

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

12-2023

Generative modelling of stochastic actions with arbitrary constraints in reinforcement learning

Changyu CHEN

Ramesha KARUNASENA

Thanh Hong NGUYEN

Arunesh SINHA

Pradeep VARAKANTHAM

Singapore Management University, pradeepv@smu.edu.sg

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

CHEN, Changyu; KARUNASENA, Ramesha; NGUYEN, Thanh Hong; SINHA, Arunesh; and VARAKANTHAM, Pradeep. Generative modelling of stochastic actions with arbitrary constraints in reinforcement learning. (2023). *Proceedings of the 37th Annual Conference on Neural Information Processing Systems, New Orleans, USA, 2023 December 10-16.*

Available at: https://ink.library.smu.edu.sg/sis_research/8589

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Generative Modelling of Stochastic Actions with Arbitrary Constraints in Reinforcement Learning

Chen Changyu¹, Ramesha Karunasena¹, Thanh Hong Nguyen²,
Arunesh Sinha³, Pradeep Varakantham¹

Singapore Management University¹, University of Oregon², Rutgers University³
cychen.2020@phdcs.smu.edu.sg, rameshak@smu.edu.sg,
thanhhng@cs.uoregon.edu, arunesh.sinha@rutgers.edu, pradeepv@smu.edu.sg

Abstract

Many problems in Reinforcement Learning (RL) seek an optimal policy with large discrete multidimensional yet unordered action spaces; these include problems in randomized allocation of resources such as placements of multiple security resources and emergency response units, etc. A challenge in this setting is that the underlying action space is categorical (discrete and unordered) and large, for which existing RL methods do not perform well. Moreover, these problems require validity of the realized action (allocation); this validity constraint is often difficult to express compactly in a closed mathematical form. The allocation nature of the problem also prefers stochastic optimal policies, if one exists. In this work, we address these challenges by (1) applying a (state) conditional normalizing flow to compactly represent the stochastic policy — the compactness arises due to the network only producing one sampled action and the corresponding log probability of the action, which is then used by an actor-critic method; and (2) employing an invalid action rejection method (via a valid action oracle) to update the base policy. The action rejection is enabled by a modified policy gradient that we derive. Finally, we conduct extensive experiments to show the scalability of our approach compared to prior methods and the ability to enforce arbitrary state-conditional constraints on the support of the distribution of actions in any state¹.

1 Introduction

Adaptive resource allocation problems with multiple resource types (e.g., fire trucks, ambulances, police vehicles) are ubiquitous in the real world [13, 23, 32]. One example is allocating security resources and emergency response vehicles to different areas depending on incidents [23]. There are many other similar problems in aggregation systems for mobility/transportation, logistics etc [13]. In this paper, we are interested in addressing the multiple difficult challenges present in such adaptive resource allocation problems: (a) Combinatorial action space, as the number of resource allocations is combinatorial; (b) *Categorical* action space, as there is no ordering of resource allocations with respect to the overall objective (as increasing or decreasing different resource types at different locations can have different impact on overall objective); (c) Constraints on feasible allocations and switching between allocations; and (d) Finally, uncertainty in demand for resources.

Existing research in RL for constrained action spaces has considered resource allocation problems with a single type of resource, thereby introducing order in action space [1]. In such ordered action spaces, actions can be converted into continuous space and this allows for the usage of continuous action RL methods (e.g., DDPG). In problems of interest in this paper, we are interested in problems

¹Our implementation is available at <https://github.com/cameron-chen/flow-iar>.

with multiple resource types (e.g., fire truck, ambulance, police). These problems have a large action space that is discrete and unordered (categorical) and there are constraints on feasible allocations (e.g., no two police vehicles can be more than 3 km away, cannot move vehicles too far away from time step to time step). Furthermore, such constraints are easy to validate by a validity oracle given any allocation action, but are hard to represent as mathematical constraints on the support of the distribution of actions (in each state) as they often require exponentially many inequalities [4, 35]. An important consideration in allocation problems is randomized (stochastic) allocation arising from issues of fair division of indivisible resources so that an allottee is not starved of resources forever [4]. Thus, we aim to output stochastic optimal policies, if one exists.

Towards addressing such resource allocation problems at scale, we propose to employ generative policies in RL. Specifically, we propose a new approach that incorporates discrete normalizing flow policy in an actor-critic framework to explore and learn in the aforementioned constrained, categorical and adaptive resource allocation problems. Prior RL approaches for large discrete multidimensional action space include ones that assume a factored action space with independent dimensions, which we call as the factored or marginal approach, since independence implies that any joint distribution over actions can be represented as the product of marginal distributions over each action dimension. Other approaches convert the selection of actions in multiple dimensions into a sequential selection approach. Both these approaches are fundamentally limited in expressivity, which we reveal in detail in our experiments. Next, we provide a formal description of the problem that we tackle.

Problem Statement: A Markov Decision Process (MDP) is represented by the tuple $\langle \mathcal{S}, \mathcal{A}, P, r, \gamma, b_0 \rangle$, where an agent can be in any state $s_t \in \mathcal{S}$ at a given time t . The agent takes an action $a_t \in \mathcal{A}$, causing the environment to transition to a new state s_{t+1} with a probability $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \mapsto [0, 1]$. Subsequently, the agent receives a reward $r : \mathcal{S} \times \mathcal{A} \mapsto \mathbb{R}$. In the infinite-horizon setting, the discounted factor is $0 < \gamma < 1$. The distribution of the initial state is b_0 .

In our work, we focus on a categorical action space, \mathcal{A} . Categorical action spaces consist of discrete, unordered actions with no inherent numerical relationship between them. We assume that for any state s , there is a set of valid actions, denoted by $\mathcal{C}(s) \subseteq \mathcal{A}$. There is an oracle to answer whether an action $a \in \mathcal{C}(s)$, but the complex constraint over categorical space cannot be expressed succinctly using closed form mathematical formula. Note that the constraint is *not* the same for every state. Our objective is to learn a stochastic policy, $\pi(\cdot|s)$, which generates a *probability distribution over actions for state s with support only over the valid actions $\mathcal{C}(s)$ in state s* . We call the set of such stochastic policies as valid policies $\Pi^{\mathcal{C}}$ given the per state constraint \mathcal{C} . We aim to maximize the long-term expected rewards over valid policies, that is, $\max_{\pi \in \Pi^{\mathcal{C}}} J(\pi)$, where J is as follows:

$$J(\pi) = \mathbb{E}_{s \sim b_0} [V(s; \pi)] \text{ where } V(s; \pi) = \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t) \mid s_0 = s; \pi \right] \quad (1)$$

In addition, we also consider settings with *partial observability* where the agent observes $o_t \in \mathcal{O}$ at time t where the observation arises from the state s_t with probability $O : \mathcal{O} \times \mathcal{S} \mapsto [0, 1]$. In this case, the optimal policy is a stochastic policy, $\pi(\cdot|h)$, where h is the history of observations till current time. For partial observability, we consider an unconstrained setting, as the lack of knowledge of the true state results in uncertainty about which constraint to enforce, which diverges from the focus in this work but is an interesting future research direction to explore. Thus, with partial observability, we search over stochastic policies Π that maximize the long term return, that is, $\max_{\pi \in \Pi} J(\pi)$. $J(\pi)$ is the same as stated in Equation 1 but where the expectation is also over the observation probability distribution in addition to the standard transition and stochastic policy probability distributions.

Contribution: We propose two key innovations to address the problem above. *First*, we present a conditional Normalizing Flow-based [26] Policy Network, which leverages Argmax Flow [16] to create a minimal representation of the policy for policy gradient algorithms. To the best of our knowledge, this is the first use of discrete normalizing flow in RL. *Second*, we demonstrate how to train the flow policies within the A2C framework. In particular, we need an estimate of the log probability of the action sampled from the stochastic policy but Argmax Flow provides only a biased lower bound via the evidence lower bound (ELBO). Thus, we design an effective *sandwich estimator* for the log probability that is sandwiched between the ELBO lower bound and an upper bound based on χ^2 divergence. *Third*, we propose a policy gradient approach which is able to reject invalid actions (that do not satisfy constraints) referred to as Invalid Action Rejection Advantage Actor-Critic (IAR-A2C). IAR-A2C queries the constraint oracle and ensures validity of actions in every state (in fully observable setting) by rejecting all invalid actions. We derive a new policy gradient estimator

for IAR-A2C. Figure 1 provides an overview of our architecture. Finally, our extensive experimental results reveal that our approach outperforms prior baselines in different environments and settings.

