# **Policy Manifold Search**

Exploring the Manifold Hypothesis for Diversity-based Neuroevolution

Nemanja Rakicevic Imperial College London n.rakicevic@imperial.ac.uk Antoine Cully Imperial College London a.cully@imperial.ac.uk Petar Kormushev Imperial College London p.kormushev@imperial.ac.uk

## ABSTRACT

Neuroevolution is an alternative to gradient-based optimisation that has the potential to avoid local minima and allows parallelisation. The main limiting factor is that usually it does not scale well with parameter space dimensionality. Inspired by recent work examining neural network intrinsic dimension and loss landscapes, we hypothesise that there exists a low-dimensional manifold, embedded in the policy network parameter space, around which a high-density of diverse and useful policies are located. This paper proposes a novel method for diversity-based policy search via Neuroevolution, that leverages learned representations of the policy network parameters, by performing policy search in this learned representation space. Our method relies on the Quality-Diversity (QD) framework which provides a principled approach to policy search, and maintains a collection of diverse policies, used as a dataset for learning policy representations. Further, we use the Jacobian of the inversemapping function to guide the search in the representation space. This ensures that the generated samples remain in the high-density regions, after mapping back to the original space. Finally, we evaluate our contributions on four continuous-control tasks in simulated environments, and compare to diversity-based baselines.

## **CCS CONCEPTS**

• Computing methodologies  $\rightarrow$  Dimensionality reduction and manifold learning; Evolutionary robotics; *Neural networks*;

## **KEYWORDS**

Neuroevolution, Manifold Learning, Quality-Diversity, MAP-Elites

#### ACM Reference Format:

Nemanja Rakicevic, Antoine Cully, and Petar Kormushev. 2021. Policy Manifold Search: Exploring the Manifold Hypothesis for Diversity-based Neuroevolution. In 2021 Genetic and Evolutionary Computation Conference (GECCO '21), July 10–14, 2021, Lille, France. ACM, Lille, France, 9 pages. https://doi.org/10.1145/3449639.3459320

## **1** INTRODUCTION

In recent years, we have seen significant progress in tackling continuous control tasks by gradient-free approaches based on Evolutionary Algorithms, such as Neuroevolution [40, 43, 44]. Compared to

GECCO '21, July 10-14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery. ACM ISBN 978-1-4503-8350-9/21/07...\$15.00

https://doi.org/10.1145/3449639.3459320



Figure 1: The proposed Policy Manifold Search approach components. Grey boxes at the top are used within the *manifold search* phase which adds policies to the collection, while the *manifold learning* is shown in green and uses all the data from the policy collection ( $C_{\theta}$ ). Blue arrows show policy network parameter ( $\theta$ ) flow, grey arrows evaluation information ( $\theta^*$ , BD), green arrow is the current autoencoder parameters, and red arrow depicts decoder Jacobian scaling.

Deep Reinforcement Learning (RL) approaches relying on gradientbased policy network optimisation [41], Neuroevolution achieves similar performance while allowing for greater parallelisation. However, these performance-based methods aim to find a unique set of controller parameters, although there are cases where a diversity of solutions is needed and this becomes an issue. An example includes environments that can accommodate multiple tasks, and that different controller policies are necessary to solve and adapt to a dramatically changing environment or recover from damage [11]. Quality-Diversity (QD) framework [12, 35], has been introduced as an elegant approach to maintaining a diversity of controller policies. A popular QD approach, MAP-Elites [33], maintains a diversity of policies by having each distinct policy assigned to the policy collection based on the behaviour descriptor, i.e. low-dimensional representation of the corresponding policy's behaviour in the environment [12]. New solutions are then generated by modifying existing ones via mutation operators.

Applications of MAP-Elites are often limited to low-dimensional or open-loop controllers [33], while Deep RL and Neuroevolution have shown to be capable of learning deep neural networks for continuous control, increasing the agent's capabilities. Still, the curse of dimensionality makes evolution quite challenging, particularly in MAP-Elites, and prevents it to scale with the parameter dimensionality [9]. A popular approach to addressing this issue is

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

representation learning [3, 6], which has been studied extensively in the context of input spaces such as environment observation vectors, images, graphs etc [25, 34]. Representation learning relies on the *manifold hypothesis* which states that high-dimensional data tends to lie in the vicinity of a low dimensional manifold. However, this notion has not been investigated much in the context of neural network parameter spaces. Recent work on examining neural network parameters space properties, gives us insights on the possible existence of low-dimensional manifolds in the network parameter space [14, 16, 20, 30, 32, 37]. Moreover, using a notion of a manifold via the hypervolume of elites within MAP-Elites, showed promising results [46]. These findings suggest a natural question whether there exists a manifold embedded in the high-dimensional network parameter space, and what are the potential benefits for policy search via Neuroevolution, using MAP-Elites?

Inspired by the manifold hypothesis, we hypothesise that there exists a lower-dimensional, non-linear manifold, embedded in the high-dimensional policy network parameter space, which contains a high density of solutions for a particular task. In this paper, we propose a novel approach, Policy Manifold Search (PoMS), that learns a mapping to a manifold in the policy network parameter space, which is used for Neuroevolution-based policy search. We expand the MAP-Elites framework, by using its policy collection as a dataset for manifold learning via an Autoencoder (AE). This is an important component as the quality of the dataset dictates the quality and generalisability of the learned latent representations. Focusing the search in the manifold with a high solution density, improves the sample complexity and the diversity of discovered behaviours. When mapping back from the manifold to the original parameter space, we take into account the additional transformation incurred, due to the imperfect mapping of the AE's decoder function. This leads to a consistent parameter search regardless of the local structure of the latent representation, by accounting for the distortions induced by the decoder. Finally, we focus solely on diversity search and do not use performance as a selection criteria.

In order to evaluate our hypothesis, we conduct ablation studies of the algorithm components, as well as comparisons to the stateof-the-art MAP-Elites approaches. These experiments are executed on four different continuous control tasks. As a performance indicator, we use the *behaviour coverage* metric, typically used in QD literature [12], to quantify the diversity of policies achieved within the policy collection. The results suggest that learning a policy manifold explicitly and using it for policy search works well, but an additional regularisation of the search process using the decoder Jacobian is necessary in order to have a well-behaved policy search.

## 2 RELATED WORK

In this section, we present relevant work on policy behaviour diversity, as well as parameter representation and generation.

