
Autoregressive Policy Optimization for Constrained Allocation Tasks

David Winkel* Niklas Strauß*
Maximilian Bernhard Zongyue Li Thomas Seidl Matthias Schubert
Munich Center for Machine Learning, LMU Munich
{winkel,strauss,bernhard,li,seidl,schubert}@dbs.ifi.lmu.de

Abstract

Allocation tasks represent a class of problems where a limited amount of resources must be allocated to a set of entities at each time step. Prominent examples of this task include portfolio optimization or distributing computational workloads across servers. Allocation tasks are typically bound by linear constraints describing practical requirements that have to be strictly fulfilled at all times. In portfolio optimization, for example, investors may be obligated to allocate less than 30% of the funds into a certain industrial sector in any investment period. Such constraints restrict the action space of allowed allocations in intricate ways, which makes learning a policy that avoids constraint violations difficult. In this paper, we propose a new method for constrained allocation tasks based on an autoregressive process to sequentially sample allocations for each entity. In addition, we introduce a novel de-biasing mechanism to counter the initial bias caused by sequential sampling. We demonstrate the superior performance of our approach compared to a variety of Constrained Reinforcement Learning (CRL) methods on three distinct constrained allocation tasks: portfolio optimization, computational workload distribution, and a synthetic allocation benchmark. Our code is available at: <https://github.com/niklasdbs/paspo>.

1 Introduction

Continuous allocation tasks are a class of problems where an agent needs to distribute a limited amount of resources over a set of entities at each time step. Many complex real-world problems are formulated as allocation tasks, and state-of-the-art solutions rely on using Reinforcement Learning (RL) to learn effective policies [6, 26, 3, 20, 27]. Notable examples include portfolio allocation tasks, where portfolio managers must allocate the available financial resources among various assets [27], or allocation tasks of computational workloads to a set of compute instances in data centers [3]. In many cases, allocation tasks come with allocation constraints [6, 20, 27, 26], such as investing at most 30 % of the portfolio into a specific subset of the assets or to restrict the maximum workload to certain servers in a data center. Formally, allocation constraints are expressed as linear constraints and form a system of linear inequalities, geometrically describing a convex polytope. Each point in this polytope describes a possible allocation and each dimension corresponds to one of the entities. Allocation tasks often require hard constraints, i.e., constraints that are explicitly given and must be satisfied at any point in time. However, most of the existing CRL literature focuses on soft constraints that are not explicitly given [2, 29, 31, 14, 25]. These approaches typically cannot guarantee constraint satisfaction and tend to have many constraint violations during training. The majority of these methods approximate the cumulative costs of constraint violations and optimize the cumulative reward while trying to adhere to the maximum cumulative costs. While less explored, there exist

*Both authors contributed equally.

several techniques that ensure the satisfaction of hard constraints [18, 6, 11, 10, 20]. These approaches might generate actions that do not satisfy the constraints but utilize a correction mechanism to map the actions back into the valid action space. In addition, most of these approaches are restricted to off-policy algorithms [11, 20]. In another line of research, solutions tailored for constrained allocation tasks have been proposed [27, 28]. However, these solutions are severely limited since they can only handle a specific subset of linear constraints and cannot handle more than two.

In this paper, we propose Polytope Action Space Policy Optimization (PASPO), a novel RL-based method that, firstly, decomposes the action space into several dependent sub-problems and, secondly, autoregressively computes the allocations step-by-step for each entity individually. In contrast to previous methods for hard constraints, we directly generate an action within the action space. This makes the correction of invalid actions unnecessary and, thus, avoids potential sampling bias introduced by the correction. Our new decomposition approach is implemented in a neural network-based policy function, which can be employed in on-policy and off-policy RL algorithms. We show that initialization bias can prevent proper exploration in early training which leads to premature convergence. Thus, we propose a de-biasing mechanism to stabilize exploration in early training stages.

We evaluate our approach against various baselines on three distinct allocation tasks: portfolio optimization, distributing computational workload in data centers, and a synthetic benchmark. These experiments demonstrate that our approach can outperform existing methods consistently and show the importance of the proposed de-biasing mechanism.

To summarize, the main contributions of our paper are:

- A new autoregressive stochastic policy function applicable to arbitrary convex polytope action spaces of constrained allocation tasks.
- A new de-biasing mechanism to prevent premature convergence to a sub-optimal policy.
- An empirical evaluation that optimizes our new policy function using PPO [22] and demonstrates improved results compared to state-of-the-art CRL methods.

The remainder of the paper is structured as follows: In Section 2, we provide an overview of the related work in CRL, constrained allocation tasks, and autoregressive policy functions. Afterward, we formalize constrained allocation tasks in Section 3 and present our novel approach in Section 4. Section 5 describes the results of our experimental evaluation, Section 6 briefly discusses limitations and future work before Section 7 concludes the paper.

2 Related Work

Resource allocation tasks are a widely researched area with numerous applications spanning logistics, power distribution, computational load balancing, security screening, and finance [6, 26, 3, 20, 27]. We identify three key research directions that are particularly important when discussing resource allocation tasks.

Safe Reinforcement Learning The majority of work in CRL addresses soft constraints, a setting often referred to as Safe RL. We will provide a brief overview of the most important methods in this field. For a more comprehensive examination of Safe RL, we direct readers to the survey papers by [15, 12]. A common technique in Safe RL is the use of Lagrangian relaxation [4, 15]. Several works employ primal-dual optimization to leverage the Lagrangian duality, including [9, 23, 13]. Another frequently used approach involves different penalty terms [25, 14, 31]. The authors of IPO [14] propose to use logarithmic barrier functions. CPO [2] extends TRPO [21] to ensure near-constraint satisfaction with each update. Additionally, two-step approaches such as FOCOPS [32] and CUP [30] are popular in the field. However, unlike our method, these approaches do not guarantee strict constraint satisfaction, particularly during training.

Hard Constraints Although less studied than Safe RL, several works address hard instantaneous constraints on actions to ensure full constraint satisfaction at any time step. Most of these approaches employ mechanisms to correct infeasible actions, i.e., those that violate constraints, into feasible actions [18, 6, 20, 11]. In contrast, our method always generates feasible actions without the need for correction. OptLayer [18] is one of the most prominent examples in this field, which employs OptNet [5] to map infeasible actions to the nearest feasible action. Similarly, [20] propose a more

efficient projection on the polytope action space than OptLayer. The authors of [6] focus on resource allocation with hierarchical allocation constraints by proposing a faster approximate version of OptLayer. In [11], the authors propose an off-policy algorithm based on the generalized reduced gradient method [1] to handle non-linear hard constraints by projecting infeasible actions. In contrast, our method is not limited to off-policy algorithms.

In [27, 28], the action space is decomposed into independent subspaces. However, these approaches can only handle up to two allocation constraints. Furthermore, they are only applicable to binary allocation constraints. In contrast, our approach can handle an arbitrary number of constraints as well as any type of linear allocation constraints.

