
General Characterization of Agents by States they Visit

Anssi Kanervisto

University of Eastern Finland
anssk@cs.uef.fi

Tomi Kinnunen

University of Eastern Finland

Ville Hautamäki

University of Eastern Finland
National University of Singapore

Abstract

Behavioural characterizations (BCs) of decision-making agents, or their policies, are used to study outcomes of training algorithms and as part of the algorithms themselves to encourage unique policies, match expert policy or restrict changes to policy per update. However, previously presented solutions are not applicable in general, either due to lack of expressive power, computational constraint or constraints on the policy or environment. Furthermore, many BCs rely on the actions of policies. We discuss and demonstrate how these BCs can be misleading, especially in stochastic environments, and propose a novel solution based on what states policies visit. We run experiments to evaluate the quality of the proposed BC against baselines and evaluate their use in studying training algorithms, novelty search and trust-region policy optimization. The code is available at <https://github.com/miffyli/policy-supervectors>.

1 Introduction

While creating or training autonomous agents, whether it is via manual coding, *reinforcement learning* (RL) [49] or *evolution strategy* (ES) [18] algorithms, one wishes to compare the solutions to find out which work and which do not. A common approach is comparing the performance of agents' policies in the given task [20] or by studying how policy behaves [39]. In addition to these explicit comparisons, RL and ES algorithms include implicit comparisons to encourage finding novel solutions [29, 8, 12, 37], to match the behaviour of an expert policy [24, 35] or to limit changes on policy's behaviour to avoid catastrophic failure [45, 46].

To facilitate such comparisons, *behavioural characterizations* (BCs) [29] aim to capture policy's behaviour in a fixed representation accompanied by a distance metric. A general BC could be used to generalize previously proposed methods [8, 36]. However, BCs in the previous work are often domain-specific or depend on policy's actions which are not descriptive of behaviour in general, as we will discuss in Section 3.

Recent work attempted to generalize the notion of a *behavioural embedding* [36], considering policies as distributions over trajectories, and providing a principled mechanism to compare policies in the behavioural space. However, this approach relies on "behavioural embedding mappings", which map trajectories to structures that are believed to describe the trajectory (e.g., reward-to-go, terminal state). It remains unclear how the choice of this mapping affects the results.

In this work, we propose and evaluate a more general BC that can be applied to different domains. Our work provides the following contributions: (1) We summarize and compare BCs used in the previous work. (2) We discuss and demonstrate the shortcomings of action-based BCs which have been actively used in the previous work. (3) Propose a novel BC method based on which states the agent visits, borrowing techniques from the field of speaker recognition [28]. (4) We use the proposed

Table 1: Comparison of different BCs. Required compute describes CPU and memory requirements for comparing a large set of policies, excluding the sampling of the environment. Environment agnostic methods do not pose requirements for the environment, and action agnostic methods can compare policies across action spaces.

Description and reference	Required compute	Describes	Expressiveness	Agnostic	
				Env.	Action
Returns [20]	Low	Policy	Low	✓	✓
Policy parameters [13]	Low	Policy	Low	✓	✗
Comparing actions [36, 19]	Low	Policy	Low	✓	✗
Termination state [8]	Low	Policy	Low	✗	✓
Transition matrix [32]	V.High	Policy	High	✗	✗
Trajectories [8]	High	Trajectory	High	✗	✓
States + Discriminator [12]	N/A	Skill [12]	High	✓	✓
States + Gaussian [3]	Low	Policy	Low	✓	✓
Trajectory encoder [52, 43]	High	Trajectory	High	✓	✗
Aggregate state-actions [17, 50]	High	Policy	High	✓	✗
States + Discriminator (Adapted)	Low	Policy	High	✓	✓
Trajectory encoder (Adapted)	High	Policy	High	✓	✓
States + GMM (Proposed)	Low	Policy	High	✓	✓

BC to study and visualize different training algorithms and also explore their use in novelty search [29] and trust-region policy optimization [45].

2 Preliminaries and definitions

We model environments as stochastic *Markov decision processes* (MDPs) [49], where agent acts on states $s \in \mathcal{S}$. The environment begins in an initial state $s_0 \sim p(s_0)$. The policy $\pi \in \Pi$ provides an action $a \in \mathcal{A}$ as a stochastic function $a \sim \pi(s)$. After executing an action, the environment evolves to the next state according to a stochastic function $s' \sim p(s'|s, a)$ (transition dynamics), and agent is provided with a reward $r \sim p(r|s, a, s')$. This process repeats until environment lands into a terminal state $s_T \in \mathcal{S}$. A single episode begins from an initial states and ends to a terminal state. A trajectory is a tuple of states and actions experienced during one episode. An environment is defined by tuple $(\mathcal{S}, \mathcal{A}, p(s_0), p(s'|s, a), p(r|s, a, s'), T)$.

Behaviour characterization is a function $b : \Pi \rightarrow \mathcal{B}$, where \mathcal{B} is space of all behaviours, accompanied with a *pseudo-distance* function $d : \mathcal{B} \times \mathcal{B} \rightarrow \mathbb{R}^+$. This need not be a proper distance metric, as long it is non-negative and symmetric. Ideally, the more different behaviour of two policies is, the larger this distance should be. Note that b may be a stochastic function if, for example, it relies on randomly sampled data to map policy to a behaviour.

3 Choosing behavioural characterization

Table 1 compares various BCs used in the previous research, which have been either used explicitly for visualization (returns, terminal state) or implicitly in the training algorithm (e.g., novelty search [47], diversity [12]). Table 1 also includes methods which model trajectories or skills (e.g., fixed policy conditioned on a latent vector), rather than policies.

While we could use policy’s parameters to compare agents, this does not generalize over different types of algorithms nor can we tell if a change in parameter vector is truly meaningful to the agent’s behaviour [8, 19]. Episodic returns are the *de facto* approach for measuring quality of policies, but policies with distinct behaviour can achieve same returns (see Figure 1) [1]. Using a termination state, such as the final coordinates of the robot, has been successful in novelty search but is limited to environments where such heuristics can be used. Element-wise comparisons of trajectories require either fixed-length episodes or heuristics to combine varying length episodes. Finally, one can estimate transition dynamics by sampling the environments [32], but this has a high computational

cost that limits the number of policies compared. With these options discarded, we are left with BCs which focus on actions policy takes or states they visit.

3.1 Actions and stochastic environments

Comparing policies by what actions they chose is intuitively a sound solution, as one may argue the policy’s actions are what define its behaviour. Action-based BCs are used in gradient-based RL algorithms to restrict changes to policy’s behaviour [45, 46] or to encourage diversity [37]. This is done by comparing actions of policies they would take in a fixed set of states (an “off-policy embedding” [36]).

However, using actions alone for BC would ignore the transition dynamics of the environment. As discussed by Pacchiano et al. [36], a small change to actions could lead the agent to wildly different states or a large change in actions could have no effect.¹ To illustrate this, Figure 1 shows three scenarios where the behaviour of two policies are represented with the average episodic return, distribution of the taken actions and distribution of the visited states.

Action-based BC is agnostic to environment stochasticity. As stochasticity of the environment increases, the policies tend towards random behaviour, and action-based BC describes them as two very different policies while they both behave the same (random agents). Another weakness is the insensitivity to “doorway” scenarios, where a single action can lead to different states. Action-based BC shows the two policies are very similar, but the green policy never reaches the other side of the world until we increase stochasticity. In real environments, doorway scenarios may manifest as literal doorways, critical points in grasping an object or pressing a button in a video game. Finally, even if some of the states used to compare action-differences were unreachable, they affect the results. In the right-most figure of Figure 1, both policies behave exactly the same in the area where they can traverse, yet the action-based distance is high.

In summary, while action-based BCs are useful for RL training, they can not be relied upon as a general BC: without accounting for the environment dynamics, one can not say if the difference in actions is meaningful.

3.2 Describing policies by states they visit

A common alternative [12, 3] to action-based BCs is comparing *what states policies visit*. This requires sampling states for each policy, but in turn, it captures the environment dynamics. Referring to Figure 1, where action-based BC provided misleading descriptions, state-based BC captures correctly the increasing stochasticity and the effect of the doorway and unreachable states. Compared to returns, it also captures the difference between two optimal but distinct policies.

A simple state-based BC is to fit a multivariate Gaussian on the states visited by the policy [3]. Both fitting and comparing policies is fast (e.g., KL-divergence, which can be evaluated in a closed-form), but limits the description to a unimodal distribution. The true distribution of states can be complex even for a simple environment (see scatter plots in Figure 2). Instead of a single Gaussian, a neural network discriminator can be trained to measure the probability of a state coming from a given policy, as done by Eysenbach et al. [12] and Ni et al. [35]. Alternatively, one can train an auto-encoder to encode varying-length trajectories into fixed length [52]. Alas, the BCs of these works are tightly integrated with the training loop, and can not be used with policies trained by other algorithm. Instead,

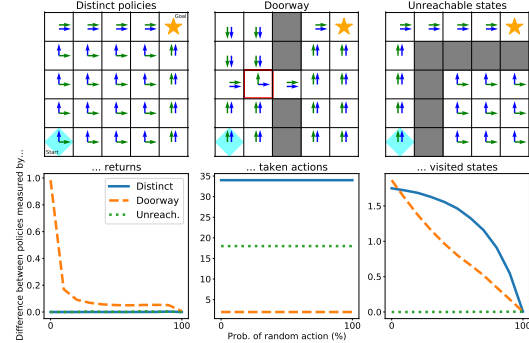


Figure 1: Difference between the two policies, blue and green, as measured by different BCs. State BC is obtained with ten thousand trajectories per policy, and distance is measured as the sum of the absolute errors over states. Action BC difference is the sum of the absolute difference between action distributions. Details are available in Appendix A.

