
Direct Behavior Specification via Constrained Reinforcement Learning

Julien Roy^{‡12} Roger Girgis^{‡12} Joshua Romoff³ Pierre-Luc Bacon¹⁴⁵ Christopher Pal¹²⁴⁶⁷

Abstract

The standard formulation of Reinforcement Learning lacks a practical way of specifying what are admissible and forbidden behaviors. Most often, practitioners go about the task of behavior specification by manually engineering the reward function, a counter-intuitive process that requires several iterations and is prone to reward hacking by the agent. In this work, we argue that constrained RL, which has almost exclusively been used for safe RL, also has the potential to significantly reduce the amount of work spent for reward specification in applied RL projects. To this end, we propose to specify behavioral preferences in the CMDP framework and to use Lagrangian methods to automatically weigh each of these behavioral constraints. Specifically, we investigate how CMDPs can be adapted to solve goal-based tasks while adhering to several constraints simultaneously. We evaluate this framework on a set of continuous control tasks relevant to the application of Reinforcement Learning for NPC design in video games.

1. Introduction

Reinforcement Learning (RL) has shown rapid progress and lead to many successful applications over the past few years (Mnih et al., 2013; Silver et al., 2017; Andrychowicz et al., 2020). The RL framework is predicated on the simple idea that all tasks could be defined as a single scalar function to maximise, an idea generally referred to as the reward hypothesis (Sutton & Barto, 2018; Silver et al., 2021;

Abel et al., 2021). This idea has proven very useful to develop the theory and concentrate research on a single theoretical framework. However, it can be significantly limiting when translating a real-life problem into an RL problem, since the question of where the reward function comes from is completely ignored (Singh et al., 2009). In practice, human-designed reward functions often lead to unforeseen behaviors and represent a serious obstacle to the reliable application of RL in the industry (Amodei et al., 2016).

Concretely, for an engineer working on applying RL methods to an industrial problem, the task of reward specification implies to: (1) characterise the desired behavior that the system should exhibit, (2) write in a computer program a reward function for which the optimal policy corresponds to that desired behavior, (3) train an RL agent on that task using one of the methods available in the literature and (4) evaluate whether the agent exhibits the expected behavior. Multiple design iterations of that reward function are generally required, each time accompanied by costly trainings of the policy (Hadfield-Menell et al., 2017; Dulac-Arnold et al., 2019). This inefficient design loop is exacerbated by the fact that current Deep RL algorithms cannot be guaranteed to find the optimal policy (Sutton & Barto, 2018), meaning that the reward function could be correctly specified but still fail to lead to the desired behavior. The design problem thus becomes “What reward function would lead SAC (Haarnoja et al., 2018) or PPO (Schulman et al., 2017) to give me a policy that I find satisfactory?”, a difficult puzzle that every RL practitioner has had to deal with.

Most published work on Reinforcement Learning focuses on point (3) i.e. improving the reliability and efficiency with which these algorithms can yield a near-optimal policy for a *given* reward function. This line of work is crucial to allow RL to tackle difficult problems. However, as agents become more and more capable of solving the tasks we present them with, our ability to (2) correctly specify these reward functions will only become more critical (Dewey, 2014).

Constrained Markov Decision Processes (Altman, 1999) offer an alternative framework for sequential decision making. The agent still seeks to maximise a single reward function,

[‡]Work conducted while interning at Ubisoft La Forge.

¹Institut d’intelligence artificielle du Québec (Mila). ²École Polytechnique de Montréal. ³Ubisoft La Forge. ⁴Université de Montréal. ⁵Facebook CIFAR AI Chair. ⁶ServiceNow. ⁷Canada CIFAR AI Chair.. Correspondence to: Julien Roy <julien.roy@mila.quebec>.




Figure 1. Depictions of our setup to evaluate direct behavior specification using constrained RL; Arena environment (left); OpenWorld environment (right). For videos see: <https://sites.google.com/view/behaviorspecificationviacr/home>.

but must do so while respecting a set of constraints defined by additional cost functions. While it is generally recognised that this formulation has the potential to allow for an easier task definition from the end user (Ray et al., 2019), most work on CMDPs focuses on the safety aspect of this framework i.e. that the constraint-satisfying behavior be maintained throughout the entire exploration process (Achiam et al., 2017; Zhang et al., 2020; Turchetta et al., 2020; Marchesini et al., 2022). In this paper we specifically focus on the benefits of CMDPs relating to behavior specification. We make the following contributions: (1) we show experimentally that reward engineering poorly scales with the complexity of the target behavior, (2) we propose a solution where a designer can directly specify the desired frequency of occurrence of some events, (3) we develop a novel algorithmic approach capable of jointly satisfying many more constraints and (4) we evaluate this framework on a set of constrained tasks illustrative of the development cycle required for deploying RL in video games.

2. The problem with reward engineering

In this section, we motivate the impracticality of using reward engineering to shape behavior. We consider a navigation task in which the agent has to reach a goal location while being subject to additional behavioral constraints. These constraints are (1) looking at a visible marker 90% of the time, (2) avoiding forbidden terrain 99% of the time and (3) avoiding to run out of energy also 99% of the time. The environment is depicted in Figure 1 (left) and the details are presented in Appendix B. The reward function for this task is of the form:

$$R'(s, a) = R(s, a) - \mathbf{1} * w_{\text{not-looking}} - \mathbf{1} * w_{\text{in-lava}} - \mathbf{1} * w_{\text{no-energy}} \quad (1)$$

where $R(s, a)$ gives a small shaping reward for progressing towards the goal and a terminal reward for reaching the goal, and the $\mathbf{1}$'s are indicator functions which are only active if their corresponding behavior is exhibited.

The main challenge for an RL practitioner is to determine the correct values of the weights $w_{\text{not-looking}}$, $w_{\text{in-lava}}$ and $w_{\text{no-energy}}$ such that the agent maximises its performance on the main task while respecting the behavioral requirements, a problem often referred to as reward engineering. Setting these weights too low results in an agent that ignores these requirements while setting them too high distracts the agent from completing the main task. In general, knowing how to scale these components relatively to one another is not intuitive and is often performed by trial and error across the space of reward coefficients w_k . To illustrate where the desired solutions can be found for this particular problem, we perform 3 grid searches on 7 different values for each of these weights, ranging from 0.1 to 10 times the scale of the main reward function, for the cases of 1, 2 and 3 behavioral constraints. The searches thus respectively must go through 7, 49 and 343 training runs. Figure 2 (and Figure 7 in Appendix D) show the results of these experiments. We can see that a smaller and smaller proportion of these trials lead to successful policies as the number of behavioral constraints grows. For an engineer searching to find the right trade-off, they find themselves cornered between two undesirable solutions: an ad-hoc manual approach guided by intuition or to run a computationally demanding grid-search. While expert knowledge or other search strategies can partially alleviate this burden, the approach of reward engineering clearly does not scale as the control problem grows in complexity.

It is important to note that whether or not it is the case that all tasks can in principle be defined as a single scalar function to maximise i.e. the reward hypothesis (Sutton & Barto, 2018), this notion should not be seen as a restrictive design principle when translating a real-life problem into an RL problem. That is because it does not guarantee that this reward function admits a simple form. Rich and multi-faceted behaviors may only be specifiable through a complex reward function (Abel et al., 2021) beyond the reach of human intuition. In the next sections we present a practical framework in which CMDPs can be used to provide a more intuitive and human-centric interface for behavioral specification.

