JMLR*, **10**, 1633–1685, (2009).
- [49] Chen Tessler, Daniel J. Mankowitz, and Shie Mannor, ‘Reward Constrained Policy Optimization’, in *ICLR*, (2019).
- [50] Brijen Thananjeyan, Ashwin Balakrishna, Suraj Nair, Michael Luo, Krishnan Srinivasan, Minh Hwang, Joseph E Gonzalez, Julian Ibarz, Chelsea Finn, and Ken Goldberg, ‘Recovery RL: Safe Reinforcement Learning With Learned Recovery Zones’, *IEEE Robotics and Automation Letters*, **6**(3), 4915–4922, (2021).
- [51] Matteo Turchetta, Andrey Kolobov, Shital Shah, Andreas Krause, and Alekh Agarwal, ‘Safe Reinforcement Learning via Curriculum Induction’, in *NeurIPS*, pp. 12151–12162, (2020).
- [52] Giulia Vezzani, Abhishek Gupta, Lorenzo Natale, and Pieter Abbeel. Learning latent state representation for speeding up exploration, 2019. [arXiv:1905.12621](https://arxiv.org/abs/1905.12621).
- [53] Zhaoxing Xie, Patrick Clary, Jeremy Dao, Pedro Morais, Jonathan W. Hurst, and Michiel van de Panne, ‘Learning Locomotion Skills for Cassie: Iterative Design and Sim-to-Real’, in *CoRL*, pp. 317–329, (2019).
- [54] Qisong Yang, Thiago D. Simão, Nils Jansen, Simon H. Tindemans, and Matthijs T. J. Spaan. Reinforcement Learning by Guided Safe Exploration, 2023. [arXiv:2307.14316](https://arxiv.org/abs/2307.14316).
- [55] Qisong Yang, Thiago D. Simão, Simon H. Tindemans, and Matthijs T. J. Spaan, ‘WCSAC: Worst-Case Soft Actor Critic for Safety-Constrained Reinforcement Learning’, in *AAAI*, pp. 10639–10646, (2021).
- [56] Qisong Yang and Matthijs T. J. Spaan, ‘CEM: Constrained Entropy Maximization for Task-Agnostic Safe Exploration’, in *AAAI*, pp. 10798–10806, (2023).
- [57] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J. Ramadge, ‘Projection-Based Constrained Policy Optimization’, in *ICLR*, (2020).
- [58] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J. Ramadge, ‘Accelerating safe reinforcement learning with constraint-mismatched baseline policies’, in *ICML*, pp. 11795–11807, (2021).
- [59] Tom Zahavy, Brendan O’Donoghue, André Barreto, Sebastian Flennerhag, Volodymyr Mnih, and Satinder Singh. Discovering Diverse Nearly Optimal Policies with Successor Features, 2021. [arXiv:2106.00669](https://arxiv.org/abs/2106.00669).
- [60] Moritz A. Zanger, Karam Daaboul, and J. Marius Zöllner, ‘Safe Continuous Control with Constrained Model-Based Policy Optimization’, in *IROS*, pp. 3512–3519, (2021).
- [61] Jesse Zhang, Brian Cheung, Chelsea Finn, Sergey Levine, and Dinesh Jayaraman, ‘Cautious Adaptation For Reinforcement Learning in Safety-Critical Settings’, in *ICML*, p. 11055–11065, (2020).
- [62] Zhuangdi Zhu, Kaixiang Lin, Anil K. Jain, and Jiayu Zhou, ‘Transfer learning in deep reinforcement learning: A survey’, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–20, (2023).
- [63] Brian D Ziebart, *Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy*, Ph.D. dissertation, CMU, 2010.