¹If an agent takes an action and it has no effect on the environment, did the agent take an action from an outside observer’s point of view?

we propose using a mixture of Gaussian models to model complex distributions of states, which is detailed in the next section.

4 Policy supervectors

We consider a setup with N policies π_i , $i \in \{1 \dots N\}$ which we want to compare. For each policy, we play M episodes and store all encountered states s to a per-policy dataset $B_i = (s_{i,1}, s_{i,2}, \dots)$. Note that the stored states may differ from states the policy acts upon; stored states can be a separate piece of information believed to describe the behaviour, such as x-y coordinates of a maze robot [39]. In the case of high-dimensional environments, one can train an encoder to turn images into more compact latent codes [3] or to preprocess them into low-dimensional representations [9].

Gaussian mixture models (GMMs) [33] can be used to model multi-modal data, and with enough components they can model almost any continuous distribution to an arbitrary accuracy [4]. However, measuring KL-divergence requires expensive sampling. While approximations exist, they come with different assumptions and drawbacks [22].

Instead, we draw inspiration from the field of speaker verification where GMMs have been used to model speaker characteristics [28]. Instead of training one GMM per speaker, Reynolds et al. [44] presented *universal background models* (GMM-UBMs, later UBM) to model a distribution of speaker-independent features by fitting a large GMM on a large pool of speaker data. Speaker-specific model can then be adapted with *Maximum a Posteriori* (MAP) adaptation of GMM components [15] using individual speaker's data. The acquired parameters can then be concatenated into a long vector, called a *supervector* [7], describing the speaker with a fixed-length vector.

We apply this method to policies instead of speakers. We first fit a GMM on state datasets from all policies and then MAP-adapt a supervector for each policy, which we call a *policy supervector*. Under the Bayesian view, the UBM represents an informative prior while the adapted GMM is a point estimate of the posterior, where we use the same data as part of the prior and adapting the posterior. We treat UBM as a common feature extractor, and by including data from all policies we ensure the prior covers all states visited by all of the policies. This process is illustrated in Figure 2.

Formally, policies' datasets are pooled together into an UBM-training set $B_{\text{ubm}} = (B_1, \dots, B_N)$. This set is then used to fit an UBM with K components and parameters $\mu_{\text{ubm}} \in \mathbb{R}^{K \times d}$, $\Sigma_{\text{ubm}} \in \mathbb{R}^{K \times d \times d}$ and $w_{\text{ubm}} \in \mathbb{R}^K$, $w_k \geq 0$, $\sum_k w_k = 1$ with EM-algorithm until convergence. We can then perform MAP-adaptation on a per-policy dataset B to obtain adapted mean $\hat{\mu}_k$ of k^{th} component with [44]

$$p(k|s_t) = \frac{w_k \mathcal{N}(\mu_k, \Sigma_k)}{\sum_{l=1}^K w_l \mathcal{N}(\mu_l, \Sigma_l)} \quad (1)$$

$$n_k = \sum_{t=1} p(k|s_t) \quad (2)$$

$$E_k(B) = \frac{\sum_{t=1} p(k|s_t) s_t}{n_k} \quad (3)$$

$$\alpha_k = \frac{n_k}{n_k + r} \quad (4)$$

$$\hat{\mu}_k = \alpha_k E_k(B) + (1 - \alpha_k) \mu_k, \quad (5)$$

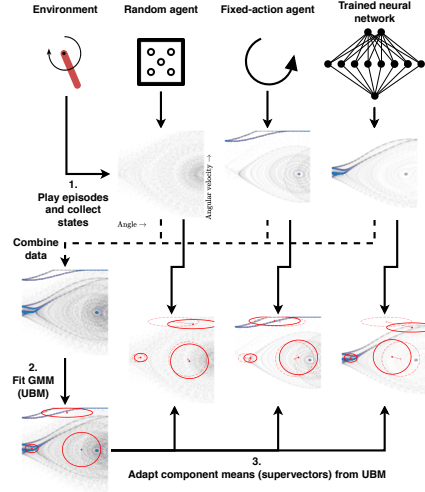


Figure 2: Computing policy supervectors for three policies in the Pendulum environment, using real data and policies. Each faint dot represents a single state of the Pendulum environment, and red circles represent Gaussian components of the GMMs.

where r , known as a *relevance factor*, is a control parameter that impacts how much new data can affect the adapted mean.

To measure the distance between two policy supervectors, Campbell et al. [7] show that an upper-bound for KL-divergence of means of two adapted GMMs μ_i and μ_j is

$$d_{\text{KL}}(\mu_i, \mu_j) \leq \frac{1}{2} \sum_{k=1}^K w_{\text{ubm}}^k (\mu_i^k - \mu_j^k) \Sigma_{\text{ubm}}^{-1} (\mu_i^k - \mu_j^k). \quad (6)$$

Unlike KL-divergence, this upper-bound is symmetric [7]. In this work, we will only adapt the means of the UBM to allow the use of the above metric, but this method can be extended by adapting covariances and weights (see [44] equations (11)-(13)). We use this upper bound as a distance metric for policies.

Limitations. Policy supervectors require each policy to sample data from the environment which can be expensive in slow environments. In highly stochastic environments we need more samples to accurately describe policies, which will make the training phase of these methods more demanding. Policy supervectors also require the environment to be the same for all policies, which can be difficult for real-world robotics, for example.

5 Adapted baselines from related work

Given the number of work towards characterising policies (Table 1), we adapt two of the previous work to form two baseline solutions. These methods share the same limitations as supervectors discussed above.

5.1 Discriminator as a state-density estimator

Motivated by the earlier use of discriminator networks as a density estimator [12, 35], we train one network per policy to distinguish between states this policy encounters versus all other policies. Specifically, we train a discriminator neural network D_π by ascending the loss

$$\mathcal{L}_i = \mathbb{E}_{s \sim B_i} [\log D_i(s)] + \mathbb{E}_{s \sim B_{j \neq i}} [\log(1 - D_i(s))] , \quad (7)$$

where i is the index of the policy we are about to compare to others. Optimal discriminator satisfies properties [16, 35]

$$D^*(s) = \frac{p_i(s)}{p_i(s) + p_{j \neq i}(s)} \quad (8)$$

$$\Rightarrow \frac{p_i(s)}{p_{j \neq i}(s)} = \frac{D^*(s)}{1 - D^*(s)} = r_i(s) , \quad (9)$$

where $p_i(s)$ is shorthand for $p_{B_i}(s)$. This value $r_i(s)$ represents how likely it is the state was observed by π_i than by the rest of the policies. With this in mind, we define a pseudo-distance between two policies,

$$f_i(s) = e^{-\log r_i(s)} \quad (10)$$

$$d(i, j) = \mathbb{E}_{s \sim B_i} [f_j(s)] + \mathbb{E}_{s \sim B_j} [f_i(s)] , \quad (11)$$

summing distance both ways for symmetry. Essentially, for each policy, we measure how likely it is that their data is also contained in others' datasets and vice versa. In practice, the expectations are computed as averages over the datasets. We opt for training one discriminator per policy (policy's states versus all others) rather than one per comparison (policy X's states versus policy Y's) to avoid the quadratic explosion of discriminator training operations required. Training details can be found in Appendix B.4.

5.2 Trajectory encoder

Wang et al. [52] proposed to train a *variational autoencoder* (VAE) [27] to reconstruct trajectories. The encoder turns a sequence of states (trajectory) into an embedding z , and the decoder reconstructs original state-action pairs from this embedding. We adapt this method for policies by modelling

the distribution of trajectory embeddings. We train the system on data from all policies, then use the encoder to sample one embedding per trajectory, and then fit a multivariate Gaussian on these embeddings. The distance between two policies is then symmetric KL-divergence of these two Gaussians, where we sum the KL-divergences both ways. Further details are available in Appendix B.5.

Compared to the two other proposed methods, this method can capture temporal dependencies between states, while the other two methods discard this information. However, this could lead to posterior collapse where the decoder predicts future states purely based on previous states, ignoring the embedding. If the environment and/or policy are complex, we need a complex model to encode the whole policy into a fixed embedding. Finally, out of the three methods, training the encoder is the most computationally expensive.

