BUILDING A SUBSPACE OF POLICIES FOR SCALABLE CONTINUAL LEARNING

Anonymous authors
Paper under double-blind review

ABSTRACT

The ability to continuously acquire new knowledge and skills is crucial for autonomous agents. Existing methods are typically based on either fixed-size models that struggle to learn a large number of diverse behaviors, or growing-size models that scale poorly with the number of tasks. In this work, we aim to strike a better balance between an agent's *size* and *performance* by designing a method that grows adaptively depending on the task sequence. We introduce Continual Subspace of Policies (CSP), a new approach that incrementally builds a subspace of policies for training a reinforcement learning agent on a sequence of tasks. The subspace's high expressivity allows CSP to perform well for many different tasks while growing sublinearly with the number of tasks. Our method does not suffer from forgetting and displays positive transfer to new tasks. CSP outperforms a number of popular baselines on a wide range of scenarios from two challenging domains, Brax (locomotion) and Continual World (manipulation). Interactive visualizations of the subspace can be found at `csp`.

1 INTRODUCTION

Developing autonomous agents that can continuously acquire new knowledge and skills is a major challenge in machine learning, with broad application in fields like robotics or dialogue systems. In the past few years, there has been growing interest in the problem of training agents on sequences of tasks, also referred to as continual reinforcement learning (CRL, Khetarpal et al. (2020)). However, current methods either use *fixed-size* models that struggle to learn a large number of diverse behaviors (Hinton et al., 2006; Rusu et al., 2016a; Li & Hoiem, 2018; Bengio & LeCun, 2007; Kaplanis et al., 2019; Traoré et al., 2019; Kirkpatrick et al., 2017; Schwarz et al., 2018; Mallya & Lazebnik, 2018), or *growing-size* models that scale poorly with the number of tasks (Berseth et al., 2018; Cheung et al., 2019; Wortsman et al., 2020). In this work, we introduce an *adaptive-size* model which strikes a better balance between **performance** and **size**, two crucial properties of continual learning systems (Veniat et al., 2020), thus scaling better to long task sequences.

Taking inspiration from the mode connectivity literature (Garipov et al., 2018; Gaya et al., 2021), we propose **Continual Subspace of Policies** (**CSP**), a new CRL approach that incrementally builds a subspace of policies (see Figure 1 for an illustration of our method). Instead of learning a single policy, CSP maintains an entire subspace of policies defined as a convex hull in parameter space. The vertices of this convex hull are called anchors, with each anchor representing the parameters of a policy. This subspace captures a large number of diverse behaviors, enabling good performance on a wide range of settings. At every stage of the CRL process, the best found policy for a previously seen task is represented as a single point in the current subspace (*i.e.* unique convex combination of the anchors), which facilitates cheap storage and easy retrieval of prior solutions.

If a new task shares some similarities with previously seen ones, a good policy can often be found in the current subspace without increasing the number of parameters. On the other hand, if a new task is very different from previously seen ones, CSP extends the current subspace by adding another anchor, and learns a new policy in the extended subspace. In this case, the pool of candidate solutions in the subspace increases, allowing CSP to deal with more diverse tasks in the future. *The size of the subspace is increased only if this leads to performance gains larger than a given threshold, allowing users to specify the desired trade-off between performance and size* (*i.e.* number of parameters or memory cost).

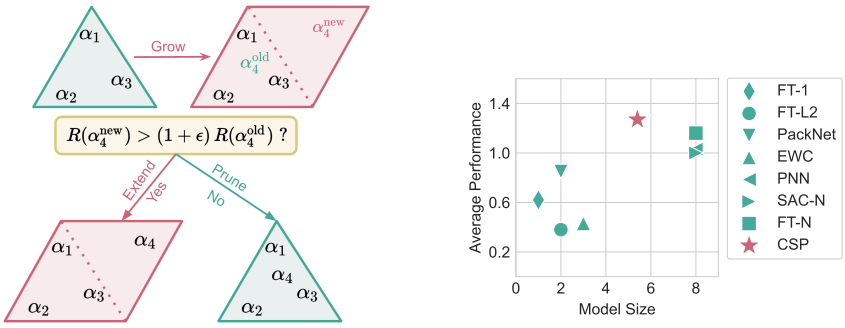

(a) Continual Subspace of Policies (CSP)

(b) Performance-Size Trade-Off

Figure 1: (a) **Continual Subspace of Policies (CSP)** iteratively learns a subspace of policies in the continual RL setting. At every stage during training, the subspace is a simplex defined by a set of anchors (*i.e.* vertices). Any policy (*i.e.* point) in this simplex can be represented as a convex combination α of the anchor parameters. α_i defines the best policy in the subspace for task i. When the agent encounters a new task, CSP tentatively *grows* the subspace by adding a new anchor. If the new task i is very different from previously seen ones, a better policy α_i^{new} can usually be learned in the new subspace. In this case, CSP *extends* the subspace by keeping the new anchor. If the new task bears some similarities to previously seen ones, a good policy α_i^{old} can typically be found in the old subspace. In this case, CSP *prunes* the subspace by removing the new anchor. The subspace is extended only if it improves improves performance relative to the old subspace by at least some threshold ϵ. (b) Trade-off between model performance and size for a number of methods, on a sequence of 8 tasks from HalfCheetah (see Table 6 for trade-offs on other scenarios).

We evaluate our approach on 18 CRL scenarios from two different domains, locomotion in Brax and robotic manipulation in Continual World, a challenging CRL benchmark (Wołczyk et al., 2021). We also compare CSP with a number of popular CRL baselines, including both fixed-size and growing-size methods. As Figure 1b shows, CSP *is competitive with the strongest existing methods while maintaining a smaller size*. CSP does not incur any forgetting of prior tasks and displays positive transfer to new ones. We demonstrate that by increasing the threshold parameter, the model size can be significantly reduced without substantially hurting performance. In addition, our qualitative analysis shows that the subspace captures diverse behaviors and even combinations of previously learned skills, allowing transfer to many new tasks without requiring additional parameters.

