



Policy search in continuous action domains: An overview

Olivier Sigaud, Freek Stulp

► To cite this version:

Olivier Sigaud, Freek Stulp. Policy search in continuous action domains: An overview. Neural Networks, 2019, 113, pp.28-40. 10.1016/j.neunet.2019.01.011 . hal-02182466

HAL Id: hal-02182466

<https://hal.sorbonne-universite.fr/hal-02182466>

Submitted on 12 Jul 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Policy Search in Continuous Action Domains: an Overview

Olivier Sigaud

INRIA Bordeaux Sud-Ouest, *équipe FLOWERS*
Sorbonne Université, CNRS UMR 7222,
Institut des Systèmes Intelligents et de Robotique, F-75005 Paris, France
olivier.sigaud@isir.upmc.fr +33 (0) 1 44 27 88 53

Freek Stulp

German Aerospace Center (DLR), Institute of Robotics and Mechatronics, Wessling, Germany
freek.stulp@dlr.de

Abstract

Continuous action policy search is currently the focus of intensive research, driven both by the recent success of deep reinforcement learning algorithms and the emergence of competitors based on evolutionary algorithms. In this paper, we present a broad survey of policy search methods, providing a unified perspective on very different approaches, including also Bayesian Optimization and directed exploration methods. The main message of this overview is in the relationship between the families of methods, but we also outline some factors underlying sample efficiency properties of the various approaches.

Keywords: policy search, sample efficiency, deep reinforcement learning, deep neuroevolution

1. Introduction

Autonomous systems are systems which know what to do in their domain without external intervention. Generally, their behavior is specified through a *policy*. The policy of a robot, for instance, is defined through a controller which determines actions to take or signals to send to the actuators for any state of the robot in its environment.

Robot policies are often designed by hand, but this manual design is only viable for systems acting in well-structured environments and for well-specified tasks. When those conditions are not met, a more appealing alternative is to let the system find its own policy by exploring various behaviors and exploiting those that perform well with respect to some predefined *utility* function. This approach is called *policy search*, a particular case of *reinforcement learning* (RL) (Sutton and Barto, 1998) where actions are vectors from a continuous space. More precisely, the goal of policy search is to optimize a policy where the function relating behaviors to their utility is *black-box*, i.e. no analytical model or gradient of the utility function is available. In

practice, a policy search algorithm runs the system with some policies to generate *rollouts* made of several state and action *steps* and gets the utility as a return (see Figure 1). These utilities are then used to improve the policy, and the process is repeated until some satisfactory set of behaviors is found. In general, policies are represented with a parametric function, and policy search explores the space of *policy parameters*. For doing so, rollout and utility data are processed by policy improvement algorithms as a set of *samples*.




Figure 1: Visualization of one episode, the information contained in a rollout, and the definition of the episode utility (which is also known as the *episode return* when the utility is a reward).

In the context of robotics, *sample efficiency* is a key concern. There are three aspects to sample ef-

efficiency: (1) data efficiency, i.e. extracting more information from available data (definition taken from (Deisenroth and Rasmussen, 2011)), (2) sample choice, i.e. obtaining data in which more information is available and (3) sample reuse, i.e. improving a policy several times by using the same samples more than once through *experience replay*.

In this paper, we provide a broad overview of policy search algorithms under the perspective of these three aspects.

1.1. Scope and Contributions

Three surveys about policy search for robotics have been published in recent years (Deisenroth et al., 2013; Stulp and Sigaud, 2013; Kober et al., 2013). With respect to these previous surveys, we cover a broader range of policy search algorithms, including optimization without a utility model, Bayesian optimization (BO), directed exploration methods, and deep RL. The counterpart of this breadth is that we do not give a detailed account of the corresponding algorithms nor their mathematical derivation. To compensate for this lack of details, we refer the reader to (Deisenroth et al., 2013) for the mathematical derivation and description of most algorithms before 2013, and we provide carefully chosen references as needed when describing more recent algorithms.

Furthermore, we focus on the case where the system is learning to solve a *single task*. That is, we do not cover the broader domain of *lifelong*, *continual* or *open-ended* learning, where a robot must learn how to perform various tasks over a potentially infinite horizon (Thrun and Mitchell, 1995). Additionally, though a subset of policy search methods are based on RL, we do not cover recent work on RL with discrete actions such as DQN and its successors (Mnih et al., 2015; Hessel et al., 2017). Finally, we restrict ourselves to the case where samples are the unique source of information for improving the policy. That is, we do not consider the interactive context where a human user can provide external guidance (Najar et al., 2016), either through feedback, shaping or demonstration (Argall et al., 2009).

1.2. Perspective and structure of the survey

The main message of this paper is as follows. In optimization, when the utility function to be optimized is known and convex, efficient convex methods can be applied (Gill et al., 1981). If the function is known but not convex, a local optimum can

be found using gradient descent, iteratively moving from the current point towards a local optimum by following the direction provided by the derivative of the function at this point. If the function is *black-box*, neither the function nor its analytic gradient are known. In policy gradient methods, policy parameters are only related to their utility *indirectly* through an intermediate set of observed behaviors. Given that policy search corresponds to this more difficult context, we consider five solutions:

1. searching for high utility policy parameters without building a utility model (Section 2),
2. learning a model of the utility function in the space of policy parameters and performing stochastic gradient descent (SGD) using this model (Section 3),
3. defining an arbitrary *outcome space* and using directed exploration of this outcome space for finding high utility policy parameters (Section 4),
4. doing the same as in Solution 2 in the state-action space (Section 5),
5. learning a model of the transition function of the system-environment interaction that predicts the next state given the current state and action, to generate samples without using the system, and then applying one of the above solutions based on the generated samples.

An important distinction in the policy search domain is whether the optimization method is episode-based or step-based (Deisenroth et al., 2013)¹. The first three solutions above are episode-based, the fourth is step-based and the fifth can be applied to all others.

Resulting from the above perspective, this survey is structured following the organization of methods depicted in Figure 2. The rest of the paper describes the different nodes in the trees and highlights some of their sample efficiency factors. A table giving a quick reference to the main paper for each algorithm is given in the end of each section.

In Sections 2 to 5, we present Solutions 1 to 4 above in more detail, showing how the corresponding methods are implemented in various policy search algorithms. We do not cover Solution 5 and refer readers to (Chatzilygeroudis et al., 2017) for a recent presentation of these *model-based policy*

¹This distinction exactly matches the phylogenetic RL versus ontogenetic RL distinction in (Togelius et al., 2009).

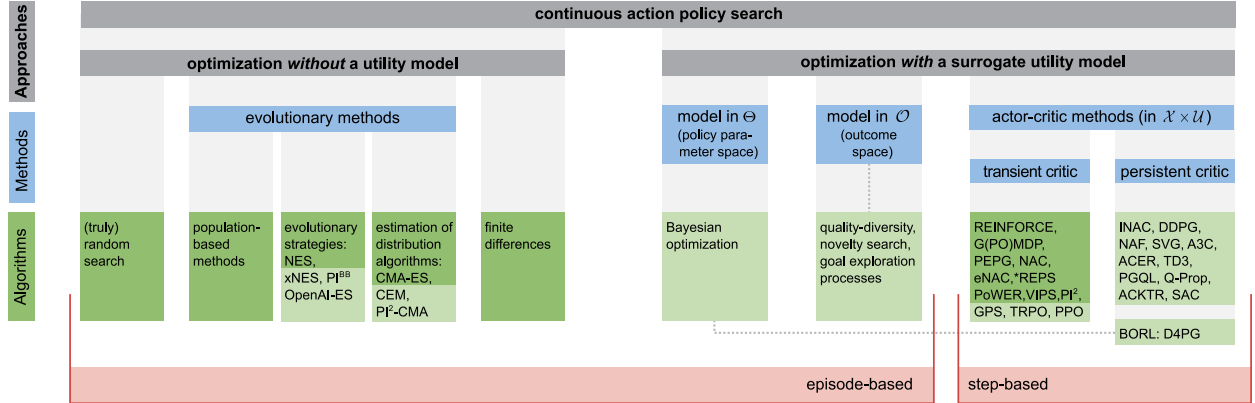


Figure 2: Simplified classification of the algorithms covered in the paper. Θ is the space of policy parameters (see Section 2), \mathcal{O} is an outcome space (see Section 4) and $\mathcal{X} \times \mathcal{U}$ is the state and action space (see Section 5). Algorithms not covered in (Deisenroth et al., 2013) have a lighter (green) background. References to the main paper for each of these algorithms is given in a table in the end of each section. From the left to the right, algorithms are grossly ranked in order of increasing sample reuse, but methods using a utility model in Θ and \mathcal{O} show better sample choice, resulting in competitive sample efficiency.

search methods. Then, in Section 6, we discuss the different elementary design choices that matter in terms of sample efficiency. Finally, Section 7 summarizes the paper and provides some perspectives about current trends in the domain.