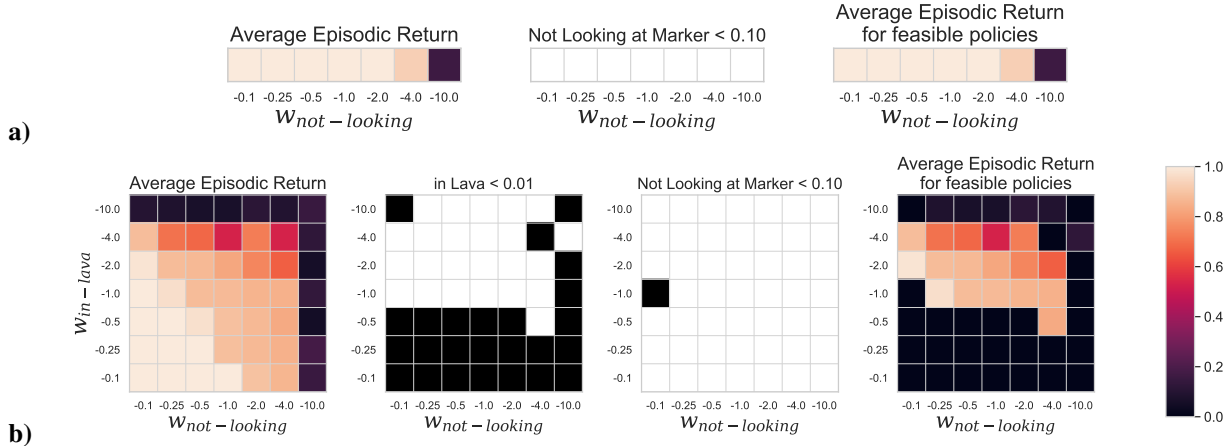


Figure 2. Enforcing behavioral constraints using reward engineering. Each grid represents a different metric. Within each grid, each square represents the final performance (according to that metric) of an agent trained for 3M steps using the reward function in Equation 1 parameterised as given by the grid coordinates. Performance is obtained by evaluating the agent on 1000 episodes. The leftmost column indicates the episodic return of the trained policies, the middle columns indicates whether or not the agent respects the behavioral constraint(s) and the rightmost column indicates the average return for these feasible policies only. **a)** The “looking-at marker” behavior does not affect too much the main task and, consequently, all chosen weights allow to satisfy the constraint (looking at marker 90% of the time) and many of them also lead to good performance on the main navigation task ($-0.1 \geq w \geq -2$). **b)** When also enforcing the “Not in Lava” behavior, which is much more in the way of the main task, most of the resulting policies do not respect the constraint or perform poorly on the navigation task, highlighting the difficulty of choosing the correct penalty weights ahead of time. On 49 experiments, only two yielded good performing feasible policies: $(-0.10, -2.0)$ and $(-0.25, -1.0)$. On the largest search with 3 behavioral constraints, none of the 343 experiments found a good performing feasible policy (see Figure 7 in Appendix D).

3. Background

Markov Decision Processes (MDPs) (Sutton & Barto, 2018) are formally defined through the following four components: $(\mathcal{S}, \mathcal{A}, P, R)$. At timestep t , an agent finds itself in state $s_t \in \mathcal{S}$ and picks an action $a_t \in \mathcal{A}(s_t)$. The transition probability function P encodes the conditional probability $P(s_{t+1}|s_t, a_t)$ of transitioning to the next state s_{t+1} . Upon entering the next state, an immediate reward is generated through a reward function $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. In this paper, we restrict our attention to stationary randomized policies of the form $\pi(a|s)$ – which are sufficient for optimality in both MDPs and CMDPs (Altman, 1999). The interaction of a policy within an MDP gives rise to trajectories $(s_0, a_0, r_0, \dots, s_T, a_T, R_T)$ over which can be computed the sum of rewards which we call the *return*. Under the Markov assumption, the probability distribution over trajectories is of the form:

$$p_\pi(\tau) := P_0(s_0) \prod_{t=0}^T P(s_{t+1}|s_t, a_t) \pi(a_t|s_t) \quad (2)$$

where P_0 is some initial state distribution. Furthermore, any such policy induces a marginal distribution over state-action pairs referred to as the *visitation distribution* or *state-action occupation measure*:

$$x_\pi(s, a) := \frac{1}{Z(\gamma, T)} \sum_{t=0}^T \gamma^t p_{\pi, t}(S_t = s, A_t = a) \quad , \quad (3)$$

where $Z(\gamma, T) = \sum_{t=0}^T \gamma^t$ is a normalising constant.

In this paper, it is useful to extend the notion of return to any function $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ over states and actions other than the reward function of the MDP itself. The expected discounted sum of f then becomes:

$$J_f(\pi) := \mathbb{E}_{\tau \sim p_\pi} \left[\sum_{t=0}^T \gamma^t f(s_t, a_t) \right] \quad (4)$$

where $\gamma \in [0, 1]$ is a discount factor. While this idea is the basis for much of the work on General Value Functions (GVFs) (White, 2015) for predictive state representation (Sutton et al., 2011), our focus here is on problem of behavior specification and not that of prediction.

Finally, in the MDP setting, a policy is said to be optimal under the expected discounted return criterion if $\pi^* = \arg \max_{\pi \in \Pi} J_R(\pi)$, where Π is the set of possible policies.

Constrained MDPs (CMDPs) (Altman, 1999) is a framework that extends the notion of optimality in MDPs to a scenario where multiple cost constraints need to be satisfied in addition to the main objective. We write $C_k : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ to denote such a cost function whose expectation must remain bounded below a specified threshold $d_k \in \mathbb{R}$. The set of

feasible policies is then:

$$\Pi_C = \{\pi \in \Pi : J_{C_k}(\pi) \leq d_k, k = 1, \dots, K\}. \quad (5)$$

Optimal policies in the CMDP framework are those of maximal expected return among the set of feasible policies:

$$\pi^* = \arg \max_{\pi \in \Pi} J_R(\pi), \text{ s.t. } J_{C_k}(\pi) \leq d_k, k = 1, \dots, K \quad (6)$$

While it is sufficient to consider the space of stationary deterministic policies in searching for optimal policies in the MDP setting, this is no longer true in general with CMDPs (Altman, 1999) and we must consider the larger space of stationary randomized policies.

Lagrangian methods for CMPDs. Several recent works have found that the class of Lagrangian methods for solving CMDPs is capable of finding good feasible solutions at convergence (Achiam et al., 2017; Ray et al., 2019; Stooke et al., 2020; Zhang et al., 2020). The basis for this line of work stems from the saddle-point characterisation of the optimal solutions in nonlinear programs with inequality constraints (Uzawa et al., 1958; Polyak, 1970; Korpelevich, 1976). Intuitively, these methods combine the main objective J_R and the constraints into a single function \mathcal{L} called the Lagrangian. The relative weight of the constraints are determined by additional variables λ_k called the Lagrange multipliers. Applied in our context, this idea leads to the following min-max formulation:

$$\begin{aligned} & \max_{\pi} \min_{\lambda \geq 0} \mathcal{L}(\pi, \lambda) \\ \mathcal{L}(\pi, \lambda) &= J_R(\pi) - \sum_{k=1}^K \lambda_k (J_{C_k}(\pi) - d_k) \end{aligned} \quad (7)$$

where we denoted $\lambda := \{\lambda_k\}_{k=1}^K$ for conciseness. Uzawa et al. (1958) proposed to find a solution to this problem iteratively by taking gradient ascent steps of the Lagrangian \mathcal{L} in the variable π and descent ones in λ . This is also the same gradient ascent-descent (Lin et al., 2020) procedure underpinning many learning algorithms for Generative Adversarial Networks (Goodfellow et al., 2014).

The maximization of the Lagrangian over the policy variables can be carried out by applying any existing unconstrained policy optimization methods to the new reward function $L : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ where:

$$L(s, a) = R(s, a) - \sum_{k=1}^K \lambda_k C_k(s, a). \quad (8)$$

For the gradient w.r.t. the Lagrange multipliers λ , the term depending on π cancels out and we are left with $\nabla_{\lambda_k} \mathcal{L}(\pi, \lambda) = -(J_{C_k}(\pi) - d_k)$. The update is followed by a projection onto $\lambda_k \geq 0$ using the max-clipping operator.