2 Continual Reinforcement Learning

A continual RL problem is defined by a sequence of N tasks denoted $t_1, ..., t_N$. Task i is defined by a Markov Decision Process (MDP) $\mathcal{M}_i = \langle \mathcal{S}_i, \mathcal{A}_i, \mathcal{T}_i, r_i, \gamma \rangle$ with a set of states $\mathcal{S}_i$, a set of actions $\mathcal{A}_i$, a transition function $\mathcal{T}_i : \mathcal{S}_i \times \mathcal{A}_i \to \mathcal{P}(\mathcal{S}_i)$, and a reward function $r_i : \mathcal{S}_i \times \mathcal{A}_i \to \mathbb{R}$. Here, we consider that all the tasks have the same state and action space. We also define a global policy $\Pi : [1..N] \times \mathcal{Z} \to (\mathcal{S} \to \mathcal{P}(\mathcal{A}))$ which takes as input a task id i and a sequence of tasks $\mathcal{Z}$, and outputs the policy which should be used when interacting with task i. $\mathcal{P}(\mathcal{A})$ is a probability distribution over the action space. We consider that each task i is associated with a *training budget* of interactions b_i. When the system switches to task t_{i+1} it no longer has access to transitions from the $\mathcal{M}_i$. We do not assume access to a replay buffer with transitions from prior tasks since this would lead to a significant increase in memory cost; instead, at each stage, we maintain a buffer with transitions from the current task, just like SAC (Haarnoja et al., 2018b).

3 Continual Subspace of Policies (CSP)

3.1 Subspace of Policies

Our work builds on Gaya et al. (2021) by leveraging a subspace of policies. However, instead of using the subspace to train on a single environment and adapt to new ones at test time, we use it to efficiently learn tasks sequentially in the continual RL setting. This requires designing a new approach for learning the subspace, as detailed below.

Let $\theta \in \mathbb{R}^d$ represent the parameters of a policy, which we denote as $\pi(a|s, \theta)$. A subspace of policies is a simplex in $\mathbb{R}^d$ defined by a set of anchors $\Theta = \{\theta_1, ..., \theta_k\}$, where any policy in this simplex can

be represented as a convex combination of the anchor parameters. Let the weight vector $\alpha \in \mathbb{R}_+^k$ with $\|\alpha\|_1 = 1$ denote the coefficients of the convex combination. Therefore, each value of α uniquely defines a policy with $\pi(a|s, [\alpha, \Theta]) = \pi(a|s, \sum \alpha_i \theta_i)$. Throughout the paper, we sometimes refer to the weight vector α and its corresponding policy $\pi(a|s, [\alpha, \Theta])$ interchangeably.

3.2 Learning Algorithm

We propose to adaptively construct a subspace of policies for the continual RL setting in which an agent learns tasks sequentially. We call our method **Continual Subspace of Policies** (**CSP**).

Our model builds a sequence of subspaces $\Theta_1, ..., \Theta_N$, one new subspace after each training task, with each subspace extending the previous one. Note that each subspace Θ_j is a collection of at most j anchors (or neural networks) *i.e.* $|\Theta_j| \leq j, \ \forall \ j \in 1...N$ since the subspace grows sublinearly with the number of tasks (as explained below). Hence, the number of anchors m of a subspace Θ_j is not the same as the number of tasks j used to create Θ_j. The learned policies are represented as single points in these subspaces. At each stage, CSP maintains both a set of anchors defining the current subspace, as well as the weights α corresponding to the best policies found for all prior tasks. The best found policy for task i after training on the first j tasks, $\forall \ i \leq j$, is denoted as $\pi_i^j(a|s)$ and can be represented as a point in the subspace Θ_j with a weight vector denoted α_i^j such that $\pi_i^j(a|s) = \pi(a|s, [\alpha_i^j, \Theta_j])$, where $|\alpha_i^j| = |\Theta_j| \leq j$.

Given a set of anchors Θ_j and a set of previously learned policies $\{\alpha_1^j, \dots, \alpha_j^j\}$, updating our model on the new task t_{j+1} produces a new subspace Θ_{j+1} and a new set of weights $\{\alpha_1^{j+1}, \dots, \alpha_{j+1}^{j+1}\}$. There are two possible cases. One possibility is that the current subspace already contains a good policy for t_{j+1}, so we just need to find the weight vector corresponding to a policy which performs well on t_{j+1}. The other possibility is that the current subspace does not contain a good policy for t_{j+1}. In this case, the algorithm produces a new subspace by adding one anchor to the previous one (see next section), and converts the previous policies to be compatible with this new subspace.

To achieve this, CSP operates in two phases:

1. *Grow* the current subspace by adding a new anchor and learning the best possible policy for task $j + 1$ in this subspace (where the previous j anchors are frozen).
2. Compare the quality of this policy with the best possible policy expressed in the previous subspace. Based on this comparison, decide whether to *extend* the subspace to match the new one or *prune* it back to the previous one.

We now describe in detail the two phases of the learning process, namely how we grow the subspace and how we decide whether to extend or prune it (see Algorithm 1 for a pseudo-code).

3.3 Grow the Subspace

Given the current subspace Θ_j composed of $m \leq j$ anchors (with j being the number of tasks seen so far), a new subspace $\tilde{\Theta}_{j+1}$ is built as follows. First a new anchor denoted θ_{j+1} is added to the set of anchors such that $\tilde{\Theta}_{j+1} = \Theta_j \bigcup \{\theta_{j+1}\}$. With all previous anchors frozen, we train the new anchor by sampling α values from a Dirichlet distribution parameterized by a uniform vector with size $j + 1$, $Dir\big(\mathcal{U}(j+1)\big)$. The new anchor θ_{j+1} is updated by maximizing the expected return obtained by interacting with the new task $j + 1$ using policies defined by the sampled α's:

$$\theta_{j+1} = \arg\max_{\theta} \mathbb{E}_{\alpha \sim Dir, \, \tau \sim \pi(a|s, [\alpha, \Theta_j \bigcup \{\theta\}])}[R_{j+1}(\tau)], \tag{1}$$

where $R_{j+1}(\tau)$ is the return obtained on task $j + 1$ throughout trajectory τ which was generated using policy $\pi(a|s, [\alpha, \tilde{\Theta}_{j+1})]$. Note that the anchor is trained such that not only one but all possible values of α tend to produce a good policy. The resulting subspace thus aims at containing as many good policies as possible for task $j + 1$.

