

Multi-objective Model-based Policy Search for Data-efficient Learning with Sparse Rewards

Rituraj Kaushik

Inria, CNRS, Université de Lorraine
rituraj.kaushik@inria.fr

Konstantinos Chatzilygeroudis

Inria, CNRS, Université de Lorraine
konstantinos.chatzilygeroudis@inria.fr

Jean-Baptiste Mouret

Inria, CNRS, Université de Lorraine
jean-baptiste.mouret@inria.fr

Abstract: The most data-efficient algorithms for reinforcement learning in robotics are model-based policy search algorithms, which alternate between learning a dynamical model of the robot and optimizing a policy to maximize the expected return given the model and its uncertainties. However, the current algorithms lack an effective exploration strategy to deal with sparse or misleading reward scenarios: if they do not experience any state with a positive reward during the initial random exploration, it is very unlikely to solve the problem. Here, we propose a novel model-based policy search algorithm, Multi-DEX, that leverages a learned dynamical model to efficiently explore the task space and solve tasks with sparse rewards in a few episodes. To achieve this, we frame the policy search problem as a multi-objective, model-based policy optimization problem with three objectives: (1) generate maximally novel state trajectories, (2) maximize the cumulative reward and (3) keep the system in state-space regions for which the model is as accurate as possible. We then optimize these objectives using a Pareto-based multi-objective optimization algorithm. The experiments show that Multi-DEX is able to solve sparse reward scenarios (with a simulated robotic arm) in much lower interaction time than VIME, TRPO, GEP-PG, CMA-ES and Black-DROPS.

Keywords: Model-based Policy Search, Exploration, Sparse Reward

1 Introduction

Reinforcement Learning algorithms (RL) now allow artificial agents to learn to play many of the Atari 2600 games directly from pixels [28] or to beat the world’s best players at Go and chess with minimal human knowledge [41]. Unfortunately, these impressive successes rely on very long *interaction times* between the algorithm and the system: for example, 38 days of play (real time) for Atari 2600 games [28], 4.8 million games for Go [41], or about 100 hours of simulation time (much more for real time) for a 9-DOF mannequin that learns to walk [20]. This makes these algorithms well suited to simulated environments, but challenging to use with real robots, for which it is often not materially possible to perform more than few dozen of trials of a few seconds/minutes.

At the other end of the spectrum, model-based policy search algorithms can learn policies in a few minutes of interaction time, at the expense of a significant computation time between episodes. For instance, PILCO [11] and Black-DROPS [5, 4] can learn a policy to balance a cart-pole or a policy to control a 5-DOF manipulator in less than 40-60 seconds of interaction time. However, these algorithms implicitly assume that most states can be associated to a positive or negative reward, as this is the case when the objective is to balance a pole or to follow a trajectory with a manipulator. In other words, these algorithms are essentially greedy and mostly exploiting.

While this assumption often holds in control tasks, rewards are much more sparse in many interesting learning scenarios: typically, we would like to reward the robot when it successfully achieves the




Figure 1: Overview of the Multi-DEX algorithm: The robot interacts with the environment with a policy in an episodic scheme; after each episode, Multi-DEX learns a GP model of the dynamics of the system and then performs multi-objective policy search to obtain a set of Pareto-optimal policies based on novelty, cumulative reward, and model variance; last, a policy is selected from the Pareto set based on an exploration parameter ε and executed on the robot to collect new data.

task, and not for all the intermediate steps that we think should lead to success. For example, a robot might need to open a drawer and get rewarded by how much the drawer is open: in most of the state space, the reward is zero because the robot does not even touch the handle, meaning that PILCO or Black-DROPS need to open the drawer purely by chance to start learning a policy. In most realistic scenarios, this is unlikely to happen.

When there is no reward available, the robot should be “curious” and search for “interesting stepping stones”. This question is central in reinforcement learning [22, 2, 27, 1, 14], developmental robotics [34, 18], and evolutionary robotics [25, 32, 31, 13]. Among the most representative ideas, Novelty Search [25, 31] attempts to find behaviors that are as different as possible from what has been tested before; Self-Adaptive Goal Generation Robust Intelligent Adaptive Curiosity (SAGG-RIAC) selects random target states and searches for policies/trajectories to reach them; or Adaptive Curiosity pushes the robot towards situations in which it maximizes its learning progress [34, 35]. Unfortunately, all these approaches require thousands (often more) learning episodes; in addition, they often focus solely on the exploration and often disregard that the robot might have a task to fulfill (and not only explore its body and the environment).

Our main insight is that we can leverage ideas from Novelty Search and developmental robotics but exploit them *in a learned dynamical model* to keep the interaction time as low as possible. In essence, we would like to fuse the “novelty-driven” approaches with model-based policy search. At their core, modern model-based policy search algorithms alternate between learning a probabilistic dynamical model and optimizing a policy according to the model (and the uncertainty of the model) [11, 5]: In this paper, we propose to extend this idea to not only optimize for policies that are likely to work on the robot, but also for those that are predicted to be good for exploration. Our second insight is that we can optimize for novelty and cumulative reward using a Pareto-based Multi-Objective Evolutionary Algorithm (MOEA) [9], which simultaneously searches for all the Pareto-optimal trade-offs (instead of a single optimal solution when using an aggregated objective, *Suppl. Mat. Sec 2*).

More concretely, we propose a novel model-based policy search algorithm, Multi-DEX (Multi-objective Data-Efficient eXploration), that efficiently explores the task space by framing the policy search problem as a multi-objective optimization problem with three objectives (Fig. 1): (1) generate maximally novel state trajectories (2) maximize the cumulative reward and (3) keep the system in state-space regions for which the model is as accurate as possible. This optimization is performed with a Pareto-based multi-objective algorithm [9] after each episode, once the dynamical model is updated with the new data. Thanks to the learned probabilistic dynamical model, we show that our proposed approach is highly data-efficient and typically needs only a few minutes of interaction time with the actual system to solve tasks with rare state transitions and sparse rewards.

We compare the performance of Multi-DEX with recent algorithms for sparse reward scenarios (VIME [21] and GEP-PG [6]), a state-of-the-art model-free policy search algorithm (TRPO [40]), a state-of-the-art black-box optimizer (CMA-ES [19]), and a model-based policy search algorithm (Black-DROPS [5]) in a sequential goal reaching task and in a drawer opening task; we show that Multi-DEX solves both tasks in order of magnitude less interaction time than the others.