2 Background

Normalizing Flows: Normalizing flows are a family of generative models that can provide both efficient sampling and density estimation. Their main idea is to construct a series of invertible and differentiable mappings that allow transforming a simple probability distribution into a more complex one. Given $\mathcal{V} = \mathcal{Z} = \mathbb{R}^d$ with densities p_V and p_Z respectively, normalizing flows [26] aim to learn a bijective and differentiable transformation $f : \mathcal{Z} \rightarrow \mathcal{V}$. This deterministic transformation allows us to evaluate the density at any point $v \in \mathcal{V}$ based on the density of $z \in \mathcal{Z}$, as follows:

$$p_V(\mathbf{v}) = p_Z(\mathbf{z}) \cdot \left| \det \frac{d\mathbf{z}}{d\mathbf{v}} \right|, \quad \mathbf{v} = f(\mathbf{z}) \quad (2)$$

In this context, p_Z can be any density, though it is typically chosen as a standard Gaussian and f is represented by a neural network. Consequently, normalizing flows offer a powerful tractable framework for learning complex density functions. However, the density estimation presented in Equation 2 is limited to continuous probability distributions. To enable the learning of probability mass functions (P) on categorical discrete data, such as natural language, Argmax Flow [16] proposed to apply the argmax operation on the output of continuous flows. Let’s consider $\mathbf{v} \in \mathbb{R}^{D \times M}$ and $\mathbf{x} \in \{1, \dots, M\}^D$. The argmax operation is interpreted as a surjective flow layer $\mathbf{v} \mapsto \mathbf{x}$, which is deterministic in one direction ($x_d = \arg \max_m v_d$, written compactly as $\mathbf{x} = \arg \max \mathbf{v}$) and stochastic in the other ($\mathbf{v} \sim q(\cdot|\mathbf{x})$). With this interpretation, the argmax operation can be considered a probabilistic right-inverse in the latent variable model expressed by:

$$P(\mathbf{x}) = \int P(\mathbf{x}|\mathbf{v})p(\mathbf{v})d\mathbf{v}, \quad P(\mathbf{x}|\mathbf{v}) = \delta(\mathbf{x} = \arg \max(\mathbf{v})) \quad (3)$$

where $\arg \max$ is applied in the last dimension of \mathbf{v} . In this scenario, the density model $p(\mathbf{v})$ is modeled using a normalizing flow. The learning process involves introducing a variational distribution $q(\mathbf{v}|\mathbf{x})$, which models the probabilistic right-inverse for the argmax surjection, and optimizing the evidence lower bound (ELBO), which is the RHS of the following inequality:

$$\log P(\mathbf{x}) \geq \mathbb{E}_{\mathbf{v} \sim q(\cdot|\mathbf{x})} [\log P(\mathbf{x}|\mathbf{v}) + \log p(\mathbf{v}) - \log q(\mathbf{v}|\mathbf{x})] = \mathbb{E}_{\mathbf{v} \sim q(\cdot|\mathbf{x})} [\log p(\mathbf{v}) - \log q(\mathbf{v}|\mathbf{x})] = \mathcal{L}$$

The last equality holds under the constraint that the support of $q(\mathbf{v}|\mathbf{x})$ is enforced to be only over the region $\mathcal{S} = \{\mathbf{v} \in \mathbb{R}^{D \times M} : \mathbf{x} = \arg \max \mathbf{v}\}$ which ensures that $P(\mathbf{x}|\mathbf{v}) = 1$. From standard variational inference results, $\log P(\mathbf{x}) - \mathcal{L} = \text{KL}(q(\mathbf{v}|\mathbf{x})||p(\mathbf{v}|\mathbf{x}))$, which also approaches 0 as the approximate posterior $q(\mathbf{v}|\mathbf{x})$ comes closer to the true posterior $p(\mathbf{v}|\mathbf{x})$ over the training time.

χ^2 **Upper Bound:** Variational inference involves proposing a family of approximating distributions and finding the family member that is closest to the posterior. Typically, the Kullback-Leibler (KL) divergence $\text{KL}(q||p)$ is employed to measure closeness, where $q(\mathbf{v}|\mathbf{x})$ represents a variational family. This approach yields ELBO of the evidence $\log P(\mathbf{x})$ as described above.

Instead of using KL divergence, the authors in [8] suggest an alternative of χ^2 -divergence to measure the closeness. As a result, they derived an upper bound of the evidence, known as CUBO: $\log P(\mathbf{x}) \leq \frac{1}{2} \log \mathbb{E}_{\mathbf{v} \sim q(\cdot|\mathbf{x})} \left[\left(\frac{p(\mathbf{x}, \mathbf{v})}{q(\mathbf{v}|\mathbf{x})} \right)^2 \right]$. Similar to the ELBO in Argmax Flow, CUBO can be further simplified under the constraint that the support of $q(\cdot|\mathbf{x})$ is restricted to the region \mathcal{S} :

$$\log P(\mathbf{x}) \leq \frac{1}{2} \log \mathbb{E}_{\mathbf{v} \sim q(\cdot|\mathbf{x})} \left[\left(\frac{p(\mathbf{v})}{q(\mathbf{v}|\mathbf{x})} \right)^2 \right] = \mathcal{L}_{\chi^2} \quad (4)$$

Also, $\mathcal{L}_{\chi^2} - \log P(\mathbf{x}) = \frac{1}{2} \log(1 + D_{\chi^2}(p(\mathbf{v}|\mathbf{x})||q(\mathbf{v}|\mathbf{x})))$, hence the gap between the ELBO and CUBO approaches 0 as the approximate posterior $q(\cdot|\mathbf{x})$ becomes closer to the true posterior $p(\cdot|\mathbf{x})$.

3 Flow-based Policy Gradient Algorithm with Invalid Action Rejection

We first present the Flow-based Policy Network, which leverages Argmax Flow to create a minimal representation of the policy for policy gradient algorithms, and then construct flow policies within

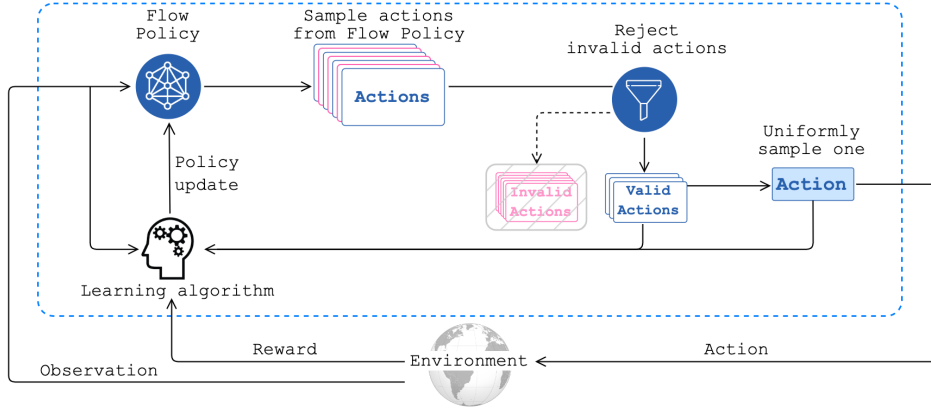


Figure 1: **Our IAR-A2C framework.** At each time step, an initial batch of action samples, along with their log probabilities, are generated using the Flow Policy. Invalid actions from this batch are rejected using an oracle. A single action is then uniformly sampled from the remaining valid ones, and executed. This selected action and the valid action set are stored along with the resulting state and reward. This collective experience is subsequently utilized to update the Flow Policy.

the A2C framework. Our exposition will present policies conditioned on state, but the framework works for partial observability also by using the sequence of past observations as state. After this, for the fully observable setting only, we introduce our novel policy gradient algorithm called IAR-A2C that enforces state dependent constraints.

3.1 Flow-based Policy Network

In policy gradient algorithms for categorical actions, the standard approach is to model the entire policy, denoted as $\pi(\cdot|s)$, which allows us to sample actions and obtain their corresponding probabilities. The policy is parameterized using a neural network where the network generates logits for each action, and then converts these logits into probabilities. The size of the output layer matches that of the action space, which can be prohibitively large for resource allocation problems of interest in this paper. We posit that we can have a better policy representation, as it is sufficient to require samples from the support set of the policy, represented by $a^{(i)} \sim \pi(\cdot|s)$ and probability value $\pi(a^{(i)}|s) > 0$.




Figure 2: Composition of a conditional flow $p_{\rho,w}(z_K|s)$ and argmax transformation resulting in the policy $\pi_{\rho,w}(a|s)$. The flow maps from a base distribution $p_{\rho}(z_0|s)$ by using a bijection F_w . The diagram is adapted from Hoogetboom et al. [16].