2. Policy search without a utility model

When the function to be optimized is available but has no favorable property, the standard optimization method known as Gradient Descent consists in iteratively following the gradient of this function towards a local optimum. When the same function is only known through a model built by regression from a batch of samples, one can also do the same, but computing the gradient requires evaluations over the whole batch, which can be computationally expensive. An alternative known as Stochastic Gradient Descent (SGD) circumvents this difficulty by taking a small subset of the batch at each iteration (Bottou, 2012). Before starting to present these methods in Section 3, we first investigate a family of methods which perform policy search without learning a model of the utility function at all. They do so by sampling the policy parameter space Θ and moving towards policy parameters θ of higher utility $J(\theta)$.

2.1. Truly random search

At one extreme, the simplest black-box optimization (BBO) method randomly searches Θ until it stumbles on a good enough utility. We call this

method “Truly random search” as the name “random search” is used in the optimization community to refer to gradient-free methods (Rastrigin, 1963). Its distinguishing feature is in its sample choice strategy: the utility of the previous θ has no impact on the choice of the next θ .

Quite obviously, this sample choice strategy is not efficient, but it requires no assumption at all on the function to be optimized. Therefore, it is an option when this function does not show any *regularity* that can be exploited. All other methods rely on the implicit assumption that $J(\theta)$ presents some smoothness around optima θ^* , which is a first step towards using a gradient.

So globally, this method provides a proof of concept that an agent can obtain a better utility without estimating any gradient at all. Recently, other forms of gradient-free methods called *random search* though they are not truly random have been shown to be competitive with deep RL (Mania et al., 2018).

The next three families of methods, population-based optimization, evolutionary strategies and estimation of distribution algorithms, are all variants of *evolutionary* methods. An overview² of these methods is depicted in Figure 3.

²A blog with dynamical visualizations and more technical details can be found at <http://blog.otoro.net/2017/10/29/visual-evolution-strategies/>.




Figure 3: One iteration of evolutionary methods. (a) Population-based methods (b) Evolutionary Strategies (c) EDAs. Blue: current generation and sampling domain. Full blue dots: samples with a good evaluation. Dots with a red cross: samples with a poor evaluation. Green: new generation and sampling domain, empty dots are not evaluated yet. Red dots: optimum guess. In population-based methods, the next generation are offspring from several elite individuals of the previous generation. In ES, it is obtained from an optimum guess and sampling from fixed Gaussian noise. In EDAs, Gaussian noise is tuned using Covariance Matrix Adaptation.

2.2. Population-based optimization

Population-based BBO methods manage a limited population of individuals, and generate new individuals randomly in the vicinity of the previous *elite* individuals. There are several families of population-based optimization methods, the most famous being Genetic Algorithms (GAs) (Goldberg, 1989), Genetic Programming (GP) (Koza, 1992), and the more advanced NEAT framework (Stanley and Miikkulainen, 2002). In these methods, the parameter θ corresponding to an individual is often called its *genotype* and the corresponding utility is called its *fitness*, see (Back, 1996) for further reading. These methods have already been combined with neural networks giving rise to *neuroevolution* (Floreano et al., 2008) but, up to recently, these methods were mostly applied to small to moderate size policy representations. However, the availability of modern computational resources have made it possible to apply them to large and deep neural network representations, defining the emerging domain of *deep neuroevolution* (Petroski Such et al., 2017). Among other things, it was shown that, given large enough computational resources, methods as simple as GAs can offer a competitive alternative to deep RL methods presented in Section 5, mostly due to their excellent parallelization capabilities (Petroski Such et al., 2017; Conti et al., 2017).

2.3. Evolutionary strategies

Evolutionary strategies (ES) can be seen as specific population-based optimization methods where only one individual is retained from one generation

to the next. More specifically, an optimum guess is computed from the previous samples, then the next samples are obtained by adding Gaussian noise to the current optimum guess.

Moving from one optimum guess to the next implements a form of policy improvement similar to SGD, but where the gradient is approximated by averaging over samples instead of being analytically computed. Hence this method is more flexible but, since gradient approximation uses a random exploration component, it is less data efficient. However, data efficiency can be improved by reusing samples between one generation and the next when their sampling domain overlaps, a method called *importance mixing* (Sun et al., 2009). An improved version of importance mixing was recently proposed in (Pourchot et al., 2018), showing a large impact on sample efficiency, but not large enough to compete with deep RL methods on this aspect. Further results about importance mixing can be found in (Pourchot and Sigaud, 2018), showing that more investigations are necessary to better understand in which context this mechanism can be most useful.

The correlation between the direction of the gradient given by SGD and the same direction for ES depends on the evolutionary algorithm. Interestingly, good ES performance can be obtained even when the correlation is not high, though this result still needs to be confirmed in the case of policy search (Zhang et al., 2017).

A specific ES implementation of deep neuroevolution where constant Gaussian noise is used at each generation was shown to compete with deep RL methods on standard benchmarks (Salimans et al.,

2017). This simple implementation generated an insightful comparison with methods based on SGD depending on various gradient landscapes, showing under which conditions ES can find better optima than SGD (Lehman et al., 2017).

Finally, instead of approximating the vanilla gradient of utility, NES (Wierstra et al., 2008) and xNES (Glasmachers et al., 2010) approximate its natural gradient (Akimoto et al., 2010), but for doing so they have to compute the inverse of the Fisher Information Matrix, which is prohibitively expensive in large dimensions (Grondman et al., 2012). We refer the reader to (Pierrot et al., 2018) for a detailed presentation of *natural gradient* and other advanced gradient descent concepts.

2.4. Estimation of Distribution Algorithms

The standard perspective about EDAs is that they are a specific family of ES using a covariance matrix Σ (Larrañaga and Lozano, 2001). This covariance matrix defines a multivariate Gaussian function over Θ , hence its size is $|\Theta|^2$. Samples at the next iteration are drawn with a probability proportional to this Gaussian function. Along iterations, the ellipsoid defined by Σ is progressively adjusted to the top part of the hill corresponding to the local optimum θ^* .

The role of Σ is to control exploration. The exploration policy can be characterized as *uncorrelated* when it only updates the diagonal of Σ and *correlated* when it updates the full Σ (Deisenroth et al., 2013). The latter is more efficient in small parameter spaces but computationally more demanding and potentially inaccurate in larger spaces as more samples are required. In particular, it cannot be applied in the deep neuroevolution context where the order of magnitude of the size of θ is between thousands and millions.

Various instances of EDAs, such as CEM, CMA-ES, PI^2 -CMA, are covered in (Stulp and Sigaud, 2012a,b, 2013). Among them, the CMA-ES algorithm is also shown to approximate the natural gradient (Arnold et al., 2011). By contrast, the PI^{BB} algorithm, also described in (Stulp and Sigaud, 2013), is a simplification of PI^2 -CMA where covariance matrix adaptation has been removed. Thus it should be considered an instance of the former ES category.

2.5. Finite difference methods

In *finite difference* methods, the gradient of utility with respect to θ is estimated as the first order

approximation of the Taylor expansion of the utility function. This estimation is performed by applying local perturbations to the current input. Thus these methods are derivative-free and we classify them as using no model, even if they are based on a local linear approximation of the gradient.

In finite difference methods, gradient estimation can be cast as a standard regression problem, but perturbations along each dimension of Θ can be treated separately, which results in a very simple algorithm (Riedmiller et al., 2008). The counterpart of this simplicity is that it suffers from a lot of variance, so in practice the methods are limited to deterministic policies.

2.6. Reference to the main algorithms

Algorithm	Main paper
CMA-ES	(Hansen and Ostermeier, 2001)
CEM	(Rubinstein and Kroese, 2004)
finite diff.	(Riedmiller et al., 2008)
NES	(Wierstra et al., 2008)
xNES	(Glasmachers et al., 2010)
PI^{BB}	(Stulp and Sigaud, 2012b)
PI^2 -CMA	(Stulp and Sigaud, 2012a)
OpenAI-ES	(Salimans et al., 2017)
Random Search	(Mania et al., 2018)

Table 1: Main gradient-free algorithms. Above the line, they were studied in (Deisenroth et al., 2013), below they were not.

2.7. Sample efficiency analysis

In all gradient-free methods, sampling a vector of policy parameters θ provides an exact information about its utility $J(\theta)$. However, the J function can be stochastic, in which case one value of $J(\theta)$ only contains partial information about the value of that θ . Anyways, sample reuse can be implemented by storing an archive of the already sampled pairs $\langle \theta, J(\theta) \rangle$. Each time an algorithm needs the utility $J(\theta)$ of a sample θ , if this utility is already available in the archive, it can use it instead of sampling again. In the deterministic case, using the stored value is enough. In the stochastic case, the archive may provide a distribution over values $J(\theta)$, and the algorithm may either draw a value from this distribution or sample again, depending on accuracy requirements.

Message 1: Policy search without a utility model is generally less data efficient than Stochastic Gradient Descent (SGD). Though sample reuse is technically possible without a utility model, in practice it is seldom used. Despite their lower sample efficiency in comparison to SGD, some of these methods are highly parallelizable and offer a viable alternative to deep RL provided enough computational resources.