6 Related work

The *de facto* method for comparing policies is by evaluating their performance in a task [25], or if the environment permits it, we can compare the different coordinates policy visits [39, 8]. Hernandez et al. [21] make use of dimensional reduction to plot fixed-length trajectories of different self-play algorithms to study the evolution of policies, and Matusch et al. [32] compare discretized transition matrices to study which metrics correlate with human behaviour. Others propose decompressing compact representations of policies back into functional policies, such as from genomes [13] or random generator seeds [48].

Trust-region policy optimization [45] limits per-update change to policy’s behaviour to ensure improvement in returns, but similar methods can also be used to encourage novel behaviour [29, 8] or diverse behaviour [12, 37]. An RL agent can be encouraged to visit unseen states [2], where visit counts can be estimated by prediction error [6]. Framing imitation learning as a task of matching state or state-action distributions has also been successful [24, 35, 30]. Compact policy embeddings have also been used to create generalized value functions [43, 50, 19] or to improve imitation learning [52].

7 Experiments and results

For an empirical assessment of the proposed BC (policy supervectors), we first evaluate it against the baseline BCs and then explore their use in various tasks. We aim to answer the following research questions (RQs). **RQ1** Does the BC separate truly different policies (expressivity)? **RQ2** How sensitive are the BCs to the random sampling of the environment? If BC does not describe the policy, there is no reason to use it. If it produces wildly different results even with a lot of data, its results may not be trusted. Note that we have not included action-based BCs because of the reasons discussed in Section 3.1 (their measurements can not be trusted). However, we will compare to them in trust-region experiments later (Section 7.5).

We use five classic control environments from the OpenAI Gym library [5] which include low dimensional environments (Cartpole, Pendulum and Acrobot) as well as two higher-dimensional environments (Lunarlander and Bipedal-walker). We opt for these simple environments to allow scaling the experiments. We train three *proximal policy optimization* (PPO) [46] agents per environment and store 100 versions of the policy during training. For each policy, we then collect a set of trajectories, which we use to compute the state-based BCs and finally measure distances between these 100 policies. We repeat this data collection and distance measurement three times. Distance matrices are min-max normalized to $[0, 1]$ to allow comparison between different repetitions. Technical details, source code and further results described can be found in the Appendix B.

7.1 Baselines

In addition to adapted baselines (Section 5), we include two baseline solutions.

Discretization. Similar to Matusch et al. [32], we discretize each state dimension to ten equally spaced bins, count occurrences of each bin and divide by the number of states visited. The distance between two policies is then $d(p_i, p_j) = \frac{1}{2} \sum_s |p_i(s) - p_j(s)|$. We opt to use this sum over KL-divergence due to the prevalence of zeros in the distributions. While simple, this method’s computational requirements explode with the increasing number of dimensions.

Single Gaussian. We fit a multivariate Gaussian on the sampled states and define distance between two policies as $d_{\text{KL}}(p_i||p_j) + d_{\text{KL}}(p_j||p_i)$ for symmetry. This has been used previously in tracking what states policy has visited [3].

7.2 Evaluation metrics

Correlation with return-distance.

To measure how well BC separates different policies, we measure the Pearson correlation between the absolute difference between average returns and distances measured by the BC. While returns are limited in their expressiveness, it still indicates the policies are doing something different if the returns differ. This especially applies to the environments we use, as the reward signal is tied to the states agent visits.

Distance error. Given a “ground truth” distance between policies under the same BC, we measure the average relative error between this ground truth and predicted distances. We select one of the repetitions with the highest amount of data (100 trajectories) to represent this ground truth. This metric aims to measure how sensitive the method is to the amount of data we collect.

Distance variance. As behavioural distances between a fixed set of policies should stay the same relative to each other, we study the variance in the results by computing *coefficient of variation* (CV) [11] σ/μ over the repetitions, where σ is the standard deviation of the distance over repetitions and μ is the sample average of distances. A lower value indicates more similar results over repetitions.

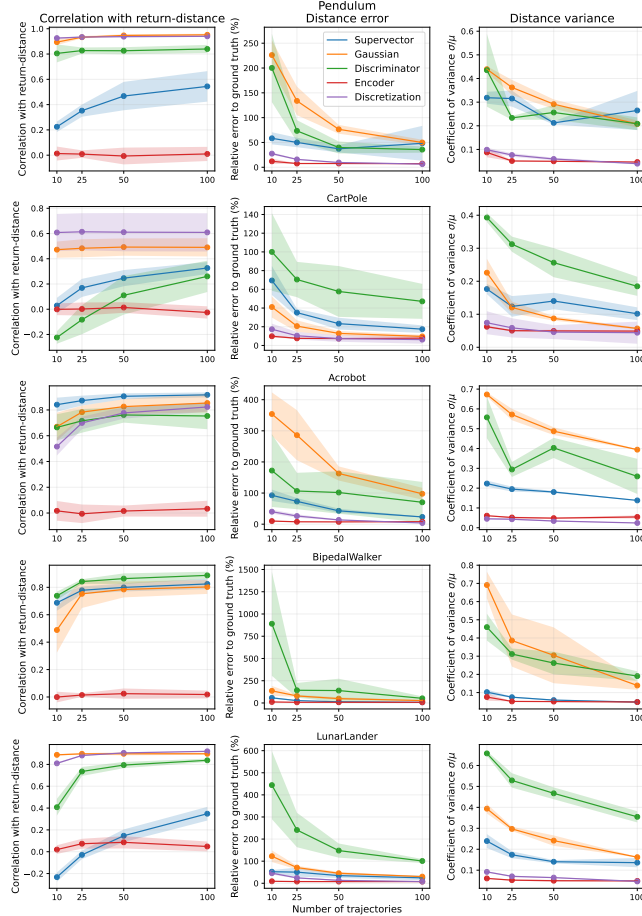


Figure 3: Evaluation results of different BCs. Averaged over three repetitions. The shaded region is plus/minus one standard deviation.

7.3 Evaluation results

Figure 3 shows the results for each environment separately. The single Gaussian and discriminator methods correlate with return differences, but require many trajectories per policy to stabilize results. Policy supervectors with 64 Gaussian components offer stabler results even at a lower number of trajectories, which we believe is due to the use of data from all policies in the UBM training. By sweeping over the different number of components, we find a strong connection between environments and the optimal number of components: in some environments one to four components provides the stablest results (Appendix B).

The discretization method appears to capture the behaviour as well as remain stable, but this method is computationally limited: experiments in the BipedalWalker environment ran out of memory (64GB of system memory) when trying to build the transition matrix over multiple dimensions. Trajectory encoders were similarly limited, which we found to be slow to train due to long sequences. We believe the low performance of encoder method is due to the difficulty of predicting future states. In summary,




Figure 4: t-SNE plot of the evolution of policies under different training algorithms, where each point represents a single policy. Plots under the same environment share the same plot scales. Rewards are scaled according to minimal and maximal attainable reward per environment.

we find policy supervectors a scalable and functioning state-based BC, with discretization a viable option in low-dimensional environments but not generally applicable due to memory constraints.

7.4 Studying evolution of policies under training algorithms

Evolution-based methods such as *neuroevolution of augmenting topologies* (NEAT) [47] and *covariance matrix adaptation ES* (CMA-ES) explore by modifying promising solutions and testing which location of the parameter space works better [18], while gradient-based RL methods like *advantage actor critic* (A2C) [34] and PPO update a single policy towards higher episodic reward. To better understand how these training algorithms evolve policies, we train agents with them and store checkpoints during training. We then extract policy supervectors of these checkpoints (with 64 components) and plot the resulting points with t-SNE [31] dimensionality reduction using the adapted distance (6). The hypothesis is that evolution-based methods cover a wider area of policies (random mutation of the parameters leads to different behaviours), while RL algorithms remain in a small region per run. We are not comparing which of these algorithms is better, rather we aim to understand if they explore different behaviours and how they evolve the policies.

Figure 4 shows the results with five A2C/PPO runs and one NEAT/CMA-ES run. In total, roughly 30,000 policies are compared against each other, with a varying number of policies from different algorithms depending on the settings (see Appendix C for details and remaining plots). NEAT and CMA-ES cover a wide area of behaviours, as expected. RL solutions cover smaller areas, with PPO forming distinguishable “worms” in more complicated environments (BipedalWalker), while A2C forms small clouds. This suggests the trust-region restriction of PPO is visible in policy supervectors as well, where the change in behaviour between successive policies is small, whereas A2C updates may change behaviour considerably.

Indeed, we find a positive correlation between total distance travelled by the trained policy and PPO ratio-clip value (Appendix C), which controls the size of the trust region. We also find a negative correlation between distance travelled after an update and average return, indicating that the initial learning steps change the agent’s behaviour the most. These results concur with observations of Engstrom et al. [10], where the KL-distance between successive policies first increases and then decreases as training moves on. The results with *behavioural cloning* [40] indicate that distance between the expert policy and trained policy decreases as training progresses, with change to behaviour slowing down as training reaches closer to the expert’s performance.