2 Problem Formulation

We consider the following dynamical systems of the form:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + f(\mathbf{x}_t, \mathbf{u}_t) + \mathbf{w} \quad (1)$$

with continuous-valued states $\mathbf{x} \in \mathbb{R}^E$ and controls $\mathbf{u} \in \mathbb{R}^F$, i.i.d. Gaussian system noise \mathbf{w} , and unknown transition dynamics f . Our goal is to find a *policy* $\pi(\mathbf{u}|\mathbf{x}, \boldsymbol{\theta})$ parameterized by $\boldsymbol{\theta} \in \mathbb{R}^\Theta$ that maximizes the *expected return* when following the policy $\pi_{\boldsymbol{\theta}}$ for T time steps:

$$J_r(\boldsymbol{\theta}) = \mathbb{E} \left[\sum_{t=1}^T r(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}) \middle| \boldsymbol{\theta} \right] \quad (2)$$

where $r(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1}) \in \mathbb{R}$ is the reward for being in state \mathbf{x}_t , taking action \mathbf{u}_t , and reaching state \mathbf{x}_{t+1} . We assume that r can be sparse or may have plateaus; *i.e.*, it can be zero for most values of $(\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{t+1})$ or have large regions with constant value.

3 Related Work

3.1 Encouraging Exploration in RL and Other Fields

Exploration in trial-and-error learning has been investigated almost independently in several fields, with different inspirations and metaphors: in “traditional” reinforcement learning [22, 23], in deep reinforcement learning [1, 38, 14, 6], in developmental robotics [16, 15], and in evolutionary robotics [25, 8].

In RL with discrete state-action spaces, Bayesian RL [22] and PAC-MDP methods [23] offer exploration with formal guarantees. When the state/action space is continuous, as this is usually the case in robotics, most successful approaches encourage exploration by performing *action perturbation* [40] or on policy *parameter perturbation* [19]. *Directed exploration* approaches try to cover the whole state-action space as uniformly as possible. In this context, approaches such as [2, 27] try to drive the agent towards novel states. Hindsight Experience Replay (HER) [1] exploits the fact that when experiencing a state-action trajectory, the goal being pursued influences the agent’s actions but not the dynamics and therefore trajectories can be replayed with an arbitrary goal in an off-policy RL algorithm. Recently, [14] proposed a “reverse learning” approach where the agent starts learning from a state nearby to the goal and gradually learns from more difficult initial states.

In order to explore large and continuous state-action spaces efficiently, many approaches such as [36, 21, 39, 29] give the agent an additional reward called intrinsic reward that encourages the agent to perform actions that reduce the uncertainty in the agent’s prediction ability. In this direction, VIME (Variational information Maximizing Exploration) [21] was able to outperform (faster convergence and better solutions) many state-of-the-art RL algorithms such as TRPO [40] and ERWR [24] that use more basic exploration strategies.

In developmental and evolutionary methods, approaches such as novelty search [25], quality-diversity [8, 32] and curiosity-driven learning [16] aim at uniformly covering an user-defined space called “outcome space”, “goal space”, or “behavior space”. Approaches based on novelty-search encourage the agent to actively search for policies that produce novel outcomes. In this direction, [31] has shown that by framing novelty and quality improvement of the policies as a multi-objective Pareto optimization problem, it is possible to evolve high quality policies compared to single objective counterparts. In curiosity-driven learning, the agent tries to accomplish goals progressively, based on increasing complexity. In this regime, IMGEP [15] tries to discover policies that produce diverse outcomes by stochastically sampling different goals. Taking inspiration from IMGEP, GEP-PG [6] fills the replay buffer of a DDPG algorithm [26] with the data from initial exploratory episodes and then switches to a normal DDPG run.

Unfortunately, all the above-mentioned approaches need thousands of trials to find a working policy. For example, VIME needs more than 10000 trials for the cart-pole swing-up task with sparse reward or HER [1] needs around 80000 trials to find a policy such that a manipulator pushes an object to a desired location. This makes them unsuitable for learning on real robotic systems. A few model-based approaches tried to address this problem by using an optimism-driven exploration in a model predictive control scheme [30, 42]. Nevertheless, they are either too computationally intensive for online control [30] or require prior knowledge (e.g., dynamics equations) of the system [42].

So far, to the best of our knowledge, no RL algorithm can efficiently (*i.e.*, in little interaction time) explore in *sparse-reward scenarios without the need of prior knowledge about the system morphology and its dynamics*.

3.2 Model-based Policy Search

The most data-efficient algorithms for RL in robotics are model-based policy search algorithms (MB-PS), which alternate between learning a dynamical model of the robot and optimizing a policy to maximize the expected return given the model and its uncertainties [5, 11, 12]. Among the few proposed approaches, PILCO [11] and Black-DROPS [5] exploit a probabilistic dynamical model to achieve high data-efficiency with minimal assumptions about the underlying system. For example, both algorithms require around 20 seconds of interaction time to solve the cart-pole swing-up task (at least 1-2 order of magnitude less than model-free policy search algorithms [12]).

PILCO uses gradient-based optimization by analytically computing the gradients of the reward with respect to the policy parameters and the probabilistic model (by using moment matching). Black-DROPS exploits the parallelization and the *implicit averaging* properties of population-, rank-based optimizers (like CMA-ES [19]) to efficiently perform sampling-based evaluation of the trajectories in the model while accounting for the modelling uncertainties. Both algorithms are conceptually close but Black-DROPS is computationally faster when several cores are available and does not impose any constraint on the type of the policy and reward functions [5].

Nevertheless, PILCO and Black-DROPS do not have an explicit exploration strategy besides a few random initial policies; they essentially *exploit* whatever knowledge is captured in the model so far. As a result, when the reward function is sparse or deceptive [17], these algorithms struggle to find working policies or converge to sub-optimal ones: if high-reward states are associated with one or more *rare state transitions* that are not observed *by chance* in previous episodes, then these algorithms will fail to find the optimal policy because the rare transitions are not part of the model.

The main idea of the present paper is to combine the data-efficiency of MB-PS approaches with the creativity offered by novelty-based approaches to propose *an algorithm that requires as little interaction time as possible even in scenarios where the reward is sparse or deceptive*.

4 Approach

4.1 Learning system dynamics and reward model

Gaussian Processes (GPs) are widely used for learning system dynamics because of their generalization ability and their ability to compute uncertainty along with the predictions [12, 5, 11]. A GP is an extension of multivariate Gaussian distribution to an infinite-dimension stochastic process for which any finite combination of dimensions will be a Gaussian distribution [37].