3. Policy search with a model of utility in the space of policy parameters

As outlined in the introduction, the utility of a vector of policy parameters is only available by observing the corresponding behavior. Although no model that relates policy parameters to utilities is given, one may approximate the utility function in Θ from these observations, by collecting samples consisting of (policy parameters, utility) pairs and using regression to infer a model of the corresponding function (see e.g. (Stulp and Sigaud, 2015)). Such a model could either be deterministic, giving one utility per policy parameters vector, or stochastic, giving a distribution over utility values.

Once such a model is learned, one could perform gradient descent on this model. These steps could be performed sequentially (first model learning and then gradient descent) or incrementally (improving the model and performing gradient descent after each new utility observation). In the latter case, the model is necessarily *persistent*: it evolves from iteration to iteration given new information, in contrast with the sequential case where it could be *transient*, that is recomputed from scratch at each iteration.

3.1. Bayesian Optimization

Though the above approach seems appealing, we are not aware of any algorithm performing what is described above in the deterministic case. A good reason for this is that utility functions are generally stochastic in Θ . Thus, algorithms which learn a model $J(\theta)$ have to learn a distribution over such models. This is exactly what Bayesian optimization (BO) does. The *distribution* over models is updated through Bayesian inference. It is initialized with a *prior*, and each new sample, considered as some new *evidence*, helps adjusting the model distribution towards a peak at the true value, whilst keeping track

of the variance over models. By estimating the uncertainty over the distribution of models, BO methods are endowed with active learning capabilities, dramatically improving their sample efficiency at the cost of a worse scalability.

A BO algorithm comes with a *covariance function* that determines how the information provided by a new sample influences the model distribution around this sample. It also comes with an *acquisition function* used to choose the next sample given the current model distribution. A good acquisition function should take into account the value and the uncertainty of the model over the sampled space.

By quickly reducing uncertainty, BO implements a form of active learning. As a consequence, it is very sample efficient when the parameter space is small enough, and it searches for the global optimum, rather than a local one. However, given the necessity to optimize globally over the acquisition function, it scales poorly in the size of the parameter space. For more details, see (Brochu et al., 2010).

The ROCK* algorithm is an instance of BO that searches for a local optimum instead of a global one (Hwangbo et al., 2014). It uses CMA-ES to find the optimum over the model function. By doing so, it performs natural rather than vanilla gradient optimization, but it does not use the available model of the utility function, though this could improve sample efficiency.

Bayesian optimization algorithms generally use Gaussian kernels to efficiently represent the distribution over models. However, some authors have started to note that, in the specific context of policy search, BO was not using all the information available in elementary steps of the agent. This led to the investigation of more appropriate data-driven kernels based on the Kullback-Leibler divergence between rollout density generated by two policies (Wilson et al., 2014).

Using BO in the context of policy search is an emerging domain (Lizotte et al., 2007; Calandra et al., 2014; Metzen et al., 2015; Martinez-Cantin et al., 2017). Furthermore, recent attempts to combine BO with reinforcement learning approaches, giving rise to the Bayesian Optimization Reinforcement Learning (BORL) framework, are described in Section 5.

3.2. Reference to the main algorithms

Algorithm	Main paper
Bayes. Opt.	(Pelikan et al., 1999)
ROCK*	(Hwangbo et al., 2014)

Table 2: Main Bayesian Optimization algorithms

3.3. Sample efficiency analysis

Learning a model of the utility function in Θ should be more sample efficient than trying to optimize without a model, as the gradient with respect to the model can be used to accelerate parameter improvement. However, learning a deterministic model is not enough for most cases, as the true utility function is generally stochastic in Θ , and learning a stochastic model comes with an additional computational cost which impacts the scalability of the approach.

Message 2: Bayesian Optimization is BBO managing a distribution over models in the policy parameter space. Its sample efficiency benefits from active choice of samples. But as it performs global search, it does not scale well to large policy parameter spaces. Thus, it is difficult to apply to deep neural network representations.

4. Directed exploration methods

Directed exploration methods are particularly useful in tasks with *sparse rewards*, i.e. where large parts of the search space have the same utility signal. These methods have two main features. First, instead of searching directly in the policy parameter space Θ , they search in a smaller *outcome space* O (also called *descriptor space* or *behavioral space*) and learn an invertible mapping between Θ and O . Second, they all optimize a task-independent criterion called novelty or diversity which is used to efficiently cover the outcome space.

The outcome itself corresponds to properties of the observed behavior. The general intuition is that if the outcome space is properly covered by known policy parametrizations, and if utility can be easily related to outcomes, then it should be easy to find policy parameters with a high utility, even when

the utility function is null for most policy parameters. Figure 4 visualizes why it is generally more efficient to perform the search for novel solutions in a dedicated outcome space and learning a mapping from Θ to O than performing this search directly in Θ (Baranes et al., 2014).

So, for the method to work, the outcome space has to be defined in such a way that determining the utility corresponding to an outcome is straightforward. Generally, the outcome space is defined by an external user to meet this requirement. Nevertheless, using representation learning methods to let the agent autonomously define its own outcome space is an emerging topic of interest (Pere et al., 2018; Laversanne-Finot et al., 2018).




Figure 4: A standard mapping between a policy parameter space Θ and an outcome space O . Most often, many policy parameters result in the same outcome (for instance, in the case of a robotic arm which must move a ball around, if the policy defines arm movements and the outcome space is defined as ball positions, most policy parameters will result into a static ball). In that case, sampling directly in Θ works poorly: you have to sample in such a way to efficiently cover O .

Directed exploration methods can be split into *novelty search* (NS) (Lehman and Stanley, 2011), *quality-diversity* (QD) (Pugh et al., 2015) and *goal exploration processes* (GEPs) (Baranes and Oudeyer, 2010; Forestier and Oudeyer, 2016; Forestier et al., 2017). The first two derive from evolutionary methods, whereas GEPs come from the developmental learning and intrinsic motivation literature.

An important distinction between them is that NS and GEPs are designed to optimize diversity only³, thus they do not use the utility function at all, whereas QD methods rely on multi-objective optimization methods to optimize diversity and utility at the same time.

³Hence the dotted line in Figure 2.

The NS approach arose from the realization that optimizing utility as a single objective is not the only option (Doncieux and Mouret, 2014). In particular, in the case of sparse or deceptive reward problems, it was shown that seeking novelty or diversity is an efficient strategy to obtain high utility solutions, even without explicitly optimizing this utility (Lehman and Stanley, 2011). The GEP approach was more inspired by thoughts on intrinsic motivations, where the goal was to have an agent achieve its own goal without an external utility signal (Forestier et al., 2017). However, researchers in evolutionary methods also realized that diversity and utility can be optimized jointly (Cuccu and Gomez, 2011), giving rise to more advanced NS and QD algorithms (Pugh et al., 2015; Cully and Demiris, 2017).

All these methods share a lot of similarities. They all start with a random search stage and, when they evaluate a policy parameter vector θ resulting in a point o in the outcome space O , they store the corresponding $\langle \theta, o \rangle$ pair in an archive. Because they use this archive for policy improvement, they all implement a form of *lazy learning*, endowing them with interesting sample efficiency properties (Aha, 1997). The archive itself can be seen as a stochastic model of the function relating Θ to O , as made particularly obvious in the *MAP-Elites* algorithm (Cully et al., 2015).

In more details, the main differences between these methods lie in the way they cover the outcome space O . NS and QD methods perform undirected variations to the elite θ vectors present in the archive. More precisely, in NS, the resulting solution is just added to the archive, whereas in QD the new solution replaces a previous one if it outperforms it both in terms of diversity and utility. By contrast, GEPs choose a desired outcome o^* and modify a copy of the θ leading to the closest outcome in the archive. The choice of a desired outcome o^* can be performed randomly or using curriculum learning or learning progress concepts (Baranes and Oudeyer, 2013; Forestier et al., 2017). Similarly, the modification of θ can be performed using undirected Gaussian noise or in more advanced ways. For instance, some GEP methods build a local linear model of the mapping from Θ to O to efficiently invert it, so as to find the θ^* corresponding to the desired outcome o^* (Baranes and Oudeyer, 2013).

Thus directed exploration methods all learn a stochastic and invertible mapping between Θ and

O . When they also learn a model of $J(o)$, this model is stochastic with respect to Θ , which makes them similar to BO methods. In that case, the outcome space is an intermediate space between Θ and utilities: policy parameters are first projected into the outcome space, and then a model of the utility function in this outcome space can be learned.

Learning a model of the utility in O shares some similarities with learning a critic in the state action space $\mathcal{X} \times \mathcal{U}$, as presented in Section 5. From this perspective, these methods can be seen as providing an intermediary family between evolutionary, BO and reinforcement learning methods. However, we shall see soon that learning a critic in the state action space $\mathcal{X} \times \mathcal{U}$ benefits from additional properties related to temporal difference learning, which limits the use of the above unifying perspective.