Action Space Decomposition/Factorization The decomposition or factorization of multi-dimensional action spaces has been examined in several works [24, 16, 19]. A notable example is [24], in which the authors discretize a continuous action space into several independent action branches, each parameterized by individual network branches. In [16], a variant of DQN [17] that discretizes a continuous action space into multiple discrete dimensions is proposed. These dimensions are sequentially parameterized, conditional on the previous sub-actions. Similarly, [19] propose an autoregressive factorization of an unconstrained action space into dependent sub-problems. Unlike our approach, these methods focus either on decomposing continuous action spaces into discrete action spaces or decomposing unconstrained action spaces. However, the decomposition of arbitrary convex polytope action spaces into tractable sub-action spaces remains a non-trivial challenge that our approach addresses.

3 Problem Description

An allocation task can be described as a finite-horizon Markov decision process (MDP) (S, A, T, R, γ) , where S represents the state space, A the action space, $T : S \times A \times S \rightarrow [0, 1]$ the state transition function, R the reward function, and $\gamma \in [0, 1]$ a discount factor. The goal of this task is to find a policy π maximizing the expected cumulative reward $J_R^\pi = \mathbb{E}_\pi [\sum_{t=1}^n \gamma^t R(s_t, \pi(s_t), s_{t+1})]$.

The action a is an allocation $a = \{a_1, \dots, a_n\} \in A$ over a set of n entities $E = \{e_1, \dots, e_n\}$ at each time step. Each element a_i of the action vector a represents the proportion allocated to entity e_i . Furthermore, allocation tasks require a complete allocation, i. e., $\sum_{i=1}^n a_i = 1$ and allocations cannot be negative ($a_i \geq 0$). Thus, the action space of unconstrained allocation tasks forms an n -dimensional standard simplex. A visualization of an unconstrained allocation action space is provided in Figure 1a.


Allocation tasks frequently include constraints, such as allocating at most 30% to a subset of the entities. An example of a constrained action space is visualized in Figure 1b. Formally, an allocation constraint can be expressed as a linear inequality $\sum_{i=1}^n c_i a_i \leq b$, where c_i denotes the weighting of the allocation variable a_i of entity e_i and $b \in \mathbb{R}$ denotes the corresponding constraint limit. For the sake of readability and simplicity, we only define \leq constraints since $a \geq b$ can be transformed into $-a \leq -b$ and $a = b$ can be rewritten as $a \leq b$ and $-a \leq -b$.

The action space A of constrained allocation tasks can be easily expressed by a set of linear inequalities, defining a polytope $A = \{a \in [0, 1]^n | Ca \leq b\}$, where


$$C \in \mathbb{R}^{m \times n} = \begin{bmatrix} c_{11} & \dots & c_{1n} \\ \vdots & \ddots & \vdots \\ c_{m1} & \dots & c_{mn} \end{bmatrix} \quad (1)$$

is a matrix of coefficients for the m constraints, including those linked to the simplex constraints $\sum_{i=1}^n a_i = 1$ and $a_i \geq 0 \forall i \in \{1, \dots, n\}$, as well as all coefficients for additional allocation constraints. Let $a \in [0, 1]^n$ represent an allocation vector and $b \in \mathbb{R}^m$ is the vector of constraint limits.

Alternatively, constrained allocation tasks can be defined using the framework of constrained Markov decision processes (CMDPs). A CMDP extends the standard MDP by a number of cost functions to incorporate the constraints. The goal is to maximize the expected cumulative reward while satisfying m constraints on the expected cumulative costs. The expected cumulative costs for the k -th cost function CF_k are defined as $J_{CF_k}^\pi = \mathbb{E}[\sum_{t=0}^T \gamma^t CF_k(s_t, a_t)]$. The m constraints to be satisfied in the CMDP are then stated as $J_{CF_k}^\pi \leq d_k$, where d_k denotes the cost limit with



(a) Unconstrained standard simplex



(b) Constrained simplex with $a_3 \leq 0.6$ and $a_2 \leq 0.7$

Figure 1: **Examples of 3-dimensional allocation action spaces** (a) unconstrained and (b) constrained (valid solutions as red area).

$k \in \{1, \dots, m\}$. To formulate constrained allocation tasks using CMDPs, the cost functions can be defined as $CF_k(s, a) = \max\{0, (Ca)_k - b_k\}$ to measure any allocation constraint violation as a cost. In addition, strict adherence to all allocation constraints at any point in time is required, i.e., $d_k = 0$. By formulating constraint allocation tasks using CMDPs, it becomes possible to use existing methods from Safe RL for soft constraints. However, these methods cannot guarantee constraint satisfaction at all times [15]. Let us note that our method does not use cost functions, instead it samples actions directly from the constrained action space.

4 Polytope Action Space Policy Optimization (PASPO)

Our approach PASPO autoregressively computes the allocation to every single entity in an iterative process until all allocations are fixed. We will later show that this step-wise decomposition allows for a tractable parametrization of the action space.

4.1 Autoregressive Polytope Decomposition

PASPO starts by determining the feasible interval $[a_1^{min}, a_1^{max}]$ for allocations into the first entity e_1 . Then, we sample the first allocation a_1 from this interval. The details of the sampling process will be further discussed in Section 4.2. Fixing an allocation impacts the shape of the remaining action space. Thus, we have to compute the shape of the polytope $A^{(2)}$ described by $C^{(2)}$ and $b^{(2)}$ before we can sample the next allocation a_2 .

Each iteration i starts with determining the interval $[a_i^{min}, a_i^{max}]$ of all feasible values for a_i . Geometrically, this interval is bounded by the minimum and the maximum value of the remaining polytope $A^{(i)}$ in the i -th dimension associated with the allocation a_i . To determine a_i^{min} , we solve the following linear program:

$$\begin{aligned} & \text{minimize} && a_i \\ & \text{s.t.} && C^{(i)} a^{(i)} \leq b^{(i)} \end{aligned}$$

where $C^{(i)}$ are the constraint coefficients for the entities e_i, \dots, e_n , $b^{(i)}$ are the adjusted constraint limits, and $a^{(i)}$ describes the unfixed allocations. We determine a_i^{max} by solving the respective maximization problem. For the first iteration $i = 1$, we define $C^{(1)} = C$, $b^{(1)} = b$ and $a^{(1)} = a$. After sampling an allocation a_i from the interval $[a_i^{min}, a_i^{max}]$. The resulting polytope $A^{(i+1)}$ for the next iteration $i + 1$ is described by the following inequality system:

$$\underbrace{\begin{bmatrix} c_{1,i+1} & \cdots & c_{1,n} \\ \vdots & \ddots & \vdots \\ c_{m,i+1} & \cdots & c_{m,n} \end{bmatrix}}_{C^{(i+1)}} \underbrace{\begin{bmatrix} a_{i+1} \\ \vdots \\ a_n \end{bmatrix}}_{a^{(i+1)}} \leq \underbrace{\begin{bmatrix} b_1^{(i)} \\ \vdots \\ b_m^{(i)} \end{bmatrix}}_{b^{(i+1)}} - a_i \underbrace{\begin{bmatrix} c_{1,i} \\ \vdots \\ c_{m,i} \end{bmatrix}}_{c^{(i)}} \quad (2)$$