We use tuples of states \mathbf{x}_t and actions \mathbf{u}_t (*i.e.*, $\tilde{\mathbf{x}}_t = (\mathbf{x}_t, \mathbf{u}_t) \in \mathbb{R}^{E+F}$) as training inputs for learning the system dynamics. The difference between the current and the next state vector, $\Delta_{\mathbf{x}_t} = \mathbf{x}_{t+1} - \mathbf{x}_t \in \mathbb{R}^E$, is used as corresponding training targets. Then, E independent GPs are used to model each dimension of the difference vector $\Delta_{\mathbf{x}_t}$. Let $D_{1:t} = \{f(\tilde{\mathbf{x}}_1), \dots, f(\tilde{\mathbf{x}}_t)\}$ is a set of observations, we can query the GP at a new input point $\tilde{\mathbf{x}}_*$: $p(\hat{f}(\tilde{\mathbf{x}}_*) | D_{1:t}, \tilde{\mathbf{x}}_*) = \mathcal{N}(\mu(\tilde{\mathbf{x}}_*), \sigma^2(\tilde{\mathbf{x}}_*))$. The mean and variance predictions of this GP are computed using a kernel vector $\mathbf{k} = k(D_{1:t}, \tilde{\mathbf{x}}_*)$, and a kernel matrix K , with entries $K_{ij} = k(\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j)$:

$$\mu(\tilde{\mathbf{x}}_*) = \mathbf{k}^T K^{-1} D_{1:t} \quad (3)$$

$$\sigma^2(\tilde{\mathbf{x}}_*) = k(\tilde{\mathbf{x}}_*, \tilde{\mathbf{x}}_*) - \mathbf{k}^T K^{-1} \mathbf{k} \quad (4)$$

In this paper, we use the exponential kernel with automatic relevance determination [37] and its hyper-parameters as well as the data noise hyper-parameter of the GP are found through Maximum Likelihood Estimation [37]. We use the Limbo C++11 library for GP regression [7]. In many cases, the immediate reward function might be unknown to the algorithm (*i.e.*, not directly computable from the state/action vector). In such situations, we use Random Forest regression [3] to learn the immediate reward function.

4.1.1 Learning system dynamics with sparse transitions

When the system dynamics have some sparse transitions that strongly affect the cumulative reward, learning the model in a naïve way (*i.e.*, using all the data points) can lead to suboptimal models and the policy search will struggle to find good policies. For example, in our sequential goal reaching task (Sec. 5.1) the state includes a Boolean switch that indicates whether the arm passed through the first way-point or not; because only one data point will have transition from *false* to *true* in any episode if the arm passes through the switch, learning a model with the full data (that mostly have no points with the switch transition) will most probably ignore these transition data points because they are likely to be (rightfully) considered as outliers. In other words, rare transitions have little impact on the likelihood or the fit of the dynamical model (as we additionally optimize the data noise hyper-parameter of the GP), whereas they are critical to leverage the dynamical model to learn a policy.

The intuition here is to have a balanced blend of ordinary trajectories and trajectories with rare transitions (leading to high reward) for model learning. Learning a dynamics model this way will retain information not only about the rare and high rewarding transitions but also about the regions where no reward was observed. As a result, this model can efficiently be used for exploration as well as exploitation. To do this, we maintain two fixed sized experience buffers, \mathbb{P} and \mathbb{H} , where we keep the trajectory data of the episodes for model learning. For each new trajectory executed on the robot, we insert a new observed trajectory into \mathbb{P} in a FIFO fashion if no reward is observed; otherwise, we insert it into \mathbb{H} . Note that, we keep the maximum buffer size of \mathbb{P} equal to the data size of \mathbb{H} to have a uniform blend of “high rewarding” and “no/low rewarding” trajectories for model learning.

4.2 Exploration-Exploitation Objectives

The intuition behind Multi-DEX is to combine the data-efficiency of model-based policy search with the ability of diversity in novelty-based search algorithms to encourage exploration in policy search in a data-efficient way. Thus, we frame exploration as an additional objective that promotes policies that will most likely produce novel state trajectories in the real system. To keep the system within the more certain regions of the dynamics model so that prediction error can be avoided as much as possible, we set an additional objective to reduce the prediction variance in the trajectory. The three objectives are typically antagonistic, as the uncertainty is likely to be higher for more novel trajectories, and more novel trajectories are likely to not have an cumulative reward as high as those with the highest cumulative rewards (*Suppl. Mat.* Sec 3).

Cumulative Reward: To compute the cumulative reward for a policy, we propagate through the learned GP model (with the mean prediction of the GPs, f_μ) of the system for T time steps using the policy, and observe the immediate reward in each step. Then, we aggregate all the immediate rewards to compute the cumulative reward (given the dynamical model) for the episode:

$$\hat{J}_r(\theta) = \sum_{t=1}^T r(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \mathbf{x}_{t-1} + f_\mu(\mathbf{x}_{t-1}, \mathbf{u}_{t-1})) \quad (5)$$

where $\mathbf{u}_{t-1} = \pi(\mathbf{x}_{t-1}|\theta)$

As it is evaluated in the model, optimizing this objective means exploiting the learned dynamics to find a policy that is likely to improve the cumulative reward.

Novelty: We compute the novelty score of a policy by comparing its expected state trajectory with the expected state trajectories of already executed policies (*Suppl. Mat.* Sec 5). To keep the computation time tractable, we sub-sample them into s_n equally spaced time-steps. We concatenate these state samples into one vector that we call *state trajectory vector*, β which represents the “expected

state trajectory” of a policy. These vectors are archived in a set \mathbb{B} of fixed size b_n and we re-compute them every time the GP model is updated. When \mathbb{B} is full, we drop the least novel policy. We define the novelty score for any policy π_θ as the minimum Euclidean distance of β_θ to the state trajectory vectors in \mathbb{B} :

$$\hat{J}_n(\theta) = \min(\|\beta_\theta - \beta\|^2)_{\forall \beta \in \mathbb{B}} \quad (6)$$

This objective is promoting exploration policies that are likely to lead to state trajectories that are different from those already observed.

Cumulative model-variance: We define the cumulative model-variance for a policy π_θ as the negative mean of the step-by-step model prediction variances:

$$\hat{J}_{\sigma^2}(\theta) = -\frac{1}{T} \sum_{t=1}^T (\|\sigma_{x_t}\|^2) \quad (7)$$

where x_t is given by applying the policy π_θ on the model

This objective tries to keep the system close to regions where the model uncertainty is low and thereby avoids prediction error on the real system.

4.3 Multi-Objective Optimization

Most algorithms aggregate objectives using a weighted sum. However, this means that the weights have to be chosen beforehand, which requires a time-consuming trial-and-error process; moreover, some optimal trade-offs cannot be found using a weighted sum [9]. Here, we rely on the multi-objective optimization literature [9] and the Pareto-optimality concept:

Definition 1 A solution x_1 is said to dominate another solution x_2 , if:

1. the solution x_1 is not worse than x_2 with respect to all objectives,
2. the solution x_1 is strictly better than x_2 at least one objective.