If the constraint is violated ($J_{C_k}(\pi) > d_k$), taking a step in the opposite direction of the gradient will increase the corresponding multiplier λ_k , thus increasing the relative importance of this constraint in $J_L(\pi)$. Inversely, if the constraint is respected ($J_{C_k}(\pi) < d_k$), the update will decrease λ_k , allowing the optimisation process to focus on the other constraints and the main reward function R .

4. Proposed Framework

In Reinforcement Learning, the reward function is often assumed to be provided apriori. For example, in most RL benchmarking environments this is indeed the case and researchers can focus on improving current algorithms at finding better policies, faster and more reliably. In industrial applications however, several desiderata are often required for the agent’s behavior, and balancing these components into a single reward function is highly non-trivial. In the next sections, we describe a framework in which CMDPs can be used for efficient behavior specification.

4.1. Indicator cost functions

The difficulty of specifying the desired behavior of an agent using a single reward function stems from the need to tune the relative scale of each reward component. Moreover, finding the most appropriate ratio becomes more challenging as the number of reward components increases (see Section 2). While the prioritisation and saturation characteristics of CMDPs help factoring the behavioral specification problem (Ray et al., 2019), there remains important design challenges. First, the CMDP framework allows for arbitrary forms of cost functions, again potentially leading to unforeseen behaviors. Second, specifying the appropriate thresholds d_k can be difficult to do solely based on intuition. For example, in the mujoco experiments performed by Zhang et al. (2020), the authors had to run an unconstrained version of PPO (Schulman et al., 2017) to first estimate the typical range of values for the cost infringements and then run their constrained solver over the appropriately chosen thresholds.

We show here that this separate phase of threshold estimation can be avoided completely if we consider a subclass of CMDPs that allows for a more intuitive connection between the chosen cost functions C_k and their expected returns J_{C_k} . More specifically, we restrict our attention to CMDPs where the cost functions are defined as indicators of the form:

$$C_k(s, a) = I(\text{behavior } k \text{ is met in } (s, a)) \quad (9)$$

which simply expresses whether an agent showcases some particular behavior k when selecting action a in state s . An interesting property of this design choice is that, by rewriting the expected discounted sum of these indicator cost functions as an expectation over the visitation distribution

of the agent, we can interpret this quantity as a re-scaled probability that the agent exhibits behavior k at any given time during its interactions with the environment:

$$J_{C_k}(\pi) = \mathbb{E}_{\tau \sim p_\pi} \left[\sum_{t=0}^T \gamma^t C_k(s_t, a_t) \right] \quad (10)$$

$$= Z(\gamma, T) \mathbb{E}_{(s,a) \sim x_\pi(s,a)} [C_k(s, a)] \quad (11)$$

$$= Z(\gamma, T) \mathbb{E}_{(s,a) \sim x_\pi(s,a)} [I(\text{behavior } k \text{ met in } (s, a))] \quad (12)$$

$$= Z(\gamma, T) Pr(\text{behavior } k \text{ met in } (s, a)), (s, a) \sim x_\pi \quad (13)$$

Dividing each side of $J_{C_k}(\pi) \leq d_k$ by $Z(\gamma, T)$, we are left with \tilde{d}_k , a normalized constraint threshold for the constraint k which represents the desired rate of encountering the behavior designated by the indicator cost function C_k . In practice, we simply compute the average cost function across the batch to give equal weighting to all state-action pairs regardless of their position t in the trajectory:

$$\tilde{J}_{C_k}(\pi) := \frac{1}{N} \sum_{i=1}^N C_k(s_i, a_i) \quad (14)$$

where i is the sample index from the batch. We also train the corresponding critic $Q^{(k)}$ using a discount factor $\gamma_k < 1$ for numerical stability.

While the class of cost functions defined in Equation 9 still allows for modelling a large variety of behavioral preferences, it has the benefit of informing the user on the range of appropriate thresholds – a probability $\tilde{d}_k \in [0, 1]$ – and the semantics is clear regarding its effect on the agent’s behavior (assuming that the constraint is binding and that a feasible policy is found). This effectively allows for minimal to no tuning behavior specification (or “zero-shot” behavior specification).

Finally, indicator cost functions also have the practical advantage of allowing to capture both desired and undesired behaviors without affecting the termination tendencies of the agent. Indeed, when using an arbitrary cost function, it could be tempting to simply flip its sign to enforce the opposite behavior. However, as noted in previous work (Kostrikov et al., 2018), the choice of whether to enforce behaviors through bonuses or penalties should instead be thought about with the termination conditions in mind. A positive bonus could cause the agent to delay termination in order to accumulate more bonuses while negative penalties could shape the agent behavior such that it seeks to trigger the termination of the episode as soon as possible. Indicator cost functions are thus very handy in that they offer a straightforward way to enforce the opposite behavior by simply inverting the indicator function $Not(I(s, a)) = 1 - I(s, a)$ without affecting the sign of the constraint (penalties v.s. bonuses).

4.2. Multiplier normalisation

When the constraint k is violated, the multiplier λ_k associated with that constraint increases to put more emphasis on that aspect of the overall behavior. While it is essential for the multipliers to be able to grow sufficiently compared to the main objective, a constraint that enforces a behavior which is long to discover can end up reaching very large multiplier values. It then leads to very large policy updates and destabilizes the learning dynamics.

To maintain the ability of one constraint to dominate the policy updates when necessary while keeping the scale of the updates bounded, we propose to normalize the multipliers. This can be readily implemented by using a softmax layer:

$$\lambda_k = \frac{\exp(z_k)}{\exp(a_0) + \sum_{k'=1}^K \exp(z_{k'})}, \quad k = 1, \dots, K \quad (15)$$

where z_k are the base parameters for each one of the multipliers and a_0 is a dummy variable used to obtain a normalized weight $\lambda_0 := 1 - \sum_{k=1}^K \lambda_k$ for the main objective $J_R(\pi)$. The corresponding min-max problem becomes:

$$\begin{aligned} & \max_{\pi} \min_{z_{1:K} \geq 0} \mathcal{L}(\pi, \lambda) \\ & \mathcal{L}(\pi, \lambda) = \lambda_0 J_R(\pi) - \sum_{k=1}^K \lambda_k (J_{C_k}(\pi) - d_k) \end{aligned} \quad (16)$$

4.3. Bootstrap Constraint

In the presence of many constraints, one difficulty that emerges with the above multiplier normalisation is that the coefficient of the Lagrangian function that weighs the main objective is constrained to be $\lambda_0 = 1 - \sum_{k=1}^K \lambda_k$, which leaves very little to no traction to improve on the main task while the process is looking for a feasible policy. Furthermore, as more constraints are added, the optimisation path becomes discontinuous between regions of feasible policies, preventing learning progress on the main task objective.

A possible solution is to grant the main objective the same powers as the behavioral constraints that we are trying to enforce. This can be done by defining an additional function $S_{K+1}(s, a)$ which captures some measure of success on the main task. Indeed, many RL tasks are defined in terms of such sparse, clearly defined success conditions, and then often only augmented with a dense reward function to guide the agent toward these conditions (Ng et al., 1999). A so-called *success constraint* of the form $J_{S_{K+1}}(\pi) \geq \tilde{d}_{K+1}$ can thus be implemented using an indicator cost function as presented above and added to the existing constraint set $\{J_{C_k}(\pi) \leq \tilde{d}_k\}_{k=1}^K$. While the use of a success constraint alone can be expected to aid learning of the main task, it is only a sparse signal and could be very difficult to discover if the main task is itself challenging. Since the success

function S_{K+1} is meant to be highly correlated with the reward function R , by going a step further and using the success constraint multiplier λ_{K+1} in place of the reward multiplier λ_0 , we can take full advantage of the density of the main reward function when enforcing that constraint. However, to maintain a true maximisation objective over the main reward function, we still need to keep using λ_0 when other constraints are satisfied, so that the most progress can be made on $J_R(\pi)$. We thus take the largest of these two coefficients for weighing the main objective $\tilde{\lambda}_0 := \max(\lambda_0, \lambda_{K+1})$ and replace λ_0 with $\tilde{\lambda}_0$ in Equation 16. Here we say that constraint $K + 1$ is used as a *bootstrap constraint*.