**Manifold Learning** aims to obtain a lower-dimensional representation, i.e. a manifold, embedded in the high-dimensional input space. It is based on the *manifold hypothesis* that assumes a high density of useful datapoints located in the vicinity of such a manifold. This notion has been thoroughly explored under *representation learning* [3], and applied in RL [1, 8, 15], with important insights

on how to exploit the structure of the manifold to improve an algorithm's robustness [38, 39]. However, this body of work has focused on manifold learning in the context of input spaces such as images, environment observations, etc, which have different structural properties compared to neural network parameter spaces. Recent work examines the concept of a manifold in the context of network parameter spaces, and show that the local optima of deep networks can be connected by continuous paths which lie in low loss value regions of the network parameter space [14, 17, 20]. This is related to the redundant nature of networks as multiple configurations could lead to the same behaviour [31]. Moreover, there are some insights on the existence of an intrinsic dimension of the network parameter space, that is dependant on the particular task's complexity and is sufficient for solving such a task [13, 19, 30, 37, 48]. Another line of work investigates the lottery ticket hypothesis, which states that deep networks contain subnetworks which can reach the same performance as the original network, when trained in isolation [16, 32]. These findings give us inspiration that such a manifold could be found and used for efficient policy search in QD.

Parameter-generating networks have been proposed as an approach to few-shot adaptation of the (controller) network parameters, conditioned on the observations or latent variables as inputs, in the context of fast/slow weights [2, 22], meta-learning [1], dynamic filter networks [47] or hypernetworks [24]. The abovementioned approaches do not explicitly learn a parameter manifold from which the controller network parameters are generated, rather a differentiable mapping from the observation space which can be updated via gradient descent. Conversely, in recent work a manifold is learned from a collection of converged neural network policies [7], or simple motion primitive based controllers [27]. They use learned representations for analysing and assessing new generated policy network parameters, rather than for policy search. To the best of our knowledge, Chang et al [7] is the only work besides ours, that learns policy network parameter representation. However, they use a set of fully converged policy networks to train the AE, which is costly to produce and could lead to a less generalisable representation.

Quality Diversity (QD) algorithms [12, 35] have been recently introduced as a framework that generates a collection of highperforming and diverse solutions. Most popular approaches are Multi-dimensional Archive of Phenotypic Elites (MAP-Elites) [11, 33] and Novelty Search with Local Competition [29], which differ in how they select and maintain a collection of controller policies based on their exhibited behaviours. In this work, we expand on MAP-Elites, due to its simplicity of implementation and proven performance in various applications [11, 23, 45]. Although MAP-Elites usually uses simple low-dimensional parameterised controllers [11], recent work proposed scaling MAP-Elites to Neuroevolution and applied it to more complex environments [9], by combining MAP-Elites with Evolutionary Strategies to effectively search the highdimensional network parameter space. Two recently introduced MAP-Elites-based approaches present ideas which are close to the notion of controller parameter manifolds. MAP-Elites via Elite Hypervolumes [46] uses a mix of isotropic and directional Gaussian operators allowing for an adaptive search that implicitly explores the hypervolume of the elites, which is similar to the learned manifold in PoMS, but in the original parameter space. Data-driven encoding

MAP-Elites [18], combines the line mutation with a reconstructioncrossover operator, based on an AE trained on the policies contained in the collection, similarly to our approach.

**Task-conditioned policies** are a widely used framework for achieving a diversity of policies in RL. These policies are conditioned on a task context identifier, sampled from a discrete [15, 25] or continuous [42] distribution. Diversity is All You Need (DIAYN) [15] focuses purely on discovering diverse skills, i.e. behaviours, and does not consider extrinsic environment rewards. While QD and task-conditioned policies both aim to promote behaviour diversity, they differ in how the diversity is maintained and in the definition of skills, i.e. behaviours. QD requires certain level of domain knowledge, still, it is able to consider a significantly larger number of diverse behaviours.

## **3 POLICY MANIFOLD SEARCH**

In this section, we describe the proposed Policy Manifold Search (PoMS) algorithm (Fig 1). PoMS is an iterative algorithm, where each iteration consists of the *parameter manifold learning* and *parameter manifold search* phases (Algorithm 1). The former consists of finding a latent representation of the parameter space and the corresponding mapping functions, while the latter performs search in the learned latent space to generate new diverse policies. Each iteration potentially adds new policies to the collection, following the MAP-Elites framework [11], which are then used to refine the latent representation in the next iteration. The goal of PoMS is to focus the search in the manifold which contains a high density of solutions, in order to improve the sample complexity and the behaviour coverage.

#### 3.1 Preliminaries

Let us consider a typical RL setting, with an environment defined as a Markov Decision Process (MDP) ( $S, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma$ ), with the state space S, action space A, and the deterministic transition function  $\mathcal{T}: \mathcal{S} \times \mathcal{A} \to \mathcal{S}$ . The reward function  $\mathcal{R}$  and discount factor and y are not used in this work. We define a deterministic policy  $\pi_{\theta}$ , parameterised as a deep neural network, that maps the current state to the action to be taken at that state  $a_t = \pi_{\theta}(s_t)$ . The policy parameters  $\theta$  are a P-dimensional set of network weights and biases,  $\theta \in \mathbb{R}^{P}$ , such that each point in the policy parameter space defines a unique policy. During an episode of length *T*, an agent interacts with the environment, using the policy  $\pi_{\theta}$ , thus generating a trajectory  $\tau = \{s_i, a_i\}_1^T$ . We want to distinguish how a certain deterministic policy  $\pi_{\theta_k}$  interacts with the environment in a quantifiable way. To this end, we use the concept of a Behaviour Descriptor (BD) from the QD literature [12] which aims to uniquely describe an episode rollout. The BD is formalized as a mapping from a state-trajectory  $\tau$  space, to a *b*-dimensional behaviour space  $\mathcal{BD} : T \to \mathcal{B}$ .

The MAP-Elites framework maintains a collection of multiple policy parameters, in a multi-dimensional cell-grid  $C_{\theta}$ , which is indexed by the behaviour index  $b \in \mathcal{B}$  obtained using the BD. Different policies can produce the same b, but a specific deterministic policy will map to a unique behaviour (surjective mapping). To resolve the surjective mapping, usually each cell of the grid is populated by a single highest-performing policy based on some performance metric. The aim of MAP-Elites is to fill all the cells of

Algorithm 1: Policy	y Manifold Search	(PoMS)	1
---------------------	-------------------	--------	---