Instead of searching for a single, optimal policy, a multi-objective optimization algorithm searches for the set of non-dominated policies

in the search space, called the Pareto set: these policies are Pareto-optimal trade-offs between the objectives. For instance, a novel policy π_1 with a low cumulative return is as interesting as less novel policy π_2 but with a better cumulative return; however, a policy π_3 which is both less novel than π_1 and lower-performing is not interesting and is not Pareto-optimal.

While the set of non-dominated solutions is a very small subset of the search space, there are still many policies to choose from at the end of the optimization process. Multi-DEX executes on the robot the policy with the maximum reward from the Pareto set with a probability $1 - \varepsilon$ (exploitation) and choose a random policy from the Pareto set with a probability ε (exploration). As a consequence, if two policies have the same cumulative return, then Multi-DEX will always choose the most novel and less uncertain one; and if two policies are equally novel, it will always choose the one with the best cumulative return and uncertainty; and among the Pareto-optimal policies, it will select either the one with the best cumulative return or a random policy from the Pareto set. To find the Pareto set according to the dynamical model, we use the NSGA-II algorithm [10] (using the Spheres2 [33] C++ library), which is a widely used, state-of-the-art algorithm (see *Suppl. Mat.* Sec 1 for details about NSGA-II).

Algorithm 1 Multi-DEX

```

1: procedure MULTI-DEX
2:   Define policy  $\pi : \mathbf{x} \times \theta \rightarrow \mathbf{u}$ 
3:   Define  $\varepsilon \in [0, 1]$   $\triangleright$  User defined exploration parameter
4:    $D = \emptyset$ 
5:   for  $i = 1 \rightarrow N_R$  do  $\triangleright N_R$  random episodes
6:      $\theta_{rand} = \text{random\_policy}()$ 
7:     Execute  $\pi_{\theta_{rand}}$  on the real system
8:     Collect  $(\mathbf{x}_t, \mathbf{u}_t) \rightarrow \mathbf{x}_{t+1}$  and update  $\mathbb{P}, \mathbb{H}$  and  $D$ 
9:   end for
10:  while  $\text{task} \neq \text{solved}$  do  $\triangleright$  Learning episodes
11:    Train dynamics model with GPs using collected data  $D$ 
12:    Update  $\mathbb{B}$  using the trained GPs
13:    Get Pareto set of policies  $\pi_\theta$  using NSGA-II according to
         $\hat{J}_r, \hat{J}_n, \hat{J}_{\sigma^2}$ 
14:     $\pi_{\theta^*} =$  policy in Pareto front that has maximum  $\hat{J}_r$ 
15:    With probability  $\varepsilon$ , replace  $\pi_{\theta^*}$  by randomly selecting
        from the Pareto front
16:    Execute  $\pi_{\theta^*}$  on the real system
17:    Collect  $(\mathbf{x}_t, \mathbf{u}_t) \rightarrow \mathbf{x}_{t+1}$  and update  $\mathbb{P}, \mathbb{H}$  and  $D$ 
18:  end while
19: end procedure

```




Figure 2: (a) Setup of the sequential goal reaching task; the goal of the 2-DOF simulated arm is to reach with its end-effector the green goal while first passing through the blue region. (b) Best reward found per trial (20 replicates). The lines are median values and the shaded regions the 25th and 75th percentiles. Multi-DEX significantly outperforms all the other approaches and finds working policies in about only 7 minutes of interaction (around 100 episodes/trials).

Overall, Multi-DEX starts with random policy and executes it on the robot (Algo. 1). Then it enters an iterative loop, where Multi-DEX (1) learns a model of the dynamics of the system and a model of the reward function (if needed) from the data collected from previous episodes, (2) finds a set of non-dominated policies, using NSGA-II, that maximizes the cumulative reward (Eq. 5), novelty (Eq. 6) and minimizes the model prediction variance (Eq. 7), and (3) selects the policy that corresponds to maximum reward (with a probability ε , it replaces it with a randomly selected policy from the Pareto front) and executes it on the robot. This iterative process continues until the task is solved (Algo. 1).

5 Experiments

We compare to several state-of-the-art approaches in a sequential goal reaching task and in a drawer opening task: (1) Black-DROPS [5], a model-based policy search algorithm; (2) TRPO [40], a model-free policy gradient approach; (3) TRPO with the VIME exploration strategy; (4) CMA-ES [19], a black-box optimizer effective for direct policy search, and (5) GEP-PG, a curiosity-driven model-free approach. In all the tasks, we give a budget of 2500 trials to all the model-free approaches and 250 trials to the model-based ones (see *Suppl. Mat.* Sec 6 and Sec 7 for further experimental details and additional experiments). The Multi-DEX source code and the supplementary video are available online.¹²

5.1 Sequential goal reaching with a robotic arm

This simulated task consists of a robotic arm with 2 degrees of freedom (Fig. 2a). The goal of the arm is to reach the green goal with its end-effector while first passing through the blue region. The length of each episode is 4 seconds with a control frequency of 10Hz. All the algorithms use feed-forward neural network policies. In more detail:

- **State:** $\mathbf{x}_{\text{seq}} = [\theta_0, \theta_1, \gamma] \in \mathbb{R}^3$, where θ_0, θ_1 are joint angles and γ is a boolean value which is set to 1 for rest of the episode if the end-effector has touched the blue region.
- **Actions:** $\mathbf{u}_{\text{seq}} = [v_0, v_1] \in \mathbb{R}^2$, $-2 \leq v_0, v_1 \leq 2 \text{ rad/s}$ are joint velocity commands.
- **Reward:** The reward function depends only on the current state and is unknown to the algorithm as it is not a *direct* function of observable state:

$$r(\mathbf{x}_{\text{seq}}) = \begin{cases} 0, & \text{if } \gamma \neq 1 \text{ or } \Delta \geq 0.1 \\ \exp(-\alpha \Delta^2), & \text{otherwise.} \end{cases} \quad (8)$$

where Δ is the Euclidean distance between green goal and the end-effector of the arm and α is a user-defined constant ($\frac{0.5}{0.2 \times 0.2}$ in our experiments). After every episode, Multi-DEX learns the reward function from the collected data using Random Forest regression [3].

¹Multi-DEX source code: https://github.com/resibots/kaushik_2018_multi-dex

²Video: <https://youtu.be/XOBWq7mkYho>




Figure 3: (a) Setup of the drawer opening task; the goal of the 2-DOF arm is to open the drawer and go back to the up-right position. (b) Best reward found per trial (20 replicates). The lines are median values and the shaded regions the 25th and 75th percentiles. Multi-DEX outperforms all the other approaches and finds working policies in about only 14 minutes of interaction (200 trials).