Our method of encoding a success criterion in the constraint set can be seen as a way of relaxing the behavioral constraints during the optimisation process without affecting the convergence requirements. For example, in previous work, Calian et al. (2020) tune the learning rate of the Lagrange multipliers to automatically turn some constraints into soft-constraints when the agent is not able to satisfy them after a given period of time. Instead, the bootstrap constraint allows to start making some progress on the main task without turning our hard constraints into soft constraints.

5. Related Work

Constrained Reinforcement Learning. CMDPs (Altman, 1999) have been the focus of several previous work in Reinforcement Learning. Lagrangian methods (Borkar, 2005; Tessler et al., 2018; Stooke et al., 2020) combine the constraints and the main objective into a single function and seek to find a saddle point corresponding to feasible solutions to the maximisation problem. Projection-based methods (Achiam et al., 2017; Chow et al., 2019; Yang et al., 2020; Zhang et al., 2020) instead use a projection step to try to map the policy back into a feasible region after the reward maximisation step. While most of these works focus on the single-constraint case (Zhang et al., 2020; Dalal et al., 2018; Calian et al., 2020; Stooke et al., 2020) and seek to minimize the total regret over the cost functions throughout training (Ray et al., 2019), we focus on the potential of CMDPs for precise and intuitive behavior specification and work on satisfying many constraints simultaneously.

Reward Specification. Imitation Learning (Zheng et al., 2021) is largely motivated by the difficulty of designing reward functions and instead seeks to use expert data to define the task. Other approaches introduce a human in the loop to either guide the agent towards the desired behavior (Christiano et al., 2017) or to prevent it from making catastrophic errors while exploring the environment (Saunders et al., 2017). While our approach of using CMDPs for behavior specification also seeks to make better use of

human knowledge, we focus on the idea of providing this knowledge by simply specifying thresholds and indicator functions rather than requiring expert demonstrations or constant human feedback. Another line of work studies whether natural language can be used as a more convenient interface to specify the agent’s desired behavior (Goyal et al., 2019; MacGlashan et al., 2015). While this idea presents interesting perspectives, natural language is inherently ambiguous and prone to reward hacking by the agent. Moreover such approaches generally come with the added complexity of having to learn a language-to-reward model. Finally, others seek to solve reward mis-specification through Inverse Reward Design (Hadfield-Menell et al., 2017; Mindermann et al., 2018; Ratner et al., 2018) which treats the provided reward function as a single observation of the true intent of the designer and seeks to learn a probabilistic model that explains it. While this approach is interesting for adapting to environmental changes, we focus on behavior specification in fixed-distribution environments.

RL in video games. Video games have been used as a benchmark for Deep RL for several years (Shao et al., 2019; Berner et al., 2019; Vinyals et al., 2019). However, examples of RL being used in a video game production are limited due to a variety of factors which include the difficulty of shaping behavior, interpretability, and compute limitations at runtime (Jacob et al., 2020; Alonso et al., 2020). Still, there has been a recent push in the video game industry to build NPCs (Non Player Characters) using RL, for applications including navigation (Alonso et al., 2020; Devlin et al., 2021), automated testing (Bergdahl et al., 2020; Gordillo et al., 2021), play-style modeling (de Woillemont et al., 2021) and content generation (Gisslén et al., 2021).

6. Experiments

To evaluate the proposed framework, we train SAC agents (Haarnoja et al., 2018) to solve navigation tasks with up to 5 constraints imposed on their behavior. Many of these constraints interact with the main task and with one another which significantly restricts the space of admissible policies. We conduct most of our experiments in the Arena environment (see Figure 1, left)¹ where we seek to verify the capacity of the proposed framework to allow for easy specification of the desired behavior and the ability of the algorithm to deal with a large number of constraints simultaneously. We also perform an experiment in the OpenWorld environment (see Figure 1, right), a much larger and richer map generated using the GameRLand map generator (Beeching et al., 2021), where we seek to verify the scalability of that approach and whether it fits the needs

¹The algorithm is presented in Appendix A. The code for the Arena environment experiments is available at: <https://github.com/ubisoft/DirectBehaviorSpecification>




Figure 3. The multiplier normalisation keeps the learning dynamics stable when discovering a constraint-satisfying behavior takes a large amount of time. To simulate such a case, an impossible constraint is set for 7.5M steps and then replaced by a feasible one for the last 2.5M steps. The method using unnormalized multipliers (red) keeps taking larger and larger steps in policy space leading to the divergence of its learning dynamics and complete collapse of its performance.




Figure 4. Each column presents the results for an experiment in which the agent is trained for 3M steps with a *single* constraint enforced on its behavior. Training is halted after every 20,000 environment steps and the agent is evaluated for 10 episodes. All curves show the average over 5 seeds and envelopes show the standard error around that mean. The top row shows the average return, the bottom row shows the average behavior rate on which the constraint is enforced. The black dotted lines mark the constraint thresholds.




Figure 5. Each row presents the results of an experiment in which an agent is trained for 10M steps. Training is halted after every 20,000 environment steps and the agent is evaluated for 10 episodes. All curves show the average over 5 seeds and envelopes show the standard error around that mean. (a) Unconstrained SAC agent; none of the behavioral preferences are enforced and consequently improvement on performance is very fast but none of the constraints are satisfied. (b) SAC-Lagrangian with the 5 behavioral constraints enforced. While each constraint was successfully dealt with when imposed one by one (see Figure 4), maximising the main objective when subject to all the constraints *simultaneously* proves to be much harder. The agent does not find a policy that improves on the main task while keeping the constraints in check. (c) By using an additional success constraint (that the agent should reach its goal in 99% of episodes), the agent can cut through infeasible policy space to start improving on the main task and optimise the remaining constraints later on. (d) By using the success constraint as a bootstrap constraint (bound to the main reward function) improvement on the main task is much faster as the agent benefits from the dense reward function to improve on the goal-reaching task.




Figure 6. A SAC-Lagrangian agent trained to solve the navigation problem in the OpenWorld environment while respecting four constraints and imposing the bootstrap constraint. Results suggest that our SAC-Lagrangian method using indicator cost functions, normalised multipliers and bootstrap constraint scales well to larger and more complex environments.

of agent behavior specification for the video game industry. See Appendices B and C for a detailed description of both experimental setups.

6.1. Experiment in the Arena environment

Multiplier Normalization Our first set of experiments showcases the effect of normalizing the Lagrange multipliers. For illustrative purposes, we designed a simple scenario where one of the constraints is not satisfied for a long period of time. Specifically, the agent is attempting to satisfy an impossible constraint of never touching the ground. Figure 3 (in red) shows that the multiplier on the unsatisfied constraint endlessly increases in magnitude, eventually harming the entire learning system; the loss on the critic diverges and the performance collapses. When using our normalization technique, Figure 3 (in blue) shows that the multiplier and critic losses remain bounded, avoiding such instabilities.

Single Constraint satisfaction We use our framework to encode the different behavioral preferences into indicator functions and specify their respective thresholds. Figure 4 shows that our SAC-Lagrangian with multiplier normalisation can solve the task while respecting the behavioral requirements when imposed with constraints individually. We note that the different constraints do not affect the main task to the same extent; while some still allow to quickly solve the navigation task, like the behavioral requirement to avoid jumping, others make the navigation task significantly more difficult to solve, like the requirement to avoid certain types of terrain (lava).