3.4 Extend or Prune the Subspace

To decide if the new anchor is kept, we propose to simply compare the best possible policy for task $j + 1$ in the new subspace with the best possible policy for the same task in the previous subspace (*i.e.* without using the new anchor). Each policy could be evaluated via Monte-Carlo (MC) estimates by doing additional rollouts in the environment and recording the average performance. However,

Algorithm 1: Continual Subspace of Policies (CSP)

Input: $\theta_1,\ldots,\theta_j$ (previous anchors), ϵ (threshold)
Initialize: W_ϕ (subspace critic), $\mathcal{B}uf$ (replay buffer)

Initialize $\theta_{j+1} \leftarrow \frac{1}{j}\sum_{i=1}^{j}\theta_i$ (new anchor) ; `// Grow the Subspace`
for $i = 1, ..., B$ **do**
  Sample $\alpha \sim Dir\big(\mathcal{U}(j+1)\big)$
  Set policy parameters $\theta_\alpha \leftarrow \sum_{i=1}^{j+1}\alpha_i\theta_i$
  for $l = 1, ..., K$ **do**
    Collect and store (s, a, r, s', α) in $\mathcal{B}uf$ by sampling $a \sim \pi_{\theta_\alpha}(\cdot|s)$
  end
  if *time to update* **then**
    Update $\pi_{\theta_{j+1}}$ and W_ϕ using the SAC algorithm and the replay buffer $\mathcal{B}uf$
  end
end
Use $\mathcal{B}uf$ and W_ϕ to estimate: ; `// Extend or Prune the Subspace`

$$\alpha^{\text{old}} \leftarrow \underset{(\alpha,0)\text{ with }\alpha\in\mathbb{R}_+^m,\|\alpha\|_1=1}{\arg\max} W_\phi(\alpha)$$
$$\alpha^{\text{new}} \leftarrow \underset{\alpha\in\mathbb{R}_+^{m+1},\|\alpha\|_1=1}{\arg\max} W_\phi(\alpha)$$

if $W_\phi(\cdot,\alpha^{\text{new}}) > (1+\epsilon)\cdot W_\phi(\cdot,\alpha^{\text{old}})$ **then**
  Return: $\theta_1,\ldots,\theta_j,\theta_{j+1},\alpha^{\text{new}}$; `// Extend`
else
  Return: $\theta_1,\ldots,\theta_j,\alpha^{\text{old}}$; `// Prune`
end

this typically requires a large number (*e.g.* millions) of interactions which may be impractical with a limited budget. Thus, we propose an alternative procedure to make this evaluation sample efficient.

For each task j and corresponding subspace Θ_j, our algorithm also learns a Q-function $Q(s,a,\alpha)$ which is trained to predict the expected return on task j for all possible states, actions, and all possible α's in the corresponding subspace. Our algorithm is based on SAC (Haarnoja et al., 2018b), so for each task, we collect a replay buffer of interactions $\mathcal{B}uf$ which contains all states and actions seen by the agent while training on that task. Thus, the Q-function $Q(s,a,\alpha)$ can help us directly estimate the quality $W(\alpha)$ of the policy represented by the weight vector α in the new subspace which can be computed as the average over all states and actions in the replay buffer:

$$W(\alpha) = \mathbb{E}_{s,a\sim\mathcal{B}uf}Q(s,a,\alpha). \tag{2}$$

It is thus possible to compute the value of α corresponding to the best policy in the extended subspace (denoted $\alpha^{\text{new}} \in \mathbb{R}_+^{m+1}$, $\|\alpha^{\text{new}}\|_1 = 1$):

$$\alpha^{\text{new}} = \underset{\alpha\in\mathbb{R}_+^{m+1},\|\alpha\|_1=1}{\arg\max} W(\alpha), \tag{3}$$

as well as the value of α corresponding to the best policy in the previous subspace (denoted $\alpha^{\text{old}} \in \mathbb{R}_+^m$, $\|\alpha^{\text{old}}\|_1 = 1$):

$$\alpha^{\text{old}} = \underset{(\alpha,0)\text{ with }\alpha\in\mathbb{R}_+^m,\|\alpha\|_1=1}{\arg\max} W(\alpha). \tag{4}$$

In practice, α^{new} and α^{old} are estimated by uniformly sampling a number of α's in the corresponding subspace as well as a number of states and actions from the buffer.

The quality of the new subspace and the previous one can thus be evaluated by comparing $W(\alpha^{\text{new}})$ and $W(\alpha^{\text{old}})$. If $W(\alpha^{\text{new}}) > (1+\epsilon)\cdot W(\alpha^{\text{old}})$, the subspace is extended to the new subspace (*i.e.* the one after the grow phase): $\Theta_{j+1} = \tilde{\Theta}_{j+1}$. Otherwise, the subspace is pruned back to the old subspace (*i.e.* the one before the growth phase): $\Theta_{j+1} = \Theta_j$. Note that, if the subspace is extended, the previously learned policies have to be mapped in the new subspace such that $\alpha_i^{j+1} \in \mathbb{R}_+^{j+1}$ *i.e.* $\forall\ i \leq j,\ \alpha_i^{j+1} := (\alpha_i^j, 0)$ and $\alpha_{j+1}^{j+1} := \alpha^{\text{new}}$. If the subspace is not extended, then old values can be kept *i.e.* $\forall\ i \leq j,\ \alpha_i^{j+1} := \alpha_i^j$ and $\alpha_{j+1}^{j+1} := \alpha^{\text{old}}$.

After finding the best α for the current task, the replay buffer and Q-function are reinitialized. Hence, memory cost for training the Q-function is constant and there is no need to store data from previous tasks unlike other CRL approaches (Rolnick et al., 2019). The policy $\pi_{\theta_{j+1}}$ and value function W_ϕ are updated using SAC (Haarnoja et al., 2018b). See Appendix C for more details.

3.5 Scalability of CSP

By having access to an large number of policies, the subspace is highly expressive so it can capture a wide range of diverse behaviors. This enables positive transfer to many new tasks without the need for additional parameters. As a consequence, *the number of parameters increases sublinearly with the the number of tasks*. The final size of the subspace depends on the sequence of tasks, with longer and more diverse sequences requiring more anchors. The speed of growth is controlled by the threshold ϵ which trades-off performance gains for memory efficiency (the higher the ϵ the more performance losses are tolerated to reduce memory costs). See Figure 2a and Appendix C.1.

4 Experiments

4.1 Environments

We evaluate CSP on 18 CRL scenarios containing 35 different RL tasks, from two continuous control domains, *locomotion* in **Brax** (Freeman et al., 2021) and *robotic manipulation* in **Continual World** (CW, Wołczyk et al. (2021)), a challenging CRL benchmark. For CW, we run experiments on both the proposed sequence of 10 tasks (CW10), and on all 8 triplet sequences (CW3). Each task in these sequences has a *different reward* function. The goal of these experiments is to compare our approach with popular CRL methods on a well-established benchmark.