To define the new coefficient matrix $C^{(i+1)}$ (red), we remove the first column of the coefficient matrix of the previous iteration $C^{(i)}$. To calculate the new vector $b^{(i+1)}$ of constraint limits, we subtract

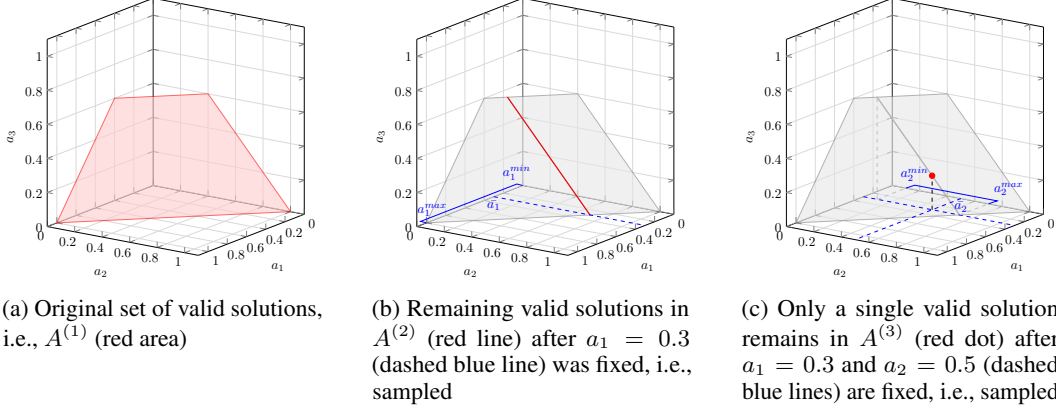


Figure 2: **Example of sampling process** of an action $a = (a_1, a_2, a_3)$ in a 3-dimensional constrained allocation task.

the removed column (blue) scaled by the fixed allocation a_i from the previous constraint limits $b^{(i)}$ (yellow). We iterate over all entities until we determine a_{n-1} . Allocation a_n is already determined as soon as the allocations a_1, \dots, a_{n-1} are fixed because of the simplex constraint $\sum_{i=1}^n a_i = 1$. Sampling an allocation using this approach always guarantees constraint satisfaction and it is possible to sample any action in the constrained action space. A formal proof of these guarantees can be found in Appendix D.

Figure 2 displays a visualization of the process for a 3-dimensional case. The set of valid solutions before any allocations have been fixed is shown in Figure 2a. Figure 2b depicts the first iteration after $a_1 = 0.3$ has been determined, and the resulting new polytope $A^{(2)}$, i.e., a set of valid solutions, shrinks to the red line is shown. Figure 2c shows the second iteration after also $a_2 = 0.5$ has been determined. It can be seen that the new polytope $A^{(3)}$ contains only a single valid solution represented as a red dot, making a third iteration unnecessary since the only remaining solution is to allocate $a_3 = 0.2$, resulting in a final allocation of $a = (0.3, 0.5, 0.2)$.

4.2 Parameterizable Policy Process

Our goal is to define a learnable stochastic policy function over the action space. For unconstrained allocation tasks, a Dirichlet distribution can be used to parameterize the action space [26, 28]. Unfortunately, to the best of our knowledge, there is no known parameterizable, closed-form distribution function over arbitrary convex polytopes as in our setting. In fact, even uniform sampling over a convex polytope is an active research problem [8].

We sequentially constructed an action a from the polytope action space A in the previous section. Now, we describe how to utilize this process to define a parameterizable policy function over the action space A . We model the distribution for allocating each individual entity using a beta distribution that is normalized to the range $[a_i^{min}, a_i^{max}]$. This distribution is also known as the four-parameter beta distribution [7]. Its probability density function is defined as:

$$p(x; \alpha, \beta, a_i^{min}, a_i^{max}) = \frac{(x - a_i^{min})^{\alpha-1} (a_i^{max} - x)^{\beta-1}}{(a_i^{max} - a_i^{min})^{\alpha+\beta-1} B(\alpha, \beta)},$$

where $B(\alpha, \beta)$ is the beta function. It is important to note that any other parameterizable distributions with bounded support in the range $[a_i^{min}, a_i^{max}]$ can be used, such as a squashed Gaussian distribution. However, our preliminary experiments indicated that the beta distribution performs particularly well.

To optimize the policy $\pi_\theta(s)$ over the complete allocations, we follow the approach of [19] for training an autoregressively dependent series of sub-policies. A fixed but arbitrary order of entities is used for sampling the allocations a_i . The sub-policy $\pi_\theta^i(a_i | a_1, \dots, a_{i-1})$ is conditional on the previous allocations a_1, \dots, a_{i-1} . Using this autoregressive dependence structure, the policy is defined as: $\pi_\theta(a|s) = \pi_\theta^1(a_1|s) \cdot \pi_\theta^2(a_2|s, a_1) \dots \pi_\theta^n(a_n|s, a_1, \dots, a_{n-1})$. This policy can be jointly optimized. We parameterize each sub-policy using a neural network that receives an embedding of the state and the previously selected actions as input.




Figure 3: **The impact of initialization in an unconstrained simplex.** (a) Mean allocations a_i to each entity in a seven entity setup when sampling each individual allocation using the **uniform distribution** (red) vs. **our initialization** (blue). (b,c) Distribution of 2500 allocations in a three entity setup when sampling each individual allocation uniformly (b) or using beta distributions with parameters set according to our initialization (c).

Algorithm 1 Maximum likelihood estimation of parameter de-biasing terms

Input: Polytope $A = \{a \in \mathbb{R}_{0,+}^n | Ca \leq b\}$, number of samples k

Output: Beta distribution shape parameters $\hat{\alpha}_i, \hat{\beta}_i$ as de-biasing terms

- 1: Let $A_j^{(1)} = A$ for $j = \{1, \dots, k\}$
 - 2: Sample k allocations $\{a_{(j)}\}_{j=1\dots k}$ uniformly within A via rejection sampling
 - 3: **for** each dimension i in $\{1, \dots, n\}$ **do**
 - 4: **for** each sample j in $\{1, \dots, k\}$ **do**
 - 5: Calculate interval $[a_{(j)i}^{min}, a_{(j)i}^{max}]$ using LP based on $A_j^{(i)}$
 - 6: Normalize sampled allocation to support $[0,1]$ of beta distribution: $a_{(j)i}^{norm} = \frac{a_{(j)i} - a_{(j)i}^{min}}{a_{(j)i}^{max} - a_{(j)i}^{min}}$
 - 7: Compute polytope $A_j^{(i+1)}$ from $A_j^{(i)}$ using sampled allocation $a_{(j)i}$ (see Eq. 2)
 - 8: **end for**
 - 9: ML estimation of beta distribution parameters $\hat{\alpha}_i, \hat{\beta}_i$ using $\{a_{(j)i}^{norm}\}_{j=1\dots k}$
 - 10: **end for**
-

An entropy term is often used to encourage exploration. However, our policy does not have a closed-form solution for entropy. Therefore, we follow [19] to empirically estimate the entropy:

$$H_{\text{emp}}(\pi_{\theta}(\cdot|s)) = \mathbb{E}_{a \sim \pi_{\theta}(\cdot|s)} \left[\sum_{i=1}^n H(\pi_{\theta}^i(\cdot|s, a_1, \dots, a_{i-1})) \right] \quad (3)$$

Here, H denotes the entropy of the beta distribution. We compute the expectation within each training batch to estimate the entropy of the complete policy function over the entropies of single actions. Let us note that when using an off-policy algorithm, the actions must be resampled using the current policy. As the current policy might have a significantly different parametrization than the sampling policy, we have to generate actions based on the current policy to estimate the entropy properly.