Multiple Constraints Satisfaction In Figure 5 we see that when imposed with all of the constraints simultaneously, the agent learns a feasible policy but fails at solving the main task entirely. The agent effectively settles on a trivial behavior in which it only focuses on satisfying the constraints, but from which it is very hard to move away without breaking the constraints. By introducing a success constraint, the agent at convergence is able to satisfy all of the constraints as well as succeeding in the navigation task. This additional incentive to traverse infeasible regions of the

policy space allows to find feasible but better performing solutions. Our best results are obtained when using the success constraint as a bootstrap constraint, effectively lending λ_{K+1} to the main reward while the agent is still looking for a feasible policy.

6.2. Experiment in the OpenWorld environment

In the OpenWorld environment, we seek to verify that the proposed solution scales well to more challenging and realistic tasks. Contrarily to the Arena environment, the OpenWorld contains uneven terrain, buildings, and interactable objects like jump-pads, which brings this evaluation setting much closer to an actual RL application in the video game industry. For this experiment, we trained a SAC-Lagrangian agent to solve the navigation problem with four constraints on its behavior: *On-Ground*, *Not-In-Lava*, *Looking-At-Marker* and *Above-Energy-Limit*. The SAC component uses the same hyper-parameters as in Alonso et al. (2020). The results are shown in Figure 6. While training the agent in this larger and more complex environment now requires up to 50M environment steps, the agent still succeeds at completing the task and respecting the constraints, favourably supporting the scalability of the proposed framework for direct behavior specification.

7. Discussion

Our work showed that CMDPs offer compelling properties when it comes to task specification in RL. More specifically, we developed an approach where the agent’s desired behavior is defined by the frequency of occurrence for given indicator events, which we view as constraints in a CMDP formulation. We showed through experiments that this methodology is preferable over the reward engineering alternative where we have to do an extensive hyper-parameter search over possible reward functions. We evaluated this framework on the many constraints case in two different environments. Our experiments showed that simultaneously satisfying a large number of constraints is difficult and can systematically prevent the agent from improving on the main task. We addressed this problem by normalizing the

constraint multipliers, which resulted in improved stability during training and proposed to bootstrap the learning on the main objective to avoid getting trapped by the composing constraint set. This bootstrap constraint becomes a way for practitioners to incorporate prior knowledge about the task and desired result – if the threshold is strenuous, a high success is prioritized – if the threshold is lax, it will simply be used to exit the initialisation point and the other constraints will quickly takeover. Our overall method is easy to implement over existing policy gradient code bases and can scale across domains easily.

We hope that these insights can contribute to a wider use of Constrained RL methods in industrial application projects, and that such adoption can be mutually beneficial to the industrial and research RL communities.

Acknowledgments

We wish to thank Philippe Marcotte, Maxim Peter, Rémi Labory, Pierre Le Pelletier De Woillemont, Julien Varnier, Pierre Falticska, Gabriel Robert, Vincent Martineau, Olivier Pomarez, Tristan Deleu and Paul Barde as well as the entire research team at Ubisoft La Forge for providing technical support and insightful comments on this work. We also acknowledge funding in support of this work from Fonds de Recherche Nature et Technologies (FRQNT), Mitacs Accelerate Program, Institut de valorisation des données (IVADO) and Ubisoft La Forge.

References

Abel, D., Dabney, W., Harutyunyan, A., Ho, M. K., Littman, M., Precup, D., and Singh, S. On the expressivity of markov reward. *Advances in Neural Information Processing Systems*, 34, 2021.

Achiam, J., Held, D., Tamar, A., and Abbeel, P. Constrained policy optimization. In *International Conference on Machine Learning*, pp. 22–31. PMLR, 2017.

Alonso, E., Peter, M., Goumard, D., and Romoff, J. Deep reinforcement learning for navigation in aaa video games. *arXiv preprint arXiv:2011.04764*, 2020.

Altman, E. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., and Mané, D. Concrete problems in ai safety. *arXiv preprint arXiv:1606.06565*, 2016.

Andrychowicz, O. M., Baker, B., Chociej, M., Jozefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., et al. Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20, 2020.

Beeching, E., Peter, M., Marcotte, P., Debangoye, J., Simonin, O., Romoff, J., and Wolf, C. Graph augmented deep reinforcement learning in the gamerland3d environment. *arXiv preprint arXiv:2112.11731*, 2021.

Bergdahl, J., Gordillo, C., Tollmar, K., and Gisslén, L. Augmenting automated game testing with deep reinforcement learning. In *2020 IEEE Conference on Games (CoG)*, pp. 600–603. IEEE, 2020.

Berner, C., Brockman, G., Chan, B., Cheung, V., Debiak, P., Dennison, C., Farhi, D., Fischer, Q., Hashme, S., Hesse, C., et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

Borkar, V. S. An actor-critic algorithm for constrained markov decision processes. *Systems & control letters*, 54(3):207–213, 2005.

Calian, D. A., Mankowitz, D. J., Zahavy, T., Xu, Z., Oh, J., Levine, N., and Mann, T. Balancing constraints and rewards with meta-gradient d4pg. *arXiv preprint arXiv:2010.06324*, 2020.

Chow, Y., Nachum, O., Faust, A., Duenez-Guzman, E., and Ghavamzadeh, M. Lyapunov-based safe policy optimization for continuous control. *arXiv preprint arXiv:1901.10031*, 2019.

Christiano, P., Leike, J., Brown, T. B., Martic, M., Legg, S., and Amodei, D. Deep reinforcement learning from human preferences. *arXiv preprint arXiv:1706.03741*, 2017.

Dalal, G., Dvijotham, K., Vecerik, M., Hester, T., Paduraru, C., and Tassa, Y. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.

de Woillemont, P. L. P., Labory, R., and Corruble, V. Configurable agent with reward as input: A play-style continuum generation. In *2021 IEEE Conference on Games (CoG)*, pp. 1–8. IEEE, 2021.

Devlin, S., Georgescu, R., Momennejad, I., Rzepecki, J., Zuniga, E., Costello, G., Leroy, G., Shaw, A., and Hofmann, K. Navigation turing test (ntt): Learning to evaluate human-like navigation. *arXiv preprint arXiv:2105.09637*, 2021.

Dewey, D. Reinforcement learning and the reward engineering principle. In *2014 AAAI Spring Symposium Series*, 2014.

Dulac-Arnold, G., Mankowitz, D., and Hester, T. Challenges of real-world reinforcement learning. *arXiv preprint arXiv:1904.12901*, 2019.