For Brax, we create new CRL scenarios based on 3 subdomains: **HalfCheetah**, **Ant**, and **Humanoid**. Each scenario has 8 tasks and each task has *specific dynamics*. We use a budget of 1M interactions for each task. The goal of these experiments is to perform an in-depth study to separately evaluate capabilities specific to CRL agents such as *forgetting, transfer, robustness, and compositionality* (see Appendix B.2 and Tables 5a, 5b for more information). For each of the 4 capabilities, we create 2 CRL scenarios, one based on HalfCheetah and one based on Ant. To further probe the effectiveness of our approach, we also create one CRL scenario with 4 tasks on the challenging Humanoid domain. Here we use a budget of 2M interactions for each task.

The CRL scenarios are created following the protocol introduced in Wołczyk et al. (2021) which proposes a systematic way of generating task sequences that test CRL agents along particular axes. While Wołczyk et al. (2021) focus on transfer and forgetting (see Appendix B.3), we also probe robustness (*e.g.* adapting to environment perturbations such as action reversal), and compositionality (*e.g.* combining two previously learned skills to solve a new task). Each task in a sequence has different dynamics which are grounded in quasi-realistic situations such as increased or decreased gravity, friction, or limb lengths. See Appendix B.1 for details on how these tasks were designed.

4.2 Baselines

We compare CSP with a number of popular CRL baselines such as PNN (Rusu et al., 2016b), EWC (Kirkpatrick et al., 2017), PackNet (Mallya & Lazebnik, 2018), FT-1 which finetunes a single model on the entire sequence of tasks, and FT-L2 which is like FT-1 with an additional L_2 regularization applied during finetuning. We also compare with SAC-N which trains one model for each task from scratch. While SAC-N avoids forgetting, it cannot transfer knowledge across tasks. Finally, we compare with a method called FT-N which combines the best of both SAC-N and FT-1. Just like SAC-N, it stores one model per task after training on it and just like FT-1, it finetunes the previous model to promote transfer. However, FT-N and SAC-N scale poorly (*i.e.* linearly) with the number of tasks in terms of both memory and compute, which makes them unfeasible for real-world applications. Note that our method is not directly comparable with CLEAR (Rolnick et al., 2019) since we assume no access to data from prior tasks. Storing data from all prior tasks (as CLEAR does) is unfeasible for long task sequences due to prohibitive memory costs. All methods use SAC (Haarnoja et al., 2018b) as a base algorithm. The means and standard deviations are computed over 10 seeds unless otherwise noted. See Appendices A.1, A.2, and A.4 for more details about our protocol, baselines, and hyperparameters, respectively.

Table 1: Aggregated results across all Brax scenarios from HalfCheetah, Ant, and Humanoid. These scenarios were designed to test forgetting, transfer, compositionality, and robustness. CSP performs as well as or better than the strongest baselines, while having a much lower model size and thus memory cost. CSP's performance is also not too far from that of CSP-ORACLE indicating that it can use the critic to find good policies in the subspace without requiring millions of interactions.

	HalfCheetah (4 scenarios)		**Ant** (4 scenarios)		**Humanoid** (1 scenario)	
Method	Performance	Model Size	Performance	Model Size	Performance	Model Size
FT-1	0.62 ± 0.29	1.0	0.52 ± 0.26	1.0	0.71 ± 0.07	1.0
FT-L2	0.38 ± 0.15	2.0	0.78 ± 0.20	2.0	0.68 ± 0.28	2.0
PACKNET	0.85 ± 0.14	2.0	1.08 ± 0.21	2.0	0.96 ± 0.21	2.0
EWC	0.43 ± 0.24	3.0	0.55 ± 0.24	3.0	0.94 ± 0.01	3.0
PNN	1.03 ± 0.14	8.0	0.98 ± 0.31	8.0	0.98 ± 0.26	4.0
SAC-N	1.00 ± 0.15	8.0	1.00 ± 0.38	8.0	1.00 ± 0.29	4.0
FT-N	1.16 ± 0.20	8.0	0.97 ± 0.20	8.0	0.65 ± 0.46	4.0
CSP (ours)	$\mathbf{1.27 \pm 0.27}$	5.4 ± 1.3	$\mathbf{1.11 \pm 0.17}$	3.9 ± 0.8	$\mathbf{1.76 \pm 0.19}$	3.4 ± 0.3
CSP-ORACLE	1.88 ± 0.19		1.24 ± 0.07		1.98 ± 0.22	

4.3 ABLATIONS

We also perform a number of ablations to understand how the different components of CSP influence performance. We first compare CSP with CSP-ORACLE which selects the best policy by sampling a large number of policies in the subspace and computing Monte-Carlo estimates of their returns. These estimates are expected to be more accurate than the critic's, so CSP-ORACLE can be considered an upper bound to CSP. However, CSP-ORACLE is less efficient than CSP since it requires singificantly more interactions with the environment to find a policy for each task (*i.e.* millions).

We also want to understand how much performance we lose, if any, by not adding one anchor per task. To do this, we run CSP-LINEAR which always extends the subspace by adding one anchor for each new task. In addition, we vary the threshold ϵ used to decide whether to extend the subspace based on how much performance is gained by doing so. This analysis can shed more light on the trade-off between performance gain and memory cost as the threshold varies.

4.4 METRICS

Agents are evaluated across a range of commonly used CRL metrics (Wołczyk et al., 2021; Powers et al., 2021). For brevity, following the notation in Section 2, we will use $\pi_i^j \coloneqq \Pi(i, [t_1, \ldots, t_j])$ to denote the policy selected by Π for task i obtained after training on the sequence of tasks $t_1, ..., t_j$ (in this order with j included). Note that Π can also be defined for a single task *i.e.* $\pi_i^{t_i} \coloneqq \Pi(i, [t_i])$ is the policy for task i after training the system only on task t_i (for the corresponding budget b_i). We also report forgetting and forward transfer (see Appendix A.3 for their formal definitions).