Multi-DEX is able to find high-performing policies in less than 100 trials on the target system (around 7 minutes of interaction), whereas none of the other approaches are able to achieve similar rewards even after 2500 trials (Fig. 2b — 20 replicates). As this task has a highly sparse reward space, CMA-ES, Black-DROPS and TRPO show almost zero median reward.

5.2 Drawer opening task with a robotic arm

The goal of this simulated task is to open a drawer with a 2-DOF robotic arm and go back to the up-right position (Fig. 3a). Similarly to the previous task, the length of each episode is 4 seconds (10Hz control) and all the algorithms use neural network policies. In more detail:

- **State:** $\mathbf{x}_{\text{drawer}} = [\theta_0, \theta_1, \delta] \in \mathbb{R}^3$, where θ_i are joint angles and δ the drawer displacement.
- **Actions:** $\mathbf{u}_{\text{drawer}} = [v_0, v_1] \in \mathbb{R}^2$, $-1 \leq v_0, v_1 \leq 1 \text{ rad/s}$ are joint velocity commands.
- **Reward:** In this task, the reward depends only on the current state and is known to the algorithm as it is a function of the observable state. The total reward is given by:

$$r(\mathbf{x}_{\text{drawer}}) = r_{\text{return}}(\mathbf{x}_{\text{drawer}}) + \delta \quad (9)$$

$$r_{\text{return}}(\mathbf{x}_{\text{drawer}}) = \begin{cases} 0, & \text{if } \delta \leq 0.2 \\ \exp(-\theta_0^2 - \theta_1^2), & \text{otherwise.} \end{cases} \quad (10)$$

In this task, the reward space is moderately sparse compared to the previous one; however, it has a misleading reward space because of the composition of two different rewards. Any algorithm without efficient exploration will converge to the reward associated with the drawer displacement only. Multi-DEX is able to find policies that complete the task in just 200 trials (around 14 minutes of interaction), whereas all the other approaches are deceived by the reward space (Fig. 3b — 20 replicates). Black-DROPS and TRPO, without any directed exploration, fall in the sub-optimal solution (*i.e.*, just opening the drawer quickly). CMA-ES, TRPO-VIME and GEP-PG either fail to converge or need more than 2500 trials to reach the same quality of solutions as Multi-DEX.

6 Conclusion

It is well established that learning a dynamical model can make policy search highly-data efficient for control tasks [5, 12] when (1) the state-space is low-dimensional enough to learn a model, (2) the model is probabilistic so that the uncertainty of the predictions can be taken into account, and (3) it is possible to rely on a large computation time between episode (for the policy optimization stage). In this paper, we showed that the same model can be leveraged to improve exploration for tasks in which the reward is sparse or deceptive: compared to state-of-the-art approaches to exploration for policy search, which are all model-free, our model-based policy search algorithm requires at least an order of magnitude fewer episodes and reach higher rewards in all our experiments.

Acknowledgments

This work received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (GA no. 637972, project “ResiBots”), and the European Commission through the project H2020 AnDy (GA no. 731540).

References

- [1] Marcin Andrychowicz et al. Hindsight experience replay. In *Proc. of NIPS*, 2017.
- [2] Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In *Proc. of NIPS*, 2016.
- [3] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [4] Konstantinos Chatzilygeroudis and Jean-Baptiste Mouret. Using Parameterized Black-Box Priors to Scale Up Model-Based Policy Search for Robotics. In *Proc. of ICRA*, 2018.
- [5] Konstantinos Chatzilygeroudis, Roberto Rama, Rituraj Kaushik, Dorian Goepp, Vassilis Vassiliades, and Jean-Baptiste Mouret. Black-Box Data-efficient Policy Search for Robotics. In *Proc. of IROS*, 2017.
- [6] Cédric Colas, Olivier Sigaud, and Pierre-Yves Oudeyer. GEP-PG: Decoupling Exploration and Exploitation in Deep Reinforcement Learning Algorithms. In *Proc. of ICML*, 2018.
- [7] A. Cully, K. Chatzilygeroudis, F. Allocati, and J.-B. Mouret. Limbo: A Flexible High-performance Library for Gaussian Processes modeling and Data-Efficient Optimization. *The Journal of Open Source Software*, 3(26):545, 2018.
- [8] Antoine Cully and Yiannis Demiris. Quality and diversity optimization: A unifying modular framework. *IEEE Trans. on Evolutionary Computation*, 22(2):245–259, 2018.
- [9] Kalyanmoy Deb and Deb Kalyanmoy. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [10] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMT Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. on Evolutionary Computation*, 6(2):182–197, 2002.
- [11] Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Trans. Pattern Anal. Mach. Intell.*, 37(2):408–423, 2015.
- [12] Marc Peter Deisenroth, Gerhard Neumann, and Jan Peters. A survey on policy search for robotics. *Foundations and Trends in Robotics*, 2(1):1–142, 2013.
- [13] Stéphane Doncieux and Jean-Baptiste Mouret. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, 7(2):71–93, 2014.
- [14] Carlos Florensa, David Held, Markus Wulfmeier, and Pieter Abbeel. Reverse curriculum generation for reinforcement learning. In *Conference on Robot Learning*, 2017.
- [15] Sébastien Forestier, Yoan Mollard, and Pierre-Yves Oudeyer. Intrinsically motivated goal exploration processes with automatic curriculum learning. *arXiv preprint arXiv:1708.02190*, 2017.
- [16] Sébastien Forestier and Pierre-Yves Oudeyer. Curiosity-driven development of tool use precursors: a computational model. In *Proc. of COGSCI*, 2016.
- [17] David E Goldberg. Simple genetic algorithms and the minimal, deceptive problem. *Genetic algorithms and simulated annealing*, pages 74–88, 1987.
- [18] Jacqueline Gottlieb, Pierre-Yves Oudeyer, et al. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11):585–593, 2013.
- [19] Nikolaus Hansen. *The CMA Evolution Strategy: A Comparing Review*. Springer, 2006.
- [20] Nicolas Heess et al. Emergence of locomotion behaviours in rich environments. *arXiv preprint arXiv:1707.02286*, 2017.
- [21] Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Vime: Variational information maximizing exploration. In *Proc. of NIPS*, 2016.