4.3 Policy Network Architecture

We create an embedding of the state using an MLP, denoted as $f_{\theta}(s) = MLP(s)$. We parameterize the probability distribution over allocations for each entity using an MLP $\pi_{\theta_i}^i(s) = MLP(f_{\theta}(s), a_1, \dots, a_{i-1})$, which receives the latent encoding of the state and the previously sampled allocations a_1, \dots, a_{i-1} as input. Note that each of the MLPs $\pi_{\theta_i}^i$ has its own parameters. For further details, we refer to the Appendix.

4.4 De-biasing Mechanism

A drawback of generating actions by an autoregressive process is that a random initialization of the beta distributions leads to a sampling bias towards the entities selected earlier in the process. The

effect is caused by the autoregressive dependency structure of our process. To sample an allocation of 80% for entity a_i the cumulated fixed allocations for earlier entities $\sum_{j=1}^{i-1} a_j$ must be at most 20%. However, this is rather unlikely if we initialize the distribution for all entities in a similar way. This effect can be observed in the red bars of Figure 3a. The red bars correspond to the average allocation for each dimension when uniformly drawing from an unconstrained seven-dimensional simplex with our autoregressive process for each entity e_i . As expected, the mean for the first dimension is 0.5, which is the mean of a uniform distribution over the interval $[0, 1]$. Correspondingly, the mean is decreased by half for any successive further entity until entity 6, which has the same mean as entity 7 due to the simplex constraint. Even though the bias is more complex for constrained allocation spaces, a similar effect can be expected.

For policy gradient methods, such a bias in the initialization of the policy function can lead to convergence to poorly performing policies or long training times. As the initial policy is crucial for ensuring sufficient exploration of the state-action space, a biased initial distribution leads to underexplored regions in the state-action space. Consequently, well-performing actions might not be discovered. To counter this effect, we propose a de-biasing mechanism that adjusts the initial parameters of beta distributions to estimate a uniform sampling over the joint action space. During learning, the amount of required exploration decreases, and the parameters of the policy function are optimized to increase the cumulative rewards. Thus, the impact of our de-biasing mechanism should diminish over time. We achieve this effect by adding a de-biasing term to the linear layers' initial bias terms, predicting α_i and β_i for entity i . As the default initialization of the bias terms has zero means, the first iterations use α and β values close to the de-biasing terms.

To determine suitable initial values for each iteration step, we proceed as described in Algorithm 1. We start by uniformly sampling n data points from the complete action space A . We do this by rejection sampling, i.e., we sample over the standard simplex and reject the samples outside the action polytope A . To determine the parameters corresponding to the acquired uniform sample, we project any allocation for each entity to a standard interval between $[0, 1]$. However, for this step, we have to determine the interval $[a_i^{min}, a_i^{max}]$ for each entity following the above process. We determine the relative position in this interval, corresponding to the position in the named standard interval. After collecting relative values for each sample and entity, we employ the standard maximum likelihood estimator to generate an empirical estimate of the α_i and β_i for each entity e_i . The blue bars in Figure 3a correspond to the results on the unconstrained seven-dimensional unconstrained simplex. Figure 3b shows autoregressive sampling based on uniform distribution, whereas Figure 3c displays the result of our initialization for a three-dimensional example. It can be seen that the result of our initialization of the autoregressive process closely resembles a uniform distribution over the complete action space.

5 Experiments

In this section, we provide an extensive experimental evaluation of our approach in various scenarios demonstrating its ability to handle various allocation tasks and constraints. We use two real-world tasks: Portfolio optimization [27] and compute load distribution [3]. Additionally, we create a synthetic benchmark with a reward surface generated by a randomly initialized MLP. Each of these tasks comes with a different set of allocation constraints. We will briefly describe each setting in the following and refer the reader to the Appendix for more details.

Portfolio Optimization Portfolio optimization is a prominent constrained allocation task. In this task the agent has to allocate its wealth over 13 assets at each time step. We use the environment of [27]. Each investment period contains 12 months and the investor needs to reallocate the portfolio each month. This environment is highly stochastic since each trajectory is sampled from a hidden Markov model fitted on real-world NASDAQ-100 data. After every 5120 environment steps, we run eight parallel evaluations on 200 fixed trajectories. Constraints in this setting define minimum and maximum allocation to groups of assets. Additionally, we add constraints where the constraint coefficients in C correspond to portfolio measures like a minimum dividend yield or a maximum on the CO2 intensity.

Compute Load Distribution The environment is based on the paper of [3] and simulates a data center in which computational jobs need to be split into sub-jobs to enable parallel processing across nine servers. Here, we use five constraints that are randomly sampled as follows: First, we sample

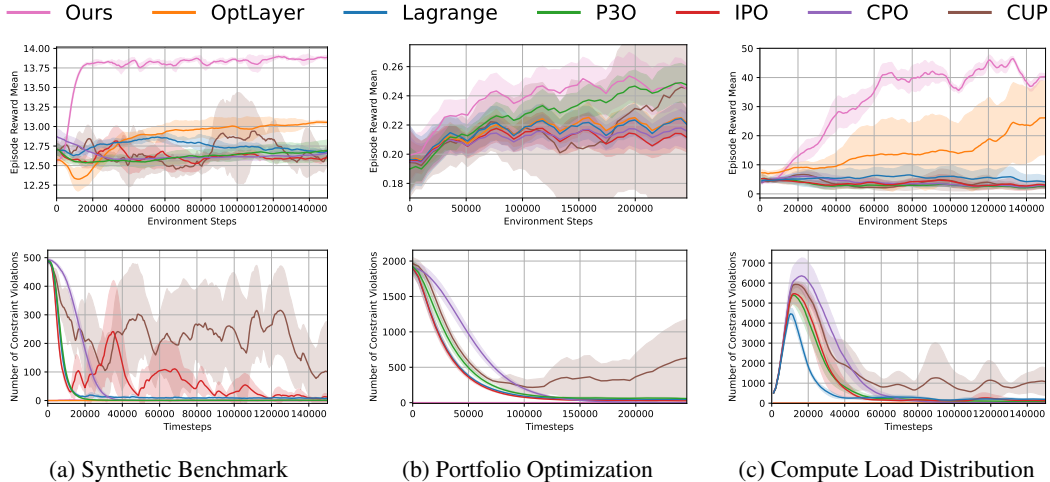


Figure 4: **Learning curves of all methods in three environments.** The x-axis corresponds to the number of environment steps. The y-axis is the average episode reward (first row), and the number of constraint violations during every epoch (second row). For portfolio optimization (b) we report the performance running eight evaluation on 200 fixed market trajectories. This is because in training, every trajectory is different which makes comparisons hard. Curves smoothed for visualization.

the number of affected entities for each constraint. We then sample the constraint coefficients from the range $[0, 1]$.