Average Performance is the average performance of the method across all tasks, after training on the entire sequence of tasks. Let $P_i(\pi)$ be the performance of the system on task i using policy π defined as $P_i(\pi) \coloneqq \mathbb{E}_{\pi, \mathcal{T}_i}\left[\sum r_i\left(s, a\right)\right]$, where $r_i(s, a)$ is the reward on task i when taking action a in state s. Then, the final average performance across all tasks can be computed as $P \coloneqq \frac{1}{N}\sum_{i=1}^{N} P_i(\pi_i^N)$.

Model Size is the final number of parameters of a method after training on all tasks, divided by the number of parameters of FT-1 (which is a single policy trained on all tasks, with size equal to that of an anchor). For example, in the case of CSP, the model size is equivalent with the number of anchors defining the final subspace $|\Theta_N|$. Similarly, in the case of FT-N and SAC-N, the model size is always equal to the number of tasks N. Following Rusu et al. (2016b), we cap PNN's model size to N since otherwise it grows quadratically with the number of tasks.

5 RESULTS

5.1 PERFORMANCE ON BRAX

Table 1 shows the aggregated results across all scenarios from HalfCheetah, Ant, and Humanoid. The results are normalized using SAC-N's performance. See Appendix D.1 for individual results.

Table 2: Results on the CW10 benchmark for CSP and other popular baselines including the state-of-the-art PACKNET (results taken from Wołczyk et al. (2021)). CSP performs almost as well as PACKNET while using about half the number of heads, and thus having a much lower memory cost.

Method	CSP (ours)	PACKNET	EWC	FT-L2	FT-1
Performance	0.81 ± 0.06	$\mathbf{0.83 \pm 0.02}$	0.66 ± 0.03	0.48 ± 0.05	0.10 ± 0.01
# Heads	$\mathbf{5.3 \pm 1.6}$	10	10	10	10

CSP performs as well as or better than the strongest baselines which grow linearly with the number of tasks (*i.e.* PNN, SAC-N, FT-N), while maintaining a much smaller size and thus memory cost. In fact, CSP's size grows sublinearly with the number of tasks. At the same time, CSP is significantly better than the most memory-efficient baselines that have a relatively small size (*i.e.* FT-1, FT-L2, EWC, PACKNET). Naive methods like FT-1 have low memory costs and good transfer, but suffer from catastrophic forgetting. In contrast, methods that aim to reduce forgetting such as FT-L2 or EWC do so at the expense of transfer. The only competitive methods in terms of performance are the ones where the number of parameters increases at least linearly with the number of tasks, such as PNN, SAC-N, and FT-N. These methods have no forgetting because they store the models trained on each task. SAC-N has no transfer since it trains each model from scratch, while FT-N promotes transfer as it finetunes the previous model. However, due to their poor scalability, these methods are unfeasible for more challenging CRL scenarios with long and varied task sequences.

In contrast, CSP doesn't suffer from forgetting since the best found policies for all prior tasks can be cheaply stored in the form of convex combinations of the subspace's anchors (*i.e.* vectors rather than model weights). In addition, due to having access to a large number of policies (*i.e.* all convex combinations of the anchors), CSP has good transfer to new tasks (that share some similarities with prior ones). This allows CSP to achieve strong performance while its memory cost grows sublinearly with the number of tasks. Nevertheless, CSP's performance is not too far from that of CSP-ORACLE, indicating that the critic can be used to find good policies in the subspace without requiring millions of interactions. See Appendix C.2 for more details about the use of the critic.

5.2 PERFORMANCE ON CONTINUAL WORLD

Table 2 shows results on CW10 (Wołczyk et al., 2021), a popular CRL benchmark. The baselines used in Wołczyk et al. (2021) use a separate linear layer for each task on top of a common network. The authors recognize this as a limitation since it scales poorly with the number of tasks. For a fair comparison, we implement CSP using a similar protocol where the anchors are represented by linear heads, so the number of heads grows adaptively depending on the performance threshold ϵ. Instead of model size, we compare the final number of heads, since this now determines a method's scalability with the number of tasks. See Appendix A for more details about the experimental setup.

CSP performs almost as well as PACKNET which is the current state-of-the-art on CW10 (Wołczyk et al., 2021), and is significantly better than all other baselines. At the same time, CSP uses about half the number of heads, thus being more memory efficient especially as the number of tasks increases. Note that PACKNET suffers from a major limitation, namely that it requires prior knowledge of the total number of tasks in order to allocate resources. If this information is not correctly specified, PACKNET is likely to fail due to either not being expressive enough to handle many tasks or being too inefficient while learning only a few tasks (Wołczyk et al., 2021). This makes PACKNET unfeasible for real-world applications where agents can face task sequences of varying lengths (including effectively infinite ones). In contrast, CSP grows adaptively depending on the sequence of tasks, so it can handle both short and long sequences without any modification to the algorithm. See Appendix D.2 for additional results on Continual World, including CW3.

To summarize, CSP *is competitive with the strongest CRL baselines, while having a lower memory cost since its size grows sublinearly with the number of tasks. Thus,* CSP *maintains a good balance between model size and performance, which allows it to scale well to long task sequences.*

5.3 ANALYSIS OF CSP

Varying the Threshold. Figure 2a shows how performance and size vary with the threshold ϵ used to decide whether to extend the subspace or not. Both the performance and the size are normalized with respect to CSP-LINEAR which always extends the subspace. As expected, as ϵ increases, performance decreases, but so does the size. Note that *performance is still above* 70% *even as the*

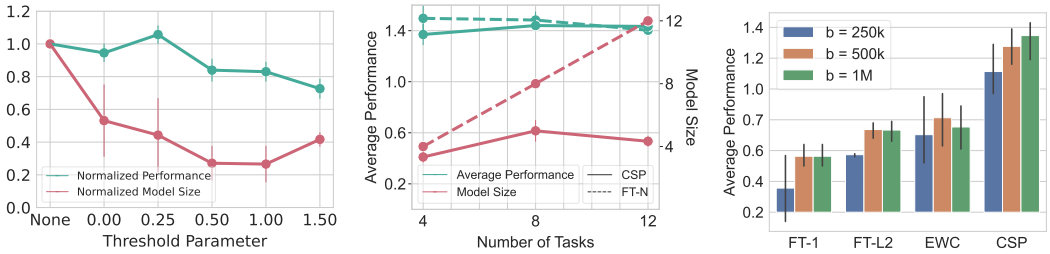

(a) Varying the Threshold (b) Scalability (c) Learning Efficiency

Figure 2: (a) shows how CSP's performance and size vary with the threshold, (b) shows how performance and size scale with the number of tasks for CSP and FT-N, and (c) compares the performance of multiple methods using different budgets. See Appendix D.3 for more details.