7.5 Applying supervectors to trust-region policy optimization

The above insights on trust-region optimization and BC distances suggest that state-based BCs could also be used in trust-region optimization. To test this, we construct a N -dimensional grid world environment where the agent can move to one of N directions per grid, where one direction is correct (+1 reward) and rest either reset the episode or do nothing. This environment is constructed to reflect

the doorway scenario: some actions have a large impact, while others do not. Agents are trained with a modified PPO where we disable the policy ratio clip and instead check if state-based BC distance is larger than the threshold after every network update. We compare state-based BCs against no constraint and *total variation divergence* (Max TV) $d(\pi, \pi') = \max_s \frac{1}{2} \sum_a |\pi(a|s) - \pi'(a|s)|$ [45]. We sweep over threshold values and present results for the method with the largest area under the learning curve. Further details are available in Appendix E.

Results in Figure 5a indicate that action-based BC is best suited for this task, but both Gaussian and supervector approaches outperform no constraint. Gaussian BC slowly improves to the highest score overall, while supervector quickly reaches peak and then drops. We believe this is due to the non-gaussian nature of the state distribution, where supervector BC is able to move individual components on different modalities, allowing it to learn fast in the beginning but later become unstable. In the current form, state-based BCs are not practical as part of the training loop, but their advantage lies in the analysis of policies (Section 7.4) and when encouraging novelty (Section 7.6).

7.6 Applying supervectors to novelty search

Novelty search encourages policies to find new behaviours to approach the task [29], which can take the form of maximizing BC distance between policies [8]. While successful, previous work has used domain-specific heuristics like terminal states as BCs to describe the behaviour. Supervectors, or Gaussians, could be used instead as a more generalizable alternative. To test this, we construct a continuous 2D point environment similar to Pacchiano et al. [36], where policy is rewarded for travelling to the positive y-direction. A wall prevents the policy from directly moving in this direction, and the policy has to learn to go around it. We use the methods and code of Conti et al. [8], where novelty encouragement is combined with fitness (NSR-ES) or agent is trained only for fitness (ES). Further details are in Appendix D.

Results in Figure 5b show that both single Gaussian and policy supervector help policy to overcome the obstacle. Using terminal state yields slightly better (but not significantly so) results, likely due to it being an “aligned” [8] BC with the objective, as walking away from previous locations can improve fitness. This demonstrates that supervectors (or even simple Gaussians) could be used as state-based BCs to encourage novel behaviour.

8 Discussion and conclusions

As demonstrated, state-based BCs, and especially policy supervectors based on GMMs, can be used regardless of the environment while also scaling to thousands of policies. Compared to using returns or actions to define behaviour, state-based BC methods are more descriptive and capture dynamics of the environment into a behavioural description. These BCs can be used in the study of policies under different training algorithms, or to encourage exploration/novelty or impose limits on how much behaviour can change per policy update.

While policy supervectors are applicable to studying policies as an outside observer, they are not yet practical as a part of the training loop: they require expensive sampling of the environment (as opposed to reusing already collected samples). This direction could be further explored by creating more sample efficient solutions or differentiable state-based BCs to train agents. Other topics include combining supervectors with the behavioural embeddings framework [36], extending the novelty search experiments [37] and further analysis of parameters learned by policy supervectors, such as the meaning of each adapted Gaussian component.

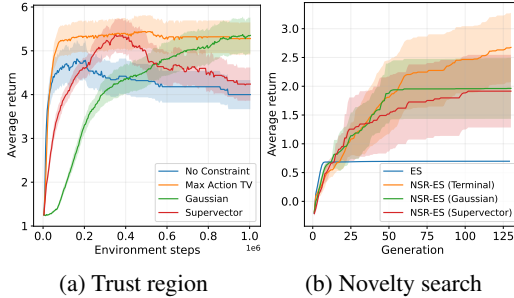


Figure 5: Results in novelty search (10 repetitions) and trust-region policy optimization (50 repetitions). Shaded region represents 95% confidence interval.

References

- [1] R. Agarwal, M. C. Machado, P. S. Castro, and M. G. Bellemare. Contrastive behavioral similarity embeddings for generalization in reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [2] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in neural information processing systems*, pages 1471–1479, 2016.
- [3] G. Berseth, D. Geng, C. Devin, C. Finn, D. Jayaraman, and S. Levine. SMiRL: Surprise minimizing rl in dynamic environments. *arXiv preprint arXiv:1912.05510*, 2019.
- [4] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [5] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. Openai gym, 2016.
- [6] Y. Burda, H. Edwards, A. Storkey, and O. Klimov. Exploration by random network distillation. In *International Conference on Learning Representations*, 2019.
- [7] W. M. Campbell, D. E. Sturim, and D. A. Reynolds. Support vector machines using gmm supervectors for speaker verification. *IEEE signal processing letters*, 13(5):308–311, 2006.
- [8] E. Conti, V. Madhavan, F. P. Such, J. Lehman, K. Stanley, and J. Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In *Advances in neural information processing systems*, pages 5027–5038, 2018.
- [9] A. Ecoffet, J. Huizinga, J. Lehman, K. O. Stanley, and J. Clune. Go-explore: a new approach for hard-exploration problems. *arXiv preprint arXiv:1901.10995*, 2019.
- [10] L. Engstrom, A. Ilyas, S. Santurkar, D. Tsipras, F. Janoos, L. Rudolph, and A. Madry. Implementation matters in deep policy gradients: A case study on PPO and TRPO. In *International Conference on Learning Representations*, 2020.
- [11] B. Everitt and A. Skrondal. *The Cambridge dictionary of statistics*, volume 106. Cambridge University Press Cambridge, 2002.
- [12] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine. Diversity is all you need: Learning skills without a reward function. *arXiv:1802.06070*, 2018.
- [13] A. Gaier and D. Ha. Weight agnostic neural networks. In *Advances in Neural Information Processing Systems*, pages 5364–5378, 2019.
- [14] A. Gaier and D. Ha. Weight agnostic neural networks. *arXiv preprint arXiv:1906.04358*, 2019.
- [15] J.-L. Gauvain and C.-H. Lee. Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains. *IEEE transactions on speech and audio processing*, 2(2):291–298, 1994.
- [16] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27:2672–2680, 2014.
- [17] A. Grover, M. Al-Shedivat, J. Gupta, Y. Burda, and H. Edwards. Learning policy representations in multiagent systems. In *International conference on machine learning*, pages 1802–1811. PMLR, 2018.
- [18] N. Hansen and A. Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [19] J. Harb, T. Schaul, D. Precup, and P.-L. Bacon. Policy evaluation networks. *arXiv preprint arXiv:2002.11833*, 2020.
- [20] P. Henderson, R. Islam, P. Bachman, J. Pineau, D. Precup, and D. Meger. Deep reinforcement learning that matters. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018.

- [21] D. Hernandez, K. Denamganai, S. Devlin, S. Samothrakis, and J. A. Walker. A comparison of self-play algorithms under a generalized framework. In *IEEE Conference on Games*, 2019.
- [22] J. R. Hershey and P. A. Olsen. Approximating the kullback leibler divergence between gaussian mixture models. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, volume 4, pages IV–317. IEEE, 2007.
- [23] A. Hill, A. Raffin, M. Ernestus, A. Gleave, A. Kanervisto, R. Traore, P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, and Y. Wu. Stable baselines. <https://github.com/hill-a/stable-baselines>, 2018.
- [24] J. Ho and S. Ermon. Generative adversarial imitation learning. In *Advances in neural information processing systems*, pages 4565–4573, 2016.
- [25] S. M. Jordan, Y. Chandak, D. Cohen, M. Zhang, and P. S. Thomas. Evaluating the performance of reinforcement learning algorithms. In *International Conference on Machine Learning*, 2020.
- [26] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014.
- [27] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations*, 2014.
- [28] T. Kinnunen and H. Li. An overview of text-independent speaker recognition: From features to supervectors. *Speech communication*, 2010.
- [29] J. Lehman and K. O. Stanley. Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336, 2008.
- [30] S. Levine, A. Kumar, G. Tucker, and J. Fu. Offline reinforcement learning: Tutorial, review, and perspectives on open problems. *arXiv preprint arXiv:2005.01643*, 2020.
- [31] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008.
- [32] B. Matusch, J. Ba, and D. Hafner. Evaluating agents without rewards. *arXiv preprint arXiv:2012.11538*, 2020.
- [33] G. J. McLachlan and K. E. Basford. *Mixture models: Inference and applications to clustering*, volume 38. M. Dekker New York, 1988.
- [34] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, 2016.
- [35] T. Ni, H. Sikchi, Y. Wang, T. Gupta, L. Lee, and B. Eysenbach. f-IRL: Inverse reinforcement learning via state marginal matching. In *Conference on Robot Learning*, 2020.
- [36] A. Pacchiano, J. Parker-Holder, Y. Tang, K. Choromanski, A. Choromanska, and M. Jordan. Learning to score behaviors for guided policy optimization. In *International Conference on Machine Learning*, pages 7445–7454. PMLR, 2020.
- [37] J. Parker-Holder, A. Pacchiano, K. Choromanski, and S. Roberts. Effective diversity in population-based reinforcement learning. In *Advances in Neural Information Processing Systems*, 2020.
- [38] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [39] S. Pitis, H. Chan, S. Zhao, B. Stadie, and J. Ba. Maximum entropy gain exploration for long horizon multi-goal reinforcement learning. In *International Conference on Machine Learning*, 2020.