Synthetic Environment In addition to the aforementioned environments, we propose a synthetic benchmark. The reward surface consists of an MLP with random weights. Each episode comprises two states. As it is completely deterministic, it provides a simple yet effective way to benchmark approaches for constrained allocation tasks. In this setting, we create the constraints by randomly sampling 30 points and use their convex hull as the polytope defining the action space. We utilize a seven-dimensional setting with 611 constraints in our experiments.

5.1 Experimental Setup

We train PASPO using PPO [22] and compare our approach to various baselines, including state-of-the-art approaches for constrained allocation tasks and Safe RL. Specifically, we compare PASPO with five representative approaches from Safe RL: CPO [2], CUP [30], IPO [14], P3O [31], and PPO with Lagrangian relaxation. Additionally, we compare our method to OptLayer [18], a popular projection-based method for linear hard constraints. To maintain a consistent and fair comparison across different methods, we use the same hyperparameters across the different methods if possible. Many Safe RL approaches have difficulties handling equality constraints [11]. Therefore, we use a Dirichlet distribution to represent the policy in the baselines, thereby ensuring satisfaction of the simplex equality constraint. We do not share the parameters between the policy and value function. We use a fully-connected MLP with two hidden layers of 32 units and ReLU non-linearities for each policy, cost, and value function. In our approach, the state encoder and each policy head consists of a two-layer MLP. The training process is run for 150,000 steps and the results are averaged over five different seeds. In the portfolio optimization task, we use ten different seeds due to the stochasticity of the financial environment and train for 250,000 steps. Given the relatively small network sizes, training is conducted exclusively on CPUs. We implement our algorithm and the baselines using RLlib and PyTorch. More details regarding the environments, training, and hyperparameters can be found in the Appendix.

5.1.1 Performance of PASPO

We visualize the performance and constraint violations of all methods across our three environments in Figure 4. A tolerance of $1e^{-3}$ is used for evaluating constraint violations and we report the total number of violations per episode. In all three environments, PASPO converges faster to a higher




Figure 5: **Ablations** in (a) show the performance of our approach with (blue) and without (orange) the de-biased initialization. In (b) depicts the impact of the allocation order. We reverse the allocation order (red).

average return compared to baselines. Additionally, while all compared soft-constraint methods display constraint violations, only the hard constraint approaches PASPO and OptLayer guarantee to permanently satisfy the constraints. Finally, we can observe that the variance of PASPO is rather low compared to other methods. However, in portfolio optimization task (b) our approach displays some variance which we attribute to the stochasticity of the environment. Overall, these results demonstrate that our approach is not only able to consistently outperform other algorithms in terms of rewards but also guarantees no constraint violations.

5.1.2 Importance of de-biased Initialization and Order

We conduct ablation studies to investigate the impact of our de-biased initialization and the order of entity allocation on our synthetic benchmark. No constraints are applied except for the simplex constraint to highlight the effects. The results, shown in Figure 5, indicate that without de-biased initialization (orange in (a)), learning is slower and converges prematurely to a sub-optimal policy. In (b), we explore the impact of allocation order by reversing it (red) and observe no significant performance difference. This indicates that our approach is robust to the allocation order due to the use of the de-biasing initialization.

6 Limitations and Future Work

While PASPO guarantees that constraints are always satisfied, it is considerably more computationally expensive than standard neural networks in allocation tasks with many entities, as the sampling of each action requires solving a series of linear programs. RL in high-dimensional continuous action spaces is a very challenging task. Our approach cannot overcome this issue and also struggles in very high-dimensional settings. For future work, we plan to extend PASPO to also incorporate state-dependent constraints. While we evaluate our approach only on benchmarks with hard constraints, it can be applied to settings with both hard and soft cumulative constraints. In these scenarios, our method for handling hard constraints can be easily combined with most Safe RL algorithms to handle soft cumulative constraints.

7 Conclusion

In this paper, we examine allocation tasks where a certain amount of a resource has to be distributed over a set of entities at every step. This problem has many applications like logistics tasks, portfolio management, and computational workload processing in distributed environments. In all these applications, the set of feasible allocations might be bound by a set of linear constraints. Formally, these restrict the action space to a convex polytope. To define a stochastic policy function that can be used with policy gradient methods in RL, we propose an autoregressive process that computes allocation sequentially. We employ linear programming to compute the range of feasible allocations for an entity given the already fixed allocations of other entities. Our policy function consists of a sequence of one-dimensional beta distributions where the shape parameters α and β are learned by neural networks. To counter the effect of initialization bias, we utilize a de-biasing mechanism to ensure sufficient exploration and prevent premature convergence to a sub-optimal policy. In our

experiments, we demonstrate that our novel method PASPO yields better results than state-of-the-art approaches while not having any constraint violations. Furthermore, we show that our initialization method yields better results than random initializations and counters the impact of the allocation order.

References

- [1] J. Abadie. Generalization of the wolfe reduced gradient method to the case of nonlinear constraints. *Optimization*, pages 37–47, 1969.
- [2] J. Achiam, D. Held, A. Tamar, and P. Abbeel. Constrained policy optimization. In *International conference on machine learning*, pages 22–31. PMLR, 2017.
- [3] L. Ale, S. A. King, N. Zhang, A. R. Sattar, and J. Skandaraniyam. D3pg: Dirichlet ddpq for task partitioning and offloading with constrained hybrid action space in mobile-edge computing. *IEEE Internet of Things Journal*, 9(19):19260–19272, 2022.
- [4] E. Altman. *Constrained Markov Decision Processes*, volume 7. CRC Press, 1999.
- [5] B. Amos and J. Z. Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pages 136–145. PMLR, 2017.
- [6] A. Bhatia, P. Varakantham, and A. Kumar. Resource constrained deep reinforcement learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 610–620, 2019.
- [7] J. Carnahan. Maximum likelihood estimation for the 4-parameter beta distribution. *Communications in Statistics-Simulation and Computation*, 18(2):513–536, 1989.
- [8] Y. Chen, R. Dwivedi, M. J. Wainwright, and B. Yu. Fast mcmc sampling algorithms on polytopes. *Journal of Machine Learning Research*, 19(55):1–86, 2018.
- [9] Y. Chow, M. Ghavamzadeh, L. Janson, and M. Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *Journal of Machine Learning Research*, 18(167):1–51, 2018.
- [10] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, and Y. Tassa. Safe exploration in continuous action spaces. *arXiv preprint arXiv:1801.08757*, 2018.
- [11] S. Ding, J. Wang, Y. Du, and Y. Shi. Reduced policy optimization for continuous control with hard constraints. *Advances in Neural Information Processing Systems*, 36, 2024.
- [12] S. Gu, L. Yang, Y. Du, G. Chen, F. Walter, J. Wang, Y. Yang, and A. Knoll. A review of safe reinforcement learning: Methods, theory and applications. *arXiv preprint arXiv:2205.10330*, 2022.
- [13] Q. Liang, F. Que, and E. Modiano. Accelerated primal-dual policy optimization for safe reinforcement learning. *arXiv preprint arXiv:1802.06480*, 2018.
- [14] Y. Liu, J. Ding, and X. Liu. Ipo: Interior-point policy optimization under constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 4940–4947, 2020.
- [15] Y. Liu, A. Halev, and X. Liu. Policy learning with constraints in model-free reinforcement learning: A survey. In *The 30th international joint conference on artificial intelligence (ijcai)*, 2021.
- [16] L. Metz, J. Ibarz, N. Jaitly, and J. Davidson. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*, 2017.
- [17] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- [18] T.-H. Pham, G. De Magistris, and R. 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.
- [19] T. PIERROT, V. Macé, J.-B. Sevestre, L. Monier, A. Laterre, N. Perrin, K. Beguir, and O. Sigaud. Factored action spaces in deep reinforcement learning, 2021. URL <https://openreview.net/forum?id=naSAkn2Xo46>.