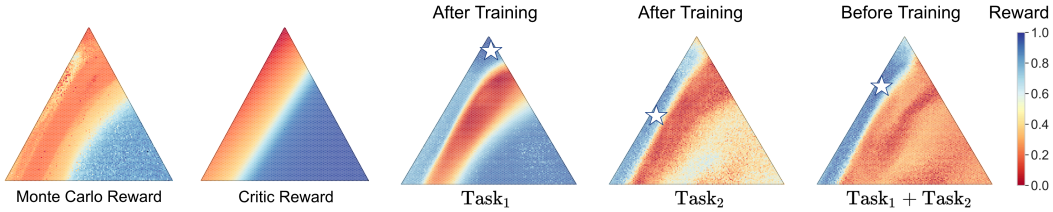

(a) Smoothness and Critic Accuracy (b) Diversity and Compositionality

Figure 3: (a) shows the value of each policy in the subspace estimated using both Monte-Carlo simulations (left) and our critic's predictions (right), demonstrating both that the subspace is smooth and that the critic learns accurate estimates of the reward, which allows CSP to find good policies in the subspace; (b) shows the subspace for three different tasks, the third being a combination of the first two (*e.g.* walk on the moon (task 1) with tiny feet (task 2)), demonstrating that the subspace already contains a policy with high reward on the compositional task before being trained on it. The star represents the best policy in the subspace for the corresponding task. The subspace contains policies covering the whole spectrum of rewards, suggesting that it captures diverse behaviors. Click here to see interactive visualizations of the subspace.

size is cut by more than 70%, relative to CSP-LINEAR which trains a new set of parameters (*i.e.* anchor) for each task. Thus, CSP *can drastically reduce memory costs without significantly hurting performance*. Practitioners can set the threshold to trade-off performance gains for memory efficiency *i.e.* the higher the ϵ the more performance losses are tolerated in order to reduce memory costs.

Scalability. Figure 2b shows how performance and size scale with the number of tasks in the sequence (on HalfCheetah's compositionality scenario) for both CSP and FT-N which is our strongest baseline on this scenario. CSP maintains both strong performance and small size (*i.e.* low memory cost) even as the number of tasks increases. In contrast, even if FT-N performs well on all these scenarios, its size grows linearly with the number of tasks, rendering it impractical for long task sequences.

Learning Efficiency. Figure 2c shows the average performance for three different budgets (*i.e.* number of interactions allowed for each task) on HalfCheetah's robustness scenario, comparing CSP with FT-1, FT-L2, and EWC. These results demonstrate that CSP can learn efficiently even with a reduced budget, while still outperforming these baselines. By keeping track of all convex combinations of previously learned behaviors, CSP enables good transfer to new tasks which in turn leads to efficient training on task sequences.

5.4 Analysis of the Subspace

Smoothness and Critic Accuracy. Figure 3a shows a snapshot of a trained subspace, along with the expected reward of all policies in the subspace, for a given task. The expected reward is computed using both Monte-Carlo (MC) rollouts, as well as our critic's predictions using Equation 2. As illustrated, the learned Q-function has a similar landscape with the MC reward *i.e.* the subspace is smooth and the critic's estimates are accurate.

Diversity and Compositionality. Figure 3b illustrates that the subspace contains behaviors composed of previously learned skills (*e.g.* walk on the moon, walk with tiny feet, walk on the moon with

tiny feet). This allows CSP to reuse previously learned skills to find good policies for new tasks without the need for additional parameters. The figure also shows that for a given task, the policies in the subspace cover the whole spectrum of rewards, thus emphasizing the diversity of the behaviors expressed by the subspace. See Appendix C.5 for more analysis of the subspace.

6 Related Work

Continual Reinforcement Learning. CRL methods aim to avoid catastrophic forgetting, as well as enable transfer to new tasks, while remaining scalable to a large number of tasks. In the past few years, multiple approaches have been proposed to address one or more of these challenges (Tessler et al., 2017; Parisi et al., 2019; Khetarpal et al., 2020; Powers et al., 2021; Wołczyk et al., 2021; Tessler et al., 2017; Berseth et al., 2022; Nagabandi et al., 2018; Xie et al., 2020; Xu et al., 2020; Berseth et al., 2021; Ren et al., 2022; Kessler et al., 2022).

One class of methods focuses on preventing catastrophic *forgetting*. Some algorithms achieve this by storing the parameters of models trained on prior tasks (Rusu et al., 2016b; Cheung et al., 2019; Wortsman et al., 2020). However, these methods scale poorly with the number of tasks (*i.e.* at least linearly) in both compute and memory, which makes them unfeasible for more realistic scenarios. Other methods maintain a buffer of experience from prior tasks to alleviate forgetting (Lopez-Paz & Ranzato, 2017; Isele & Cosgun, 2018; Riemer et al., 2019b; Rolnick et al., 2019; Caccia et al., 2020a). However, this is also not scalable as the memory cost increases significantly with the number and complexity of the tasks. In addition, many real-world applications in domains like healthcare or insurance prevent data storage due to privacy or ethical concerns. Another class of methods focuses on improving *transfer* to new tasks. Naive approaches like finetuning that train a single model on each new task provide good scalability and plasticity, but suffer from catastrophic forgetting. To overcome this effect, methods like elastic weight consolidation (EWC) (Kirkpatrick et al., 2017) alleviate catastrophic forgetting by constraining how much the network's weights change, but this can in turn reduce plasticity. Another class of methods employs knowledge distillation to improve transfer in CRL (Hinton et al., 2006; Rusu et al., 2016a; Li & Hoiem, 2018; Schwarz et al., 2018; Bengio & LeCun, 2007; Kaplanis et al., 2019; Traoré et al., 2019). However, since these methods train a single network, they struggle to capture a large number of diverse behaviors.

There is also a large body of work which leverages the shared structure of the tasks (Xu & Zhu, 2018; Pasunuru & Bansal, 2019; Lu et al., 2021; Mankowitz et al., 2018; Abel et al., 2017; Sodhani et al., 2021), meta-learning (Javed & White, 2019; Spigler, 2019; Beaulieu et al., 2020; Caccia et al., 2020b; Riemer et al., 2019a; Zhou et al., 2020; Schmidhuber, 2013), or generative models (Robins, 1995; Silver et al., 2013; Shin et al., 2017; Atkinson et al., 2021) to improve CRL agents, but these methods don't work very well in practice. Lee et al. (2020) use an iteratively expanding network but they assume no task boundaries and only consider the supervised learning setting.