// autoencoder parameter initialisation  $\xi \leftarrow$ Xavier-Glorot // collection initialisation  $C_{\theta} = \emptyset$  $\boldsymbol{\theta}^{INIT} \sim \mathcal{U}(-1,1)$  $\boldsymbol{\tau}^{INIT} = \text{environment}\_\text{eval}(\boldsymbol{\pi}_{\boldsymbol{\theta}^{INIT}})$  $\boldsymbol{b}^{INIT} = \mathcal{BD}(\boldsymbol{\tau}^{INIT})$  $C_{\boldsymbol{\theta}}[\boldsymbol{b}^{INIT}] \leftarrow \boldsymbol{\theta}^{INIT}$ for nloop in PoMS\_loops do // parameter manifold search phase for niter in MAP-Elites\_iterations do  $\theta^{SEL} \sim C_{\theta}$  $\theta^{MUT} = \text{region}_{\text{based}_{\text{search}}}(\theta^{SEL})$  $\tau^{MUT} = \text{environment}_{\text{eval}}(\pi_{\theta^{MUT}})$  $\boldsymbol{b}^{MUT} = \mathcal{BD}(\boldsymbol{\tau}^{MUT})$  $C_{\boldsymbol{\theta}}[\boldsymbol{b}^{MUT}] \leftarrow \boldsymbol{\theta}^{MUT}$ end // parameter manifold learning phase for  $\theta_{batch}$  in  $C_{\theta}$  do  $\hat{\boldsymbol{\theta}}_{\text{batch}} = f_D \circ f_E(\boldsymbol{\theta}_{\text{batch}}; \boldsymbol{\xi})$  $\mathcal{L}_{AE} \equiv \|\boldsymbol{\theta}_{\text{batch}} - \hat{\boldsymbol{\theta}}_{\text{batch}}\|_2^2$ argmin  $\mathcal{L}_{AE}$ ξ end end

 $C_{\theta}$  with the best possible policies through an iterative process. Each iteration consists of (i) randomly selecting a batch of individuals from the collection, (ii) applying a mutation and evaluating these modified individuals, (iii) based on the outcome of the evaluation, add the new individuals to the grid if the corresponding cell is vacant or if they outperform the currently occupying individual. In this study, we focus solely on policy behaviour diversity, so a cell occupancy conflict is resolved by a coin-flip. A typical mutation operator adds an isotropic Gaussian noise  $\mathcal{N}(0, \Sigma_{\Theta})$ , with a unit covariance matrix  $\Sigma_{\Theta} = \sigma_{\Theta} \mathbb{I}$ , where  $\sigma_{\Theta}$  is a hyperparameter.

#### 3.2 Parameter Manifold Learning Phase

The main insight of PoMS is learning a lower-dimensional manifold  $\mathbb{R}^M$ , embedded in  $\mathbb{R}^P$  where M << P, around which a high-density of interesting (i.e. non-degenerate) policies are located. This manifold can then serve as a smaller search space for a more efficient exploration. To find the manifold, we first generate a set of randomly distributed policy network parameters (uniformly sampled  $\theta_i \sim \mathcal{U}(-1, 1)$ ) and add them to the policy collection according to the MAP-Elites framework, which preserves diverse and unique behaviours. Then, we use the parameters that have been added to the policy collection to train a dimensionality reduction algorithm, like a deep AutoEncoder (AE) [26, 28]. The bottleneck layer of the AE defines the latent parameter representation space, which we refer to as the learned manifold. In this way, each point in the original parameter space  $\theta_i$  can be directly mapped into the corresponding point on the manifold  $z_i \in \mathbb{R}^M$  using the encoder  $(f_E)$ , and reconstructed back using the decoder  $(f_D)$ . The AE used is fully-connected and symmetrical, where  $f_E$  and  $f_D$  are parameterised by  $\xi = \{\xi_E, \xi_D\}$ . The AE parameters  $\xi$  are learned by minimising the reconstruction loss:  $\mathcal{L}_{AE} = \frac{1}{|C_\theta|} \sum_i^{|C_\theta|} \|\theta_i - \hat{\theta}_i\|_2^2$ , where  $\hat{\theta}_i = f_D \circ f_E(\theta_i; \xi)$ . In contrast with recent works [7], we do not apply any additional regularisation on the latent space, and train the AE in an unsupervised manner to reconstruct its inputs. As opposed to the common training strategy on static datasets, in the case of PoMS, it is not beneficial to normalise the training data before fitting the AE, as periodic additions to the collection lead to instability in the training.

Each iteration of PoMS is likely to add new policies to the collection. The *parameter manifold learning* phase then uses all the policy parameters from the collection and continues the training of the AE to refine the latent representation. Further details on the training procedure and hyperparameters, are presented in Appendix A.

#### 3.3 Parameter Manifold Search Phase

One of the strengths of MAP-Elites is that it constantly applies small mutations to the "elites", a subset of solutions contained in the collection, therefore only exploring around solutions that have shown to be either high-performing (i.e. exploiting) or different (i.e. exploring). In PoMS, we sample a random subset of solutions from the policy collection, as we focus on behaviour diversity rather than performance. Further, we also apply small mutations to the selected solutions, although these mutations are applied to the latent representations of the solutions in the latent space, as opposed to directly in parameter space. However, a small perturbation in the latent space can lead to a very large perturbation in the parameter space, once reconstructed, due to the complexity of the learned decoder. Therefore, there is a significant risk that applying mutations directly in the latent space (e.g. via Gaussian noise) and reconstructing, will lead to an uncontrolled mutation overall, similar to random search.

**Considering the decoder Jacobian.** To address this issue, we propose limiting the magnitude of the reconstructed samples, by making the latent parameter space search heteroscedastic. We impose that each mutated point in the latent space  $z' \sim N(z_k, \Sigma_Z)$ , when reconstructed, lands within an isotropic Gaussian distribution in the original parameter space, i.e.  $f_D(z') \sim N(f_D(z_k), \Sigma_{\Theta} = \sigma_{\Theta} \mathbb{I})$ . Meaning, the reconstructed samples are limited to a hyper-sphere region around the parameters of the selected policy. To achieve this, we use the Jacobian of the decoder, which gives us a linear approximation (first order Taylor expansion) of the transformation around a specific point in the latent space *z*, denoted as  $J_D(z)$ :

$$\mathbf{J}_D(z) = J_D(z)_{ij} = \frac{\partial f_D(z)_i}{\partial z_j} \tag{1}$$

Let us define the desired covariance matrix of an isotropic Gaussian in the original parameter space as  $\Sigma_{\Theta} = \sigma_{\Theta} \mathbb{I}$ , where  $\sigma_{\Theta}$  is the desired radius of the spherical Gaussian and  $\mathbb{I} \in \mathbb{R}^{P \times P}$  is the square identity matrix. The objective is to estimate the appropriate covariance matrix for the Gaussian noise  $\Sigma_Z$  applied in the latent space as a function of  $\Sigma_{\Theta}$  and the Jacobian. This is given by:

$$\Sigma_Z = \mathbf{J}_D^T \Sigma_{\Theta} \mathbf{J}_D \tag{2}$$

For the full derivation of Equation (2) using Taylor expansion, please refer to Appendix C.