- [22] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of artificial intelligence research*, 4:237–285, 1996.
- [23] Michael Kearns and Satinder Singh. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- [24] Jens Kober and Jan Peters. Policy search for motor primitives in robotics. *Machine Learning*, 84:171–203, 2011.
- [25] Joel Lehman and Kenneth O Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, 2011.
- [26] Timothy Paul Lillicrap et al. Continuous control with deep reinforcement learning. In *Proc. of ICLR*, 2016.
- [27] Manuel Lopes, Tobias Lang, Marc Toussaint, and Pierre-Yves Oudeyer. Exploration in model-based reinforcement learning by empirically estimating learning progress. In *Proc. of NIPS*, 2012.
- [28] Volodymyr Mnih et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- [29] Shakir Mohamed and Danilo Jimenez Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Proc. of NIPS*, 2015.
- [30] Teodor Mihai Moldovan, Sergey Levine, Michael I Jordan, and Pieter Abbeel. Optimism-driven exploration for nonlinear systems. In *Proc. of ICRA*, 2015.
- [31] Jean-Baptiste Mouret. Novelty-based Multiobjectivization. In *New Horizons in Evolutionary Robotics*, pages 139–154. Springer Berlin/Heidelberg, 2011.
- [32] Jean-Baptiste Mouret and Jeff Clune. Illuminating search spaces by mapping elites. *arxiv:1504.04909*, 2015.
- [33] Jean-Baptiste Mouret and Stéphane Doncieux. Sferes v2: Evolving in the multi-core world. In *Proc. of CEC*, 2010.
- [34] Pierre-Yves Oudeyer, Frdric Kaplan, and Verena V Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE Trans. on Evolutionary Computation*, 11(2):265–286, 2007.
- [35] Pierre-Yves Oudeyer, Frédéric Kaplan, Verena V Hafner, and Andrew Whyte. The playground experiment: Task-independent development of a curious robot. In *Proc. of the AAAI Spring Symposium on Developmental Robotics*, pages 42–47. Stanford, California, 2005.
- [36] Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proc. of ICML*, 2017.
- [37] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.
- [38] Martin Riedmiller et al. Learning by Playing-Solving Sparse Reward Tasks from Scratch. *arXiv preprint arXiv:1802.10567*, 2018.
- [39] Jürgen Schmidhuber. A Possibility for Implementing Curiosity and Boredom in Model-building Neural Controllers. In *Proc. of SAB*, 1990.
- [40] John Schulman, Sergey Levine, Philipp Moritz, Michael I Jordan, and Pieter Abbeel. Trust region policy optimization. In *Proc. of ICML*, 2015.
- [41] David Silver et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- [42] Chris Xie, Sachin Patil, Teodor Moldovan, Sergey Levine, and Pieter Abbeel. Model-based reinforcement learning with parametrized physical models and optimism-driven exploration. In *Proc. of ICRA*, 2016.

Supplementary Material

1 Non-dominant Sorting Genetic Algorithm - II (NSGA II)

We use NSGA-II for multi-objective policy optimization.

NSGA-II starts with random initialization of parent population P_0 of size N . P_0 is sorted based on non-domination and a fitness score (or rank) is assigned to each of the individuals. For example, solutions which are not dominated by any solutions in all the objectives are assigned a fitness or rank 0 (i.e 0^{th} level of domination). Similarly, solutions which are dominated in all the objectives by rank 0 solutions are assigned a fitness value or rank of 1 (i.e 1^{st} level of domination) and so on. Then $N/2$ individuals are selected with binary tournament selection based on the fitness/rank scores (lower the better) and then applied binary crossover and mutation operators to obtain a offspring population Q_0 of size N . For each generation k , the algorithm performs the following steps:

1. A combined population $R_k = P_k \cup Q_k$ of size $2N$ is formed.
2. R_k is sorted in ascending order based on the level of nondomination into sets F_0, F_1, F_2 and so on. Then solutions are selected starting from F_0 to the last possible level (say F_l) to form a new parent population P_{k+1} of size N . If last set F_l cannot be accommodated completely in P_{k+1} , then the solutions from F_l is selected based on crowded comparison operator \prec_n [6].
3. Selection, crossover (sbx) and polynomial mutation operators are applied on the new parent population P_{k+1} to produce the new offspring population Q_{k+1} of size N .
4. The algorithm continues until a specified number of generation is reached. Then it returns the set of non-dominated solutions (approximation of the true Pareto-front).

2 Why Pareto optimality should be preferred over weighted sum objectives?

When optimizing several objectives with a weighted sum, the weights determine the trade-off between the objectives (e.g., more exploration vs more exploitation) and the optimum is a particular Pareto-optimal solution. The important question for multi-objective optimisation is therefore: can any Pareto-optimal solution (i.e., any optimal trade-off) be found by selecting the right combination of weights? The answer depends on the shape of the Pareto front (the set of Pareto-optimal solutions). If it is convex, then all the Pareto-optimal solutions can be found by varying the weights; however, if it is not, then some Pareto-optimal trade-offs cannot be found using a weighted sum [5]. In the learning problems tackled in this paper, there is no formal guarantee that the Pareto front is convex, therefore it is theoretically better to use another approach to find a suitable Pareto-optimal trade-off. In other words, the best trade-off between reward, novelty, and model variance might not be expressible with a combination of weights. In addition, setting weights requires a good normalization method (since the trade-off will be affected by the scale of each objective), whereas a Pareto-based approach will usually be independent of the scale.

For completeness, we compared experimentally our Pareto-based approach with both fixed and stochastic weighted aggregation of the objectives to a single objective. For fixed weighted aggregation we used weights 0.4, 0.3 and 0.3 for novelty, variance and cumulative reward respectively. On the other hand, for stochastic weighted aggregation we used 4 sets of weights: $\{0.6, 0.3, 0.1\}$, $\{0.1, 0.3, 0.6\}$, $\{0.1, 0.45, 0.45\}$ where each element of the sets corresponds to the weights for novelty, variance and cumulative reward respectively. We randomly select a set of these weights for each episode and use it for aggregation of the objectives. The results show that weighted aggregation doesn't give consistent results on both the experiments. Although a different set of weights could give different results, in the drawer opening task Pareto-based approach out performs both the variants of weighted aggregation approach with higher mean performance and smaller variance over the replicates. In the sequential goal reaching task the stochastic weighted aggregation shows similar performance to the Pareto-based approach. However, fixed weighted aggregation approach did not solve the task at all.




Figure 4: Best reward found vs number of trials plots show the comparison between Pareto based multi-objective approach and weighted aggregation approach. For stochastic weighted aggregation approach we defined 4 set of weights with different trade-offs among the objectives. For each episode we randomly select one set of weights and use them for aggregation of the objectives into a single objective. On the other hand, for the fixed weighted aggregation approach we defined a single set of weights and used them for the aggregation of the objectives.