4.1. Reference to the main algorithms

Algorithm	Main paper
Novelty Search	(Lehman and Stanley, 2011)
Quality-Diversity	(Pugh et al., 2015)
Goal Exploration	(Forestier et al., 2017)

Table 3: Main directed exploration algorithms

4.2. Sample efficiency analysis

The defining characteristic of all directed exploration methods is their capability to widely cover the outcome space. This provides efficient exploration, which in turn critically improves sample efficiency when combined with more standard evolutionary methods mentioned in Section 2 (Conti et al., 2017) or deep RL method mentioned in Section 5 (Colas et al., 2018).

Even though our article focuses on single-task learning, it is worth mentioning that direct exploration methods may very much improve sample efficiency in multi-task learning scenarios. This is because such methods aim at covering the (interesting) outcome space, and can thus more easily adapt when facing multiple tasks, and thus potentially multiple outcomes.

Message 3: Looking for diversity only in a user defined outcome space is an efficient way to perform exploration, and can help solve sparse or deceptive reward problems, where more standard exploration would fail. Directed exploration methods are thus useful complements to other methods covered in this survey.

5. Policy search with a critic

The previous two sections have presented methods which learn mappings from policy parameter space Θ to utilities or outcomes. We now cover methods which learn a model of utility in the state-action space $\mathcal{X} \times \mathcal{U}$.

An important component in the RL formalization, the utility $U(\mathbf{x}, \mathbf{u})$ corresponds to the return the agent may expect from performing action \mathbf{u} when it is in state \mathbf{x} and then following either its current policy π_{θ} or the optimal policy π^* . This quantity may also depend on a discount factor γ and a noise parameter β .

The true utility $U(\mathbf{x}, \mathbf{u})$ can be approximated with a model $\hat{U}_{\eta}(\mathbf{x}, \mathbf{u})$ with parameters η . Such a model is called a *critic*. A key feature is that the critic can be learned from samples corresponding to *single steps* in the rollouts of the agent, either with temporal differencing or Monte Carlo methods. Methods that approximate U by \hat{U}_{η} , and determine the policy parameters θ by descending the gradient of θ with respect to \hat{U}_{η} are called *actor-critic* methods, the policy π_{θ} being the actor (Peters and Schaal, 2008b; Deisenroth et al., 2013).

This actor-critic approach can be applied to stochastic and deterministic policies (Silver et al., 2014). The space of deterministic policies being smaller than the space of stochastic policies, the latter can be advantageous because searching the former is faster than searching the latter. However, a stochastic policy might be more appropriate when Markov property does not hold (Williams and Singh, 1998; Sigaud and Buffet, 2010) or in adversarial contexts (Wang et al., 2016b).

5.1. Exploration in parameter or state-action space

As mentioned in Section 3, learning a model of the utility in the space Θ is a regression problem that is performed by sampling and exploring directly in Θ . In contrast, $\mathcal{X} \times \mathcal{U}$ cannot be sampled

directly, as one does not know in advance which policy parameters will result in visiting which states and performing which action. Exploration is therefore performed in either by adding noise to the θ (policy parameter perturbation), or adding noise to the actions the policy outputs (action perturbation). In the latter case, exploration is generally undirected and adds Gaussian noise or correlated Ornstein-Uhlenbeck noise to the actions taken by the policy. Policy parameter perturbation is done in PEPG, POWER and PI², and more recently to DDPG (Fortunato et al., 2017; Plappert et al., 2017), whereas action perturbation in the other algorithms presented in this paper.

All actor-critic algorithms iterate over the following three steps:

- A Collect new step samples from the current policy with policy parameter perturbation or action perturbation for exploration,
- B Compute a new critic \hat{U}_{η} based on these samples, by determining η through a temporal difference method,
- C Update the policy parameters θ through gradient descent with respect to the critic.

A distinction here should be made on whether 1) the critic is discarded after step C, and must thus be learned from scratch in step B in the next iteration, or 2) the critic is persistent throughout the learning, and incrementally updated in step B. We discuss the differences between these two variants – which we denote *transient critic* and *persistent critic* respectively – in more detail in the next two sections.

5.2. Transient Critic Algorithms

In methods with a transient critic, Monte Carlo sampling – running a large set of episodes and averaging over the stochastic return – is used to evaluate the current policy and generate a new set of step samples. Then, determining the optimal critic parameters given these samples can be cast as a batch regression problem.

Among these methods, one must distinguish between three families: *likelihood ratio* methods such as REINFORCE (Williams, 1992) and PEPG (Sehnke et al., 2010), *natural gradient* methods such as NAC and eNAC ⁴ (Peters and Schaal, 2008a) and *EM-*

⁴The critic is generally persistent in actor-critic methods, but this is not the case in NAC and eNAC.

based methods such as POWER⁵ (Kober and Peters, 2009) and the variants of REPS (Peters et al., 2010). All the corresponding algorithms are well described in (Deisenroth et al., 2013).

Although they derive from a different mathematical framework, likelihood ratio methods and EM-based methods are similar: they both use unbiased estimation of the gradient through Monte Carlo sampling and they are both mathematically designed so that the most rewarding rollouts get the highest probability.

The TRPO (Schulman et al., 2015) algorithm also follows an iterative approach and can use a deep neural network representation, thus it can be classified as a deep RL method. Among other things, it uses a bound on the Kullback-Leibler divergence between policies at successive iterations to ensure safe and efficient exploration. Finally, the Guided Policy Search (GPS) algorithm (Levine and Koltun, 2013; Montgomery and Levine, 2016) is another transient critic deep RL method inspired from REINFORCE, but adding guiding rollouts obtained from simpler policies.

5.3. Persistent Critic Algorithms

In contrast with transient critic algorithms, persistent critic algorithms incrementally update the critic during training. Most such algorithms use an actor-critic architecture, with the notable exception of NAF (Gu et al., 2016b), which does not have an explicit representation of the actor. To our knowledge, before the emergence of deep RL algorithms described below, the four INAC algorithms were the only representative of this family (Bhatnagar et al., 2007).

The way to compute the critic incrementally can be named a *temporal difference* (TD) method, also named a *bootstrap* method (Sutton, 1988). They compute at each step a *temporal difference error* or *reward prediction error* (RPE) between the immediate reward predicted by the current values of the critic and the actual reward received by the agent. This RPE can then be used as a loss that the critic should minimize over iterations (Sutton and Barto, 1998).

⁵Interestingly, PI^2 can also be seen as a transient critic method, though it could in principle use a persistent one and fall into Section 5.3. This is just because using batch updates make it more stable (Deisenroth et al., 2013).

5.4. Key properties of Persistent Critic Algorithms

Most mechanisms that made deep actor-critic algorithms possible were first introduced in DQN (Mnih et al., 2015). Though DQN is a discrete action algorithm which is outside the scope of this survey, we briefly review its important concepts and mechanisms before listing the main algorithms in the family of continuous action deep RL methods.

5.4.1. Accuracy and scalability: deep neural networks

By using deep neural networks as approximation functions and making profit of large computational capabilities of modern clusters of computers, all deep RL algorithms are capable of addressing much larger problems than before, and to approximate gradients with unprecedented accuracy, which makes them more stable than the previous linear architectures of NAC and POWER, hence amenable to incremental updates of a persistent critic rather than recomputing a transient one.

5.4.2. Stability: the target critic

Deep RL methods introduced a *target critic* as a way to improve stability. Standard regression is the process of fitting samples to a model so as to approximate an unknown stationary function (Stulp and Sigaud, 2015). Estimating a critic through temporal difference methods is similar to regression, but the target function is not stationary: it is itself a function of the estimated critic, thus it is modified each time the critic is modified. This can result in divergence of the critic when the target function and the estimated critic are racing after each other (Baird, 1994). To mitigate this instability, one should keep the target function stationary during several updates and reset it periodically to a new function corresponding to the current critic estimate, switching from a regression problem to another. This idea was first introduced in DQN (Mnih et al., 2015) and then modified from periodic updates to smooth variations in DDPG (Lillicrap et al., 2015).

5.4.3. Sample reuse: the replay buffer

Since they are based on *value propagation*, TD methods can give rise to more sample reuse than standard regression methods, provided that these samples are saved into a *replay buffer*. Using a replay buffer is at the heart of the emergence of modern actor-critic approaches in deep RL. Actually,

learning from the samples in the order in which they are collected is detrimental to learning performance and stability because these samples are not independent and identically distributed (i.i.d.). Stability is improved by drawing the samples randomly from the replay buffer and sample efficiency can be further improved by better choosing the samples, using *prioritized experience replay* (Schaul et al., 2015).

5.4.4. Adaptive step sizes and return lengths

Modern SGD methods provided by most machine learning libraries now incorporate adaptive step sizes, removing a difficulty with previous actor-critic algorithms such as eNAC. Another important ingredient for the success of some recent methods in the use of *n-step return*, which consists in performing temporal difference updates over several time steps, resulting in the possibility to control the bias-variance trade-off (see Section 6.3.1).