Algorithm 2: Region-based search
<b>Input:</b> $\theta^{SEL}$ ; $C_{\theta}$ , $\Sigma_{\Theta}$ , $f_E$ , $f_D$
<b>Output:</b> $\theta^{MUT}$
$\epsilon_{recn} = \frac{1}{ C_{\theta} } \sum_{i}^{ C_{\theta} } \ \theta_i - \hat{\theta}_i\ _2^2$
$\boldsymbol{\theta}^{MUT} = \boldsymbol{\emptyset}$
for $\theta_i$ in $\theta^{SEL}$ do
$\hat{\theta}_i = f_D \circ f_E(\theta_i)$
<b>if</b> $\ \theta_i - \hat{\theta}_i\ _2^2 < \epsilon_{recn}$ then
$z_i = f_E(\tilde{\theta}_i)$
$J_D = \text{calculate_jacobian_matrix}(f_D; z_i)$
$\Sigma_{Z} = \mathbf{J}_{D}^{T} \Sigma_{\Theta} \mathbf{J}_{D}$
$z_i^{MUT} = z_i + \mathcal{N}(0, \Sigma_Z)$
$\theta_i^{MUT} = f_D(z_i^{MUT})$
else
$\theta_i^{MUI} = \theta_i + \mathcal{N}(0, \Sigma_{\Theta})$
end
$\mid \boldsymbol{\theta}^{MUT} \leftarrow \boldsymbol{\theta}_{i}^{MUT}$
end

Region-based mixing strategies. Performing the parameter search in the latent space has the advantage of offering a smaller search space with a high-density of different and interesting policies. However, like in most autoregressive algorithms, the AE is unable to generalise far beyond the training set data support. To overcome this problem, we employ a region-based exploration strategy. If the reconstruction error of a selected  $\theta_i$  is below a threshold value  $\epsilon_{recn}$ , we perform the mutation in the latent space as explained above. If the reconstruction error is above the threshold, the mutation is applied directly in the parameter space using  $\mathcal{N}(\theta_i, \Sigma_{\Theta})$ . The reconstruction error threshold  $\epsilon_{recn}$ , is determined heuristically based on the average reconstruction error of all the points in the collection, achieved after the manifold learning phase. The pseudocode for region-based policy search is given in Algorithm 2. This approach allows us to use mutations in the latent space in regions that are well known by the AE (i.e. low reconstruction error), and unconstrained mutations in the parameter space otherwise. Periodically adding solutions obtained via parameter space search helps reduce overfitting, thus making the AE generalise better. This can be regarded as a type of active learning based on model uncertainty [36]. In the experiment section, when using the mixing strategy, we show its average mixing ratio (proportion of the total number of samples obtained in the parameter space).

## **4 EVALUATION**

Our experimental evaluation  $^1$  aims to find the following insights: **Q1.** Are there any benefits of using the learned manifold for policy search, as opposed to the original parameter space?

Q2. What is the effect of using the Jacobian of the decoder, for

<sup>&</sup>lt;sup>1</sup>The code for the proposed approach and evaluation environments is available online: https://sites.google.com/view/policy-manifold-search

Policy Manifold Search



Figure 2: Screenshots of the continuous control environments used for experimental evaluation. The environments (a-c) are implemented in Box2D [5] and (d) in PyBullet [10].

scaling the latent space sampling covariance matrix? Q3. Is the non-linear manifold learning necessary, or does a linear projection suffice?

**Q4.** How does PoMS compare to state-of-the-art QD approaches? In order to answer Q1, we compare the performance of the proposed PoMS approach, to MAP-Elites (MAPE-Iso) which performs search in the original parameter space. Within this analysis, we also consider two state-of-the-art baselines which are based on a notion of a manifold: MAP-Elites via Elite Hypervolumes (MAPE-IsoLineDD) [46] and Data-driven encoding MAP-Elites (DDE) [18]. In this paper, we mainly focus on the parameter versus manifold exploration. While other work exists that scale MAP-Elites to Neuroevolution using more advanced search strategies in the parameter space [9], we leave for future work the investigation on how these two families of approaches can be combined.

Further, we perform an ablation study to address Q2 and Q3. The first aim is to demonstrate the importance of scaling the latent space sampling covariance matrix, using the decoder Jacobian. The alternative to this would be a latent space search, where the latent sampling covariance matrix  $\Sigma_Z$  is determined based on the current ranges of the latent representations (PoMS-no-jacobian). The second aim is to examine the importance of non-linear manifold learning, using AE, as opposed to using a linear projection of the parameter space via PCA (PoMS-PCA) to obtain the latent space.

Regarding Q4, we focus on how PoMS compares to MAPE-IsoLineDD and DDE. Additionally, we examined Diversity is All You Need (DIAYN) [15], as a task-conditioned policy approach. However, preliminary tests showed that it is unable to scale to several thousand skills, and in the time required by MAP-Elites and PoMS to evaluate million episodes, DIAYN managed to evaluate just over 2000, without observing any promising results. Therefore, we decided to exclude this algorithm from our analysis.

**Random baselines** are included as a 'sanity check' to compare the proposed algorithm with random search. We evaluate two approaches: (ps-uniform) where each policy network samples are drawn from a uniform distribution  $\mathcal{U}(-1, 1)$ , and (ps-glorot) in which each sample is initialised according to Xavier-Glorot scaled normal distribution, a popular network weight initialiser [21].

#### 4.1 Experiments

We evaluate the methods on four deterministic simulated environments (shown in Fig. 2). The agent acting in the environments is controlled by a neural network parameterised policy, whose parameters are generated by the policy search method and stored in the policy collection. The policy network takes the full observation vector as input, and outputs a desired action vector (joint torques or velocities). The architecture of the policy network is kept the same across all experiments, and is based on the Proximal Policy Optimisation [41] policy implementation for continuous control tasks. It is as a fully-connected neural network with two hidden layers of 32 neurons each, with *tanh* activation functions. The output layer activation function is linear. The input and output size vary depending on the observation and action vector sizes, which are specific to each environment. The environments used are:

Bipedal-Walker [observation 26D, action 4D, collection size 50000] is a standard OpenAI gym [4] environment (Fig. 2a). The original observation vector has 24 elements, which include the robot hull angle, horizontal, vertical and angular velocities, joints angles and angular velocities, legs-ground contact information, and 10 lidar rangefinder measurements. We also added the absolute coordinates of the robot hull, thus creating a 26D observation vector. The 4D action vector is unaltered and provides torques for each of the leg joints. At the start of the episode, the robot is placed in the middle of the terrain so it can walk either forward or backwards, in as many diverse ways as possible. The episode is limited to 500 steps. The policy collection is a 4D grid populated using the behaviour descriptor based on the agent's absolute positions and leg-ground contacts during the episode, where each dimension index is: D1: average hull y-coordinate, limited to [4.5, 6.2] range and discretised into 5 bins; D2: final hull x-coordinate, 100 bins spanning the terrain length; D3-4: proportion of time left and right legs spent in contact with ground, respectively, normalized to [0, 1] range discretised into 10 bins each.