3 Interaction of the objectives among themselves

The first objective in multi-DEX is to find trajectories (i.e., corresponding policies) that have high reward. The second objective tries to ensure that for a particular value of reward we get as novel trajectory as possible. Now these two objectives might be antagonistic and for most of cases both cannot be improved without compromising the other. Thus we want a Pareto-optimal set of policies that gives the trade-off between these two objective, starting from zero reward to maximum reward possible on the model. However, as we are using GP mean prediction only to find trajectories, it is possible that we find policies that drive the system though highly uncertain regions on the model. If this is the case then those policies will have completely different behavior on the system from the one predicted using the model. Thus we have the third objective that tries to keep the trajectories as close as possible to the more certain regions of the model. Again this new objective might be antagonistic to previous two objectives in some cases. Therefore, we want a Pareto-optimal trade-off among all the objectives. It is to be noted here that in Multi-DEX novelty score is highly dependent upon the order in which the states are visited and not that much on the individual states. So it is possible to have a high novelty score by going through similar states that were observed before (i.e., points with less uncertainty) but in a completely different order. Once the Pareto-optimal trade-offs are identified, all of them are potentially good candidates. This is why we alternatively take a random policy from these good candidate (minimize bias in exploration) or the one with the best predicted reward (exploitation).

4 How scalable is the multi-objective optimization method (NSGA-II) with regards to the number of parameters of the policy?

Several recent papers show that evolutionary algorithms, such as NSGA-II, can be competitive to optimize deep neural network policies in reinforcement learning scenarios [11, 14]. Thus, we believe that NSGA-II should be able to scale up to higher-dimensional policy optimization, at the obvious expense of more optimization time. At any rate, it should be emphasized that what matters the most in this work is the interaction time with the system/robot, and not the policy optimization time that happens on the learned dynamical model.

5 Why Expected Trajectory for Novelty Score?

In Multi-DEX we have defined the novelty score of any policy as the minimum Euclidean distance between the expected state trajectory of that policy to all the *expected state trajectories* of the executed policies on the system. There is a strong reason for keeping expected state trajectory vectors and not the observed state trajectory vectors on the system to compute the novelty score of a policy. Because of the imperfection in the model, the model might predict unachievable state trajectories to maximize novelty, which might give a completely different state trajectory on the actual system than

the one expected by the model. Now, if we compute the novelty score for a policy by comparing with the observed state trajectory vector on the system, then the already tried policies predicting unachievable behaviors will have very high novelty scores. As a result, optimizer might return these trajectories again and again for every iteration of the algorithm.

A solution to this problem is to keep the expected state trajectory vectors instead to compute the novelty scores. As a result, the unachievable state trajectories that were already tried on the system will have low novelty scores; and hence, policies with these state trajectories will be rejected because of low novelty score by the optimizer.

6 Details on the Experimental Setup

6.1 Simulator and source code

For all the experiments we have used DART physics simulation library [7] to implement the scenarios. For TRPO, VIME and GEP-PG we used python version of the DART library (pydart2³); for CMA-ES, Black-DROPS and Multi-DEX we used the C++ version of the library for simulation.

For CMA-ES, we used Limbo’s wrapper of libcmaes library⁴. For all other algorithms (GEP-PG, TRPO, TRPO-VIME and Black-DROPS), we used the code provided by the authors and integrated our own scenarios. The source code of Multi-DEX can be found here: https://github.com/resibots/kaushik_2018_multi-dex

6.2 General and Exploration Parameters

For GEP-PG, we used 500 exploratory episodes for the “Drawer Opening” and “Sequential Goal Reaching” tasks; and 100 exploratory episodes for “Pendulum Swing-up” task. For all the experiments with GEP-PG we used 10 initial bootstrap episodes.

For TRPO-VIME we used $\eta = 0.001$ for the “Drawer Opening” task. For the “Sequential Goal Reaching” and “Pendulum Swing-up” tasks, we used $\eta = 0.001$.

For CMA-ES, we used the active-IPOP version with elitism enabled [1, 9]. We used 3 restarts and 2500 maximum function evaluations.

For Multi-DEX, we used $\varepsilon = 0.3$ for the “Sequential Goal Reaching” and “Pendulum Swing-up” tasks. We used $\varepsilon = 0.4$ for the “Drawer Opening” task. In all the experiments, to compute the novelty score, we sub-sample the expected state trajectories at 10 equally spaced points in time. In all the cases, we bootstrapped the algorithm with 5 initial random trials.

6.3 Policy and Parameter Bounds

For all the experiments, we used neural network policies. For Black-DROPS, GEP-PG and Multi-DEX in the “Sequential Goal Reaching” and “Drawer Opening” tasks, we used one hidden layer with 5 neurons and in the “Pendulum Swing-up” task we used one hidden layer with 10 neurons. For TRPO and TRPO-VIME we used two hidden layers with 10 neurons each (the implementations that we used required at least 2 hidden layers for the neural network policies).

The policy parameter bounds used in Multi-DEX are $[-1, 1]$ for “Sequential Goal Reaching”, $[-3, 3]$ for “Drawer Opening” tasks and $[-5, 5]$ for “Pendulum Swing-up” task.

6.4 NSGA-II Parameters

We used the Sferes2 C++ library’s [10] implementation of NSGA-II algorithm in Multi-DEX. The parameters used for NSGA-II are the following:

- crossover rate = 0.5
- mutation rate = 0.1

³<https://github.com/sehoonha/pydart2>

⁴<https://github.com/beniz/libcmaes>




Figure 5: Best reward found vs number of trials plot for Pendulum Swing-up Task. The plot clearly outperforms all the competing approaches and achieves very high reward (balancing the pendulum in upright position) in just 100 trials (approx 6.6 minutes of total interaction)

- $\eta_m = 15.0$
- $\eta_c = 10.0$
- mutation type: Polynomial [5]
- cross-over type: Simulated Binary Crossover [4]

We run NSGA-II for 600 generations with population sizes of 200 for the “Drawer Opening” and “Sequential Goal Reaching” tasks and 300 for the “Pendulum Swing-up” task.

For the first episode, the initial population of policies is randomly initialized. From the 2nd episode onwards, we insert some policies from the Pareto-front into the initial population. The maximum number of these inserted policies is not more than 30% of the total population.

6.5 Gaussian Process Model learning

In this work, we used the Limbo C++ library’s [3] implementation of Gaussian process regression for probabilistic dynamics model learning. We used exponential kernel with automatic relevance determination [12]. We optimize the hyper-parameters of the kernel via Maximum Likelihood Estimation using the Rprop optimizer [13, 2].

We maintain two experience buffers (of fixed size) \mathbb{P} and \mathbb{H} to keep the state-action trajectories for model learning (see the main text).