Based on this observation, our first contribution is a compact policy representation using Argmax Flow [16], which we refer to as the *Flow Policy*. Argmax Flow is a state-of-the-art discrete normalizing flow model and it has shown great capability of learning categorical data such as in sentence generation. In our context, Flow Policy will output the action (a sample of the flow policy) and its probability, instead of explicitly outputting the entire distribution and sampling from it as in prior work. Once trained, we can sample from the Argmax Flow and, more importantly, estimate the probability of the sample. A normalizing flow model transforms a base distribution,

given by a random variable z_0 , to the desired distribution. In our approach, the desired distribution is the distribution given by policy π . We adapt Argmax Flow approach for learning a stochastic policy.

Before diving into the specifics of the Flow Policy, we first discuss our rationale for choosing the normalizing flow model over other contemporary deep generative models, such as Generative Adversarial Networks (GANs) [12] and Diffusion Models [15, 29]. GANs, categorized as implicit generative models, do not allow for an estimation of the data’s log probability. While prior research has successfully devised an approach for exact log probability estimation with Diffusion Models, these

Algorithm 1: ELBO Optimization

```
1 Input: Invertible flow  $F_w = f_{w,K} \circ \dots \circ f_{w,1}$ , State encoder  $E_\rho$ , Posterior  $q_\psi$ , rollout  $\tau$ 
2 Sample  $States = \{(s^{(i)})\}_{i=1}^B \sim \tau$ 
3 for  $s \in States$  do
4   for  $j \leftarrow 1 : n\_elbo\_steps$  do
5      $z^{(j)} = F_w(z_0^{(j)})$ , where  $z_0^{(j)} \sim E_\rho(s)$ 
6      $a^{(j)} = \operatorname{argmax} z^{(j)}$ 
7   for  $j \leftarrow 1 : n\_elbo\_steps$  do
8      $z'^{(jn)} = \operatorname{threshold}_T(u^{(jn)})$ , where  $\{(u^{(jn)})\}_{n=1}^N \sim q_\psi(\cdot | a^{(j)}, s)$ 
9      $z_0'^{(jn)} = F_w^{-1}(z'^{(jn)})$ 
10    Take a gradient ascending step of ELBO w.r.t.  $\rho, w$  and  $\psi$ 
11     $\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left( \log p_\rho(z_0'^{(jn)} | s) - \sum_{k=1}^K \log \left| \det \frac{\partial f_{w,k}}{\partial z_{k-1}'} \right| - \log q_\psi(z'^{(jn)} | a^{(j)}, s) \right)$ 
```

models encounter the problem of slow sampling and expensive probability estimation, necessitating the resolution of an ordinary differential equation [29]. In contrast, normalizing flow has low cost for sampling and probability evaluation in our case, as we find the simple flow model is sufficient to learn a good policy. So normalizing flow is a natural choice to construct our policy.

We demonstrate how to construct the flow policy within the A2C framework. An illustration is provided in Figure 2. To define the policy, $\pi(a|s)$, we first encode the state that is then used to define the base distribution z_0 . In accordance with standard practice, we select z_0 as a Gaussian distribution with parameters μ, σ defined by the state encoder (a state-dependent neural network): $z_0 = \mu_\rho(s) + \sigma_\rho(s) \cdot \epsilon$ where $\epsilon \sim \mathcal{N}(0, 1)$. We write this as $z_0 \sim E_\rho(s)$, where ρ denotes the weights of the state encoder. We then apply a series of invertible transformations given as functions f_k that define the flow, followed by the argmax operation (as defined in background on Argmax flow). Consequently, the final sampled action is given by $a = \operatorname{argmax} (f_K \circ \dots \circ f_1(z_0))$. Each f_k is an invertible neural network [26] and we use $F_w = f_{w,K} \circ \dots \circ f_{w,1}$ to denote the composed function, where w, k is the weight of the network representing f_k . Thus, sampled action $a = \operatorname{argmax} (F_w(z_0))$ for $z_0 \sim E_\rho(s)$. We use shorthand $\theta = (\rho, w)$ when needed. Also, we use $z_k = f_{w,k}(z_{k-1})$ to denote the output of $f_{w,k}$ and $p_\rho(z_0|s)$ to denote the probability density of $z_0 \sim \mathcal{N}(\mu_\rho, \sigma_\rho)$.

We use a variational distribution q_ψ for the reverse mapping from a discrete action (conditional on state s) to a continuous output of $F_w(z_0)$. The corresponding estimate of the log probability of the sampled action a , denoted by \hat{l}_π , is the evidence lower bound ELBO, computed as follows:

$$\hat{l}_\pi(a|s) = \mathcal{L} = \mathbb{E}_{z_K \sim q_\psi(\cdot | a, s)} \left[\log p_\rho(z_0|s) - \sum_{k=1}^K \left(\log \left| \det \frac{\partial f_k}{\partial z_{k-1}} \right| \right) - \log q_\psi(z_K | a, s) \right] \quad (5)$$

To ensure that our approximation closely approximates the evidence, we optimize ELBO progressively, following the training scheme of Argmax flow, as shown in the subroutine Algorithm 1 used within the outer A2C framework in Algorithm 2. The target distribution in this subroutine is given by actions $a^{(j)}$ sampled from the current policy (line 6); thus, this subroutine aims to update the flow network components (ρ, w) to make the distribution of overall policy be closer to this target distribution and to improve the posterior q_ψ estimate (ELBO update in line 11), which gets used in estimating the log probability of action required by the outer A2C framework (shown later in Algorithm 2). The various quantities needed for ELBO update are obtained in lines 8, 9 by passing $a^{(j)}$ back through the invertible flow layers (using q_ψ for the discrete to continuous inverse map).

3.2 Policy Gradient with Invalid Action Rejection

We aim to use the flow-based policy within an A2C framework, illustrated in Figure 1; here we describe the same along with how to enforce arbitrary state based constraint on actions. We propose a sandwich estimator which combines ELBO and CUBO to obtain a low-bias estimation of $\log \pi(a|s)$, thus improving the convergence of policy gradient learning. Moreover, to tackle the challenge of aforementioned complex constraints, our algorithm proceeds by sampling a number of actions from

the flow-based policy and then utilizing the constraint oracle to filter the invalid actions. We then provide theoretical results showing the adaptation of the policy gradient computation accordingly.

Sandwich estimator: The ELBO \hat{l}_π is a biased estimator of $\log \pi(a|s)$, but it is known in literature [5] that \hat{l}_π is a consistent estimator (i.e., converges to $\log \pi(a|s)$ in the limit) and there are generative methods based on this consistency property, such as [33]. However, we aim to use \hat{l}_π in the policy gradient and stochastic gradient descent typically uses an unbiased gradient estimate (but not always, see [6]). Therefore, we propose to use a new technique which combines ELBO and CUBO to reduce the bias, improving the convergence of our policy gradient based learning process. In particular, we estimate an upper bound of $\log \pi(a|s)$ using the following CUBO:

$$\hat{l}_\pi^u(a|s) = \mathcal{L}_{\chi^2} = \frac{1}{2} \log \mathbb{E}_{z_K \sim q_\psi(z_K|a,s)} \left[\left(\frac{p_\rho(z_0|s)}{\prod_{k=1}^K \left| \det \frac{\partial f_{w,k}}{\partial z_{k-1}} \right| q_\psi(z_K|a,s)} \right)^2 \right] \quad (6)$$

We then use a weighted average of the upper and lower bounds as a low-bias estimate of $\log \pi(a|s)$, denoted by $\widehat{\log p_{\theta,\psi}} = \alpha \hat{l}_\pi + (1 - \alpha) \hat{l}_\pi^u$ where α is a hyperparameter. We call this the *sandwich estimator* of log probability. We observed in our experiments that an adaptive $\alpha(\hat{l}_\pi, \hat{l}_\pi^u)$ as a function of the two bounds provides better results than a static $\alpha = \frac{1}{2}$ (see Appendix B.1 for more details).

Constraints: Since the agent only has oracle access to the constraints, existing safe exploration approaches [20, 24] are not directly applicable to this setting. If the agent queries the validity of all actions, then it gains complete knowledge of the valid action set $\mathcal{C}(s)$. However, querying action validity for the entire action space at each timestep can be time-consuming, particularly when the action space is large, e.g., 1000 categories. We demonstrate this issue in our experiments.