Bipedal-Kicker [observation 30D, action 4D, collection size 10000] extends the Bipedal-Walker task by adding a ball (Fig. 2b). Therefore, the observation vector is extended with the ball x, y position and velocities, making it 30D. The action output is the same. Since the goal is to have a diversity of ball ballistic trajectories, we make the terrain flat to avoid biasing the outcomes to local valleys. In order to facilitate kicking, as the agent does not have a foot, at the start of the episode the ball is dropped from a small height so the agent can hit it. The agent is allowed to move for 100 timesteps and then stops to avoid multiple kicks, while the ball moves until it stops due to the damping effects. The behaviour descriptor, used for the 2D grid policy collection (Fig. 4), is based on the ball trajectory, as a usual way of defining a 2D ballistic trajectory: D1: final ball x-coordinate, in the right half of the terrain divided into 200 bins; D2: maximum ball y-coordinate achieved during the episode, limited to [4, 7] range and discretised into 50 bins.

**2D-Striker [observation 14D, action 3D, collection size 15300]** is a bounded air-hockey-like environment implemented in Box2D [5], with the goal of controlling the striker to hit the puck so it lands on as many diverse positions as possible (Fig. 2c). The arena is bounded by four walls to the size of 100x100 units, created to be proportional to the striker size (5x2.5 units). The puck has a radius of 2.5 units. The 14D observation vector, consists of the striker's x, y position and angle  $\phi$ , the puck's x, y position, as well as their corresponding velocities, and puck-wall distances for each of the walls. The 3D action vector controls the striker's x, y and angular



Figure 3: Behaviour coverage and mixing-ratio plots achieved by the compared methods, for four continuous control environments. The markers on the lines of approaches using a latent representation, correspond to latent representation updates.

velocities. Similarly to Bipedal-Kicker, the agent is allowed to act 100 steps before the actions are set to 0, in order to have only one interaction with the puck per episode, while the puck moves until it stops due to damping effects. The policy collection is a 3D grid, with the indices corresponding to: D1-2: final x, y position of the puck, with 30 bins per dimension; D3: index wall(s) the puck bounced off during the episode, and has 17 possible values (no wall, south, east, north, west, and 12 additional second order combinations).

**Panda-Striker [observation 24D, action 7D, collection size 15300]** has the same goal, policy collection and behaviour descriptor as 2D-Striker. However, here the striker is attached to the end-effector of a 7-DOF Franka-Panda arm implemented in PyBullet [10], which is used to hit the ball on a bounded table (Fig. 2d). The 24D observation vector contains the 7 joint positions and velocities, cartesian position and orientation of the end-effector and the ball's position and velocities. The 7D action vector contains joint velocity controls for moving the robot arm.

In contrast to Bipedal-Walker, the Bipedal-Kicker and both Striker environments contain a ball, which is an external object manipulated by the agent. This adds complexity to the task as certain elements of the observation vector can vary independently of the agent's actions. Another distinction is that Striker environments are bounded, while Bipedal ones are not. This leads to certain elements of the observation vector having different scales. Usually they can be normalised, but in unbounded environments this is not straightforward. We investigate the influence of these environment differences on the behaviour diversity via the evaluation metric.

## 4.2 Metric

We use *behaviour coverage* as a metric to quantify the behaviour diversity, in order to assess a search method's performance [12]. It quantitatively describes the number of distinct behaviours discovered by the algorithm, as a percentage of the policy collection filled. For the experiments to have statistical significance, 5 runs are executed with different parameter initialisation seeds. We show the median of the 5 seeds, 25th and 75th percentile. The x-axis of the behaviour coverage plots shows the total cumulative number of episode rollouts to have a fair comparison.

#### **5 RESULTS**

Below, we discuss the behaviour coverage achieved by the compared methods (Fig. 3), in the context of questions from Sec. 4. Implementation details of these methods are given in Appendix B.

A1. Learned latent vs original parameter space search. The first conclusion we note from the experiments, is that the standard MAP-Elites algorithm (MAPE-Iso) is competitive in highdimensional parameter space problems, which has not been sufficiently investigated in previous work. Comparing the proposed PoMS approach with MAPE-Iso, PoMS systematically achieves higher behaviour coverage across the tasks. This is especially evident with Bipedal-Kicker, where the relationship between the observation vector and the action outputs, which depend on the policy parameterisation, is highly complex. This could be attributed to the unbounded nature of the environment and the presence of an additional object. The 2D-Striker case is an exception as all methods converge to the same performance (Fig. 2c). This result can be explained by the simplicity of the task, as the policy outputs can directly influence the planar movement of the striker, and by extension the puck, while such connection is more complex within Bipedal environments and to some extent with the Panda arm. Even though MAPE-IsoLineDD operates in the parameter space, it uses the hypervolume of elites for search, which can be related to the notion of a manifold, thus improving its efficiency over MAPE-Iso. However, besides the Striker tasks where it reaches equal asymptotic performance, it converges to a lower behaviour coverage compared to PoMS. By definition of the line mutation, the MAPE-IsoLineDD usually performs well when the hypervolume of elites is convex. When this assumption does not hold, a more involved transformation is needed which PoMS realises via manifold learning. DDE uses the learned latent representation indirectly by exploiting the reconstruction inaccuracies through its mutation operator, which on average achieves higher behaviour coverage than direct policy space search, such as MAPE-Iso. Still, accounting for the decoder transformation and having a more structured latent sampling as in PoMS enables better performance. These findings validate the benefits of using a learned latent representation of the policy parameters for policy search.

Policy Manifold Search



Figure 4: Policy collections showing the behaviour coverage achieved by the compared methods for the Bipedal-Kicker.

**A2.** Contribution of decoder Jacobian scaling. By comparing PoMS and PoMS-no-jacobian, we can see that accounting for the decoder transformation of the latent space samples, via Jacobian scaling of their covariance matrix, is crucial for the algorithm's performance. As hypothesised, simple latent perturbation and reconstruction done in PoMS-no-jacobian, leads to a search that performs similarly to random search (Fig. 3b, 3c).

