

---

# From Complexity to Simplicity: Adaptive ES-Active Subspaces for Blackbox Optimization

---

**Krzysztof Choromanski\***  
Google Brain Robotics  
kchoro@google.com

**Aldo Pacchiano\***  
UC Berkeley  
pacchiano@berkeley.edu

**Jack Parker-Holder\***  
Columbia University  
jh3764@columbia.edu

**Yunhao Tang\***  
Columbia University  
yt2541@columbia.edu

## Abstract

We present a new algorithm (ASEBO) for optimizing high-dimensional blackbox functions. ASEBO adapts to the geometry of the function and learns optimal sets of sensing directions, which are used to probe it, on-the-fly. It addresses the exploration-exploitation trade-off of blackbox optimization with expensive blackbox queries by continuously learning the bias of the lower-dimensional model used to approximate gradients of smoothings of the function via compressed sensing and contextual bandits methods. To obtain this model, it leverages techniques from the emerging theory of active subspaces [8] in the novel ES blackbox optimization context. As a result, ASEBO learns the dynamically changing intrinsic dimensionality of the gradient space and adapts to the hardness of different stages of the optimization without external supervision. Consequently, it leads to more sample-efficient blackbox optimization than state-of-the-art algorithms. We provide theoretical results and test ASEBO advantages over other methods empirically by evaluating it on the set of reinforcement learning policy optimization tasks as well as functions from the recently open-sourced Nevergrad library.

## 1 Introduction

Consider a high-dimensional function  $F : \mathbb{R}^d \rightarrow \mathbb{R}$ . We assume that querying it is expensive. Examples include reinforcement learning (RL) blackbox functions taking as inputs vectors  $\theta$  encoding policies  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  mapping states to actions and outputting total (expected/discounted) rewards obtained by agents applying  $\pi$  in given environments [6]. For this class of functions evaluations usually require running a simulator. Other examples include wind configuration design optimization problems for high speed civil transport aircrafts, optimizing computer codes (e.g. NASA synthetic tool FLOPS/ENGENN used to size the aircraft and propulsion system [2]), crash tests, medical and chemical reaction experiments [37].

Evolution strategy (ES) methods have traditionally been used in low-dimensional regimes (e.g. hyperparameter tuning), and considered ill-equipped for higher dimensional problems due to poor sampling complexity [27]. However, a flurry of recent work has shown they can scale better than previously believed [33, 11, 29, 25, 7, 30, 21]. This is thanks to a couple of reasons.

First of all, new ES methods apply several efficient heuristics (filtering, various normalization techniques as in [25] and new exploration strategies as in [11]) in order to substantially improve sampling complexity. Other recent methods [29, 7] are based on more accurate Quasi Monte Carlo (MC) estimators of the gradients of Gaussian smoothings of blackbox functions with theoretical

guarantees. These approaches provide better quality gradient sensing mechanisms. Additionally, in applications such as RL, new compact structured policy architectures (such as low-displacement rank neural networks from [7] or even linear policies [14]) are used to reduce the number of policies’ parameters and dimensionality of the optimization problem.

Recent research also shows that ES-type blackbox optimization in RL leads to more stable policies than policy gradient methods since ES methods search for parameters that are robust to perturbations [19]. Unlike policy gradient methods, ES aims to find parameters maximizing expected reward (rather than just a reward) in respect to Gaussian perturbations.

Finally, pure ES methods as opposed to state-of-the-art policy optimization techniques (TRPO, PPO or ARS [32, 15, 31, 25]), can be applied also for blackbox optimization problems that do not exhibit MDP structure required for policy gradient methods and cannot benefit from state normalization algorithm central to ARS. This has led to their recent popularity for non-differentiable tasks [17, 14].