- [20] S. Sanket, A. Sinha, P. Varakantham, P. Andrew, and M. Tambe. Solving online threat screening games using constrained action space reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 2226–2235, 2020.
- [21] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.
- [22] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [23] A. Stooke, J. Achiam, and P. Abbeel. Responsive safety in reinforcement learning by pid lagrangian methods. In *International Conference on Machine Learning*, pages 9133–9143. PMLR, 2020.
- [24] A. Tavakoli, F. Pardo, and P. Kormushev. Action branching architectures for deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [25] C. Tessler, D. J. Mankowitz, and S. Mannor. Reward constrained policy optimization. In *International Conference on Learning Representations*, 2018.
- [26] Y. Tian, M. Han, C. Kulkarni, and O. Fink. A prescriptive dirichlet power allocation policy with deep reinforcement learning. *Reliability Engineering & System Safety*, 224:108529, 2022.
- [27] D. Winkel, N. Strauß, M. Schubert, Y. Ma, and T. Seidl. Constrained portfolio management using action space decomposition for reinforcement learning. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 373–385. Springer, 2023.
- [28] D. Winkel, N. Strauß, M. Schubert, and T. Seidl. Simplex decomposition for portfolio allocation constraints in reinforcement learning. In K. Gal, A. Nowé, G. J. Nalepa, R. Fairstein, and R. Radulescu, editors, *ECAI 2023 - 26th European Conference on Artificial Intelligence, September 30 - October 4, 2023, Kraków, Poland - Including 12th Conference on Prestigious Applications of Intelligent Systems (PAIS 2023)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, pages 2655–2662. IOS Press, 2023. doi: 10.3233/FAIA230573. URL <https://doi.org/10.3233/FAIA230573>.
- [29] L. Yang, J. Ji, J. Dai, L. Zhang, B. Zhou, P. Li, Y. Yang, and G. Pan. Constrained update projection approach to safe policy optimization. *Advances in Neural Information Processing Systems*, 35:9111–9124, 2022.
- [30] L. Yang, J. Ji, J. Dai, Y. Zhang, P. Li, and G. Pan. Cup: A conservative update policy algorithm for safe reinforcement learning. *arXiv preprint arXiv:2202.07565*, 2022.
- [31] L. Zhang, L. Shen, L. Yang, S. Chen, X. Wang, B. Yuan, and D. Tao. Penalized proximal policy optimization for safe reinforcement learning. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 3744–3750, 2022.
- [32] Y. Zhang, Q. Vuong, and K. Ross. First order constrained optimization in policy space. *Advances in Neural Information Processing Systems*, 33:15338–15349, 2020.

A Environments

The implementation for all three environments can be found at: <https://github.com/niklasdbs/paspo>.

A.1 Financial Environment

The financial environment used for testing our approach is based on [28]. The financial market trajectories in this environment are sampled from a hidden Markov model, which was fitted based on real-world NASDAQ-100 data from January 3rd, 2011 to December 1st, 2021. The environment offers differently sized data sets of randomly selected assets contained in the NASDAQ-100. The experiments in the financial environment for this paper are run with 13 assets, which corresponds to

the *model_parameter_data_set_G_markov_states_2_12* data set with the *include_cash_asset=true* option. We initialize the environment with *seed=2*.

Table 1 shows a list of the assets used in the experiments.

For the experiments a random combination of two types of constraints typical for financial tasks is used: (a) a randomly selected subset of assets to either stay above or below an randomly selected allocation threshold and (b) thresholds for financial or environmental portfolio measures that can be calculated as an weighted average of the assets' individual measures, i.e., as a weighted linear combination. A list of these measure can be found in Table 2.

The experiments include 5 constraints, of which the number constraints of type (a) and type (b) is randomly decided. The exact implementation can be found in our code *polytope_loader.py* (*generate_random_fin_env_polytope_rejection_sampling*). We use the seed 2 to generate the constraints.

Index	ISIN	Ticker	Name
1	-	CASH	CASH
2	US5949181045	MSFT	Microsoft Corporation
3	US0567521085	BIDU	Baidu Inc.
4	US00724F1012	ADBE	Adobe Inc.
5	US6937181088	PCAR	Paccar Inc.
6	US67066G1040	NVDA	NVIDIA Corporation
7	US8552441094	SBUX	Starbucks Corporation
8	US4612021034	INTU	Intuit Inc.
9	US0530151036	ADP	Automatic Data Processing Inc.
10	US0231351067	AMZN	Amazon.com Inc.
11	US2786421030	EBAY	eBay Inc.
12	US0311621009	AMGN	Amgen Inc.
13	US7475251036	QCOM	Qualcomm Inc.

Table 1: List of assets used in the environment.

	Est. Total Energy Use To EVIC USD in million	Est. Total CO2 Equivalent Emissions To EVIC USD in million	Est. Weighted Average Cost of Capital, (%)	Est. Dividend yield, (%)	Est. Return On Equity, (%)
CASH	0.00	0.00	0.00	0.00	0.00
MSFT	23.49	1.75	8.19	1.89	40.15
BIDU	70.80	15.17	7.12	0.00	9.01
ADBE	12.98	1.12	6.81	2.28	92.49
PCAR	83.66	10.27	7.31	3.14	-43.29
NVDA	17.17	1.58	6.20	3.00	127.85
SBUX	85.73	6.79	7.24	1.13	26.61
INTU	41.23	17.15	7.87	0.00	21.50
ADP	1.70	0.15	8.12	0.56	22.46
AMZN	4.69	0.39	8.28	0.00	44.48
EBAY	3.49	0.30	9.35	0.02	80.24
AMGN	33.87	3.29	7.09	0.73	37.76
QCOM	47.07	4.09	8.33	2.16	36.06

Table 2: KPI estimates for assets based on 2021 (final year of the used data set, source: Refinitiv); EVIC - Enterprise value including Cash

A.2 Compute Environment

The compute environment used for testing our approach is based on [3]. The agent’s task is to allocate compute jobs to a given set of servers in a data center. A reward is triggered for each job that was completed in a predetermined maximum allowed computation time. The challenge of this environment is that the agent needs to match the queue of jobs still to be allocated with the different computational capabilities of the servers as well as each server’s individual queue of jobs still to be computed. It is assumed that the compute jobs in the environment can be arbitrarily split and computed in parallel. The creation of new compute jobs is triggered by n users and follows a Poisson process. A job is defined by its *payload size*, i.e., the data to be transferred to a server, its *required CPU cycles* for the processing workload, and its maximum allowed time until the job needs to be completely processed. These attributes for the jobs that can be created by each user are randomly sampled at creation of the environment.