-
- Fujimoto, S., Hoof, H., and Meger, D. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pp. 1587–1596. PMLR, 2018.
- Gisslén, L., Eakins, A., Gordillo, C., Bergdahl, J., and Tollmar, K. Adversarial reinforcement learning for procedural content generation. *arXiv preprint arXiv:2103.04847*, 2021.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- Gordillo, C., Bergdahl, J., Tollmar, K., and Gisslén, L. Improving playtesting coverage via curiosity driven reinforcement learning agents. *arXiv preprint arXiv:2103.13798*, 2021.
- Goyal, P., Niekum, S., and Mooney, R. J. Using natural language for reward shaping in reinforcement learning. *arXiv preprint arXiv:1903.02020*, 2019.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018.
- Hadfield-Menell, D., Milli, S., Abbeel, P., Russell, S., and Dragan, A. Inverse reward design. *arXiv preprint arXiv:1711.02827*, 2017.
- Jacob, M., Devlin, S., and Hofmann, K. “it’s unwieldy and it takes a lot of time”—challenges and opportunities for creating agents in commercial games. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 16, pp. 88–94, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Korpelevich, G. An extragradient method for finding saddle points and for other problems. 1976.
- Kostrikov, I., Agrawal, K. K., Dwibedi, D., Levine, S., and Tompson, J. Discriminator-actor-critic: Addressing sample inefficiency and reward bias in adversarial imitation learning. *arXiv preprint arXiv:1809.02925*, 2018.
- Lin, T., Jin, C., and Jordan, M. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pp. 6083–6093. PMLR, 2020.
- MacGlashan, J., Babes-Vroman, M., desJardins, M., Littman, M. L., Muresan, S., Squire, S., Tellex, S., Arumugam, D., and Yang, L. Grounding english commands to reward functions. In *Robotics: Science and Systems*, 2015.
- Marchesini, E., Corsi, D., and Farinelli, A. Exploring safer behaviors for deep reinforcement learning. AAAI, 2022.
- Mindermann, S., Shah, R., Gleave, A., and Hadfield-Menell, D. Active inverse reward design. *arXiv preprint arXiv:1809.03060*, 2018.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Ng, A. Y., Harada, D., and Russell, S. Policy invariance under reward transformations: Theory and application to reward shaping. In *Icml*, volume 99, pp. 278–287, 1999.
- Polyak, B. Iterative methods using lagrange multipliers for solving extremal problems with constraints of the equation type. *USSR Computational Mathematics and Mathematical Physics*, 10(5):42–52, 1970.
- Ratner, E., Hadfield-Menell, D., and Dragan, A. D. Simplifying reward design through divide-and-conquer. *arXiv preprint arXiv:1806.02501*, 2018.
- Ray, A., Achiam, J., and Amodei, D. Benchmarking safe exploration in deep reinforcement learning. *arXiv preprint arXiv:1910.01708*, 7, 2019.
- Saunders, W., Sastry, G., Stuhlmüller, A., and Evans, O. Trial without error: Towards safe reinforcement learning via human intervention. *arXiv preprint arXiv:1707.05173*, 2017.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Shao, K., Tang, Z., Zhu, Y., Li, N., and Zhao, D. A survey of deep reinforcement learning in video games. *arXiv preprint arXiv:1912.10944*, 2019.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- Silver, D., Singh, S., Precup, D., and Sutton, R. S. Reward is enough. *Artificial Intelligence*, pp. 103535, 2021.
- Singh, S., Lewis, R. L., and Barto, A. G. Where do rewards come from. In *Proceedings of the annual conference of the cognitive science society*, pp. 2601–2606. Cognitive Science Society, 2009.

-
- Stooke, A., Achiam, J., and Abbeel, P. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pp. 9133–9143. PMLR, 2020.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Sutton, R. S., Modayil, J., Delp, M., Degris, T., Pilarski, P. M., White, A., and Precup, D. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. In *10th International Conference on Autonomous Agents and Multiagent Systems*, pp. 761–768, 2011.
- Tessler, C., Mankowitz, D. J., and Mannor, S. Reward constrained policy optimization. *arXiv preprint arXiv:1805.11074*, 2018.
- Turchetta, M., Kolobov, A., Shah, S., Krause, A., and Agarwal, A. Safe reinforcement learning via curriculum induction. *arXiv preprint arXiv:2006.12136*, 2020.
- Uzawa, H., Anow, K., and Hurwicz, L. *Studies in linear and nonlinear programming*, 1958.
- Vinyals, O., Babuschkin, I., Czarnecki, W. M., Mathieu, M., Dudzik, A., Chung, J., Choi, D. H., Powell, R., Ewalds, T., Georgiev, P., et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575 (7782):350–354, 2019.
- White, A. *Developing a predictive approach to knowledge*. PhD thesis, University of Alberta, 2015.
- Yang, T.-Y., Rosca, J., Narasimhan, K., and Ramadge, P. J. Projection-based constrained policy optimization. *arXiv preprint arXiv:2010.03152*, 2020.
- Zhang, Y., Vuong, Q., and Ross, K. W. First order constrained optimization in policy space. *arXiv preprint arXiv:2002.06506*, 2020.
- Zheng, B., Verma, S., Zhou, J., Tsang, I., and Chen, F. Imitation learning: Progress, taxonomies and opportunities. *arXiv preprint arXiv:2106.12177*, 2021.

A. Algorithm

Our implementation of the SAC-Lagrangian algorithm is presented below. The exact values of each hyper-parameter for all of our experiments are listed in Tables 1 and 2. One notable difference between an unconstrained Soft-Actor Critic (Haarnoja et al., 2018) and our constrained version is that SAC is typically updated after every environment step to maximise the sample efficiency of the algorithm. In the constrained case however, since the constraints are optimized on-policy, updating the SAC agent at every environment step would only allow for one-sample estimates of the multiplier’s objective. On the other hand, freezing the SAC-agent for as many environment steps as the Lagrange multiplier batch-size N_λ makes the overall algorithm significantly less sample efficient. One *could* disregard the “on-policyness” of the multiplier’s objective but in preliminary experiments we found that, unsurprisingly, updating the Lagrange multipliers very frequently while using a large set of samples (many of which were collected using previous versions of the policy) lead to significant overshoot and harms the ability of the multipliers to converge to a stable behavior. There is thus a tradeoff to make between the variance of the multiplier’s objective estimate, the degree to which the multipliers are updated on-policy and the sample efficiency of the overall algorithm. In practice we found that the values for M_θ and M_λ presented in Tables 1 and 2 represented good compromises between these different characteristics. Another important detail is that we use $K + 1$ separate critics to model the discounted expected sum of reward and costs. $Q^{(0)}$ is the critic that models the main objective and $Q^{(k)}$, $k = 1, \dots, K + 1$ are the critics that model the constraint components of the Lagrangian. Using separate critics allows to avoid fast changes in the scale of the objective, as seen by the critics, when the multipliers λ_k get adjusted; they can solely focus on modeling the agent’s changing behavior with respect to their respective function (reward or costs).

Algorithm 1 SAC-Lagrangian with Bootstrap Constraint

Require: learning rate β , replay buffer \mathcal{B} , entropy coefficient α and minibatch sizes N_θ and N_λ

Require: Initialise the policy π_θ and value-functions $Q_\phi^{(k)}$ randomly, $k = 0, \dots, K + 1$

Require: Initialise the Lagrange multiplier parameters z_k

Require: Collect enough transitions to fill \mathcal{B} with $\max(N_\theta, N_\lambda)$ samples

for updates $u = 1, \dots$ (until convergence) **do**

 # Data collection

 Sample from the current policy: $a \sim \pi_\theta(\cdot|s)$

 Query next state, reward and indicators $(s', r, \{e\}_{k=1}^{K+1})$ by interacting with the environment

 Append transition $(s, a, r, s', \{e\}_{k=1}^{K+1})$ to the replay buffer \mathcal{B}

 # Policy Gradient update

if $u \% M_\theta == 0$ **then**

 Sample a minibatch of N_θ transitions **uniformly** from the replay buffer

 Sample next actions: $a'_i \sim \pi_\theta(\cdot|s'_i) \quad i = 1, \dots, N_\theta$

for $k = 0, \dots, K + 1$ **do**

 Set the “rewards” to their corresponding values: $r_i^{(0)} = r_i$ and $r_i^{(k)} = e_i^{(k)}$

 Compute the Q-targets: $y_i^{(k)} = -\alpha \log \pi_\theta(a'_i|s'_i) + \min_{j \in \{1,2\}} Q_{\phi_j}^{(k)}(s'_i, a'_i)$

 Adam descent on Q-nets with: $\nabla_{\phi_j} \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} \|Q_{\phi_j}^{(k)}(s_i, a_i) - (r_i^{(k)} + (1 - done)\gamma y_i^{(k)})\|_2$

end for

 Re-sample the current actions: $a_i \sim \pi_\theta(\cdot|s_i) \quad i = 1, \dots, N_\theta$