- [40] D. A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. In *Advances in neural information processing systems*, pages 305–313, 1989.
- [41] A. Raffin. RL baselines zoo. <https://github.com/araffin/rl-baselines-zoo>, 2018.
- [42] A. Raffin, A. Hill, M. Ernestus, A. Gleave, A. Kanervisto, and N. Dormann. Stable baselines3. <https://github.com/DLR-RM/stable-baselines3>, 2019.
- [43] R. Raileanu, M. Goldstein, A. Szlam, and R. Fergus. Fast adaptation to new environments via policy-dynamics value functions. In *International Conference on Machine Learning*, pages 7920–7931. PMLR, 2020.
- [44] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn. Speaker verification using adapted gaussian mixture models. *Digital signal processing*, 10(1-3):19–41, 2000.
- [45] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz. Trust region policy optimization. In *International Conference on Machine Learning*, 2015.
- [46] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- [47] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- [48] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint arXiv:1712.06567*, 2017.
- [49] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [50] H. Tang, Z. Meng, J. Hao, C. Chen, D. Graves, D. Li, W. Liu, and Y. Yang. Represent your own policies: Reinforcement learning with policy-extended value function approximator. *arXiv preprint arXiv:2010.09536*, 2020.
- [51] S. Wang, S. Toyer, A. Gleave, and S. Emmons. The imitation library for imitation learning and inverse reinforcement learning. <https://github.com/HumanCompatibleAI/imitation>, 2020.
- [52] Z. Wang, J. S. Merel, S. E. Reed, N. de Freitas, G. Wayne, and N. Heess. Robust imitation of diverse behaviors. In *Advances in Neural Information Processing Systems*, pages 5320–5329, 2017.

A Grid world illustration setup (Section 3)

Results in Figure 1 of the original paper are the result of running the deterministic policies shown in the grid world images for ten thousand episodes. Agent receives a positive reward if it reaches the goal (higher the faster it reaches the goal with +1 being maximum), otherwise zero. Return-based distance is the absolute difference of average episodic rewards. Both action and state-based distances are computed as a sum of absolute differences between respective distributions. A policy is represented by tensor $\pi \in \mathbb{R}^{5 \times 5 \times 4}$, representing probability of choosing one of the four directions in each grid, and distance between two such matrices is defined as $d_{\text{action}}(\pi, \hat{\pi}) = \sum_{i=1, j=1, k=1}^{5, 5, 4} |\pi_{i,j,k} - \hat{\pi}_{i,j,k}|$, summing over all grid and action pairs. Grids marked as walls are ignored from this. State distribution is estimated by running policy on the environment for ten thousand episodes and dividing visits to a cell by a total number of steps done during ten thousand episodes. The distance between two state distributions is then computed as $d_{\text{state}}(p_{\pi}, p_{\hat{\pi}}) = \sum_{i=1, j=1}^{5, 5} |p_{\pi}(i, j) - p_{\hat{\pi}}(i, j)|$.

B Behavioural characterization evaluation experiment

B.1 Environments and training hyperparameters

Exact versions of the Gym environments are Acrobot-v0 (6D), Pendulum-v0 (3D), CartPole-v1 (4D), LunarLander-v2 (8D) and BipedalWalker-v3 (24D). The dimensionality refers to the observation vector dimensionality, which is also used as a state for the policy supervector extraction. We selected these environments to cover different levels of dimensionality and complexity while still remaining lightweight enough to allow us to scale experiments.

For the PPO algorithm we use the implementation from stable-baselines [23] with hyperparameters provided by the rl-zoo package [41] (see Table 3). In all environments, the agent reaches optimal reward with these settings. All agents use a neural network of two 64-unit layers with tanh-activations for value and policy estimation (separate network for both).

B.2 Evaluation metrics and data collection (Section 7.2)

We train three PPO agents per environment and store 100 checkpoints of the policy during training (a set of policies), and then collect 200 trajectories of data per policy. We then sample $\{100, 50, 25, 10\}$ trajectories from these 200, and repeat this sampling three times to simulate different runs of sampling of data. Three sets of policies per environment are used to capture stochasticity of final results over different sets of policies, while three repetitions of data sampling aim to capture stochasticity from sampling the environment. The axis over which we average for final results depends on the evaluation metric.

Correlation with return-distance. For each policy we measure the average episodic reward over the 200 trajectories and define return-distance as absolute difference between average returns $d_{\text{returns}} = |\mathbb{E}[R_{\pi_i}] - \mathbb{E}[R_{\pi_j}]|$, where R_{π} is the sum of rewards of a trajectory. We then compute the Pearson correlation between this distance and the distance measured by the BC. The shown result is averaged over three sets of policies, three repetitions and over all 100×100 distance pairs, minus symmetric duplicates.

Distance error. We compute distances between all policies inside a set. As policies remain fixed, and so does their behaviour, distances between them should stay the same relative to each other. If BC requires training, as is the case with our proposed methods, the scale of these numbers may change. For this reason, we normalize the distance matrices by min-max normalization to $[0, 1]$ over the whole distance matrix. We then select the first repetition with 100 trajectories as “ground truth”, and for rest of the repetitions we measure the relative error to this distance $|d - d_{\text{ground truth}}|/d_{\text{ground truth}}$. The report result is averaged over the three sets of policies.

Distance variance. We compute distance matrices as above, including normalization, and then compute the standard deviation over distance matrices over the three repetitions. The result shown is then an average of this over three sets of policies.

B.3 Implementation details of policy supervectors (Section 4)

See Algorithm 2 for pseudo-code of the method. We use scikit-learn to train GMMs [38]. We use relevance factor $r = 16$ in all of our experiments. GMMs use diagonal covariance matrices and are initialized with k-means before EM-training (the default setting). Figure 6 repeats these experiments with different number of components, which was used to decide the 64 components used in the main results (the lowest overall distance variance and error).

B.4 Implementation details of the discriminator method (Section 5.1)

See Algorithm 3 for pseudo-code of the method. The network consists of three layers of 256 units with tanh-activations, followed by a linear mapping to a scalar value. This value was clipped to interval $[-10, 10]$ (following Ni et al. [35]) and then fed through sigmoid function for $[0, 1]$ value. Given two datasets the discriminator is then trained to output high values for one and low values for the second one (see Section 4.2) for 30 epochs with Adam optimizer (learning rate 10^{-3}) with mini-batches of 128 samples each.

B.5 Implementation details of the trajectory encoder method (Section 5.2)

See Algorithm 4 for pseudo-code of the method. Following Wang et al. [52], the encoder is a single bi-directional LSTM layer with 256 units. For a single trajectory of states the backward output of this encoder is averaged over and then fed through a fully connected layer to parametrize a Gaussian to sample the embedding from (one fully connected layer for mean and another for variance). The decoder is a single fully connected layer that maps the sampled latent ($d = 256$) and previous state vector to a Gaussian (mean and variance), and finally the whole system is trained to maximize the log-likelihood of successive states in the dataset. In other words, the decoder needs to predict what will be the next state, given the previous state and a latent that (supposedly) describes the policy. After training a policy is represented by encoding all trajectories, sampling one latent each and finally fitting a diagonal, multivariate Gaussian on these latents, which is used to describe the policy.

Training of the system is done with Adam optimizer [26] (learning rate 10^{-3}) for five epochs over the data (data from all policies which are about to be compared), using mini-batches of eight trajectories at a time. We found five epochs to be sufficient to reach a stable loss, and eight trajectories per batch as one trajectory usually consists of hundreds of steps. Training these encoders took an order of magnitude longer than *e.g.* discriminator method above, despite the number amount of epochs and small network size. On our four-core Intel Xeon system, the bipedal-walker experiments took four days alone.

B.6 Implementation details of the discretization method

For each policy we construct a matrix of $p_\pi = \mathbb{R}^{10 \times 10 \dots D \dots \times 10}$, where D is the dimensionality of the state vector. For each state dimension we create 10 uniformly spaced bins with edges at the observed minimum and maximum over all the data we are about to compare. We then discretize each state and increment corresponding cells in the p_π , and finally divide by the total number of states to obtain probability of being in that state over all states. Distance between two policies is measured as $d_{\text{discretization}}(p_{\pi_a}, p_{\pi_b}) = 0.5 \sum_s |p_{\pi_a}(s) - p_{\pi_b}(s)|$, where s goes over all possible discretized states.

C Analysis of policies under different training algorithms (Section 7.4)

C.1 Experiment setup for t-SNE visualization

We trained five A2C agents, five PPO agents, one NEAT and one CMA-ES agent on the Gym environments used in the evaluation experiments. For A2C and PPO agents we use the same setup as above, and for NEAT/CMA-ES we use the code of Wann et al. [14] with the hyperparameters shown in Table 3. We include only one run of NEAT/CMA-ES as they generate a large number of candidates. CMA-ES agents use fixed two-layer networks with sigmoid activations, with 40 units each for BipedalWalker and 5 units for other tasks. CMA-ES training terminates when fitness does not improve, which results in a smaller cloud of points for CMA-ES in t-SNE figures.