**A3. Linear vs non-linear representations.** The difference in performance between PoMS and PoMS-PCA, speaks mostly about the intrinsic complexity of the given task control problem. Environments in which locomotion is involved have an intrinsic non-linearity in the mapping of the policy outputs and actual motions contained in the observation vector, which is used to determine behaviour descriptors. This can explain the performance gap between PoMS and PoMS-PCA in both Bipedal environments (Fig. 3a, 3b).

A4. State-of-the-art performance comparison. The methods with consistently highest behaviour coverage are PoMS, PoMS-PCA, DDE and MAPE-IsoLineDD, where the common characteristic is using a notion of a manifold, by focusing the search in the hypervolume of elites or the learned manifold. This further solidifies claims from A1. Even though DDE and MAPE-IsoLineDD were not originally used with neural network parameterised controllers, they perform well in high-dimensional parameter spaces. As we increase the environment complexity, by adding a Panda arm to Striker, and considering the Bipedal agent, we can observe that the proposed PoMS, maintains a high behaviour coverage, while other state-of-the-art methods are not as robust. This drop in performance is higher in purely parameter space search methods such as MAPE-Iso, ps-uniform and ps-normal. PoMS outperforms the next best state-of-the-art approach by a relative increase of the median coverage of 15% and 9% respectively for Bipedal-Walker and Bipedal-Kicker environments, while it is on-par in both Striker environments at convergence. A qualitative comparison of the achieved behaviour coverage at convergence, for each of the compared methods is shown in Fig. 4.

**Mixing ratio** Below the corresponding behaviour coverage graphs in Fig. 3, we show the mixing ratios for PoMS and PoMS-PCA. The mixing ratio represents the averaged ratio of samples generated in the latent, versus the parameter space, during one search iteration. Mixing ratio of 1 means that all the policy parameters are sampled directly in the parameter space, while 0 means that all are sampled in the latent space (using the scaled covariance

matrix) and then reconstructed. The first loop of the algorithm starts with a mixing ratio of 0.5 and subsequently the ratio changes based on the mean reconstruction error, as explained in Section 3. From the mixing ratios we can see that in the beginning, there is usually a spike towards parameter space sampling. This is due to the fact that initially there are many points with a high reconstruction error, because the AE is trained on a small amount of data, and needs more diversity - thus it 'explores'. The mixing ratio slowly decreases in favor of the latent space samples, with several salient 'dips' which loosely correlate to jumps in behaviour discovery. This can be interpreted as the algorithm 'exploiting' the learned latent representation. This is not evident with PCA as its representations tend to be more rigid and do not change often with new data.

#### 6 CONCLUSION

In this paper, we proposed the Policy Manifold Search algorithm, inspired by the manifold hypothesis, that aims to discover a collection of policies with diverse behaviours by performing search in the learned manifold embedded in the policy parameter space. We learn this manifold from a collection of policy parameters maintained according to the MAP-Elites framework, and use the manifold to generate novel solutions, thus augmenting the collection. Experimental evaluations of PoMS validate the benefits of using a learned manifold, coupled with the Jacobian of the decoder for guiding the search, to discover larger collections of diverse policies compared to the baselines. The benefits of considering the manifold hypothesis were further reinforced by examining two state-of-the-art methods which rely on this notion as well, and achieve good performance. These findings suggest that the policy parameter space is redundant, and learning its latent representation is useful as it contains many diverse solutions close to each other, which helps policy search.

In future work, we plan to further investigate the structure of the learned manifold and how it can be formed without using domain knowledge through behaviour descriptors. Another interesting line of work we plan to pursue is extending this approach to full network graph encoding, rather than just weight parameters.

## ACKNOWLEDGMENTS

This work was supported by the Imperial College London President's PhD Scholarship, and EPSRC (Engineering and Physical Sciences Research Council) grant EP/V006673/1 project REcoVER.

## A AUTOENCODER TRAINING DETAILS

During the *parameter manifold learning* phase of PoMS, the AE is trained by minimising the reconstruction loss  $\mathcal{L}_{AE}$  in order to find the optimal AE parameters ( $\xi$ ). We use the Adam optimiser with:  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\varepsilon = 10^{-8}$  and learning rate of  $10^{-5}$ , for  $2 \cdot 10^4$  epochs, with batch size of 64. The training hyperparameters are the same in all experiments. To improve the robustness of the optimiser, we reset the momentum variables at every loop. Moreover, 30% of each batch is used as a test set for early stopping of the training. If the slope of the line fitted to the last 100 test set values is larger than  $10^{-5}$ , the training is stopped. We found that this improves the generalisation of the AE and reduces training time.

#### **B** IMPLEMENTATION DETAILS

All algorithms based on MAP-Elites (PoMS versions, MAPE-Iso, MAPE-IsoLineDD and DDE) run 100 iterations of MAP-Elites with a budget of 200 samples. The policy collection is initialised by drawing 2000 policy samples from a uniform distribution, as usually done in the MAP-Elites literature [12]. The ps-uniform and ps-glorot methods run for the same total amount of samples as other algorithms, while the progress is displayed every 2000 samples.

**PoMS** has three main hyperparameters that need to be tuned: AE architecture, latent space dimension (LD) and  $\Sigma_{\Theta}$ . The AE is symmetric, i.e. both encoder and decoder have one hidden layer with ELU activations, and we vary the number of hidden nodes (HD). The activation function of the bottleneck layer forming the latent space is linear, same as for the output layer of the decoder. The hyperparameter values used in the experiments are:

Environment	HD	LD	$\Sigma_{\Theta}$
Bipedal-Walker	100	50	0.1
Bipedal-Kicker	100	100	0.01
2D-Striker	100	50	0.1
Panda-Striker	100	50	0.5

**PoMS-no-jacobian** keeps the same AE architectures as PoMS, in order to have an appropriate ablation study. The  $\Sigma_Z$  is not fixed, rather, dynamically updated based on the current latent space parameter ranges, per latent dimension  $r_Z$ , in order to scale a unit covariance matrix  $\Sigma_Z = r_Z^T \mathbb{I}$ . In this way, we give importance of each of the latent dimensions based on their spread. As we can see from the results, two issues arise with this approach: (i) range does not equal importance (solution density), (ii) applying the inverse transformation applies an additional distortion which can lead to undesirable values, because  $\theta' \neq f_D(z')$ , where  $z' \sim \mathcal{N}(\mu_Z, \Sigma_Z)$  and  $\theta' \sim \mathcal{N}(f_D(\mu_Z), \Sigma_{\Theta})$ , even if  $\Sigma_{\Theta} = \Sigma_Z = \mathbb{I}$ .

**PoMS-PCA** requires LD and  $\Sigma_{\Theta}$  as hyperparameters. These values are kept the same as with PoMS, for a proper ablation study.