7 Additional Experiments

7.1 Deceptive pendulum swing-up task

This simulated task consists of a pendulum powered by an underpowered torque-controlled actuator. The goal in this task is to swing the pendulum to the upright position applying torques as small as possible (*i.e.*, using minimum power) to the actuator and hold it in that position. The learning algorithm gets a constant positive reward of +10 every time-step the pendulum is in upright position (within some region). It also gets a negative reward proportional to the square of torque for every time step. Because of this two rewards, the total reward function possesses a deceptive landscape and algorithms without efficient exploration will converge to a reward of zero which is achieved when no torque is applied to the pendulum. This type of “*Gradient Cliff*” reward landscape was first proposed by [8]. The control frequency in this task is $10Hz$ and every episode is 4 seconds long. In more detail:




Figure 6: Best reward found vs number of trials plot for drawer opening task with 4-DOF simulated robotic arm. Plot shows that both the median over the replicates with Multi-DEX is able solve the task in 1000 trials contrary to CMA-ES which is able to reach only half the reward achieved by Multi-DEX)

- **State:** $\mathbf{x}_{\text{pend_sys}} = [\theta, \dot{\theta}] \in \mathbb{R}^2$, where θ is the joint angle and $\dot{\theta}$ is the joint angular velocity. The initial state of the system is $x_0 = [0, 0]$
- **Actions:** $\mathbf{u}_{\text{pend_sys}} = [\tau] \in \mathbb{R}$, $-2.0 \leq \tau \leq 2.0$, where τ joint torque command for the arm.
- **Reward:** In this task, the reward is known to the algorithm as it is a function of the observable states and applied action to the system. The total reward is given by:

$$r(\mathbf{x}_{\text{pend_sys}}, \tau) = r_1(\mathbf{x}_{\text{pend_sys}}, \tau) + r_2(\mathbf{x}_{\text{pend_sys}}, \tau) \quad (11)$$

$$r_1(\mathbf{x}_{\text{pend_sys}}, \tau) = \begin{cases} 0, & \text{if } \pi - \theta < \pi/60 \\ +10, & \text{otherwise.} \end{cases} \quad (12)$$

$$r_2(\mathbf{x}_{\text{pend_sys}}, \tau) = -0.001 * \tau^2 \quad (13)$$

The results show that Multi-DEX quickly reaches to a very high positive reward with very small variance within the budget of 100 trials (Fig. 5). On the contrary, GEP-PG and TRPO-VIME could not reach to the same quality of solutions even after 2500 trials on the system. As TRPO, BlackDROPS and CMA-ES do not have any directed exploration, they could not improve the reward and stay close to zero by minimizing the applied torque only.

7.2 Drawer opening task with 4-DOF arm

To evaluate Multi-DEX on more difficult sparse reward scenarios, we ran it on the drawer opening task mentioned in the paper with a 4-dof simulated robotic arm. We compared the result with CMA-ES and found that Multi-DEX takes around 1000 trials on the system to solve it (median over 20 replicates). On the contrary, CMA-ES was able to reach only half the reward achieved by Multi-DEX in 1000 trials (Fig. 6).

7.3 Multi-DEX on non-sparse reward task

We tested Multi-DEX on a single goal reaching task with a 5-dof simulated arm. In this task the state-space is 5 dimensional (5 joint angles of the arm) and action space is also 5 dimensional (torque inputs to the joints). The length of each episode is 4 seconds (10Hz control) and we used neural network policy with one hidden layer. The hidden layer had 10 neurons and thus the entire policy space is 115 dimensional. We chose a reward that is continuous in the state-space as given below:




Figure 7: Best reward found vs number of trials plot for single goal reaching task with a 5-DOF simulated arm. Plot shows that both the variants of Multi-DEX are competitive enough with Black-DROPS in this non-sparse reward scenario

$$r(\mathbf{x}_{\text{arm}}) = \exp\left(-\frac{0.5}{0.2 \times 0.2} \|\mathbf{x}_{\text{goal}} - \mathbf{x}_{\text{arm}}\|^2\right) \quad (14)$$

where \mathbf{x}_{goal} and \mathbf{x}_{arm} are the 3D coordinate vectors of the the goal and the end-effector. We compared two variants of Multi-DEX: with $\epsilon = 0$ and $\epsilon = 0.4$. We compared the result with standard Black-DROPS to see how efficient Multi-DEX is on non-sparse reward setup. This preliminary results show that both the variants of Multi-DEX is competitive enough with Black-DROPS (Fig. 7). Point to be noted here that we plot here the median, 25 percentile and 75 percentile of "best reward so far" for the replicates.

References

- [1] Anne Auger and Nikolaus Hansen. A restart CMA evolution strategy with increasing population size. In *Congress on Evolutionary Computation*, 2005.
- [2] Manuel Blum and Martin A Riedmiller. Optimization of Gaussian process hyperparameters using Rprop. In *Proc. of ESANN*, 2013.
- [3] A. Cully, K. Chatzilygeroudis, F. Allocati, and J.-B. Mouret. Limbo: A Flexible High-performance Library for Gaussian Processes modeling and Data-Efficient Optimization. *The Journal of Open Source Software*, 3(26):545, 2018.
- [4] Kalyanmoy Deb and Hans-Georg Beyer. *Self-adaptive genetic algorithms with simulated binary crossover*. Secretary of the SFB 531, 1999.
- [5] Kalyanmoy Deb and Deb Kalyanmoy. *Multi-Objective Optimization Using Evolutionary Algorithms*. John Wiley & Sons, Inc., New York, NY, USA, 2001.
- [6] Kalyanmoy Deb, Amrit Pratap, Sameer Agarwal, and TAMS Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. on Evolutionary Computation*, 6(2):182–197, 2002.
- [7] Jeongseok Lee, Michael X Grey, Sehoon Ha, Tobias Kunz, Sumit Jain, Yuting Ye, Siddhartha S Srinivasa, Mike Stilman, and C Karen Liu. DART: Dynamic animation and robotics toolkit. *The Journal of Open Source Software*, 2018.
- [8] Joel Lehman, Jay Chen, Jeff Clune, and Kenneth O Stanley. Es is more than just a traditional finite-difference approximator. *arXiv preprint arXiv:1712.06568*, 2017.
- [9] Ilya Loshchilov. CMA-ES with restarts for solving CEC 2013 benchmark problems. In *Congress on Evolutionary Computation*, 2013.

- [10] Jean-Baptiste Mouret and Stéphane Doncieux. Sferes v2: Evolvin' in the multi-core world. In *Proc. of CEC*, 2010.
- [11] Felipe Petroski Such, Vashisht Madhavan, Edoardo Conti, Joel Lehman, Kenneth O Stanley, and Jeff Clune. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *arXiv preprint arXiv:1712.06567*, 2017.
- [12] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. MIT Press, 2006.
- [13] Martin Riedmiller and Heinrich Braun. Rprop-a fast adaptive learning algorithm. In *Proc. of ISCIS VII*, 1992.
- [14] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.