Mode Connectivity and Neural Network Subspaces. Based on prior studies of mode connectivity (Garipov et al., 2018; Draxler et al., 2018; Mirzadeh et al., 2021), Wortsman et al. (2021); Benton et al. (2021) proposed the neural network subspaces to connect the solutions of different models. More similar to our work, Doan et al. (2022) leverages subspace properties to mitigate forgetting on a sequence of supervised learning tasks, but doesn't focus on other aspects of the continual learning problem. Our work was also inspired by Gaya et al. (2021) which learns a subspace of policies for a single task for fast adaptation to new environment. In contrast to Gaya et al. (2021), we consider the continual RL setting and adaptively grow the subspace as the agent encounters new tasks. Our work is first to demonstrate the effectiveness of mode connectivity for the continual RL setting.

7 Discussion

In this paper, we propose CSP, a new continual RL method which adaptively builds a subspace of policies to learn a sequence of tasks. CSP is competitive with the best existing methods while using significantly fewer parameters. Thus, it strikes a good balance between performance and memory cost which allows it to scale well to long task sequences. Our paper is first to use a subspace of policies for continual RL, and thus opens up many interesting directions for future work. For example, one can assume a fixed size for the subspace and update all the anchors whenever the agent encounters a new task. Another promising direction is to leverage the structure of the subspace to meta-learn or search for good policies on a given task. Finally, active learning techniques could be employed to more efficiently evaluate policies contained in the subspace.

REFERENCES

David Abel, Dilip Arumugam, Lucas Lehnert, and Michael L. Littman. Toward good abstractions for lifelong learning. 2017.

Craig Atkinson, B. McCane, Lech Szymanski, and Anthony V. Robins. Pseudo-rehearsal: Achieving deep reinforcement learning without catastrophic forgetting. *ArXiv*, abs/1812.02464, 2021.

Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization, 2016. URL https://arxiv.org/abs/1607.06450.

Shawn L. E. Beaulieu, Lapo Frati, Thomas Miconi, Joel Lehman, Kenneth O. Stanley, Jeff Clune, and Nick Cheney. Learning to continually learn. *ArXiv*, abs/2002.09571, 2020.

Yoshua Bengio and Yann LeCun. Scaling learning algorithms towards AI. In *Large Scale Kernel Machines*. MIT Press, 2007.

Gregory W. Benton, Wesley Maddox, Sanae Lotfi, and Andrew Gordon Wilson. Loss surface simplexes for mode connecting volumes and fast ensembling. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pp. 769–779. PMLR, 2021.

Glen Berseth, Cheng Xie, Paul Cernek, and Michiel Van de Panne. Progressive reinforcement learning with distillation for multi-skilled motion control. In *International Conference on Learning Representations (ICLR) 2018*, 2018.

Glen Berseth, Zhiwei Zhang, Grace Zhang, Chelsea Finn, and Sergey Levine. Comps: Continual meta policy search. *arXiv preprint arXiv:2112.04467*, 2021.

Glen Berseth, Zhiwei Zhang, Grace Zhang, Chelsea Finn, and Sergey Levine. CoMPS: Continual meta policy search. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=PVJ6j87gOHz.

Lucas Caccia, Eugene Belilovsky, Massimo Caccia, and Joelle Pineau. Online learned continual compression with adaptive quantization modules. *Proceedings of the 37th International Conference on Machine Learning*, 2020a.

Massimo Caccia, Pau Rodriguez, Oleksiy Ostapenko, Fabrice Normandin, Min Lin, Lucas Caccia, Issam Laradji, Irina Rish, Alexandre Lacoste, David Vazquez, and Laurent Charlin. Online fast adaptation and knowledge accumulation: a new approach to continual learning. *NeurIPS*, 2020b. URL https://arxiv.org/abs/2003.05856.

Brian Cheung, Alex Terekhov, Yubei Chen, Pulkit Agrawal, and Bruno A. Olshausen. Superposition of many models into one. *ArXiv*, abs/1902.05522, 2019.

Thang Doan, Seyed Iman Mirzadeh, Joelle Pineau, and Mehrdad Farajtabar. Efficient continual learning ensembles in neural network subspaces. *arXiv preprint arXiv:2202.09826*, 2022.

Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. Essentially no barriers in neural network energy landscape. In *ICML 2018: Thirty-fifth International Conference on Machine Learning*, pp. 1308–1317, 2018.

C Daniel Freeman, Erik Frey, Anton Raichuk, Sertan Girgin, Igor Mordatch, and Olivier Bachem. Brax–a differentiable physics engine for large scale rigid body simulation. *arXiv preprint arXiv:2106.13281*, 2021.

T. Garipov, Pavel Izmailov, Dmitrii Podoprikhin, Dmitry P. Vetrov, and Andrew Gordon Wilson. Loss surfaces, mode connectivity, and fast ensembling of dnns. In *NeurIPS*, 2018.

Jean-Baptiste Gaya, Laure Soulier, and Ludovic Denoyer. Learning a subspace of policies for online adaptation in reinforcement learning. *ArXiv*, abs/2110.05169, 2021.

Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications, 2018a. URL https://arxiv.org/abs/1812.05905.

Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2018b.

Geoffrey E. Hinton, Simon Osindero, and Yee Whye Teh. A fast learning algorithm for deep belief nets. *Neural Computation*, 18:1527–1554, 2006.

David Isele and Akansel Cosgun. Selective experience replay for lifelong learning. In *AAAI*, 2018.

Khurram Javed and Martha White. Meta-learning representations for continual learning. In *NeurIPS*, 2019.

Christos Kaplanis, Murray Shanahan, and Claudia Clopath. Policy consolidation for continual reinforcement learning. *ArXiv*, abs/1902.00255, 2019.

Samuel Kessler, Jack Parker-Holder, Philip Ball, Stefan Zohren, and Stephen J Roberts. Same state, different task: Continual reinforcement learning without interference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 7143–7151, 2022.

Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. Towards continual reinforcement learning: A review and perspectives. *ArXiv*, abs/2012.13490, 2020.

James Kirkpatrick, Razvan Pascanu, Neil C. Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114:3521 – 3526, 2017.