The experiments in this paper run with a setup of 9 servers and 9 users that generate compute jobs. The parameter set used is *parameter_set_9_9_id_0*. We initialize the environment with *seed=1*. The randomly sampled specifications for the nine servers can be found in Table 3 and the job attributes created by the nine users can be found in Table 4.

To generate the constraints, we first sample the number of affected entities between 2 and 8 for each constraint and randomly choose the affected entities accordingly. We then uniformly sample constraint coefficients from the interval $[0, 1]$, as well as a corresponding constraint limit between 0 and 1. We use a seed of 1 to generate 5 constraints. The implementation can be found in `polytope_loader.py` (`generate_random_polytope_rejection_sampling`).

Index	Max Compute Cycles per Second
1	2 836 258 583
2	855 913 878
3	652 109 364
4	789 819 414
5	3 187 852 760
6	974 311 629
7	2 005 143 973
8	1 481 875 307
9	2 216 715 088

Table 3: Server Specifications

User	Data Size in Bits per Job	Required Compute Cycles per Job	Average Number of Jobs Created per Interval	Interval Length in Seconds
1	587 168	1 690 694	10	0.01
2	240 447	1 092 255	10	0.01
3	257 396	867 139	10	0.01
4	364 400	819 594	10	0.01
5	387 953	3 463 247	10	0.01
6	309 269	2 300 810	10	0.01
7	44 420	1 129 119	10	0.01
8	318 062	1 092 402	10	0.01
9	490 880	1 044 736	10	0.01

Table 4: User/Job Specifications

A.3 Synthetic Benchmark

In addition to these environments, we propose a synthetic benchmark. Its reward surface consists of an MLP with random weights. An example of the reward surface in three dimensions is visualized in Figure 6 Each episode has two states. Since it is completely deterministic, it provides a simple but




Figure 6: Example of the reward surface of our synthetic benchmark in three dimensions under constraints.

effective way to benchmark approaches for constrained allocation tasks. In this setting, we create the constraints by randomly sampling 30 points and use their convex hull as the polytope defining the action space. We use a seven-dimensional setting with 611 constraints in our experiments. More specifically, the network has one hidden layer and ReLU as a non-linearity. The input layer receives the state (as a number, i.e., 0 or 1) and the action as input and has an output size of 32, the hidden layer has an input size of 32 and output size of 16. The output layer has an input size 16 and a output size of 1. The exact initialization of the neural network weights can be found in our code (synth_env.py: MLPRewardNetwork). To generate the environment we use the seed 1 in our experiments.

To generate the constraints, we sample 30 randomly from a Dirichlet distribution with concentration parameters set to 1. We then build the convex hull of these points and convert the resulting polytope into its halfspace representation, i.e., a system of linear inequalities which we use as constraints. We use the seed of 1 to generate the constraints. This results in 611 constraints. The algorithm to generate the constraints can be found in the code (random_polytope_generator.py)

B Architecture

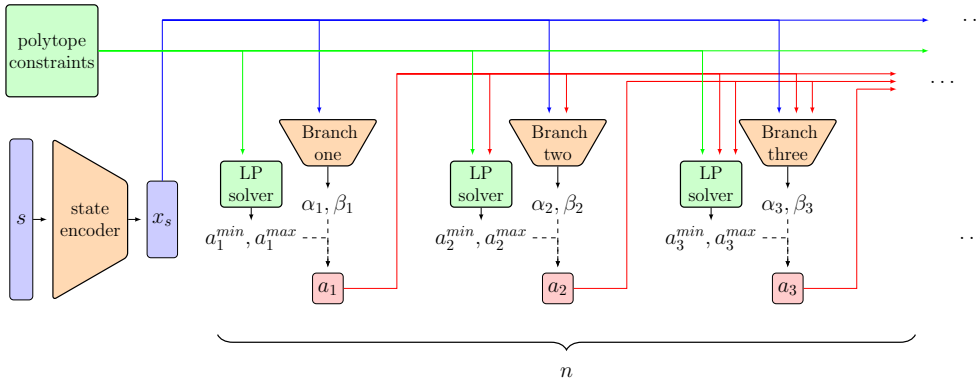


Figure 7: Architecture of PASPO

C Hyperparameters/Training

In Table 5 we list the most important parameters and hyperparameters. The full configurations used can be found in the config files (yaml/hydra based) in our code (run configs directory). We tuned hyperparameters on our synthetic benchmark with five dimensions and five constraints.

We do not train using GPUs because of the small network sizes. We used an internal CPU cluster with consumer machines and servers ranging from 8 to 90 cores and RAM between 32GB and 512GB.

Parameter	Ours	IPO	P3O	CUP	Lag.	OptLayer	CPO
Training env steps	150,000 (synth, compute), 250,000 (portfolio optimization)						
Episode/Rollout length	512 environment steps						
Number of parallel envs	8						
Learning Rate	1e-3	1e-3	1e-3	1e-3	1e-3	1e-3	–
Gradient clipping	2.0						
Minibatch size	64						
Optimizer	Adam						
GAE lambda	0.95						
Discount factor	1.0						
No. grad update it per epoch	10 (CPO only for the critic 40)						
PPO clip parameter	0.3	0.3	0.3	0.3	0.3	0.3	–
Entropy coefficient	0.01	0.01	0.01	0.01	0.01	0.01	–
Cost limit	–	1e-3	0.0	0.0	0.0	0.0	0.0

Table 5: The most important Parameters and Hyperparameters for Various Methods

D Guaranteed Constraint Satisfaction

In the following, we proof that our approach PASPO always guarantees constraint satisfaction and that our method is able to sample all possible actions from the constrained action space.

Definitions: Let A be the set of all actions that can be sampled with PASPO, and let $P = \{a \in \mathbb{R}^n | Ca \leq b\}$ be the convex polytope that corresponds to constrained action space. We define $A_*^{(i+1)} = \{a \in \mathbb{R}^n | C^{(i+1)}a^{(i+1)} \leq b_*^{(i+1)}, \forall j = 1, \dots, i : a_j = a_j^*\}$ where $b_*^{(i+1)} = b -$

$$\sum_{j=1}^i a_j^* \begin{bmatrix} c_{1j} \\ \vdots \\ c_{mj} \end{bmatrix} \text{ and } C^{(i+1)} \text{ and } a^{(i+1)} \text{ as defined in the paper.}$$

Thus, $A_*^{(i+1)}$ is the restricted action space after sampling/fixing already the allocations a_1^*, \dots, a_i^* .