To address this challenge of complex *state-dependent* constraints in the full observation setting, we propose a new policy gradient algorithm called Invalid Action Rejection Advantage Actor-Critic (IAR-A2C). IAR-A2C operates by sampling a set of actions from the current flow-based policy and then leveraging the constraint oracle to reject all invalid actions, enabling the agent to explore safely with only valid actions. We then derive the policy gradient estimator for this algorithm. Recall that $\mathcal{C}(s)$ is the set of valid actions in state s and let $\mathcal{C}^I(s)$ be the set of invalid actions. We sample an action from $\pi_\theta(a|s)$ (recall $\theta = (\rho, w)$ when using Flow Policy), rejecting any invalid action till a valid action is obtained. Clearly, the effective policy π' induced by this rejection produces only valid actions. In fact, by renormalizing we obtain $\pi'_\theta(a|s) = \frac{\pi_\theta(a|s)}{\sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s)}$ for $a \in \mathcal{C}(s)$. For the purpose of policy gradient, we need to obtain the gradient of the long term reward with π' : $J(\theta) = \mathbb{E}_{s \sim b_0} [V(s; \pi')]$. We show that:

Theorem 1 *With π' defined as above,*

$$\nabla_\theta J(\theta) = \mathbb{E}_{\pi'} \left[Q^{\pi'}(s, a) \nabla_\theta \log \pi_\theta(a|s) - Q^{\pi'}(s, a) \frac{\sum_{a \in \mathcal{C}(s)} \nabla_\theta \pi_\theta(a|s)}{\sum_{a \in \mathcal{C}(s)} \pi_\theta(a|s)} \right] \quad (7)$$

In practice, analogous to standard approach in literature [22] to reduce variance, we use the TD error (an approximation of advantage) form instead of the Q function in Theorem 1. Then, given the network π_θ , the empirical estimate of the first gradient term is readily obtainable from trajectory samples of π' . To estimate the second term, for every state s in the trajectories, we sample S actions from $\pi_\theta(a|s)$ and reject any $a \in \mathcal{C}^I(s)$ to get $l \leq S$ valid actions a_1, \dots, a_l . We then apply Lemma 1:

Lemma 1 $\frac{1}{l} \sum_{j \in [l]} \nabla_\theta \log \pi_\theta(a_j|s)$ and $\frac{1}{S} \sum_{a \in \mathcal{C}(s)} \nabla_\theta \pi_\theta(a|s)$ and $\sum_{a \in \mathcal{C}(s)} \pi_\theta(a|s)$ respectively.

The full approach is shown in Algorithm 2, with the changes from standard A2C highlighted in red. Recall that, $\theta = (\rho, w)$ is the parameter of the flow-based policy network, ϕ is the parameter of the critic, and ψ is the parameter of the variational distribution network q_ψ for ELBO.

4 Related Work

Large discrete action space. Dulac-Arnold et al. [10] attempt to solve this problem by embedding discrete actions into a continuous space; however, this approach does not work for our unordered

Algorithm 2: IAR-A2C

```
1 Initialize step count  $t \leftarrow 1$ , Initialize episode counter  $E \leftarrow 1$ 
2 repeat
3   Reset gradients:  $d(\theta, \psi) \leftarrow 0$ ,  $d\phi \leftarrow 0$ , Initialize experience set  $\tau$ ,  $t_{start} = t$ , get state  $s_t$ 
4   repeat
5     Execute  $a_t$  according to policy  $\pi'_\theta(a_t|s_t)$  // Defined in Section 3.2
6     Get a number of valid and sampled actions  $l_t$  and  $S_t$ 
7     Receive reward  $r_t$  and new state  $s_{t+1}$ 
8     Add new experience  $\tau \leftarrow \tau \cup \{s_t, a_t, s_{t+1}, r_t\}$ , and update  $t \leftarrow t + 1$ 
9   until terminal  $s_t$  or  $t - t_{start} = t_{max}$ 
10   $R = \begin{cases} 0 & \text{for terminal } s_t \\ V_\phi(s_t) & \text{for non-terminal } s_t \text{ // Bootstrap from last state} \end{cases}$ 
11  for  $i \in \{t - 1, \dots, t_{start}\}$  do
12     $R \leftarrow r_i + \gamma R$ 
13     $\widehat{\log p_{\theta, \psi}(a_i|s_i)} = \alpha \hat{l}_\pi(a_i|s_i) + (1 - \alpha) \hat{l}_\pi^u(a_i|s_i)$  // Sandwich estimator
14    Accumulate gradients w.r.t.  $(\theta, \psi)$ :  $d(\theta, \psi) \leftarrow$ 
       $d(\theta, \psi) + (R - V_\phi(s_i)) \left( \nabla_{\theta, \psi} \widehat{\log p_{\theta, \psi}(a_i|s_i)} - \frac{S_j}{|I_j^2|} \sum_{j \in [I]} \nabla_{\theta, \psi} \widehat{\log p_{\theta, \psi}(a_{ij}|s_i)} \right)$ 
15    Accumulate gradients w.r.t.  $\phi$ :  $d\phi \leftarrow d\phi + \partial(R - V_\phi(s_i))^2 / \partial\phi$ 
16  Perform update of  $\theta$  using  $d\theta$ , of  $\psi$  using  $d\psi$  and of  $\phi$  using  $d\phi$ 
17  Execute ELBO updating by running Algorithm 1 with  $\tau$  as input
18   $E \leftarrow E + 1$ 
19 until  $E > E_{max}$ 
```

discrete space as we reveal by thorough comparison in our experiments. Another type of approach relies on being able to represent a joint probability distribution over multi-dimensional discrete actions using marginal probability distribution over each dimension (also called factored representation). Tang and Agrawal [30] apply this approach to discretized continuous control problems to decrease the learning complexity. Similarly, Delalleau et al. [7] assumes independence among dimensions to model each dimension independently. However, it is well-known in optimization and randomized allocation literature [4, 35] that dimension independence or marginal representation is not valid in the presence of constraints on the support of the probability distribution. Another type of approach [34, 36] converts the choice of multi-dimensional discrete action into a sequential choice of action across the dimensions at each time step (using a LSTM based policy network), where the choice of action in any dimension is conditioned on actions chosen for prior dimensions.

Safe actions. Our approach falls into the field of constrained action space in RL. Situated in the broader safe RL literature [11, 14], these works aim to constrain actions at each step for various purposes, including safety. Notably, we *do not* have an auxiliary cost that we wish to bound, thus, frameworks based on Constrained MDP [11] cannot solve our problem. Also, our state dependent constraints if written as inequalities can be extremely large and hence Lagrangian methods [2] that pull constraints into the objective are infeasible. Most of the methods that we know for constrained actions need the constraints to be written as mathematical inequalities and even then cannot handle state dependent constraints. Although action masking [17] does offer a solution for state-dependent constraints, it necessitates supplementary information in the form of masks. Thus, they do not readily apply to our problem [20]. Some of these methods[9, 24] aim at differentiating through an optimization; these methods are slow due to the presence of optimization in the forward pass and some approaches to make them faster[27] work only with linear constraints; all these approaches scale poorly in the number of inequality constraints.

Normalizing flow policy. The application of normalizing flows in RL has been an emerging field of study with a focus on complex policy modeling, efficient exploration, and control stability [18, 21, 33]. Normalizing flow was integrated into the Soft Actor-Critic framework to enhance the modelling expressivity beyond conventional conditional Gaussian policies and achieve more efficient exploration and higher rewards [21, 33]. In robotic manipulation, Khader et al. [18] proposed to improve the control stability by a novel normalizing-flow control structure. However, these works primarily focus

on continuous actions. To the best of our knowledge, ours is the first work that has successfully incorporated a discrete flow, extending the scope of flow policies into discrete action spaces.

5 Experiments

Through our experiments, we aim to address two main research questions related to our primary contributions: (1) Is the flow-based policy effective in representing categorical actions? (2) Does IAR-A2C offer advantages in constrained action space with oracle constraints? We address these two questions in Sections 5.1 and 5.2 accordingly. Following these, we present an ablation study that examines the effectiveness of various modules in our approach.

We first describe our set-up. We evaluate IAR-A2C against prior works across a diverse set of environments, including low-dimensional discrete control tasks such as CartPole and Acrobot [3], the visually challenging Pistonball task [31] with high-dimensional image inputs and an extremely large action space (upto 59, 049 categories), and an emergency resource allocation simulator in a city, referred to as Emergency Resource Allocation (ERA). CartPole and Acrobot are well-established environments. Pistonball is also a standard environment where a series of pistons on a line needs to move (up, down, or stay) to move a ball from right to left (Figure 3). While the Pistonball environment was originally designed for a multi-agent setting with each piston controlled by a distinct agent, we reconfigure this as a single-agent control of all pistons. This modification presents a challenging task for the central controller as the action space is exponentially large. We show results for Acrobot and three versions of Pistonball: v1 with 3^5 , v2 with 3^8 , and v3 with 3^{10} actions. Results for CartPole are in the appendix.