In this paper we introduce a new adaptive sample-efficient blackbox optimization algorithm. ASEBO adapts to the geometry of blackbox functions and learns optimal sets of sensing directions, which are used to probe them, on-the-fly. To do this, it leverages techniques from the emerging theory of active subspaces [8, 10, 9, 20] in a novel ES blackbox optimization context. Active subspaces and their extensions are becoming popular as effective techniques for dimensionality reduction (see for instance: active manifolds [5] or ResNets for learning isosurfaces [36]). However, to the best of our knowledge we are the first to apply active subspace ideas for ES optimization.

ASEBO addresses the exploration-exploitation trade-off of blackbox optimization with expensive function queries by continuously learning the bias of the lower-dimensional model used to approximate gradients of smoothings of the function with compressed sensing and contextual bandits methods. The adaptiveness is what distinguishes it from some recently introduced guided ES methods such as [24] that rely on fixed hyperparameters that are hard to tune in advance (e.g. the length of the buffer defining lower dimensional space for gradient search). We provide theoretical results and empirically evaluate ASEBO on a set of RL blackbox optimization tasks as well as non-RL blackbox functions from the recently open-sourced Nevergrad library [34], showing that it consistently learns optimal inputs with fewer queries to a blackbox function than other methods.

**ASEBO versus CMA-ES:** There have been a variety of works seeking to reduce sampling complexity for ES methods through the use of metric learning. The prominent class of the covariance matrix adaptation evolution strategy (CMA-ES) methods derives state-of-the-art derivative free blackbox optimization algorithms, which seek to learn and maintain a fully parameterized Gaussian distribution. CMA-ES suffers from quadratic time complexity for each evaluation which can be limiting for high dimensional problems. As such, a series of attempts have been made to produce scalable variants of CMA-ES, by restricting the covariance matrix to the diagonal (sep-CMA-ES [28]) or a low rank approximation as in VD-CMA-ES [3] and LM-CMA-ES [22]. Two recent algorithms, Vd-CMA-ES [4] and LM-MA-ES [23], seek to combine the above ideas and have been shown to be successful in large-scale settings, including RL policy learning [26]. Although these methods are able to quickly learn and adapt the covariance matrix, they are heavily dependent on hyperparameter selection [4, 35] and lack the means to avoid learning a bias. As our experiments show, this can severely hurt their performance. The best CMA-ES variants often struggle with RL tasks of challenging objective landscapes, displaying inconsistent performance across tasks. Furthermore, they require careful hyperparameter tuning for good performance (see: analysis in Section 4, Fig. 3).

## 2 Adaptive Sample-Efficient Blackbox Optimization

Before we describe ASEBO, we explain key theoretical ideas behind the algorithm. ASEBO uses online PCA to maintain and update on-the-fly subspaces which we call *ES-active subspaces*  $\mathcal{L}_{\text{active}}^{\text{ES}}$ , accurately approximating the gradient data space at a given phase of the algorithm. The bias of the obtained gradient estimators is measured by sensing the length of its component from the orthogonal complement  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$  via compressed sensing or computing optimal probabilities for exploration (e.g. sensing from  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$ ) via contextual bandits methods [1]. The algorithm corrects its probabilistic distributions used for choosing directions for gradient sensing based on these measurements. As we show, we can measure that bias accurately using only a constant number of additional function queries, regardless of the dimensionality. This in turn determines an exploration strategy, as we explain later. Estimated gradients are then used to update parameters.




Figure 1: The motivation behind ASEBO. Two first plots: ES baseline for HalfCheetah and Swimmer tasks from the OpenAI Gym library for 212-dimensional policies - the plot shows how the dimensionality of the space capturing a given percentage of variance of approximate gradient data depends on the iteration of the algorithm. This information is never exploited by the algorithm, even though 99.5% of the variance resides in the much lower-dimensional space (100 dimensions). Two last plots: ASEBO taking advantage of this information (# of sample/sensing directions reflects the hardness of the optimization at each iteration and is strongly correlated with the PCA dimensionality).