For RL agents we store 500 checkpoints of policies per training run, and for NEAT/CMA-ES we store 25% of the population candidates every tenth generation for visualization. We then train one UBM with 64 components per environment, using at most 100 policies per algorithm and/or 10M samples in the dataset due to memory constraints. With the UBM we extract the policy supervectors and use the upper bound of KL-divergence as a distance for t-SNE, which is used to plot the results. See per-environment t-SNE visualizations in Figure 7.

C.2 Analysis of trust-region optimization

To study the effect of trust-region optimization on policy’s behaviour during training, we train PPO agents with different ratio-clip values ϵ , where the lower value should restrict updates to smaller changes in the policy. We train five PPO agents per different clip-ratio levels, where all other hyperparameters for training are the same (Table 3) except we set clip-ratio to the studied value. We also train five A2C agents with the original settings for comparison. We store 50 policies during the training, gather 100 trajectories per policy, train UBMs (64 components), extract policy-supervectors and compute the distance between successive policies experienced during training (a vector of 49 values per one training run). We then sum these distances together for “total distance travelled”, average over the five repetitions and report the result. Results can be found in Table 4, where we can see a connection between higher clip-ratio (larger trust-region) and the amount of distance travelled by the policy during training. This is especially pronounced in higher-dimensional environments, and also when comparing PPO (with trust-region) against A2C (no trust-region).

To study how much change in behaviour is required as the agent gets better, we measure the Pearson correlation between average returns and distance to the next policy. Results in Figure 8 suggest that once policy gets better in the task, the amount of change to its behaviour decreases. Table 4 shows the final results, where we see a strong negative correlation between behaviour change and amount of returns, which concurs with the original observation. Results also show that the average distance travelled by policy increases as we increase the size of the trust-region, which concurs with the intuition of trust-region optimization.

C.3 Analysis of imitation learning

To study the connection between behaviour changes and imitation learning, we train five PPO expert agents on BipedalWalkerHardcore-v3 and LunarLander-v2 (hyperparameters in Table 3), collect 10 trajectories of data and use it to train a behavioural cloning agent with 50 epochs over the data. We chose a harder version of the bipedal-walker environment for a more challenging environment. We use PPO implementation from stable-baselines3 [42] and imitation learning code from “imitation” library [51].

After each epoch, we store the current version of the trained policy. After training, we collect additional 100 trajectories for the expert agent and for each BC policy, train the UBM (64 components), extract policy-supervectors and measure the distance of all BC policies to the expert agent. The resulting learning curve and distance-to-expert are shown in Figure 9, where we see a connection between changes in returns and distance to expert policy, with an average Pearson correlation of -0.465 with a standard deviation of 0.297 . This shows that as behavioural cloning is training the policy to get a better average return (similar to expert’s), its behaviour also approaches expert’s behaviour.

D Novelty search experiment (Section 7.6)

The point-environment used in novelty search mimics the point environment used by Pacchiano et al. [36] but created in plain Python without MuJoCo. The agent lives on a 2D plane (no bounds), where observations are the current coordinates. Action space consists of a single real value $\phi \in [0, 2\pi]$ that represents the direction agent should move on to the next step (constant movement speed). Agent is rewarded based on movement on the y-axis directly with $r = y_{\text{new}} - y_{\text{old}}$. To move further in this axis, the agent has to avoid a trap by first moving either left or right on the x-axis (and slightly negative direction in the y-axis). If the agent does not avoid this trap it only gets 0.65 reward, while the optimal the solution reaches 3.7 reward.

For ES/NSR-ES implementations we used the code of Conti et al. [8] and parameters from the humanoid experiments, but lowered data per update to 100 episodes as the environment is significantly simpler than MuJoCo humanoid. We also decreased network size to two layers of 16 units, followed by tanh-activations. For novelty search, we use a population of $K = 3$ sets of parameters, with five episodes per BC (*e.g.* Gaussian/supervector trained on five episodes of data, or terminal state is averaged over five episodes as was in the original implementation). For Gaussian BC we used symmetric KL-divergence (summed KL-divergence both ways), and for supervector, we used four Gaussian components and the KL-divergence upper bound distance.

E Applying supervectors for trust-region optimization (Section 7.5)

To test the usefulness of state-based trust-region in policy optimization, the environment (dubbed “dangerous path”) is a N -dimensional grid world, where the player always starts in the same location. The environment provides coordinates of the player as observations, and at each step, the agent may choose one of N actions, which moves the player to the positive direction in that axis. For each location, there is only one valid action that moves the player forward and is rewarded with +1 reward. Two of the remaining actions are “mines” which will throw the player back to the beginning, and two actions do nothing. The assignment of these outcomes is done at random per grid. The agent is only rewarded when they move to a new grid location (*i.e.* when they choose the correct action for the first time). In the experiments we set $N = 5$ and limit episode length to 25 steps, making the maximum reward 25. We noticed the choice of the path can have a large effect on results, and for this reason, we ensure each compared method uses the same random seeds for the environment (*i.e.* with four methods and ten repetitions each, we have ten random seeds used by all four methods).

For the agent we modify the PPO implementation of stable-baselines3 [42] by removing the ratio clipping and replacing this with a constraint check: if the constraint is exceeded, we stop updating the policy with current data and move on to collect trajectories for the next policy update. The constraint is checked for after every parameter update, comparing the old policy (which collected data) to the current one (after updates). We use eight environments and collect 512 samples from each to perform policy updates to ensure a large amount of data to sample from. We do not use generalized advantage estimation and set the discount factor to zero as the agent is immediately rewarded for correct actions. The agent uses a two-layer network with 16 units and tanh-activations each. The policy is updated for 100 mini-batches of 64 items over the collected data, or until constraint prevents updates.

No constraint. Policy is always updated for 100 mini-batches.

Total variation divergence (Max TV). We measure total variation divergence of action probabilities over collected data and take the maximum of them. Essentially, this limits the maximum amount of how much probability of taking any single action (in any state) can change. Constraint value is searched from $\{0.001, 0.005, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5\}$.

Gaussian. We collect five trajectories of data and fit a multivariate, diagonal Gaussian on it, and measure symmetric KL-divergence between old and current policy (summing both ways). We add a small amount of Gaussian noise ($\sigma = 10^{-3}$) to collected states to ensure we have unique points. Constraint value is searched from $\{0.5, 1.0, 2.0, 3.0, 5.0, 10.0, 15.0, 20.0\}$.

Supervector. As above, but we fit a four-component UBM on the data from old and new policy (five trajectories from both), extract policy supervectors and compute the upper bound of KL-divergence. We use four components due to the simplicity of the environment. Constraint value is searched from $\{0.01, 0.05, 0.1, 0.15, 0.2, 0.3, 0.4, 0.5\}$

To ensure constraint violations happen we set the learning rate to relatively high 10^{-3} (for Adam optimizer, compared to default $3 \cdot 10^{-4}$). We then selected the constraint value with the largest area under the learning curve, which was averaged over ten repetitions, resulting in constraint values 0.4, 10.0 and 0.05 for Max-TV, Gaussian and supervector, respectively. With these settings each of the methods terminate the policy updates occasionally, separating them from results with no constraint. With a lower learning rate we need a higher number of allowed updates than 100, however, this started to become practically slow as Gaussian/supervector constraints constantly sampled the environment after every update. It is also worth noting that standard PPO with ratio-clipping $\epsilon = 0.2$ reached notably higher results than any of the compared results (above ten), which is due to more mini-batch updates per iteration (640 vs. 100) and guided trust-region (the constraint is part of the loss).

Table 2: Min/Max values used for reward scaling per environment.

Environment	Min reward	Max reward
Pendulum-v0	-1600	-200
Acrobot-v1	-500	-100
LunarLander-v2	-230	200
BipedalWalker-v3	-100	300
CartPole-v1	0	500

F Hardware used to run the experiments

All experiments were run on either a four-core Intel Xeon W-2125, 64GB RAM workstation or on a 16-core Intel Xeon E5-2620 v4 (dual socket), 64GB RAM server machine. The only hard requirement for our experiments is the amount of system memory (64GB or more) and mass storage for storing policy trajectories (approx. 300GB in total), and with the above systems, all of the experiments can be conducted in approximately one week when parallelized over the two systems.

Algorithm 1 Gathering data for a policy (function gatherData). Python-like pseudo-code with OpenAI Gym environments.

Input: An environment env , a policy π and number of trajectories to collect M .
#Initialize buffer.
 $B = []$
for $m \in \{1 \dots M\}$ **do**
 $s = env.reset()$
 done = False
 while not done **do**
 $B.append(s)$
 $a \sim \pi(s)$
 $s, done = env.step(a)$
 end while
end for
Output: Policy data B .

Algorithm 2 Computing distances with policy supervectors. Python-like pseudo-code.