 Adam ascent on policy with:

$$\nabla_\theta \frac{1}{N_\theta} \sum_{i=1}^{N_\theta} -\alpha \log \pi_\theta(a_i|s_i) + \max(\lambda_0, \lambda_{K+1}) \min_j Q_{\phi_j}^{(0)}(s_i, a_i) + \lambda_{K+1} \min_j Q_{\phi_j}^{(K+1)}(s_i, a_i) - \sum_{k=1}^K \lambda_k \min_j Q_{\phi_j}^{(k)}(s_i, a_i)$$

end if

 # Multipliers update

if $u \% M_\lambda == 0$ **then**

 Draw from the replay buffer a minibatch composed of **the last** N_λ transitions

for $k = 0, \dots, K + 1$ **do**

 Compute average costs: $\tilde{J}_{C_k}(\pi) = \frac{1}{N_\lambda} \sum_{i=1}^{N_\lambda} e_i^{(k)}$

 Adam descent on multipliers with: $\nabla_{z_k} \lambda_k (\tilde{J}_{C_k}(\pi) - \tilde{d}_k)$ if $k = K + 1$ else $\nabla_{z_k} \lambda_k (\tilde{d}_k - \tilde{J}_{C_k}(\pi))$

end for

end if

end for

B. Details for experiments in the Arena environment

B.1. Environment details

In the Arena Environment, the agent’s main goal is to navigate to the green tile (see Figure 1, left). The constraints that we explore in this environment are $\{On-Ground, Not-in-Lava, Looking-At-Marker, Under-Speed-Limit$ and $Above-Energy-Limit\}$. It receives as observations its XYZ position, direction and velocity, the relative XZ position of the goal, its distance to the goal, as well as an indicator for whether it is on the ground. For the looking-at constraint, it also receives the XZ vector for the direction it is looking at, its Y-angular velocity, the marker’s relative XZ position and distance, the normalised angle between the agent’s looking direction and the marker as well as an indicator for whether the marker is within its field of view (a fixed-angle cone in front of the agent). For the energy constraint, the agent receives the normalised value of its energy bar and an indicator for whether it is currently recharging. Finally for the lava constraint, the agent receives an indicator of whether it currently stands in lava as well as an indicator for 25 vertical raycast of its surrounding (0 indicating safe ground and 1 indicating lava). We also add to the agent’s observations the per-episode rates of indicator cost functions to the agent observation for each of the constraint as well a normalised representation of the remaining time-steps before reaching the time limit condition, leading to a total dimensionality of 53 for the observation vector. The action space is composed of 5 continuous actions (clamped between -1 and 1) which represent its XZ velocity and Y-angular velocity, a jump action (jump is triggered when the agent outputs a value above 0 for that dimensionality) and a recharge action (also with threshold of 0). The reward function is simply 1 when the agent reaches the goal (causing termination), 0 otherwise, and augmented with a small shaping reward function (Ng et al., 1999) based on whether the agent got closer or further away from the goal location.

B.2. Hyper-parameters

Most of the hyper-parameters are the same as in the original unconstrained Soft Actor-Critic (SAC) (Haarnoja et al., 2018). Some additional hyper-parameters emerge from the constraint enforcement aspect of our version of SAC-Lagrangian and are described in the Algorithm section above. We use the Adam optimizer (Kingma & Ba, 2014) for all parameter updates (policy, critics and Lagrange multipliers). For all experiments taking place in the Arena Environment, the policy is parameterized as a two layer neural networks that outputs the parameters of a Gaussian distribution with a diagonal covariance matrix. The hidden layers are composed of 256 units and followed by a \tanh activation function. The first hidden layer also uses layer-normalisation before the application of the \tanh function. We use $K + 1$ fully independent critic models to estimate the expected discount sum of each of the constraint and of the main reward function. The critic models are also parameterized with two-hidden-layers neural networks with the same size for the hidden layers as the policy but instead followed by relu activation functions. Table 1 shows the hyper-parameters used in our experiments conducted in the Arena environment.

Table 1. HYPER-PARAMETERS FOR EXPERIMENTS IN THE ARENA ENVIRONMENT.

GENERAL	DISCOUNT FACTOR γ	0.9
	NUMBER OF RANDOM EXPLORATION STEPS	10000
	NUMBER OF BUFFER WARMUP STEPS	2560
SAC AGENT	LEARNING RATE β	0.0003
	TRANSITIONS BETWEEN UPDATES M_θ	200
	BATCH SIZE N_θ	256
	REPLAY BUFFER SIZE	1,000,000
	INITIAL ENTROPY COEFFICIENT α	0.02
LAGRANGE MULTIPLIERS	TARGET NETWORKS SOFT-UPDATE COEFFICIENT τ	0.005
	LEARNING RATE β	0.03
	INITIAL MULTIPLIER PARAMETERS VALUE z_k	0.02
	TRANSITIONS BETWEEN UPDATES M_λ	2000
	BATCH SIZE N_λ	2000
CONSTRAINT THRESHOLDS	HAS REACHED GOAL (LOWER-BOUND)	0.99
	NOT LOOKING AT MARKER	0.10
	NOT ON GROUND	0.40
	IN LAVA	0.01
	ABOVE SPEED LIMIT	0.01
	IS UNDER THE MINIMUM ENERGY LEVEL	0.01

C. Details for experiments in the OpenWorld environment

C.1. Environment details

The OpenWorld environment is a large environment (approximately 30,000 times larger than the agent) that includes multiple multi-storey buildings with staircases, mountains, tunnels, natural bridges and lava. In addition, the environment includes 50 jump-pads that propel the agent into the air when it steps on one of them. The agent is tasked with navigating towards a goal randomly placed in the environment at the beginning of every episode. The agent controls include translation in the XY frame (2 inputs), a jumping action (1 input), a rotation action controlling where the agent is looking independent of its direction of travel (1 input), and a recharging action which allows the agent to recharge its energy level (1 input). The recharging action immobilizes the agent, i.e., it does not allow the agent to progress towards its goal. The environment also includes a look-at marker which we would like the agent to look at while it accomplishes its main navigation task.

At every timestep, the agent receives as observations its XYZ position relative to the goal as well as its normalized velocity and acceleration in the environment. In addition, it receives its relative position to the nearest jump-pad in the environment. For looking at the marker, as in the Arena environment, the agent receives the marker’s relative XZ position and distance, the normalised angle between the agent’s looking direction and the marker, as well as an indicator for whether the marker is within its field of view (a fixed-angle cone in front of the agent). For the energy-limit constraint, the agent obtains the value of its energy level, a boolean describing if it is currently recharging and a Boolean indicating if it was recharging in the previous timestep. The agent also receives a series of indicators denoting whether it is currently standing in lava, if it is touching the ground, and if the agent is currently below the minimum energy level. In order for the agent to observe lava and other elements it can collide with in the environment (e.g., buildings, doors, mountains), the agent receives 2 channels of 8×8 raycasts around the agent.

C.2. Hyper-parameters

The SAC agent in the OpenWorld environment uses the same architecture and similar hyper-parameters as in (Alonso et al., 2020). The raycasts and raw state described above are processed using two separate embedding models. For the raycasts, we employ a CNN with 3 convolutional layers, each with a corresponding ReLU layer. The raw state is processed using a separate 3-layer MLP with 1024 hidden units at each layer. The two representations are concatenated into a single vector representing the current state. The policy is parameterized by a 3-layer MLP that receives as input the concatenated representation and outputs the parameters of a Gaussian distribution with a diagonal covariance matrix. Each hidden layer is composed of 1024 hidden units and is followed by a ReLU activation function. The critic models are also parameterized by 3-layer MLP, are composed of 1024 hidden units and use ReLU activation functions. Table 2 shows some of these hyper-parameters with a focus on the constrained enforcement aspect of our version of SAC-Lagrangian.

Table 2. HYPER-PARAMETERS FOR EXPERIMENTS IN THE OPENWORLD ENVIRONMENT.