5.5. Overview of deep RL algorithms

All these favorable properties are common traits of several incremental deep RL algorithms: DDPG (Lillicrap et al., 2015), NAF (Gu et al., 2016b), PPO (Schulman et al., 2017), ACKTR (Wu et al., 2017), SAC (Haarnoja et al., 2018), TD3 (Fujimoto et al., 2018) and D4PG (Barth-maroon et al., 2018). As depicted in Figure 2, the last one, D4PG, is an instance of Bayesian Optimization Reinforcement Learning (BORL) algorithms which derive from BO but belong to the step-based category of methods described in Section 5. These algorithms result from an effort to incorporate Bayesian computations into the deep RL framework, and correspond to a very active trend in the field. Most of these works address discrete actions (Azizzadenesheli et al., 2018; Tang and Kucukelbir, 2017), but D4PG is an exception that derives from adopting a distributional perspective on policy gradient computation, resulting in more accurate estimates on the gradient and better sample efficiency (Bellemare et al., 2017).

Finally, a few algorithms such as ACER (Wang et al., 2016b), Q-PROP (Gu et al., 2016a) and PGQL (O’Donoghue et al., 2016) combine properties of transient and persistent critic methods, and are captured into the more general framework of Interpolated Policy Gradient (IPG) (Gu et al., 2017). For a more detailed description of all these algorithms, we refer the reader to the corresponding papers and to a recent survey (Arulkumaran et al., 2017).

5.6. Reference to the main algorithms

Algorithm	Main paper
REINFORCE	(Williams, 1992)
G(PO)MDP	(Baxter and Bartlett, 2001)
NAC	(Peters and Schaal, 2008a)
eNAC	(Peters and Schaal, 2008a)
POWER	(Kober and Peters, 2009)
PI ²	(Theodorou et al., 2010)
REPS	(Peters et al., 2010)
PEPG	(Sehnke et al., 2010)
VIPS	(Neumann, 2011)
iNAC	(Bhatnagar et al., 2007)
GPS	(Levine and Koltun, 2013)
TRPO	(Schulman et al., 2015)
DDPG	(Lillicrap et al., 2015)
A3C	(Mnih et al., 2016)
NAF	(Gu et al., 2016b)
ACER	(Wang et al., 2016b)
Q-PROP	(Gu et al., 2016a)
PGQL	(O’Donoghue et al., 2016)
PPO	(Schulman et al., 2017)
ACKTR	(Wu et al., 2017)
SAC	(Haarnoja et al., 2018)
TD3	(Fujimoto et al., 2018)
D4PG	(Barth-maroon et al., 2018)

Table 4: Main reinforcement learning algorithms. Algorithms below the line have not yet been covered in (Deisenroth et al., 2013).

5.7. Sample efficiency analysis

Message 4: Being step-based, deep RL methods are able to use more information from rollouts than episode-based methods. Furthermore, using a replay buffer leads to further sample reuse.

6. Discussion

In the previous sections we have presented methods which: (1) do *not* build a utility model; (2) learn a utility model: (2a) in the policy parameter space Θ , (2b) in an arbitrary outcome space O , (2c) in the state-action space $\mathcal{X} \times \mathcal{U}$. In this section, we come back to the sample efficiency properties of these different methods. We do so by descending the tree of design choices depicted in Figure 2.

6.1. Building a model or not

We have outlined that policy search methods which build a model of the utility function are generally more sample efficient than methods which do not. However, the reliance of the latter to SGD can make them less robust to local optima (Lehman et al., 2017) and it has been shown recently that methods which do not build a model of utility are still competitive in terms of final performance, due to their higher parallelization capability and distinguishing properties with respect to various gradient landscapes (Salimans et al., 2017; Petroski Such et al., 2017; Zhang et al., 2017).

6.2. Building a utility function model in the policy parameter space versus the state-action space.

Several elements speak in favor of the higher sample efficiency of learning a critic in the state-action space $\mathcal{X} \times \mathcal{U}$. First, it can give rise to more sample reuse than learning a model of the utility function in Θ . Second, learning from each step separately makes a better use of the information available from a rollout than learning from global episodes.

Furthermore, $\mathcal{X} \times \mathcal{U}$ may naturally exhibit a hierarchical structure – especially the state – which is not so obvious for Θ . As a consequence, methods that model utility in $\mathcal{X} \times \mathcal{U}$ may benefit from learning intermediate representations at different levels in the hierarchy, thus reducing the dimensionality of the policy search problem. Learning such intermediate and more compact representations is the focus of hierarchical reinforcement learning, a domain which has also been impacted by the emergence of deep RL (Kulkarni et al., 2016; Bacon et al., 2017). Hierarchical reinforcement learning can also be performed off-line, which corresponds to the perspective of the DREAM project⁶, illustrated for instance in (Zimmer and Doncieux, 2017).

Finally, an important factor of sample efficiency is the size and structure of $\mathcal{X} \times \mathcal{U}$ with respect to Θ . In both respects, the emergence of deep RL methods using large neural networks as policy representation has changed the perspective. First, in deep RL, the size of Θ can become larger than that of $\mathcal{X} \times \mathcal{U}$, which speaks in favor of learning a critic. Second, deep neural networks seem to generally induce a smooth structure between Θ and $\mathcal{X} \times \mathcal{U}$, which facilitates learning. Finally, a utility function modeled in a larger space may suffer from fewer

local minima, as more directions remain for improving the gradient (Kawaguchi, 2016).

The above conclusions might be mitigated by considering exploration. Indeed, in several surveys about policy search for robotics, policy parameter perturbation methods are considered superior to action perturbation methods (Stulp and Sigaud, 2013; Deisenroth et al., 2013). This analysis is backed-up with several mathematical arguments, but it might be true mostly when the space Θ is smaller than the space $\mathcal{X} \times \mathcal{U}$. Until recently, all deep RL methods were using action perturbation. But deep RL algorithms using policy parameter perturbation have recently been published, showing again that one can model the utility function in $\mathcal{X} \times \mathcal{U}$ while performing exploration in Θ (Fortunato et al., 2017; Plappert et al., 2017). Exploration is currently one of the hottest topics in deep RL and directed exploration methods presented in Section 4 may play a key role in this story, despite the lower data efficiency of their policy improvement mechanisms (Conti et al., 2017; Colas et al., 2018).

Message 5: There are more arguments for learning a utility model in $\mathcal{X} \times \mathcal{U}$ than in Θ , but this ultimately depends on the size of these spaces and the structure of their relationship.

6.3. Transient versus persistent critic

At first glance, having a persistent critic may seem superior to having a transient one, for three reasons. First, by avoiding to compute the critic again at each iteration, it is computationally more efficient. Second, immediate updates favor data efficiency because the policy is improved as soon as possible, which in turn helps generating better samples. Third, being based on bootstrap methods, they give rise to more sample reuse. However, these statements must be differentiated, as two factors (described below) must be taken into account.

6.3.1. Trading bias against variance

Estimating the utility of a policy in $\mathcal{X} \times \mathcal{U}$ is subject to a bias-variance compromise (Kearns and Singh, 2000). On the one hand, estimating the utility of a given policy through Monte Carlo sampling – as is generally done in transient critic approaches – is subject to a variance which grows with the

⁶<http://www.robotsthatdream.eu/>

length of the episodes. On the other hand, incrementally updating a persistent critic reduces variance, but may suffer from bias, resulting in potential sub-optimality, or even divergence. Instead of performing bootstrap updates of a critic over one step, one can do so over N steps. The larger N , the closer to Monte Carlo estimation, thus tuning N is a way of controlling the bias-variance compromise. For instance, while the transient critic TRPO algorithm is less sample efficient than actor-critic methods but more stable, often resulting in superior performance (Duan et al., 2016), its immediate successor, PPO, uses N steps return, resulting in a good compromise between both families (Schulman et al., 2017).

6.3.2. Off-policy versus on-policy updates

In on-policy methods such as SARSA, the samples used to learn the critic must come from the current policy, whereas in off-policy methods such as Q-LEARNING, they can come from any policy. In most transient critic methods, the samples are discarded from one iteration to the next and these methods are generally on-policy. By contrast, persistent critic methods using a replay buffer are generally off-policy⁷.

This on-policy versus off-policy distinction is related to the bias-variance compromise. Indeed, when learning a persistent critic incrementally, using off-policy updates is more flexible because the samples can come from any policy, but these off-policy updates introduce bias in the estimation of the critic. As a result, off-policy methods such as DDPG and NAF are more sample efficient because they use a replay buffer, but they are also more prone to sub-optimality and divergence. In that respect, a key contribution of ACER and Q-PROP is that they provide an off-policy, sample efficient update method which strongly controls the bias, resulting in more stability (Gu et al., 2016a; Wang et al., 2016b; Wu et al., 2017; Gu et al., 2017). These aspects are currently the subject of intensive research but the resulting algorithms suffer from being more complex, with additional meta-parameters.

⁷The A3C algorithm is an incremental actor-critic method which does not use a replay buffer, and it is classified as on-policy.