Input: An environment env , a set of policies we wish to study $\pi_n, n \in \{1 \dots N\}$, number of GMM components K , number of trajectories to gather per policy M .
#Datasets of states for each policy.
for $n \in \{1 \dots N\}$ **do**
 $B_n = \text{gatherData}(env, \pi_n, M)$
end for
#Concatenate all data into one big dataset.
 $B_{\text{ubm}} = \{B_1, \dots, B_N\}$
#Fit a GMM with K components (UBM).
 $(\mu_{\text{ubm}}, \Sigma_{\text{ubm}}, w_{\text{ubm}}) = \text{gmm.fit}(B_{\text{ubm}}, K)$
#Extract supervectors (adapted means)
for $n \in \{1 \dots N\}$ **do**
 #MAP-adapt mean as in (1)-(5).
 $\hat{\mu}_n = \text{mapAdapt}(\mu_{\text{ubm}}, \Sigma_{\text{ubm}}, w_{\text{ubm}}, B_n)$
end for
#Initialize distance matrix.
 $\text{policyDistances} = \mathbf{0}^{N \times N}$
for $i \in \{1 \dots N\}$ **do**
 for $j \in \{1 \dots N\}$ **do**
 #Compute KL-divergence upper bound (6).
 $\text{policyDistances}_{i,j} = \text{KLUpperBound}(\hat{\mu}_i, \hat{\mu}_j, \mu_{\text{ubm}}, \Sigma_{\text{ubm}}, w_{\text{ubm}})$
 end for
end for
Output: Distance matrix policyDistances

Algorithm 3 Computing distances with discriminators. Python-like pseudo-code.

Input: An environment env , a set of policies we wish to study $\pi_n, n \in \{1 \dots N\}$, number of GMM components K and number of trajectories to gather per policy M .
#Datasets of states for each policy.
for $n \in \{1 \dots N\}$ **do**
 $B_n = \text{gatherData}(env, \pi_n, M)$
end for
#Initialize distance matrix.
 $\text{policyDistances} = \mathbf{0}^{N \times N}$
for $i \in \{1 \dots N\}$ **do**
 #Train discriminator to separate i^{th} dataset from all others.
 $D_i = \text{trainDiscriminator}(B_i, B_{j \neq i})$
 for $j \in \{1 \dots N\}$ **do**
 #Test how likely is that j^{th} dataset came from i^{th} policy using (10).
 $\text{distance} = \frac{1}{|B_j|} \sum_{s \in B_j} f_i(s)$
 #Accumulate both ways for symmetry.
 $\text{policyDistances}_{i,j} = \text{policyDistances}_{i,j} + \text{distance}$
 $\text{policyDistances}_{j,i} = \text{policyDistances}_{j,i} + \text{distance}$
 end for
end for
Output: Distance matrix policyDistances

Algorithm 4 Computing distances with trajectory encoders. Python-like pseudo-code.

Input: An environment env , a set of policies we wish to study $\pi_n, n \in \{1 \dots N\}$, number of GMM components K and number of trajectories to gather per policy M .
#Collect trajectories for each policy (keep trajectories separate).

```

for  $n \in \{1 \dots N\}$  do
   $T_n = []$ 
  for  $m \in \{1 \dots M\}$  do
    states = []
     $s = env.reset()$ 
    done = False
    while not done do
      states.append( $s$ )
       $a \sim \pi(s)$ 
       $s, done = env.step(a)$ 
    end while
     $T_n.append(states)$ 
  end for
end for
#Train a variational autoencoder to reconstruct trajectories, using all data.
encoder = trainVAE( $T$ )
#Encode and sample latents from each trajectory and fit one diagonal Gaussian per policy.
for  $n \in \{1 \dots N\}$  do
   $Z = []$ 
  for  $t \in \{1 \dots |T_n|\}$  do
     $z \sim \text{encoder}(T_{n,t})$ 
     $Z.append(z)$ 
  end for
   $\mu_n = Z.mean(axis=0); \sigma_n = Z.std(axis=0)$ 
end for
#Initialize distance matrix.
policyDistances =  $\mathbf{0}^{N \times N}$ 
for  $i \in \{1 \dots N\}$  do
  for  $j \in \{1 \dots N\}$  do
    #Compute KL-divergence between policies' Gaussians.
    distance = DiagonalGaussianKL( $\mu_i, \mu_j, \sigma_i, \sigma_j$ )
    #Accumulate both ways for symmetry.
    policyDistances $_{i,j}$  = policyDistances $_{i,j}$  + distance
    policyDistances $_{j,i}$  = policyDistances $_{j,i}$  + distance
  end for
end for
Output: Distance matrix policyDistances

```

Table 3: Agent training hyperparameters. “*” marks parameters that are linearly decayed to zero over training. "Bipedal.H." refers to BipedalWalkerHardcore, used in the imitation learning experiments with PPO expert agent.

PPO parameters	Bipedal.H.	Bipedal.	LunarLan.	Acrobot	CartPole	Pendulum
Environment steps	100 000 000	5 000 000	1 000 000	1 000 000	100 000	2 000 000
Number of envs.	16	16	16	16	8	8
Rollout size	2048	2048	1024	256	32	2048
Training epochs	10	10	4	4	20	10
Batch size	1024	1024	512	512	256	512
Entropy weight	0.001	0.001	0.01	0.0	0.0	0.0
Policy ratio clip	0.2*	0.2	0.2	0.2	0.2*	0.2
Value ratio clip	0.2*	0.2	0.2	0.2	0.2*	0.2
Learning rate	0.00025*	0.00025	0.00025	0.00025	0.001*	0.0003
GAE λ	0.95	0.95	0.98	0.94	0.8	0.95
Discount factor	0.99	0.99	0.999	0.99	0.98	0.99
A2C parameters						
Environment steps	-	5 000 000	200 000	500 000	500 000	2 000 000
Number of envs	-	16	8	16	8	8
N-steps	-	5	5	5	5	5
Entropy weight	-	0.0	0.00001	0.0	0.0	0.0
Learning rate	-	0.0007*	0.00083*	0.0007	0.0007	0.0007
Discount factor	-	0.99	0.995	0.99	0.99	0.95
NEAT parameters						
Number of gens.	-	1024	1024	256	256	256
Population size	-	192	128	128	128	128
CMA-ES parameters						
Number of gens.	-	1024	1024	256	256	256
Population size	-	192	128	128	128	128