**DDE** hyperparameters consist of the AE architecture and mutation operator specific hyperparameters. The former is kept the same as with PoMS, for each of the experiments, while the latter are as in the original paper [18]. Instead of running a fixed window for the multi-armed bandit upper confidence bound operator selector, we maintain a moving average.

**MAPE-IsoLineDD** has two hyperparameters related to the weighing of the isometric and directional components of the mutation operator, and they are kept the same as in the original paper [46].

**MAPE-Iso** needs only  $\Sigma_{\Theta}$ , which we set to  $\Sigma_{\Theta}$ =0.1 as this achieved the best performance for MAPE-Iso across the experiments.

## C DERIVATION OF JACOBIAN SCALING

Let us assume  $\theta \in \Theta \subset \mathbb{R}^P$  and  $z \in Z \subset \mathbb{R}^M$  to be Gaussians in the parameter and latent space respectively, such that  $\theta \sim \mathcal{N}(\mu_{\Theta}, \Sigma_{\Theta})$  and  $z \sim \mathcal{N}(\mu_Z, \Sigma_Z)$ . We further assume a non-linear, vector-valued function  $f_D : \mathbb{R}^M \to \mathbb{R}^P$  which maps the latent space to the original parameter space. Then, we can get a linear approximation  $\hat{f}_D$  in a point  $\mu_Z$  by performing a first-order Taylor expansion:

$$\begin{split} \hat{f}_{D}(z) &\approx f_{D}(\mu_{z}) + \sum_{i}^{M} \left. \frac{\partial f_{D}(z)}{\partial z_{i}} \right|_{z=\mu_{Z}} (z_{i} - \mu_{Zi}) \\ &= f_{D}(\mu_{Z}) + \begin{bmatrix} \nabla f_{D}(z)_{1} \\ \nabla f_{D}(z)_{2} \\ \dots \\ \nabla f_{D}(z)_{P} \end{bmatrix} (z - \mu_{Z}) \end{split}$$
(3)
$$\\ &= \mu_{\Theta} + \mathbf{J}_{D}(z - \mu_{Z}) \end{split}$$

where **J**<sub>D</sub> is the Jacobian matrix of  $f_D$  at  $\mu_Z$ :

$$\mathbf{J}_D(\mu_Z) = J_D(\mu_Z)_{ij} = \frac{\partial f_D(\mu_Z)_i}{\partial \mu_{Z_j}} \tag{4}$$

where indices *i* and *j* refer to the corresponding elements of the reconstructed or latent parameter vector, respectively. Further, if  $\hat{\theta} = \hat{f}_D(z)$  its expected value can be obtained as:

$$\mathbb{E}[\hat{\theta}] = \mathbb{E} \left[ \mu_{\Theta} + \mathbf{J}_D(z - \mu_Z) \right]$$
  
=  $\mathbb{E}[\mu_{\Theta}] + \mathbb{E} \left[ \mathbf{J}_D(z - \mu_Z) \right]$   
(expected values of a sum is the sum of expected values)  
=  $\mu_{\Theta} + \mathbb{E} \left[ \mathbf{J}_D z \right] - \mathbb{E} \left[ \mathbf{J}_D \mu_Z \right]$   
(expectation of a constant is a constant)  
=  $\mu_{\Theta} + \mathbf{J}_D \mathbb{E}[z] - \mathbf{J}_D \mathbb{E}[\mu_Z]$   
=  $\mu_{\Theta} + \mathbf{J}_D \mu_Z - \mathbf{J}_D \mu_Z$   
=  $\mu_{\Theta}$ 

We can obtain  $\Sigma_{\Theta}$  based on  $\Sigma_Z$ . We start with the standard equation for covariance:

$$\Sigma_{\Theta} = \mathbb{E}\left[ (\hat{\theta} - \mathbb{E}[\hat{\theta}])(\hat{\theta} - \mathbb{E}[\hat{\theta}])^{T} \right]$$
  

$$= \mathbb{E}\left[ (\mu_{\Theta} + \mathbf{J}_{D}(z - \mu_{Z}) - \mu_{\Theta})(\mu_{\Theta} + \mathbf{J}_{D}(z - \mu_{Z}) - \mu_{\Theta})^{T} \right]$$
  
(using Equations 3 and 4)  

$$= \mathbb{E}\left[ (\mathbf{J}_{D}(z - \mu_{Z}))(\mathbf{J}_{D}(z - \mu_{Z}))^{T} \right]$$
  

$$= \mathbb{E}\left[ \mathbf{J}_{D}(z - \mu_{Z})(z - \mu_{Z})^{T} \mathbf{J}_{D}^{T} \right]$$
  

$$= \mathbf{J}_{D}\mathbb{E}\left[ (z - \mu_{Z})(z - \mu_{Z})^{T} \right] \mathbf{J}_{D}^{T}$$
(6)

(covariance definition for z)

$$= \mathbf{J}_D \Sigma_Z \mathbf{J}_D^T$$

By rearranging the previous equation we finally get:

$$\Sigma_Z = \mathbf{J}_D^T \Sigma_\Theta \mathbf{J}_D \tag{7}$$

(5)