Figure 3: **Pistonball**, goal is to move the ball to the left border by operating the pistons.

Finally, our custom environment, named ERA, simulates a city that is divided into different districts, represented by a graph where nodes denote districts and edges signify connectivity between them. An action is to allocate a limited number of resources to the nodes of the graph. A tuple including graph, current allocation, and the emergency events is the state of the underlying MDP. The allocations change every time step but an allocation action is subject to *constraints*, namely that a resource can only move to a neighboring node and resources must be located in proximity (e.g. within 2 hops on the graph) as they collaborate to perform tasks. Emergency events arise at random on the nodes, and the decision maker aims to minimize the costs associated with such events by attending to them as quickly as possible. Moreover, we explore a partially observable scenario (with no constraints) in which the optimal allocation is randomized, thus, the next node for a resource is sampled from the probability distribution over neighboring nodes that the stochastic policy represents (see appendix for set-up details). We show results for five versions of ERA: ERA-Partial with 9 actions and partial observability in unconstrained scenario, while v2 with 7^3 actions, v3 with 8^3 actions, v4 with 9^3 actions, and v5 with 10^3 actions in constrained scenario. The results for ERA-v1 are in Appendix C.

Benchmarks: In the unconstrained setting, we compare our approach to Wol-DDPG [10], A2C [22], factored representation (Factored, discussed in Section 4), and autoregressive approach (AR, discussed in Section 4). Wol-DDPG is chosen as it is designed to handle large discrete action spaces without relying on the dimension independence assumption. In the oracle constraints setting, we compare our method to action masking (MASK) [17], which determines the action mask by querying all actions within the action space. As far as we know action masking is currently the only existing approach for constraining action with oracle constraints, we also include a comparison with IAR-augmented AR (AR+IAR) which is able to handle the constraints by utilizing our invalid action rejection technique, as well as a comparison with Wol-DDPG to demonstrate the performance of a method that cannot enforce the constraints.

5.1 Learning in Categorical Action Space without Constraints

In the Acrobot environments, we use A2C as the benchmark because this task is relatively simple and A2C can easily learn the optimal policy. For the Pistonball environment, we consider both Wol-DDPG and A2C as benchmarks. The results are displayed in Figure 4. We also conduct experiments on CartPole and ERA-v5 with all constraints removed, which can be found in Appendix C.




Figure 4: Learning curves over time in (a) **Acrobot** and (b) **Pistonball-v[1-3]** (without constraints). All approaches are trained with 5 seeds, except for Wol-DDPG, trained with 3 seeds due to its expensive training cost on Pistonball.




Figure 5: Learning curves in **ERA-v[2-5]** (with constraints). All settings are trained with 5 seeds.

The results of Acrobot demonstrate that the flow-based policy exhibits comparable performance to the optimal policy (A2C), which also highlights the effectiveness of our sandwich estimator for $\log \pi(a|s)$. In more challenging environments with high-dimensional image observations and extremely large action spaces, our model has comparable performance to Factored and AR, while significantly outperforming A2C and Wol-DDPG. Interestingly, we observe that Wol-DDPG struggles to learn even in the simplest Pistonball environment, despite being designed for large discrete action spaces. We hypothesize that Wol-DDPG might function properly if the actions are discrete yet ordered, as demonstrated in the quantized continuous control task presented in the Wol-DDPG paper [10].

On ERA-Partial, we demonstrate the known fact that the optimal policy in environments with partial observability may be a stochastic one [28]. We compare with the factored approach (Figure 6). In this environment, the optimal stochastic policy is a joint distribution that is not factorizable into a product of independent distributions over each dimension. Thus, the factored approach cannot effectively represent the joint distribution due to its independence assumption among dimensions. The experiment results show our approach significantly outperforms the factored approach. We further present a smaller example in the appendix which shows that the inability to represent an arbitrary joint distribution makes the factored approach extremely suboptimal in partial observable settings.




Figure 6: Learning curves in **ERA-Partial**. Our approach obtains higher return than Factored approach.

5.2 Learning in Categorical Action Space with State-Dependent Constraints

We now address the second question about constraint enforcement by IAR-A2C. The results are in Figure 5. We observe that our approach demonstrates better or comparable performance compared to the benchmarks. AR is known as a strong baseline for environments with large discrete action space, but surprisingly, it performs poorly. We hypothesize this is due to the case that the autoregressive model does not have a sense of the constraints of the remaining dimensions when it outputs the action for the first dimension, thereby producing first dimension actions that may be optimal without

constraints but are suboptimal with constraints. Detailed analysis and experimental evidence to support our hypothesis are provided in Appendix D.

Also, action masking achieves its performance by querying all actions of the entire action space, whereas our approach only requires querying a batch of actions, which is substantially smaller (e.g., 64 versus 1,000 for ERA-v5). Thus, while Figure 5 shows IAR-A2C taking more iteration for convergence, the same figure when drawn with the x-axis as wall clock time in Figure 7 (shown only for v4 here, others are in appendix C) shows an order of magnitude faster convergence in wall clock time. Another critical property of our approach is the guaranteed absence of constraint violations, similar to the action masking method. However, while action masking demands the full knowledge of the validity of all actions, our method merely requires the validity of the sampled actions within a batch. Note that Wol-DDPG can and does violate constraints during the learning process. Further details regarding the action violation of Wol-DDPG are provided in the appendix C.




Figure 7: Learning curves for wall clock time. Our approach converges much faster than action masking.

5.3 Ablation Study




Figure 8: (a) Ablation of gradient correction and sandwich estimator (b) Ablation of posterior type.

the effects if we use only the ELBO estimator for log probability of action instead of our sandwich estimator. We find that the ELBO estimator is also prone to a reduction in valid actions in Figure 8a and unstable learning as a consequence, similar to the observation when using the standard policy gradient.

Posterior Type: The posterior $q(z|a, s)$ can be modeled by a conditional Gaussian or a normalizing flow. In our experiments, we discover that modelling q with a flow posterior is crucial for learning, as it can approximate the true posterior more effectively than a Gaussian, as seen in Figure 8b.

6 Conclusion and Limitations

We introduced a novel discrete normalizing flow based architecture and an action rejection approach to enforce constraints on actions in order to handle categorical action spaces with state dependent oracle constraints in RL. Our approach shows superior performance compared to available baselines, and we analyzed the importance of critical modules of our approach. A limitation of our method is in scenarios where the fraction of valid actions out of all actions is very small, and hence our sampling based rejection will need a lot of samples to be effective, making training slower. This motivates future work on improved sampling; further, better estimation of the log probability of actions is also a promising research direction for improved performance.

We conduct ablation studies of IAR-A2C on ERA-v4 to investigate the effect of various choices of modules and estimators in our approach.

Policy Gradient: We compare the performance of approaches using the policy gradient provided by Theorem 1 (gradient correction) and the original policy gradient of A2C (standard policy gradient), while still applying invalid action rejection. We observe in Figure 8a that the number of valid actions in the batch decreases rapidly, and the program may crash if no valid actions are available.

Sandwich Estimator: We examine

Acknowledgement

This research is supported by the National Research Foundation Singapore under its AI Singapore Programme (Award Number: AISG2-RP-2020-016). Dr. Nguyen was supported by grant W911NF-20-1-0344 from the US Army Research Office.