Soochan Lee, Junsoo Ha, Dongsu Zhang, and Gunhee Kim. A neural dirichlet process mixture model for task-free continual learning. *arXiv preprint arXiv:2001.00689*, 2020.

Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:2935–2947, 2018.

David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In *NIPS*, 2017.

Kevin Lu, Aditya Grover, P. Abbeel, and Igor Mordatch. Reset-free lifelong learning with skill-space planning. *ArXiv*, abs/2012.03548, 2021.

Song Duong Jean-Baptiste Gaya Pierre-Alexandre Kamienny Daniel H. Thompson Ludovic Denoyer, Alfredo de la Fuente. Salina: Sequential learning of agents. https://gitHub.com/facebookresearch/salina, 2021.

Arun Mallya and Svetlana Lazebnik. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pp. 7765–7773, 2018.

Daniel Jaymin Mankowitz, Augustin Zídek, André Barreto, Dan Horgan, Matteo Hessel, John Quan, Junhyuk Oh, H. V. Hasselt, David Silver, and Tom Schaul. Unicorn: Continual learning with a universal, off-policy agent. *ArXiv*, abs/1802.08294, 2018.

Seyed Iman Mirzadeh, Mehrdad Farajtabar, Dilan Gorur, Razvan Pascanu, and Hassan Ghasemzadeh. Linear mode connectivity in multitask and continual learning. *ArXiv*, abs/2010.04495, 2021.

Anusha Nagabandi, Chelsea Finn, and Sergey Levine. Deep online learning via meta-learning: Continual adaptation for model-based rl. *arXiv preprint arXiv:1812.07671*, 2018.

German Ignacio Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks : the official journal of the International Neural Network Society*, 113:54–71, 2019.

Ramakanth Pasunuru and Mohit Bansal. Continual and multi-task architecture search. *ArXiv*, abs/1906.05226, 2019.

Sam Powers, Eliot Xing, Eric Kolve, Roozbeh Mottaghi, and Abhinav Kumar Gupta. Cora: Benchmarks, baselines, and metrics as a platform for continual reinforcement learning agents. *ArXiv*, abs/2110.10067, 2021.

Hang Ren, Aivar Sootla, Taher Jafferjee, Junxiao Shen, Jun Wang, and Haitham Bou-Ammar. Reinforcement learning in presence of discrete markovian context evolution. *arXiv preprint arXiv:2202.06557*, 2022.

Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to learn without forgetting by maximizing transfer and minimizing interference. *ArXiv*, abs/1810.11910, 2019a.

Matthew Riemer, Tim Klinger, Djallel Bouneffouf, and Michele M. Franceschini. Scalable recollections for continual lifelong learning. In *AAAI*, 2019b.

Anthony V. Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connect. Sci.*, 7:123–146, 1995.

David Rolnick, Arun Ahuja, Jonathan Schwarz, Timothy P. Lillicrap, and Greg Wayne. Experience replay for continual learning. In *NeurIPS*, 2019.

Andrei A. Rusu, Sergio Gomez Colmenarejo, Çaglar Gülçehre, Guillaume Desjardins, James Kirkpatrick, Razvan Pascanu, Volodymyr Mnih, Koray Kavukcuoglu, and Raia Hadsell. Policy distillation. *CoRR*, abs/1511.06295, 2016a.

Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive neural networks. *ArXiv*, abs/1606.04671, 2016b.

Jürgen Schmidhuber. Powerplay: Training an increasingly general problem solver by continually searching for the simplest still unsolvable problem. *Frontiers in Psychology*, 4, 2013.

Jonathan Schwarz, Wojciech M. Czarnecki, Jelena Luketina, Agnieszka Grabska-Barwinska, Yee Whye Teh, Razvan Pascanu, and Raia Hadsell. Progress & compress: A scalable framework for continual learning. *ArXiv*, abs/1805.06370, 2018.

Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In *NIPS*, 2017.

Daniel L. Silver, Qiang Yang, and Lianghao Li. Lifelong machine learning systems: Beyond learning algorithms. In *AAAI Spring Symposium: Lifelong Machine Learning*, 2013.

Shagun Sodhani, Franziska Meier, Joelle Pineau, and Amy Zhang. Block contextual mdps for continual learning. *arXiv preprint arXiv:2110.06972*, 2021.

Giacomo Spigler. Meta-learnt priors slow down catastrophic forgetting in neural networks. *ArXiv*, abs/1909.04170, 2019.

Chen Tessler, Shahar Givony, Tom Zahavy, Daniel Mankowitz, and Shie Mannor. A deep hierarchical approach to lifelong learning in minecraft. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.

Kalifou René Traoré, Hugo Caselles-Dupré, Timothée Lesort, Te Sun, Guanghang Cai, Natalia Díaz Rodríguez, and David Filliat. Discorl: Continual reinforcement learning via policy distillation. *ArXiv*, abs/1907.05855, 2019.

Tom Veniat, Ludovic Denoyer, and Marc'Aurelio Ranzato. Efficient continual learning with modular networks and task-driven priors. *arXiv preprint arXiv:2012.12631*, 2020.

Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. *ArXiv*, abs/2006.14769, 2020.

Mitchell Wortsman, Maxwell Horton, Carlos Guestrin, Ali Farhadi, and Mohammad Rastegari. Learning neural network subspaces. *ArXiv*, abs/2102.10472, 2021.

Maciej Wołczyk, Michal Zajkac, Razvan Pascanu, Lukasz Kuci'nski, and Piotr Milo's. Continual world: A robotic benchmark for continual reinforcement learning. *ArXiv*, abs/2105.10919, 2021.

Annie Xie, James Harrison, and Chelsea Finn. Deep reinforcement learning amidst lifelong non-stationarity. *arXiv preprint arXiv:2006.10701*, 2020.

Ju Xu and Zhanxing Zhu. Reinforced continual learning. In *NeurIPS*, 2018.

Mengdi Xu, Wenhao Ding, Jiacheng Zhu, Zuxin Liu, Baiming Chen, and Ding Zhao. Task-agnostic online reinforcement learning with an infinite mixture of gaussian processes. *Advances in Neural Information Processing Systems*, 33:6429–6440, 2020.

Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019.

Wei Zhou, Yiying Li, Yongxin Yang, Huaimin Wang, and Timothy M. Hospedales. Online meta-critic learning for off-policy actor-critic methods. *ArXiv*, abs/2003.05334, 2020.