GENERAL	DISCOUNT FACTOR γ	0.99
	NUMBER OF RANDOM EXPLORATION STEPS β	200
	NUMBER OF BUFFER WARMUP STEPS β	2560
SAC AGENT	LEARNING RATE β	0.0001
	BATCH SIZE N_θ	2560
	REPLAY BUFFER SIZE	4,000,000
	INITIAL ENTROPY COEFFICIENT α	0.005
	TARGET NETWORKS SOFT-UPDATE COEFFICIENT τ	0.005
LAGRANGE MULTIPLIERS	LEARNING RATE β	0.00005
	INITIAL MULTIPLIER PARAMETERS VALUE z_k	0.02
	TRANSITIONS BETWEEN UPDATES	EVERY TIMESTEP
	BATCH SIZE N_λ	5000
CONSTRAINT THRESHOLDS	HAS REACHED GOAL (LOWER-BOUND)	0.80
	NOT LOOKING AT MARKER	0.10
	NOT ON GROUND	0.40
	IN LAVA	0.001
	IS UNDER THE MINIMUM ENERGY LEVEL	0.01

D. Additional experiments on reward engineering

See Section 2 for the description of our experiments motivating against the use of reward engineering for behavior specification. Figure 7 below shows the results for the biggest of the 3 grid searches performed to showcase the difficulty of finding a reward function that fits the behavioral requirements when the number of requirements grows.

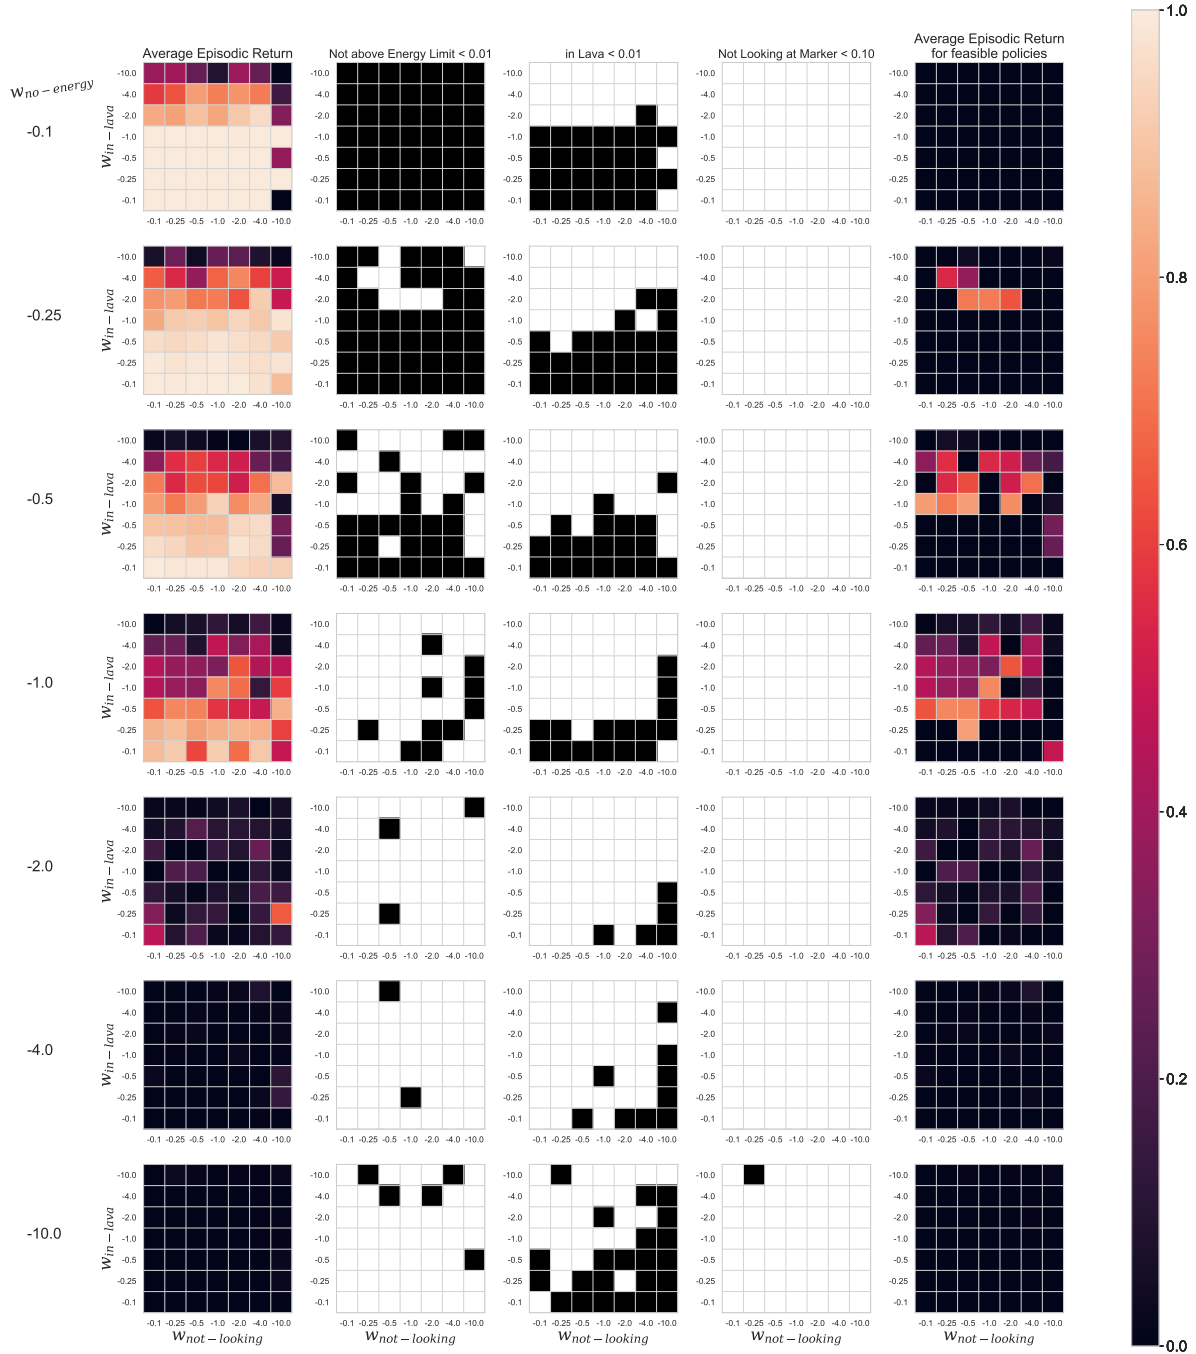


Figure 7. Also see Figure 2. When enforcing 3 behavioral requirements with reward engineering, an ever larger proportion of the experiments are wasted finding either low-performing policies or policies that do not satisfy the behavioral constraints. In this case, none of the 343 experiments yielded a feasible policy that also solves the task (success rate near 1.0), showcasing that reward engineering scales poorly with the number of constraints due to the curse of dimensionality and to the composing effect of the multiple constraints in narrowing the space of feasible policies.

E. Additional experiments on TD3

We validate that our framework can be combined with any policy optimisation algorithm by applying it to the TD3 algorithm (Fujimoto et al., 2018). This leads to a TD3-Lagrangian formulation using our indicator cost functions, normalized multipliers and bootstrap constraint. As for our experiments with SAC (Figure 5-d), our TD3-Lagrangian agent performs well and all constraints are satisfied. The results are presented in Figure 8.




Figure 8. TD3-Lagrangian agent in the Arena environment using normalised multipliers, indicator cost functions and using the success constraint as a bootstrap constraint. Training is halted after every 20,000 environment steps and the agent is evaluated for 10 episodes. All curves show the average over 5 seeds and envelopes show the standard error around that mean.