Message 6: Persistent critic methods are superior to transient critic methods in many respects, but the latter are more stable because they decorrelate the problem of estimating the utility function from the problem of descending its gradient, and they suffer from less bias.

7. Conclusion

In this paper, we have contrasted various approaches to policy search, from evolutionary methods which do not learn a model of the utility function to deep RL methods which do so in the state-action space.

In (Stulp and Sigaud, 2013), the authors have shown that policy search applied to robotics was shifting from actor-critic methods to evolutionary methods. Part of this shift was due to the use of open-loop DMPs (Ijspeert et al., 2013) as a policy representation, which favors episode-based approaches, but another part resulted from the higher stability and efficiency of evolutionary methods by that time.

The emergence of deep RL methods has changed this perspective. It should be clear from this survey that, in the context of large problems where deep neural network representations are now the standard option, deep RL is generally more sample efficient than deep neuroevolution methods, as empirically confirmed in (de Froissard de Broissia and Sigaud, 2016) and (Pourchot et al., 2018). The higher sample efficiency of deep RL methods, and particularly actor-critic architectures with a persistent critic, results from several mechanisms. They benefit from better approximation capability of non-linear critics and the incorporation of an adapted step size in SGD, they model the utility function in the state-action space, and they benefit from massive sample reuse by using a replay buffer. Using a target network has also mitigated the intrinsic instability of incrementally approximating a critic. However, it is important to acknowledge that incremental deep RL methods still suffer from significant instability⁸.

⁸As outlined at <https://www.alexirpan.com/2018/02/14/rl-hard.html>.

7.1. Future directions

The field of policy search is currently the object of an intense race for increased performance, stability and sample efficiency. We now outline what we currently consider as promising research directions.

7.1.1. More analyses than competitions

Up to now, the main trend in the literature focuses on performance comparisons (Duan et al., 2016; Islam et al., 2017; Henderson et al., 2017; Petroski Such et al., 2017) showing that, despite their lower sample efficiency, methods which do not build a model of utility are still a competitive alternative in terms of final performance (Salimans et al., 2017; Chrabaszcz et al., 2018). But stability and sample efficiency comparisons are missing and works analyzing the reasons why an algorithm performs better than another are only just emerging (Lehman et al., 2017; Zhang et al., 2017; Gangwani and Peng, 2017). By drawing an overview of the whole field and revealing some important factors behind sample efficiency, this paper was intended as a starting point towards broader and deeper analyses of the efficacy of various policy search methods.

7.1.2. More combinations than competitions

An important trend corresponds to the emergence of methods which combine algorithms from various families described above. As already noted in Section 4, directed exploration methods are often combined with evolutionary or deep RL methods (Conti et al., 2017; Colas et al., 2018). There is also an emerging trend combining evolutionary or population-based methods and deep RL methods (Jaderberg et al., 2017; Khadka and Tumer, 2018; Pourchot and Sigaud, 2018) which seem to be able to take the best of both worlds. We believe we are just at the beginning of such combinations and that this area has a lot of potential.

7.1.3. Beyond single policy improvement

Though we decided to keep *lifelong*, *continual* and *open-ended* learning outside the scope of this survey, we must mention that fast progress in policy improvement has favored an important tendency to address several tasks at the same time (Yang and Hospedales, 2014). This subfield is extremely active at the moment, with many works in multitask learning (Vezhnevets et al., 2017; Veeriah et al., 2018; Gangwani and Peng, 2018), Hierarchical Reinforcement Learning (Levy et al., 2018;

Nachum et al., 2018) and meta reinforcement learning (Wang et al., 2016a), to cite only a few.

Finally, because we focused on these elementary aspects, we have left aside the emerging topic of state representation learning (Jonschkowski and Brock, 2014; Raffin et al., 2016; Lesort et al., 2018) or using auxiliary tasks for improving deep RL (Shelhamer et al., 2016; Jaderberg et al., 2016; Riedmiller et al., 2018). The impact of these methods should be made clearer in the future.

7.2. Final word

As we have highlighted in the article, research in policy search and deep RL moves at a very high pace. Therefore, forecasting future trends, as we have done above, is risky, and even attempts to analyze the factors underlying current trends may be quickly outdated, but this also what makes this such an exciting research field.

Acknowledgments

Olivier Sigaud was supported by the European Commission, within the DREAM project, and has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement N° 640891. Freek Stulp was supported by the HGF project “Reduced Complexity Models” and received funding from the European Commission through the project H2020 AnDy (GA no. 731540). We thank David Filliat, Nicolas Perrin and Pierre-Yves Oudeyer for their feedback on this article.

References

- Aha, D. W., 1997. Editorial. In: *Lazy learning*. Springer, pp. 7–10.
- Akimoto, Y., Nagata, Y., Ono, I., Kobayashi, S., 2010. Bidirectional relation between cma evolution strategies and natural evolution strategies. In: *International Conference on Parallel Problem Solving from Nature*. Springer, pp. 154–163.
- Argall, B. D., Chernova, S., Veloso, M., Browning, B., 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 469–483.
- Arnold, L., Auger, A., Hansen, N., Ollivier, Y., 2011. Information-geometric optimization algorithms: A unifying picture via invariance principles. Tech. rep., INRIA Saclay.
- Arulkumaran, K., Deisenroth, M. P., Brundage, M., Bharath, A. A., 2017. A brief survey of deep reinforcement learning. *arXiv preprint arXiv:1708.05866*.
- Azizzadenesheli, K., Brunskill, E., Anandkumar, A., 2018. Efficient exploration through bayesian deep q-networks. *arXiv preprint arXiv:1802.04412*.

- Back, T., 1996. Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms. Oxford university press.
- Bacon, P.-L., Harb, J., Precup, D., 2017. The option-critic architecture. In: AAAI. pp. 1726–1734.
- Baird, L. C., 1994. Reinforcement learning in continuous time: Advantage updating. In: Proceedings of the International Conference on Neural Networks. Orlando, FL.
- Baranes, A., Oudeyer, P.-Y., 2010. Intrinsically motivated goal exploration for active motor learning in robots: A case study. In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010). IEEE, Taipei, Taiwan, Province Of China.
- Baranes, A., Oudeyer, P.-Y., 2013. Active learning of inverse models with intrinsically motivated goal exploration in robots. *Robotics and Autonomous Systems* 61 (1), 49–73.
- Baranes, A. F., Oudeyer, P.-Y., Gottlieb, J., 2014. The effects of task difficulty, novelty and the size of the search space on intrinsically motivated exploration. *Frontiers in neuroscience* 8, 317.
- Barth-marion, G., Hoffman, M., Budden, D., Dabney, W., Horgan, D., TB, D., Muldal, A., Heess, N., Lillicrap, T. P., 2018. Distributional policy gradient. In: ICLR. pp. 1–16.
- Baxter, J., Bartlett, P. L., 2001. Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research* 15, 319–350.
- Bellemare, M. G., Dabney, W., Munos, R., 2017. A distributional perspective on reinforcement learning. arXiv preprint arXiv:1707.06887.
- Bhatnagar, S., Sutton, R. S., Ghavamzadeh, M., Lee, M., 2007. Incremental natural actor-critic algorithms. In: Advances in Neural Information Processing Systems. MIT Press.
- Bottou, L., 2012. Stochastic gradient descent tricks. In: Neural networks: Tricks of the trade. Springer, pp. 421–436.
- Brochu, E., Cora, V. M., De Freitas, N., 2010. A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. arXiv preprint arXiv:1012.2599.
- Calandra, R., Gopalan, N., Seyfarth, A., Peters, J., Deisenroth, M. P., 2014. Bayesian gait optimization for bipedal locomotion. In: International Conference on Learning and Intelligent Optimization. Springer, pp. 274–290.
- Chatzilygeroudis, K., Rama, R., Kaushik, R., Goepp, D., Vassiliades, V., Mouret, J.-B., 2017. Black-box data-efficient policy search for robotics. arXiv preprint arXiv:1703.07261.
- Chrabaszcz, P., Loshchilov, I., Hutter, F., 2018. Back to basics: Benchmarking canonical evolution strategies for playing atari. arXiv preprint arXiv:1802.08842.
- Colas, C., Sigaud, O., Oudeyer, P.-Y., 2018. GEP-PG: Decoupling exploration and exploitation in deep reinforcement learning algorithms. arXiv preprint arXiv:1802.05054.
- Conti, E., Madhavan, V., Such, F. P., Lehman, J., Stanley, K. O., Clune, J., 2017. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. arXiv preprint arXiv:1712.06560.
- Cuccu, G., Gomez, F., 2011. When novelty is not enough. In: European Conference on the Applications of Evolutionary Computation. Springer, pp. 234–243.
- Cully, A., Clune, J., Tarapore, D., Mouret, J.-B., 2015. Robots that can adapt like animals. *Nature* 521 (7553), 503–507.
- Cully, A., Demiris, Y., 2017. Quality and diversity optimization: A unifying modular framework. *IEEE Transactions on Evolutionary Computation*.
- de Froissard de Broissia, A., Sigaud, O., 2016. Actor-critic versus direct policy search: a comparison based on sample complexity. arXiv preprint arXiv:1606.09152.
- Deisenroth, M., Rasmussen, C. E., 2011. Pilco: A model-based and data-efficient approach to policy search. In: Proceedings of the 28th International Conference on machine learning. pp. 465–472.
- Deisenroth, M. P., Neumann, G., Peters, J., et al., 2013. A survey on policy search for robotics. *Foundations and Trends® in Robotics* 2 (1–2), 1–142.
- Doncieux, S., Mouret, J.-B., 2014. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence* 7 (2), 71–93.
- Duan, Y., Chen, X., Houthoofd, R., Schulman, J., Abbeel, P., 2016. Benchmarking deep reinforcement learning for continuous control. arXiv preprint arXiv:1604.06778.
- Floreano, D., Dürr, P., Mattiussi, C., 2008. Neuroevolution: from architectures to learning. *Evolutionary Intelligence* 1 (1), 47–62.
- Forestier, S., Mollard, Y., Oudeyer, P.-Y., 2017. Intrinsically motivated goal exploration processes with automatic curriculum learning. arXiv preprint arXiv:1708.02190.
- Forestier, S., Oudeyer, P.-Y., 2016. Overlapping waves in tool use development: a curiosity-driven computational model.
- Fortunato, M., Azar, M. G., Piot, B., Menick, J., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., et al., 2017. Noisy networks for exploration. arXiv preprint arXiv:1706.10295.
- Fujimoto, S., van Hoof, H., Meger, D., 2018. Addressing function approximation error in actor-critic methods. arXiv preprint arXiv:1802.09477.
- Gangwani, T., Peng, J., 2017. Genetic policy optimization. arXiv preprint arXiv:1711.01012.
- Gangwani, T., Peng, J., 2018. Policy optimization by genetic distillation. In: ICLR 2018.
- Gill, P. E., Murray, W., Wright, M. H., 1981. Practical optimization. Academic press.
- Gasmachars, T., Schaul, T., Yi, S., Wierstra, D., Schmidhuber, J., 2010. Exponential natural evolution strategies. In: Proceedings of the 12th annual conference on Genetic and evolutionary computation. ACM, pp. 393–400.
- Goldberg, D. E., 1989. Genetic Algorithms in Search, Optimization, and Machine Learning. Addison Wesley, Reading, MA.
- Grondman, I., Busoniu, L., Lopes, G. A., Babuska, R., 2012. A survey of actor-critic reinforcement learning: Standard and natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 42 (6), 1291–1307.
- Gu, S., Lillicrap, T., Ghahramani, Z., Turner, R. E., Levine, S., 2016a. Q-prop: Sample-efficient policy gradient with an off-policy critic. arXiv preprint arXiv:1611.02247.
- Gu, S., Lillicrap, T., Ghahramani, Z., Turner, R. E., Schölkopf, B., Levine, S., 2017. Interpolated policy gradient: Merging on-policy and off-policy gradient estimation for deep reinforcement learning. arXiv preprint arXiv:1706.00387.
- Gu, S., Lillicrap, T., Sutskever, I., Levine, S., 2016b. Continuous deep q-learning with model-based acceleration. arXiv preprint arXiv:1603.00748.