#### REFERENCES

- Jakub Sygnowski Oriol Vinyals Razvan Pascanu Simon Osindero Andrei A. Rusu, Dushyant Rao and Raia Hadsell. 2019. Meta-Learning with Latent Embedding Optimization. In International Conference on Learning Representations.
- [2] Jimmy Ba, Geoffrey E Hinton, Volodymyr Mnih, Joel Z Leibo, and Catalin Ionescu. 2016. Using fast weights to attend to the recent past. In Advances in Neural Information Processing Systems.
- [3] Yoshua Bengio, Aaron Courville, and Pascal Vincent. 2013. Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis* and Machine Intelligence.
- [4] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. Openai gym. arXiv preprint arXiv:1606.01540 (2016).
- [5] Erin Catto. 2011. Box2d: A 2d physics engine for games. URL: http://www.box2d. org (2011).
- [6] Lawrence Cayton. 2005. Algorithms for manifold learning. Univ. of California at San Diego Tech. Rep (2005).
- [7] Oscar Chang, Robert Kwiatkowski, Siyuan Chen, and Hod Lipson. 2019. Agent Embeddings: A Latent Representation for Pole-Balancing Networks. In International Conference on Autonomous Agents and Multiagent Systems.
- [8] Nutan Chen, Alexej Klushyn, Alexandros Paraschos, Djalel Benbouzid, and Patrick Van der Smagt. 2018. Active learning based on data uncertainty and model sensitivity. In International Conference on Intelligent Robots and Systems.
- [9] Cédric Colas, Vashisht Madhavan, Joost Huizinga, and Jeff Clune. 2020. Scaling MAP-Elites to Deep Neuroevolution. In Genetic and Evolutionary Computation Conference.
- [10] Erwin Coumans and Yunfei Bai. 2016. Pybullet, a python module for physics simulation for games, robotics and machine learning. URL: http://pybullet.org (2016).
- [11] Antoine Cully, Jeff Clune, Danesh Tarapore, and Jean-Baptiste Mouret. 2015. Robots that can adapt like animals. *Nature*.
- [12] Antoine Cully and Yiannis Demiris. 2017. Quality and diversity optimization: A unifying modular framework. IEEE Transactions on Evolutionary Computation.
- [13] Misha Denil, Babak Shakibi, Laurent Dinh, Marc'Aurelio Ranzato, and Nando de Freitas. 2013. Predicting Parameters in Deep Learning. In Advances in Neural Information Processing Systems.
- [14] Felix Draxler, Kambis Veschgini, Manfred Salmhofer, and Fred Hamprecht. 2018. Essentially no barriers in neural network energy landscape. In *International conference on machine learning*.
- [15] Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. 2019. Diversity is All You Need: Learning Skills without a Reward Function. In International Conference on Learning Representations.
- [16] Jonathan Frankle and Michael Carbin. 2019. The Lottery Ticket Hypothesis: Finding Sparse, Trainable Neural Networks. In International Conference on Learning Representations.
- [17] C. Daniel Freeman and Joan Bruna. 2017. Topology and Geometry of Half-Rectified Network Optimization. In International Conference on Learning Representations.
- [18] Adam Gaier, Alexander Asteroth, and Jean-Baptiste Mouret. 2020. Discovering Representations for Black-box Optimization. In *Genetic and Evolutionary Computation Conference*.
- [19] Adam Gaier and David Ha. 2019. Weight agnostic neural networks. In Advances in Neural Information Processing Systems.
- [20] Timur Garipov, Pavel Izmailov, Dmitrii Podoprikhin, Dmitry P. Vetrov, and Andrew Gordon Wilson. 2018. Loss Surfaces, Mode Connectivity, and Fast Ensembling of DNNs. In Advances in Neural Information Processing Systems.
- [21] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In International Conference on Artificial Intelligence and Statistics.
- [22] Faustino Gomez and Jürgen Schmidhuber. 2005. Evolving modular fast-weight networks for control. In International Conference on Artificial Neural Networks.
- [23] Daniele Gravina, Ahmed Khalifa, Antonios Liapis, Julian Togelius, and Georgios N. Yannakakis. 2019. Procedural Content Generation through Quality Diversity. In Conference on Games.
- [24] David Ha, Andrew Dai, and Quoc V Le. 2017. HyperNetworks. In International Conference on Learning Representations.

- [25] Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin A. Riedmiller. 2018. Learning an Embedding Space for Transferable Robot Skills. In International Conference on Learning Representations.
- [26] Geoffrey E Hinton and Ruslan R Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science*.
- [27] Marija Jegorova, Stéphane Doncieux, and Timothy Hospedales. 2018. Generative Adversarial Policy Networks for Behavioural Repertoire. arXiv preprint arXiv:1811.02945 (2018).
- [28] Mark A Kramer. 1991. Nonlinear principal component analysis using autoassociative neural networks. AIChE journal.
- [29] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In *Genetic and Evolutionary Computation Conference*.
- [30] Chunyuan Li, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski. 2018. Measuring the Intrinsic Dimension of Objective Landscapes. In International Conference on Learning Representations.
- [31] Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John E. Hopcroft. 2016. Convergent Learning: Do different neural networks learn the same representations?. In International Conference on Learning Representations.
- [32] Ari S. Morcos, Haonan Yu, Michela Paganini, and Yuandong Tian. 2019. One ticket to win them all: generalizing lottery ticket initializations across datasets and optimizers. In Advances in Neural Information Processing Systems.
- [33] Jean-Baptiste Mouret and Jeff Clune. 2015. Illuminating search spaces by mapping elites. arXiv preprint arXiv:1504.04909 (2015).
- [34] Alexandre Péré, Sébastien Forestier, Olivier Sigaud, and Pierre-Yves Oudeyer. 2018. Unsupervised Learning of Goal Spaces for Intrinsically Motivated Goal Exploration. In International Conference on Learning Representations.
- [35] Justin K Pugh, Lisa B Soros, and Kenneth O Stanley. 2016. Quality diversity: A new frontier for evolutionary computation. Frontiers in Robotics and AI.
- [36] Nemanja Rakicevic and Petar Kormushev. 2019. Active learning via informed search in movement parameter space for efficient robot task learning and transfer. Autonomous Robots.
- [37] Pierre H Richemond, Arinbjörn Kolbeinsson, and Yike Guo. 2019. Biologically inspired architectures for sample-efficient deep reinforcement learning. *Deep Reinforcement Learning Workshop, NeurIPS*.
- [38] Salah Rifai, Yoshua Bengio, Yann N Dauphin, and Pascal Vincent. 2012. A generative process for sampling contractive auto-encoders. In International Conference on Machine Learning.
- [39] Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. 2011. Contractive Auto-Encoders: Explicit Invariance During Feature Extraction. In International Conference on Machine Learning.
- [40] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution strategies as a scalable alternative to reinforcement learning. arXiv preprint arXiv:1703.03864 (2017).
- [41] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347 (2017).
- [42] Archit Sharma, Shixiang Gu, Sergey Levine, Vikash Kumar, and Karol Hausman. 2020. Dynamics-Aware Unsupervised Discovery of Skills. In International Conference on Learning Representations.
- [43] Kenneth O Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. Nature Machine Intelligence.
- [44] Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O Stanley, and Jeff Clune. 2017. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. arXiv preprint arXiv:1712.06567 (2017).
- [45] Neil Urquhart, Silke Höhl, and Emma Hart. 2019. An illumination algorithm approach to solving the micro-depot routing problem. In *Genetic and Evolutionary Computation Conference*.
- [46] Vassiiis Vassiliades and Jean-Baptiste Mouret. 2018. Discovering the elite hypervolume by leveraging interspecies correlation. In Proceedings of the Genetic and Evolutionary Computation Conference.
- [47] Tinne Tuytelaars Xu Jia, Bert De Brabandere and Luc Van Gool. 2016. Dynamic Filter Networks. In Advances in Neural Information Processing Systems.
- [48] Zichao Yang, Marcin Moczulski, Misha Denil, Nando De Freitas, Alex Smola, Le Song, and Ziyu Wang. 2015. Deep fried convnets. In International Conference on Computer Vision.