Theorem 1. Let $P = \{a \in \mathbb{R}^n | Ca \leq b\} \neq \emptyset$ be the convex polytope that corresponds to a constrained action space. Let A be the set of all the points that can be generated by PASPO. It holds that $A = P$.

Proof. Well-defined: Show that $A^{(n)} \neq \emptyset$ if $P \neq \emptyset$.

Induction over i :

$$i = 1 : A_*^{(1)} = \{a \in \mathbb{R}^n | C^{(1)}a^{(1)} \leq b_*^{(i+1)}\} = \{a \in \mathbb{R}^n | Ca \leq b\} = P \neq \emptyset$$

$i \rightarrow i + 1$:

$$(i + 1 \leq n) A_*^{(i)} \neq \emptyset \Rightarrow \exists a^\uparrow, a^\downarrow \in A_*^{(i)} : a_i^\uparrow = a_i^{\min}, a_i^\downarrow = a_i^{\max}$$

Now assume an arbitrary a_i^* is sampled from $[a_i^{\min}, a_i^{\max}]$

$$\Rightarrow \exists \lambda \in [0, 1] : a_i^* = \underbrace{(\lambda a_i^\downarrow + (1 - \lambda) a_i^\uparrow)}_{:= a_i^\lambda}$$

By convexity of polytopes as solution spaces for linear inequality systems, we get:

$$\begin{bmatrix} c_{1,i} & \cdots & c_{1,n} \\ \vdots & \ddots & \vdots \\ c_{m,i} & \cdots & c_{m,n} \end{bmatrix} \begin{bmatrix} a_i^\lambda \\ \vdots \\ a_n^\lambda \end{bmatrix} \leq b - \sum_{j=1}^{i-1} a_j^* \begin{bmatrix} c_{1j} \\ \vdots \\ c_{mj} \end{bmatrix} \xLeftrightarrow_{(a_i^\lambda = a_i^*)} \begin{bmatrix} c_{1,i+1} & \cdots & c_{1,n} \\ \vdots & \ddots & \vdots \\ c_{m,i+1} & \cdots & c_{m,n} \end{bmatrix} \begin{bmatrix} a_{i+1}^\lambda \\ \vdots \\ a_n^\lambda \end{bmatrix} \leq b - \sum_{j=1}^i a_j^* \begin{bmatrix} c_{1j} \\ \vdots \\ c_{mj} \end{bmatrix} \Rightarrow a^\lambda \in A_*^{(i+1)}$$

To show that $A = P$:

$A \subseteq P$: Let $a^* \in A$. In the last step (n), a_n^* is sampled (by design) such that

$$C^{(n)} a_n^* \leq b - \sum_{j=1}^{n-1} a_j^* \begin{bmatrix} c_{1j} \\ \vdots \\ c_{mj} \end{bmatrix} \Leftrightarrow C a^* \leq b \Leftrightarrow a^* \in P$$

$A \supseteq P$: Let $a^* \in P$. $\Leftrightarrow C a^* \leq b \Leftrightarrow C^{(i)} a^* \leq b^{(i)} \forall i \Leftrightarrow a^* \in A_*^{(i)} \forall i$
 \Rightarrow We can construct a^* by sampling a_i^* in every step i . $\Rightarrow a^* \in A$

□

The intuition of why our approach can guarantee the satisfaction of constraints is based on three properties that we utilize: (1) If P_i is the set of solutions to a system of linear inequalities, then by adding further constraints to the system of linear inequalities there will be a new set of solutions P_{i+1} but always such that $P_{i+1} \subseteq P_i$. (2) For any two points a^{\min} and a^{\max} in a convex set it can be implied that there exists a point a^λ for which the following is true for its i -th dimension $\exists \lambda \in [0, 1] : a_i^* = \underbrace{(\lambda a^{\min} + (1 - \lambda) a^{\max})}_{{:=} a^\lambda}$ (3) Linear Programming can determine the upper and lower bounds for single variables in a system of linear inequalities, i.e. $[a_i^{\min}, a_i^{\max}] \forall i$.

We start with the original system of linear inequalities with the solution space $P_1 \neq \emptyset$ which is a convex polytope. We use (3) on P_1 to determine the upper and lower bounds for a_1 , i.e. $[a_1^{\min}, a_1^{\max}]$. We sample a value a_1^* from the range $[a_1^{\min}, a_1^{\max}]$. We know that the solution space P_1 must contain at least one point for which in its 1st dimension $a_1 = a_1^*$ due to (2). In the next step we add the further constraint $a_1 = a_1^*$ to the system of linear inequalities. This updated system of linear inequalities will have the solution space P_2 . Due to (2) $P_2 \neq \emptyset$, as well as $P_2 \subseteq P_1$. We then repeat the entire process and use (3) on P_2 to determine the upper and lower bounds for a_2 , i.e. $[a_2^{\min}, a_2^{\max}] \dots$

After the n -th iteration a_n^* will be determined and we then have completed the generation of point $a^* = (a_1^*, \dots, a_n^*) \in P_n \subseteq \dots \subseteq P_1$, i.e. we succeed generating a point a^* that satisfies all original constraints P_1 .

E The Impact of the Allocation Order

As already discussed in the ablations in the main paper, with our de-biased initialization the impact of the allocation order is small. However, without our de-biased initialization, the order of the allocation has a significant impact on the performance, as illustrated in Figure 8.




Figure 8: The impact of the allocation order on PASPO without de-biased initialization in the synthetic benchmark with two states, a 7-dimensional action space, and no additional allocation constraints. Blue depicts the standard allocation order (i.e., e_1, e_2, \dots, e_n) and red depicts the reversed allocation order (i.e., the entities are allocated in the reversed order). A significant difference in performance can be observed with respect to the order without our de-biased initialization. In contrast, Figure 5b in the paper shows that with the de-biased initialization the difference is not significant.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We define the scope by thoroughly describing the considered task and empirically justify our approach in the experiments.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We address the limitations of our method in Section 6.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: Our paper does not contain theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We describe our experimental setup and the used hyperparameters in the main paper and the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We disclose the code to reproduce the experiments presented in the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We describe the experimental settings in a reasonable level of detail in the main paper and provide complete information in the appendix and the submitted code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We use a 2-sigma confidence interval around the estimated mean shown in all of our plots.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide information about the compute workers used in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research does not involve human subjects and conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Constrained allocations tasks have no direct societal impacts, however, applying methods for allocation tasks in certain contexts such as finance may indirectly impact society.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: As we are not aware of direct ethical issues or direct negative consequences of our work, we abstain from putting safeguards as described above in place.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We acknowledge the authors of the Portfolio Optimization and Compute Load Distribution environments by citing the corresponding papers.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We publish the code of our implementation and benchmark environments. A documentation can be found in the Appendix and code.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The research presented in this paper does not involve crowdsourcing nor human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The research presented in this paper does not involve crowdsourcing nor human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.

- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.