- Haarnoja, T., Zhou, A., Abbeel, P., Levine, S., 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. arXiv preprint arXiv:1801.01290.
- Hansen, N., Ostermeier, A., 2001. Completely derandomized self-adaptation in evolution strategies. *Evolutionary Computation* 9 (2), 159–195.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D., 2017. Deep reinforcement learning that matters. arXiv preprint arXiv:1709.06560.
- Hessel, M., Modayil, J., Van Hasselt, H., Schaul, T., Ostrovski, G., Dabney, W., Horgan, D., Piot, B., Azar, M., Silver, D., 2017. Rainbow: Combining improvements in deep reinforcement learning. arXiv preprint arXiv:1710.02298.
- Hwangbo, J., Gehring, C., Sommer, H., Siegwart, R., Buchli, J., 2014. ROCK*: efficient black-box optimization for policy learning. In: *IEEE-RAS International Conference on Humanoid Robots*. IEEE, pp. 535–540.
- Ijspeert, A. J., Nakanishi, J., Hoffmann, H., Pastor, P., Schaal, S., 2013. Dynamical movement primitives: learning attractor models for motor behaviors. *Neural computation* 25 (2), 328–373.
- Islam, R., Henderson, P., Gomrokchi, M., Precup, D., 2017. Reproducibility of benchmarked deep reinforcement learning tasks for continuous control. In: *Proceedings of the ICML 2017 workshop on Reproducibility in Machine Learning (RML)*.
- Jaderberg, M., Dalibard, V., Osindero, S., Czarnecki, W. M., Donahue, J., Razavi, A., Vinyals, O., Green, T., Dunning, I., Simonyan, K., et al., 2017. Population based training of neural networks. arXiv preprint arXiv:1711.09846.
- Jaderberg, M., Mnih, V., Czarnecki, W. M., Schaul, T., Leibo, J. Z., Silver, D., Kavukcuoglu, K., 2016. Reinforcement learning with unsupervised auxiliary tasks. arXiv preprint arXiv:1611.05397.
- Jonschkowski, R., Brock, O., 2014. State representation learning in robotics: Using prior knowledge about physical interaction. In: *Proceedings of Robotics, Science and Systems*.
- Kawaguchi, K., 2016. Deep learning without poor local minima. In: *Advances In Neural Information Processing Systems*. pp. 586–594.
- Kearns, M. J., Singh, S. P., 2000. Bias-variance error bounds for temporal difference updates. In: *COLT*. pp. 142–147.
- Khadka, S., Tumer, K., 2018. Evolutionary reinforcement learning. arXiv preprint arXiv:1805.07917.
- Kober, J., Bagnell, J. A., Peters, J., 2013. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* 32 (11), 1238–1274.
- Kober, J., Peters, J., 2009. Learning motor primitives for robotics. In: *IEEE International Conference on Robotics and Automation*. IEEE, pp. 2112–2118.
- Koza, J. R., 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- Kulkarni, T. D., Narasimhan, K. R., Saeedi, A., Tenenbaum, J. B., 2016. Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation. arXiv preprint arXiv:1604.06057.
- Larrañaga, P., Lozano, J. A., 2001. Estimation of distribution algorithms: A new tool for evolutionary computation. Vol. 2. Springer Science & Business Media.
- Laversanne-Finot, A., Péré, A., Oudeyer, P.-Y., 2018. Curiosity driven exploration of learned disentangled goal spaces. arXiv preprint arXiv:1807.01521.
- Lehman, J., Chen, J., Clune, J., Stanley, K. O., 2017. ES is more than just a traditional finite-difference approximator. arXiv preprint arXiv:1712.06568.
- Lehman, J., Stanley, K. O., 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19 (2), 189–223.
- Lesort, T., Díaz-Rodríguez, N., Goudou, J.-F., Filliat, D., 2018. State representation learning for control: An overview. arXiv preprint arXiv:1802.04181.
- Levine, S., Koltun, V., 2013. Guided policy search. In: *Proceedings of the 30th International Conference on Machine Learning*. pp. 1–9.
- Levy, A., Platt, R., Saenko, K., 2018. Hierarchical reinforcement learning with hindsight. arXiv preprint arXiv:1805.08180.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., Wierstra, D., 2015. Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971.
- Lizotte, D. J., Wang, T., Bowling, M. H., Schuurmans, D., 2007. Automatic gait optimization with gaussian process regression. In: *IJCAI*. Vol. 7. pp. 944–949.
- Mania, H., Guy, A., Recht, B., 2018. Simple random search provides a competitive approach to reinforcement learning. arXiv preprint arXiv:1803.07055.
- Martinez-Cantin, R., Tee, K., McCourt, M., 2017. Policy search using robust Bayesian optimization. In: *Neural Information Processing Systems (NIPS) Workshop on Acting and Interacting in the Real World: Challenges in Robot Learning*.
- Metzen, J. H., Fabisch, A., Hansen, J., 2015. Bayesian optimization for contextual policy search. In: *Proceedings of the Second Machine Learning in Planning and Control of Robot Motion Workshop*, Hamburg.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., Kavukcuoglu, K., 2016. Asynchronous methods for deep reinforcement learning. arXiv preprint arXiv:1602.01783.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al., 2015. Human-level control through deep reinforcement learning. *Nature* 518 (7540), 529–533.
- Montgomery, W. H., Levine, S., 2016. Guided policy search via approximate mirror descent. In: *Advances in Neural Information Processing Systems*. pp. 4008–4016.
- Nachum, O., Gu, S., Lee, H., Levine, S., 2018. Data-efficient hierarchical reinforcement learning. arXiv preprint arXiv:1805.08296.
- Najar, A., Sigaud, O., Chetouani, M., 2016. Training a robot with evaluative feedback and unlabeled guidance signals. In: *25th IEEE International Symposium on Robot and Human Interactive Communication*. IEEE, pp. 261–266.
- Neumann, G., 2011. Variational inference for policy search in changing situations. In: *Proceedings of the 28th international conference on machine learning*. pp. 817–824.
- O’Donoghue, B., Munos, R., Kavukcuoglu, K., Mnih, V., 2016. Combining policy gradient and q-learning. arXiv preprint arXiv:1611.01626.
- Pelikan, M., Goldberg, D. E., Cantú-Paz, E., 1999. BoA: The Bayesian optimization algorithm. In: *Proceedings of the 1st Annual Conference on Genetic and Evolutionary Computation-Volume 1*. Morgan Kaufmann Publishers Inc., pp. 525–532.
- Pere, A., Forestier, S., Sigaud, O., Oudeyer, P.-Y., 2018.