References

- [1] Abhinav Bhatia, Pradeep Varakantham, and Akshat Kumar. Resource constrained deep reinforcement learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 610–620, 2019.
- [2] Steven Bohez, Abbas Abdolmaleki, Michael Neunert, Jonas Buchli, Nicolas Heess, and Raia Hadsell. Value constrained model-free continuous control. *ArXiv preprint*, abs/1902.04623, 2019. URL <https://arxiv.org/abs/1902.04623>.
- [3] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *ArXiv preprint*, abs/1606.01540, 2016. URL <https://arxiv.org/abs/1606.01540>.
- [4] Eric Budish, Yeon-Koo Che, Fuhito Kojima, and Paul Milgrom. Designing random allocation mechanisms: Theory and applications. *American economic review*, 103(2):585–623, 2013.
- [5] Yuri Burda, Roger B. Grosse, and Ruslan Salakhutdinov. Importance weighted autoencoders. In Yoshua Bengio and Yann LeCun, editors, *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*, 2016. URL <http://arxiv.org/abs/1509.00519>.
- [6] Jie Chen and Ronny Luss. Stochastic gradient descent with biased but consistent gradient estimators. *ArXiv preprint*, abs/1807.11880, 2018. URL <https://arxiv.org/abs/1807.11880>.
- [7] Olivier Delalleau, Maxim Peter, Eloi Alonso, and Adrien Logut. Discrete and continuous action representation for practical rl in video games. *ArXiv preprint*, abs/1912.11077, 2019. URL <https://arxiv.org/abs/1912.11077>.
- [8] Adji Bousso Dieng, Dustin Tran, Rajesh Ranganath, John Paisley, and David Blei. Variational inference via *chi* upper bound minimization. *Advances in Neural Information Processing Systems*, 30, 2017.
- [9] Priya L. Donti, David Rolnick, and J. Zico Kolter. DC3: A learning method for optimization with hard constraints. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=V1ZHVxJ6dSS>.
- [10] Gabriel Dulac-Arnold, Richard Evans, Hado van Hasselt, Peter Sunehag, Timothy Lillicrap, Jonathan Hunt, Timothy Mann, Theophane Weber, Thomas Degris, and Ben Coppin. Deep reinforcement learning in large discrete action spaces. *ArXiv preprint*, abs/1512.07679, 2015. URL <https://arxiv.org/abs/1512.07679>.
- [11] Javier Garcia and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16(1):1437–1480, 2015.
- [12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020.
- [13] Sergei Pavlovich Grachev, Aleksei Aleksandrovich Zhilyaev, Vladimir Borisovich Laryukhin, Dmitrii Evgen’evich Novichkov, Vladimir Andreevich Galuzin, Elena V Simonova, IV Maiy-ov, and Peter O Skobelev. Methods and tools for developing intelligent systems for solving complex real-time adaptive resource management problems. *Automation and Remote Control*, 82:1857–1885, 2021.

- [14] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, Yaodong Yang, and Alois Knoll. A review of safe reinforcement learning: Methods, theory and applications. *ArXiv preprint*, abs/2205.10330, 2022. URL <https://arxiv.org/abs/2205.10330>.
- [15] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/4c5bcfec8584af0d967f1ab10179ca4b-Abstract.html>.
- [16] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 12454–12465, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/67d96d458abdef21792e6d8e590244e7-Abstract.html>.
- [17] Shengyi Huang and Santiago Ontañón. A closer look at invalid action masking in policy gradient algorithms. In *The International FLAIRS Conference Proceedings*, volume 35, 2022.
- [18] Shahbaz Abdul Khader, Hang Yin, Pietro Falco, and Danica Kragic. Learning stable normalizing-flow control for robotic manipulation. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1644–1650. IEEE, 2021.
- [19] Dexun Li, Meghna Lowalekar, and Pradeep Varakantham. Claim: Curriculum learning policy for influence maximization in unknown social networks. In *Uncertainty in Artificial Intelligence*, pages 1455–1465. PMLR, 2021.
- [20] Jyun-Li Lin, Wei Hung, Shang-Hsuan Yang, Ping-Chun Hsieh, and Xi Liu. Escaping from zero gradient: Revisiting action-constrained reinforcement learning via frank-wolfe policy optimization. In Cassio P. de Campos, Marloes H. Maathuis, and Erik Quaeghebeur, editors, *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence, UAI 2021, Virtual Event, 27-30 July 2021*, volume 161 of *Proceedings of Machine Learning Research*, pages 397–407. AUAI Press, 2021. URL <https://proceedings.mlr.press/v161/lin21b.html>.
- [21] Bogdan Mazouze, Thang Doan, Audrey Durand, Joelle Pineau, and R Devon Hjelm. Leveraging exploration in off-policy algorithms via normalizing flows. In *Conference on Robot Learning*, pages 430–444. PMLR, 2020.
- [22] Volodymyr Mnih, Adrià Puigdomènech Badia, Mehdi Mirza, Alex Graves, Timothy P. Lillicrap, Tim Harley, David Silver, and Koray Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In Maria-Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1928–1937. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/mniha16.html>.
- [23] Ayan Mukhopadhyay, Geoffrey Pettet, Sayyed Mohsen Vazirizade, Di Lu, Alejandro Jaimes, Said El Said, Hiba Baroud, Yevgeniy Vorobeychik, Mykel Kochenderfer, and Abhishek Dubey. A review of incident prediction, resource allocation, and dispatch models for emergency management. *Accident Analysis & Prevention*, 165:106501, 2022.
- [24] Tu-Hoa Pham, Giovanni De Magistris, and Ryuki Tachibana. Optlayer-practical constrained optimization for deep reinforcement learning in the real world. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6236–6243. IEEE, 2018.
- [25] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL <http://jmlr.org/papers/v22/20-1364.html>.

- [26] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 1530–1538. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/rezende15.html>.
- [27] Sanket Shah, Arunesh Sinha, Pradeep Varakantham, Andrew Perrault, and Milind Tambe. Solving online threat screening games using constrained action space reinforcement learning. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, February 7-12, 2020*, pages 2226–2235. AAAI Press, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/5599>.
- [28] Satinder P Singh, Tommi Jaakkola, and Michael I Jordan. Learning without state-estimation in partially observable markovian decision processes. In *Machine Learning Proceedings 1994*, pages 284–292. Elsevier, 1994.
- [29] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=PxtTIG12RRHS>.
- [30] Yunhao Tang and Shipra Agrawal. Discretizing continuous action space for on-policy optimization. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 5981–5988. AAAI Press, 2020. URL <https://aaai.org/ojs/index.php/AAAI/article/view/6059>.
- [31] Justin K. Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S. Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, Niall L. Williams, Yashas Lokesh, and Praveen Ravi. Pettingzoo: Gym for multi-agent reinforcement learning. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 15032–15043, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/7ed2d3454c5eea71148b11d0c25104ff-Abstract.html>.
- [32] Ivan B Vermeulen, Sander M Bohte, Sylvia G Elkhuisen, Han Lameris, Piet JM Bakker, and Han La Poutré. Adaptive resource allocation for efficient patient scheduling. *Artificial intelligence in medicine*, 46(1):67–80, 2009.
- [33] Patrick Nadeem Ward, Ariella Smofsky, and Avishek Joey Bose. Improving exploration in soft-actor-critic with normalizing flows policies. *ArXiv preprint*, abs/1906.02771, 2019. URL <https://arxiv.org/abs/1906.02771>.
- [34] Mason Wright, Yongzhao Wang, and Michael P Wellman. Iterated deep reinforcement learning in games: History-aware training for improved stability. In *EC*, pages 617–636, 2019.
- [35] Haifeng Xu. The mysteries of security games: Equilibrium computation becomes combinatorial algorithm design. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 497–514, 2016.
- [36] Yiming Zhang, Quan Ho Vuong, Kenny Song, Xiao-Yue Gong, and Keith W Ross. Efficient entropy for policy gradient with multidimensional action space. *ArXiv preprint*, abs/1806.00589, 2018. URL <https://arxiv.org/abs/1806.00589>.

A Proofs and Derivation

A.1 Proof for Theorem 1

The following sequence of equations show that proof, relying on the fact that $\pi'(a|s) = \frac{\pi_\theta(a|s)}{\sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s)}$. We start with the standard policy gradient for any policy π' , shown in the first line below, and then replace $\pi'(a|s) = \frac{\pi_\theta(a|s)}{\sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s)}$ in the second line, followed by standard manipulation of the log function.

$$\nabla_\theta J(\theta) = \mathbb{E}_{\pi'} \left[Q^{\pi'}(s, a) \nabla_\theta \log \pi'(a|s) \right] \quad (8)$$

$$= \mathbb{E}_{\pi'} \left[Q^{\pi'}(s, a) \nabla_\theta \log \frac{\pi_\theta(a|s)}{\sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s)} \right] \quad (9)$$

$$= \mathbb{E}_{\pi'} \left[Q^{\pi'}(s, a) \nabla_\theta \log \pi_\theta(a|s) - Q^{\pi'}(s, a) \nabla_\theta \log \sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s) \right] \quad (10)$$

$$= \mathbb{E}_{\pi'} \left[Q^{\pi'}(s, a) \nabla_\theta \log \pi_\theta(a|s) - Q^{\pi'}(s, a) \frac{\sum_{a_i \in \mathcal{C}(s)} \nabla_\theta \pi_\theta(a_i|s)}{\sum_{a_i \in \mathcal{C}(s)} \pi_\theta(a_i|s)} \right] \quad (11)$$

A.2 Proof for Lemma 1

Fix state s and consider a function $F(a) = \begin{cases} \frac{\nabla_\theta \pi_\theta(a|s)}{\pi_\theta(a|s)} & \text{for } a \in \mathcal{C}(s) \\ 0 & \text{otherwise} \end{cases}$. Then,

$$\mathbb{E}_\pi[F(a)] = \sum_{a \in \mathcal{C}(s)} \nabla_\theta \pi_\theta(a|s)$$

Thus, if we obtain a sample average estimate for $\mathbb{E}_\pi[F(a)]$ then it is an unbiased estimate for $\sum_{a \in \mathcal{C}(s)} \nabla_\theta \pi_\theta(a|s)$. For S samples from π with l being valid samples, the sample average estimate for $\mathbb{E}_\pi[F(a)]$ is $\frac{1}{l} \sum_{j \in [l]} \nabla_\theta \log \pi_\theta(a_j|s)$.