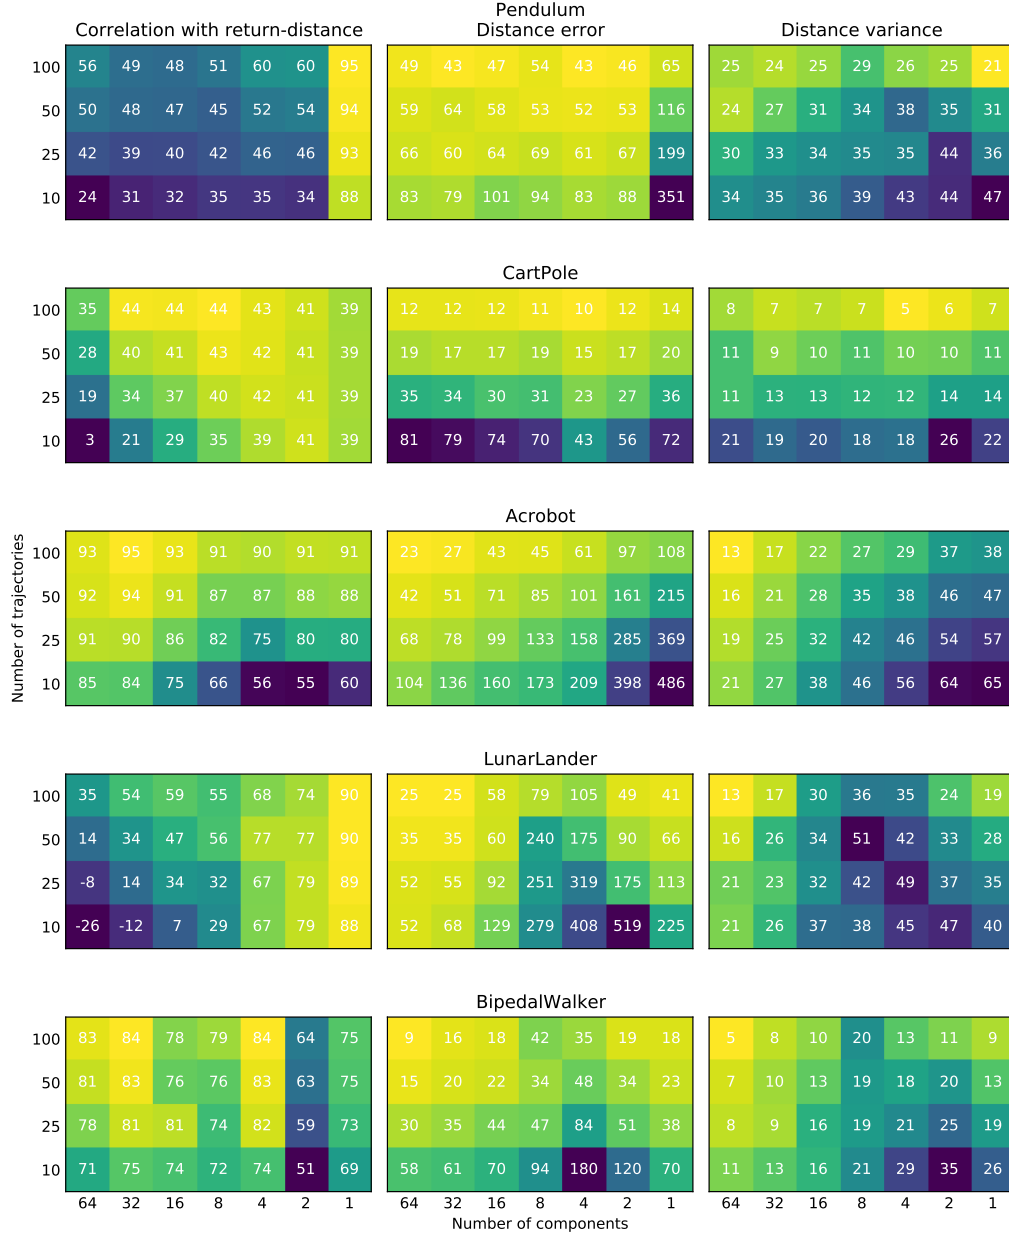


Figure 6: Brighter (yellow) colour means better. Per-environment results when analyzing the quality of final results under different amounts of data per policy and GMM components for policy supervectors. All values have been multiplied by a hundred.

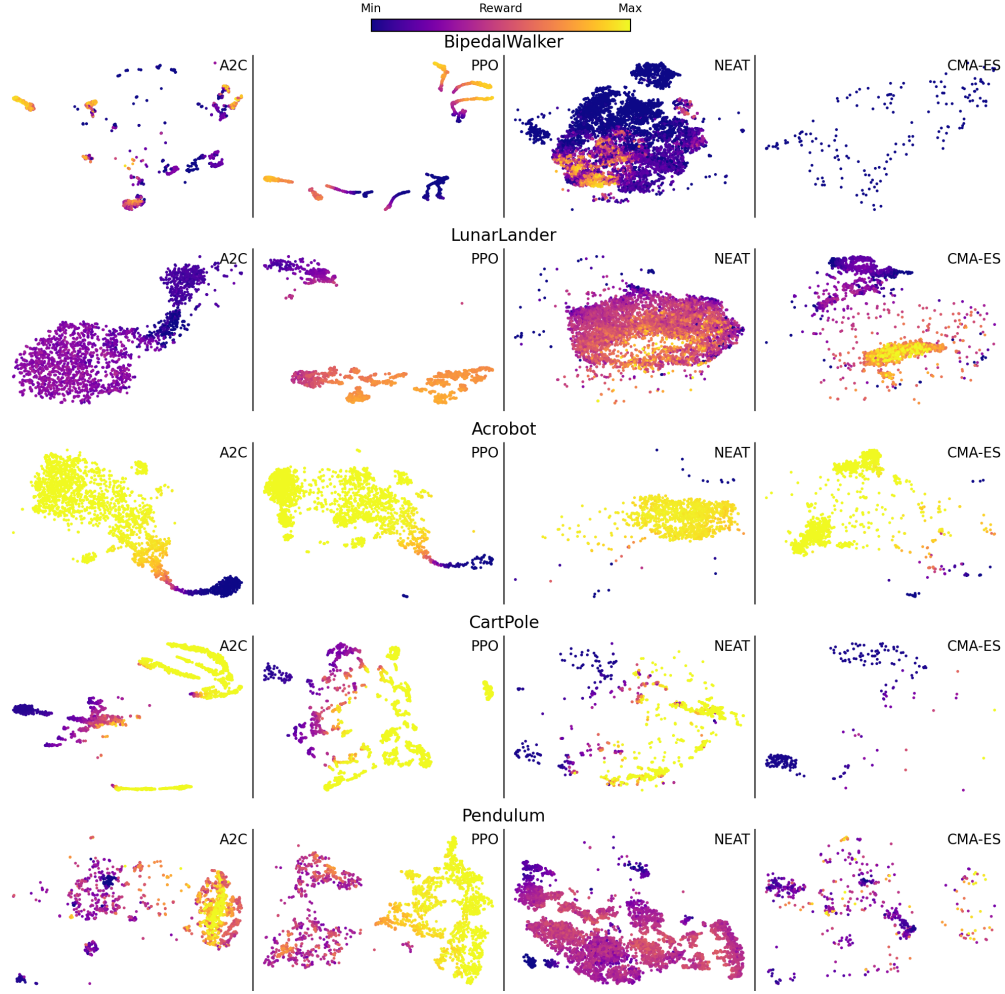


Figure 7: t-SNE plots of policies under different training algorithms in different environments. Each dot is a single policy (the behaviour characterization), where colour shows its average episodic return. A2C and PPO plots include policies from five different training runs, while CMA-ES and NEAT include policies from a single training run.




Figure 8: Learning curves with different clip-ratios ϵ when using PPO, with A2C included for comparison. Reward (dashed line) is normalized from minimum achievable reward (zero) to maximum (one). Distance refers to the distance between successive policy versions. Curves are averaged over five repetitions. Variance is omitted for visual clarity.

Table 4: Total distance travelled by the policy during PPO training with different ratio-clipping values ϵ , averaged over five training runs along with standard deviation. $\text{Corr}(R, d)$ is the Pearson correlation between average episodic reward and distance to next policy in training.

ϵ	Bipedal	Lunar	Acro	Cart	Pendulum
Distance traveled					
.01	23.6 \pm 1.1	4.5 \pm .5	1.0 \pm .1	2.0 \pm .1	2.8 \pm .4
.1	1.7 \pm 1.5	9.9 \pm .5	2.2 \pm .1	7.7 \pm 1.6	1.2 \pm .1
.2	11.3 \pm 2.0	9.7 \pm 1.1	2.6 \pm .2	13.6 \pm 1.0	1.3 \pm .2
.3	12.3 \pm 4.5	1.8 \pm .4	2.5 \pm .3	12.2 \pm 2.0	1.6 \pm .3
.4	11.9 \pm 2.3	14.7 \pm 2.3	2.4 \pm .3	14.9 \pm 1.2	2.1 \pm .5
.5	12.0 \pm 2.6	15.0 \pm .7	2.6 \pm .1	15.3 \pm 1.7	2.5 \pm .4
.75	17.7 \pm 7.4	16.7 \pm 1.0	2.5 \pm .3	15.9 \pm 3.7	3.7 \pm .5
1	25.1 \pm 1.9	18.9 \pm 2.5	2.6 \pm .2	14.0 \pm 4.0	4.1 \pm .6
2	19.0 \pm 3.4	17.3 \pm 1.2	2.7 \pm .2	17.1 \pm 4.6	3.5 \pm .8
5	25.4 \pm 7.2	22.9 \pm 4.3	2.9 \pm .3	2.2 \pm 2.6	3.4 \pm .3
10	51.5 \pm 21.1	21.5 \pm 2.6	2.8 \pm .2	14.9 \pm 4.9	4.5 \pm .7
A2C	61.8 \pm 16.6	29.4 \pm 4.3	2.2 \pm .2	6.5 \pm .7	9.8 \pm 2.4
$\text{Corr}(R, d)$					
.01	.25 \pm .18	-.06 \pm .12	.46 \pm .11	-.54 \pm .07	-.13 \pm .20
.1	-.47 \pm .19	-.25 \pm .03	-.34 \pm .23	-.14 \pm .11	-.53 \pm .04
.2	-.73 \pm .08	-.34 \pm .05	-.53 \pm .21	-.01 \pm .10	-.56 \pm .11
.3	-.61 \pm .07	-.45 \pm .11	-.43 \pm .14	-.20 \pm .14	-.33 \pm .10
.4	-.66 \pm .07	-.50 \pm .05	-.40 \pm .10	-.10 \pm .09	-.42 \pm .11
.5	-.58 \pm .12	-.43 \pm .10	-.35 \pm .13	.06 \pm .24	-.19 \pm .17
.75	-.54 \pm .11	-.32 \pm .09	-.40 \pm .23	-.14 \pm .28	.08 \pm .09
1	-.60 \pm .20	-.27 \pm .06	-.37 \pm .12	.23 \pm .30	.11 \pm .10
2	-.67 \pm .10	-.15 \pm .10	-.45 \pm .22	-.15 \pm .23	.04 \pm .28
5	-.63 \pm .14	.03 \pm .13	-.47 \pm .19	-.28 \pm .17	.19 \pm .18
10	-.58 \pm .13	.03 \pm .14	-.37 \pm .18	.04 \pm .39	.20 \pm .26
A2C	-.46 \pm .25	.40 \pm .10	.20 \pm .08	-.30 \pm .12	.20 \pm .18




Figure 9: Behavioural cloning learning curves. The reward is normalized by dividing with expert’s reward, with one equaling to performance of the expert. Distance refers to the distance to the expert policy. Averaged over five repetitions. The shaded area represents plus-minus one standard deviation.