- Unsupervised learning of goal spaces for intrinsically motivated goal exploration. In: International Conference on Learning Representations (ICLR). ArXiv preprint arXiv:1803.00781.
- Peters, J., Mülling, K., Altun, Y., 2010. Relative entropy policy search. In: AAAI. Atlanta, pp. 1607–1612.
- Peters, J., Schaal, S., 2008a. Natural actor-critic. *Neurocomputing* 71 (7-9), 1180–1190.
- Peters, J., Schaal, S., 2008b. Reinforcement learning of motor skills with policy gradients. *Neural networks* 21 (4), 682–697.
- Petroski Such, F., Madhavan, V., Conti, E., Lehman, J., Stanley, K. O., Clune, J., 2017. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. arXiv preprint arXiv:1712.06567.
- Pierrot, T., Perrin, N., Sigaud, O., 2018. First-order and second-order variants of the gradient descent: a unified framework. arXiv preprint arXiv:1810.08102.
- Plappert, M., Houthoofd, R., Dhariwal, P., Sidor, S., Chen, R. Y., Chen, X., Asfour, T., Abbeel, P., Andrychowicz, M., 2017. Parameter space noise for exploration. arXiv preprint arXiv:1706.01905.
- Pourchot, A., Perrin, N., Sigaud, O., 2018. Importance mixing: Improving sample reuse in evolutionary policy search methods. arXiv preprint arXiv:1808.05832.
- Pourchot, A., Sigaud, O., 2018. Cem-rl: Combining evolutionary and gradient-based methods for policy search. arXiv preprint arXiv:1810.01222.
- Pugh, J. K., Soros, L., Szerlip, P. A., Stanley, K. O., 2015. Confronting the challenge of quality diversity. In: Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation. ACM, pp. 967–974.
- Raffin, A., Höfer, S., Jonschkowski, R., Brock, O., Stulp, F., 2016. Unsupervised learning of state representations for multiple tasks. Tech. rep., University of Berlin.
- Rastrigin, L., 1963. The convergence of the random search method in the extremal control of a many parameter system. *Automation and Remote Control* 24 (10), 1337–1342.
- Riedmiller, M., Hafner, R., Lampe, T., Neunert, M., Degraeve, J., Van de Wiele, T., Mnih, V., Heess, N., Springenberg, J. T., 2018. Learning by playing-solving sparse reward tasks from scratch. arXiv preprint arXiv:1802.10567.
- Riedmiller, M., Peters, J., Schaal, S., 2008. Evaluation of policy gradient methods and variants on the cartpole benchmark. In: IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning (ADPRL).
- Rubinstein, R., Kroese, D., 2004. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation, and Machine Learning*. Springer-Verlag.
- Salimans, T., Ho, J., Chen, X., Sutskever, I., 2017. Evolution strategies as a scalable alternative to reinforcement learning. arXiv preprint arXiv:1703.03864.
- Schaul, T., Quan, J., Antonoglou, I., Silver, D., 2015. Prioritized experience replay. arXiv preprint arXiv:1511.05952.
- Schulman, J., Levine, S., Moritz, P., Jordan, M. I., Abbeel, P., 2015. Trust region policy optimization. CoRR, abs/1502.05477.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O., 2017. Proximal policy optimization algorithms. arXiv preprint arXiv:1707.06347.
- Sehnke, F., Osendorfer, C., Rückstieß, T., Graves, A., Peters, J., Schmidhuber, J., 2010. Parameter-exploring policy gradients. *Neural Networks* 23 (4), 551–559.
- Shelhamer, E., Mahmoudieh, P., Argus, M., Darrell, T., 2016. Loss is its own reward: Self-supervision for reinforcement learning. arXiv preprint arXiv:1612.07307.
- Sigaud, O., Buffet, O., 2010. *Markov Decision Processes in Artificial Intelligence*. iSTE - Wiley.
- Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., Riedmiller, M., 2014. Deterministic policy gradient algorithms. In: Proceedings of the 30th International Conference in Machine Learning.
- Stanley, K. O., Miikkulainen, R., 2002. Efficient evolution of neural network topologies. In: *Evolutionary Computation, 2002. CEC’02. Proceedings of the 2002 Congress on*. Vol. 2. IEEE, pp. 1757–1762.
- Stulp, F., Sigaud, O., 2012a. Path integral policy improvement with covariance matrix adaptation. In: Proceedings of the 29th International Conference on Machine Learning. Edinburgh, Scotland, pp. 1–8.
- Stulp, F., Sigaud, O., 2012b. Policy improvement methods: Between black-box optimization and episodic reinforcement learning. Tech. rep., hal-00738463.
- Stulp, F., Sigaud, O., august 2013. Robot skill learning: From reinforcement learning to evolution strategies. *Paladyn Journal of Behavioral Robotics* 4 (1), 49–61.
- Stulp, F., Sigaud, O., 2015. Many regression algorithms, one unified model: A review. *Neural Networks* 69, 60–79.
- Sun, Y., Wierstra, D., Schaul, T., Schmidhuber, J., 2009. Efficient natural evolution strategies. In: Proceedings of the 11th Annual conference on Genetic and evolutionary computation. ACM, pp. 539–546.
- Sutton, R. S., 1988. Learning to Predict by the Method of Temporal Differences. *Machine Learning* 3, 9–44.
- Sutton, R. S., Barto, A. G., 1998. *Reinforcement Learning: An Introduction*. MIT Press.
- Tang, Y., Kucukelbir, A., 2017. Variational deep q network. arXiv preprint arXiv:1711.11225.
- Theodorou, E., Buchli, J., Schaal, S., 2010. A generalized path integral control approach to reinforcement learning. *Journal of Machine Learning Research* 11, 3137–3181.
- Thrun, S., Mitchell, T. M., 1995. Lifelong robot learning. *Robotics and autonomous systems* 15 (1-2), 25–46.
- Togelius, J., Schaul, T., Wierstra, D., Igel, C., Gomez, F., Schmidhuber, J., 2009. Ontogenetic and phylogenetic reinforcement learning. *Künstliche Intelligenz* 23 (3), 30–33.
- Veeriah, V., Oh, J., Singh, S., 2018. Many-goals reinforcement learning. arXiv preprint arXiv:1806.09605.
- Vezhnevets, A. S., Osindero, S., Schaul, T., Heess, N., Jaderberg, M., Silver, D., Kavukcuoglu, K., 2017. Feudal networks for hierarchical reinforcement learning. arXiv preprint arXiv:1703.01161.
- Wang, J. X., Kurth-Nelson, Z., Tirumala, D., Soyer, H., Leibo, J. Z., Munos, R., Blundell, C., Kumaran, D., Botvinick, M., 2016a. Learning to reinforcement learn. arXiv preprint arXiv:1611.05763.
- Wang, Z., Bapst, V., Heess, N., Mnih, V., Munos, R., Kavukcuoglu, K., de Freitas, N., 2016b. Sample efficient actor-critic with experience replay. arXiv preprint arXiv:1611.01224.
- Wierstra, D., Schaul, T., Peters, J., Schmidhuber, J., 2008. Natural evolution strategies. In: *IEEE Congress on Evolutionary Computation*. IEEE, pp. 3381–3387.
- Williams, J. K., Singh, S. P., 1998. Experimental results on learning stochastic memoryless policies for partially observable markov decision processes. In: *NIPS*. pp. 1073–1080.

- Williams, R. J., May 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* 8 (3-4), 229–256.
- Wilson, A., Fern, A., Tadepalli, P., 2014. Using trajectory data to improve bayesian optimization for reinforcement learning. *The Journal of Machine Learning Research* 15 (1), 253–282.
- Wu, Y., Mansimov, E., Liao, S., Grosse, R., Ba, J., 2017. Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. *arXiv preprint arXiv:1708.05144*.
- Yang, Y., Hospedales, T. M., 2014. A unified perspective on multi-domain and multi-task learning. *arXiv preprint arXiv:1412.7489*.
- Zhang, X., Clune, J., Stanley, K. O., 2017. On the relationship between the openai evolution strategy and stochastic gradient descent. *arXiv preprint arXiv:1712.06564*.
- Zimmer, M., Doncieux, S., 2017. Bootstrapping Q-learning for robotics from neuro-evolution results. *IEEE Transactions on Cognitive and Developmental Systems*.