Similarly, for the next estimate, consider a function $G(a) = \begin{cases} 1 & \text{for } a \in \mathcal{C}(s) \\ 0 & \text{otherwise} \end{cases}$. Clearly, then $\mathbb{E}_\pi[G(a)] = \sum_{a \in \mathcal{C}(s)} \pi_\theta(a|s)$ and a sample average estimate of $\mathbb{E}_\pi[G(a)]$ is $\frac{l}{S}$.

A.3 Soft Threshold Function

In Argmax Flow [16], a threshold function was introduced to enforce the argmax constraints, i.e. the variational distribution $q(\mathbf{v}|\mathbf{x})$ should have support limited to $\mathcal{S}(\mathbf{x}) = \{\mathbf{v} \in \mathbb{R}^{D \times K} : \mathbf{x} = \arg \max \mathbf{v}\}$. The thresholding-based $q(\mathbf{v}|\mathbf{x})$ was defined by Alg. 3 in Argmax Flow. However, the formula to evaluate $\det d\mathbf{v}/d\mathbf{u}$ is not given, which is essential when estimating ELBO \hat{l}_π and CUBO \hat{l}_π^u . We derive it here.

Let's follow the notations in Alg. 3 of Argmax Flow. Suppose index i is the one that we want to be largest (i is a fixed index). The soft threshold function is given by

$$v_j = u_i - \log(1 + e^{u_i - u_j})$$

Note that the threshold $T = v_x = u_x$ (we cannot use v_x to define v itself, so T is u_x). Then,

- if $j = i$ then $v_j = u_i$, $\frac{\partial v_j}{\partial u_k} = 1$
- if $k \neq j$ or $k \neq i$ then $\frac{\partial v_j}{\partial u_k} = 0$
- if $k = i$ then $\frac{\partial v_j}{\partial u_k} = 1 - \frac{1}{1 + e^{u_i - u_j}} \times e^{u_i - u_j}$
- if $k = j$ then $\frac{\partial v_j}{\partial u_k} = \frac{1}{1 + e^{u_i - u_j}} \times e^{u_i - u_j}$

Table A.1: Various Types of Alpha α

Types of α	Remark
Static α	α is fixed to be 0.5
Trainable α	α is a trainable parameter, updated by the policy gradient
Adaptive $\alpha(\hat{l}_\pi, \hat{l}_\pi^u)$	α is conditioned on the ELBO and the CUBO


$\det dv/du$ is a $K \times K$ determinant, where only the elements on the diagonal and on the column i is non-zero, other elements are zero. We can unfold the determinant by the i -th row. Finally, we have

$$\det dv/du = \prod_{j=1, j \neq i}^K \frac{\partial v_j}{\partial u_j} = \prod_{j=1, j \neq i}^K \text{sigmoid}(u_i - u_j)$$

B Experimental Details

B.1 Effect of the Sandwich Estimator’s Weight α

The value of α plays a crucial role in estimating the log probability, which in turn impacts the performance of the model. Various types of α are presented in Table A.1. Through our experiments, we have observed that an adaptive α yields greater stability and, in certain environments, leads to improved performance. This is illustrated in Figure A.1, where the adaptive α exhibits superior performance compared to the static value of $\alpha = 0.5$.

Figure A.1: Performance with adaptive α and static α in **ERA-v4**.

B.2 Effect of Invertible Functions F_w

We have explored different types of invertible function F_w (called latent flow model) in our study, including affine coupling bijections², as well as more advanced models such as AR Flow and Coupling Flow, as described in section B.1 of Argmax Flow [16]. The AR Flow and Coupling Flow methods offer enhanced capabilities for modeling complex distributions, and they have been successfully applied to language modeling tasks within the Argmax Flow framework. However, through our experiments, we have observed that even a simple latent flow model is sufficient for achieving good performance and exhibits faster convergence. We attribute this finding to the fact that the increased parameterization in AR Flow requires a larger amount of training data to effectively learn.

Table A.2: Optimization details

Environment	n_envs	lr	Optimizer	batch size
CartPole	8	3e-4	RMSprop	256
Acrobot	8	3e-4	RMSprop	256
ERA-v5 w/o cstr	64	3e-4	RMSprop	512
Pistonball-v[1-3] w/o cstr	64	3e-4	RMSprop	512
ERA-v[1-5] w/ cstr	64	3e-4	RMSprop	256
Pistonball-v[1-2] w/ cstr	64	3e-4	RMSprop	512

B.3 Optimization Details

The total timesteps of training for each environment are determined based on the convergence of our model and benchmarks. Typically, we train each setting using 5 different random seeds, unless

²We exploit the implementation from: https://github.com/didriknielsen/survae_flows




Figure A.2: **(a)** Learning curves in **CartPole** and **ERA-v5** (without constraints); **(b)** Learning curves in **ERA-v1** (with constraints).




Figure A.3: Learning curves in **Toy-Partial**.

otherwise specified. When evaluating the model’s performance at a specific timestep with a specific seed, we employ a separate set of 10 testing environments and report the mean return over these environments. Further details can be found in Tables A.2. Note that `n_envs` denotes the number of environments running in parallel, `lr` denotes the learning rate, and `batch_size` refers to the batch size when we execute ELBO updating (Algorithm 1 in the main paper). Furthermore, we will make the code used to reproduce these results publicly available.

B.4 Range of Considered hyperparameters

We conducted experiments varying the number of samples used for estimating $\log \pi(a|s)$, specifically considering the values $\{1, 2, 4, 8\}$, as well as the inclusion of reward normalization. We find that using 2 or 4 samples generally leads to good performance across most of our experiments.

B.5 Network structure

Our implementation is built on Stable-Baseline3 [25]. In different environments, different state encoders were exploited. We used MLP encoder for discrete control tasks and CNN encoder for Pistonball task. In ERA environment, a customized state encoder was applied to handle the graph state based on the implementation from [19].

B.6 Computational resources

Experiments were run on NVIDIA Quadro RTX 6000 GPUs, CUDA 11.0 with Python version 3.8.13 in Pytorch 1.11.

C Additional Experiment

In this section, we present additional experimental results obtained from our study.




Figure A.4: Learning curves for wall clock time.

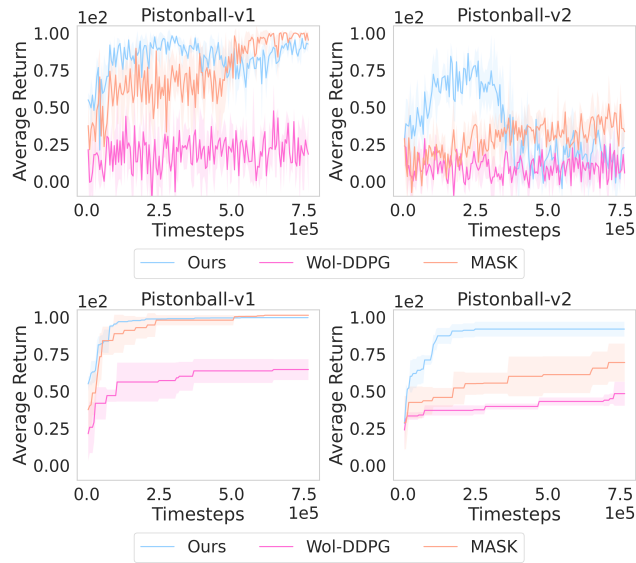


Figure A.5: Top: Learning curves in **Pistonball-v1** and **v2** (with constraints); Bottom: Best-till-now returns in **Pistonball-v1** and **v2**.




Figure A.6: Valid action ratio of Wol-DDPG during training.

In the **non-constrained setting**, we conduct experiments on CartPole, ERA-v5, and a toy environment with partial observability - Toy-Partial. For ERA-v5, we remove all constraints, allowing for the allocation of resources in any nodes. Figure A.2a illustrates that our approach achieves comparable performance to the best benchmark when the action space is not excessively large. Toy-Partial is built exactly on the example described in Section 3.1 of [28], in which the optimal stochastic policy can be arbitrarily better than any deterministic policy. We only modify the setting by augmenting the dimensionality of actions. There are two actions A and B in their example, while we use a 2-dimensional representation ((0, 1) representing action A, (1, 0) representing action B, (1, 1) and (0, 1) staying at the current state.) to simulate the case in multi-dimensional action space. Figure A.3 shows that our approach can perform significantly better than Factored approach in Toy-Partial.

In the **constrained setting**, we perform additional experiments on ERA-v1, Pistonball-v1, and Pistonball-v2. In the Pistonball environment, we introduce a constraint that restricts the upward movement of pistons on the left side to ensure the ball continues rolling to the left. Our experiments demonstrate that this constraint presents significant difficulty when the action space is large.

Regarding ERA-v1, our model outperforms the benchmark, as depicted in Figure A.2b. Additionally, we analyze the average return over wall clock time for ERA-v1, v2, v3, and v5, and our model exhibits an order of magnitude faster convergence, as illustrated in Figure A.4.

For Pistonball, our model demonstrates comparable performance in the smaller environment (Pistonball-v1) and superior performance in the larger environment (Pistonball-v2) where the benchmarks struggle to effectively learn. These findings are presented in Figure A.5, which includes the average return over timesteps in the top plot and the best-till-now returns in the bottom plot. By best-till-now we mean the best evaluated return till the current timestamp, which is a commonly used metric particularly when the return does not increase monotonically over time steps and hence the best model might be an intermediate one.

We also observed that our approach can learn a stochastic optimal policy in ERA environment, which corresponds to our motivation that a stochastic policy is preferred in many resource allocation problems. For example, we have observed that in ERA-v4, a stochastic optimal allocation was learned by our approach in a given state (shown as a distribution of 200 sampled actions (Table A.3)).

Finally, we investigate the **constraint violation of Wol-DDPG** across ERA-v1 to v5 in Figure A.6. Our observations reveal that the constraint violation of Wol-DDPG can be significant, with a valid action ratio reaching only 15% in the worst-case scenario. Note that the valid action ratio mentioned here specifically pertains to the ratio of valid actions generated by the agent (the action output by the policy at each timestep) to all the actions output in a single episode, which differs from the valid action rate discussed in the ablation study in the main paper. In the ablation study shown in Figure 8a, the valid action rate refers to the fraction (l/S) of valid actions l within the S samples at a particular timestep in IAR framework shown in Figure 1. Thus, valid action rate is metric specific to our framework; note that valid action ratio for our framework is always 1 as IAR always outputs valid actions.

D AR in Constrained Scenario

We have observed that AR approach does not perform well in the constrained scenario. We give our analysis and experimental evidence here.

First, we describe our specific ERA set-up. We are aiming to allocate 3 resources to 9 areas with the 9 areas lying on a graph (we do not show the graph here). The constraints are that in any allocation, R2 and R3 must be within 2 hops (inclusive) of each other.

Each allocation of resource R1, R2, R3 is given as a vector, e.g., (4, 0, 3) is the allocation of R1 to area 4, R2 to area 0, R3 to area 3. In an AR approach there is a dimension dependency in the policy network. Here allocation of R2 (output by a neural network, which we call R2 network) depends on allocation of R1 and allocation of R3 (output of R3 network) depends on allocation of R2 and R1. The R1 network, which outputs location of R1, is conditionally independent (note that R1, R2, R3 networks use the shared weights).

Shown below is an optimal allocation learned by **our approach** in state 1 (current allocation (8, 2, 8)), shown as a distribution of 200 sampled actions (Table A.3).

Table A.3: Optimal policy (ours) in state 1. Showing actions with top-3 probability.

Action	Number	Distribution
(4, 0, 3)	100	0.5
(4, 0, 6)	100	0.5
(0, 0, 0)	0	0.0
...

Shown below is the best allocation learned by the **AR approach** in state 1, shown as a distribution of 200 sampled actions (Table A.4).

Table A.4: AR policy in state 1.

Action	Number	Distribution
(3, 6, 6)	173	0.865
(3, 3, 6)	17	0.085
(6, 3, 6)	5	0.025
...

However, the above is not optimal in state 1. But, if we fix R1 to be in area 4 and R2 to be in area 0 and provide that as forced input to the R3 network, then we get the Table A.5 from the AR network.

Table A.5: AR policy in state 1 after fixing R1 in area 4 and R2 in area 0.

Action	Number	Distribution
(4, 0, 6)	164	0.82
(4, 0, 3)	36	0.18
(0, 0, 0)	0	0.00
...

This shows that R3 network has learned a good policy, but the R1 network is unable to independently produce the area 4 that could trigger the optimal output. We conjecture this could be due to fact that R1 is not aware of constraints since the constraint can be solely enforced by R2 network and R3 network. Hence, R1 might be outputting areas that can be optimal if there were no constraints. We have further evidence of the same happening in another state:

Shown below is an optimal allocation learned by **our approach** in state 2 (current allocation (1, 7, 5)), shown as a distribution of 200 sampled actions (Table A.6).

Table A.6: Optimal policy (ours) in state 2.

Action	Number	Distribution
(4, 0, 3)	200	1.0
(4, 0, 6)	0	0.0
(0, 0, 0)	0	0.0
...

Shown in Table A.7 is the best allocation learned by the **AR approach** in state 2, shown as a distribution of 200 sampled actions. Again, if we fix R1 to be in area 4 and R2 to be in area 0 and provide that as forced input to the R3 network, then we get the Table A.8 from the AR network. This supports our claim that the R1 network finds it difficult to reason about constraints.

While this is not a complex example and may be resolved through another ordering, we wish to highlight that constraints can also be complex and they will always introduce issues with ordering dependent approaches. To handle such issues, one may need knowledge on the dimension dependency, or make efforts on trying different orders of generating the allocations.

Table A.7: AR policy in state 2.

Action	Number	Distribution
(3, 6, 3)	81	0.405
(3, 6, 6)	81	0.405
(6, 6, 3)	16	0.008
...

Table A.8: AR policy in state 2.

Action	Number	Distribution
(4, 0, 3)	197	0.985
(4, 0, 1)	2	0.010
(4, 0, 2)	1	0.005
...

E Details of Emergency Resource Allocation Environment

We describe the details of our custom environment ERA in this section. This environment simulates a city that is divided into different districts, represented by a graph where nodes denote districts and edges signify connectivity between them. An action is to allocate a limited number of resources to the nodes of the graph. A tuple including graph, current allocation, and the emergency events is the state of the underlying MDP. The allocations change every time step but an allocation action is subject to constraints, namely that a resource can only move to a neighboring node and resources must be located in proximity (e.g. within 2 hops on the graph) as they collaborate to perform tasks. Emergency events arise at random on the nodes and the decision maker aims to minimize the costs associated with such events by attending to them as quickly as possible. Finally, the optimal allocation is randomized, thus, the next node for a resource is sampled from the probability distribution over neighboring nodes that the stochastic policy represents.

Each version of the ERA environment is characterized by an adjacency matrix that defines the connectivity of districts within the simulated city and a cost matrix that quantifies the expenses associated with traversing from one node to another. The agent’s performance is evaluated based on the successful resolution of emergency events, leading to rewards, while penalties are incurred for failure to address the emergency. The agent’s utility at each timestep encompasses the reward (or penalty) and the negative moving costs. To increase the complexity of the task, we introduce different types of emergency events. Each event type follows a fixed distribution over the nodes, and the event type itself is determined by a categorical distribution. For example, a distribution of [0.3, 0.7] means that 30% of the events are of type 1 and 70% are of type 2. The specifics of each version are presented in Table A.9, where the columns # rsc, # nodes, and hops represent the number of resources, the number of nodes in the graph, and the maximal allocation distance between two resources, respectively.

ERA-Partial is one setting with partial observability, which has three states but one possible observation. The RL agent is encouraged to change its (unobserved) state by obtaining a reward, otherwise it obtains a penalty. To perform well, the RL agent needs to perform stochastically. The ERA implementation, including configuration file that consists of the adjacency matrix, cost matrix and other relevant parameters, will be made available for reproducibility purposes.

Table A.9: Specifics for ERA

Version	# rsc	# nodes	hops
partial	2	3	1
v1	3	6	1
v2	3	7	1
v3	3	8	2
v4	3	9	2
v5	3	10	3