## 2.1 Preliminaries

Consider a blackbox function  $F : \mathbb{R}^d \rightarrow \mathbb{R}$ . We do not assume that  $F$  is differentiable. The *Gaussian smoothing* [27]  $F_\sigma$  of  $F$  parameterized by smoothing parameter  $\sigma > 0$  is given as:

$F_\sigma(\theta) = \mathbb{E}_{\mathbf{g} \in \mathcal{N}(0, \mathbf{I}_d)}[F(\theta + \sigma \mathbf{g})] = (2\pi)^{-\frac{d}{2}} \int_{\mathbb{R}^d} F(\theta + \sigma \mathbf{g}) e^{-\frac{\|\mathbf{g}\|^2}{2}} d\mathbf{g}$ . The gradient of the Gaussian smoothing of  $F$  is given by the formula:

$$\nabla F_\sigma(\theta) = \frac{1}{\sigma} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_d)}[F(\theta + \sigma \mathbf{g})\mathbf{g}]. \quad (1)$$

Formula 1 on  $\nabla F_\sigma(\theta)$  leads straightforwardly to several unbiased Monte Carlo (MC) estimators of  $\nabla F_\sigma(\theta)$ , where the most popular ones are: the *forward finite difference* estimator [7] defined as:  $\widehat{\nabla}_{\text{MC}}^{\text{FD}} F_\sigma(\theta) = \frac{1}{k\sigma} \sum_{i=1}^k (F(\theta + \sigma \mathbf{g}_i) - F(\theta))\mathbf{g}_i$ , and an *antithetic ES gradient estimator* [30] given as:  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta) = \frac{1}{2k\sigma} \sum_{i=1}^k (F(\theta + \sigma \mathbf{g}_i) - F(\theta - \sigma \mathbf{g}_i))\mathbf{g}_i$ , where typically  $\mathbf{g}_1, \dots, \mathbf{g}_k$  are taken independently at random from  $\mathcal{N}(0, \mathbf{I}_d)$  or from more complex joint distributions for variance reduction (see: [7]). We call samples  $\mathbf{g}_1, \dots, \mathbf{g}_k$  the *sensing directions* since they are used to sense gradients  $\nabla F_\sigma(\theta)$ . The antithetic formula can be alternatively rationalized as giving the renormalized gradient of  $F$  (if  $F$  is smooth), if not taking into account cubic and higher-order terms of the Taylor expansion  $F(\theta + \mathbf{v}) = F(\theta) + \nabla F^\top \mathbf{v} + \frac{1}{2} \mathbf{v}^\top H(\theta) \mathbf{v}$  (where  $H(\theta)$  stands for the Hessian of  $F$  in  $\theta$ ).

Standard ES methods apply different gradient-based techniques such as SGD or Adam, fed with the above MC estimators of  $\nabla F_\sigma$  to conduct blackbox optimization. The number of samples  $k$  per iteration of the optimization procedure is usually of the order  $O(d)$ . This becomes a computational bottleneck for high-dimensional blackbox functions  $F$  (for instance, even for relatively small RL tasks with policies encoded by compact neural networks we still have  $d > 100$  parameters).

## 2.2 ES-active subspaces via online PCA with decaying weights

The first idea leading to the ASEBO algorithm is that in practice one does not need to estimate the gradient of  $F$  accurately (after all ES-type methods do not even aim to compute the gradient of  $F$ , but rather focus on  $\nabla F_\sigma$ ). Poor scalability of ES-type blackbox optimization algorithms is caused by high-dimensionality of the gradient vector. However, during the optimization process the space spanned by gradients may be locally well approximated by a lower-dimensional subspace  $\mathcal{L}$  and sensing the gradient in that subspace might be more effective. In some recent papers such as [24] such a subspace is defined simply as  $\mathcal{L} = \text{span}\{\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_i), \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_{i-1}), \dots, \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_{i-s+1})\}$ , where  $\{\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_i), \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_{i-1}), \dots, \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_{i-s+1})\}$  stands for the batch of last  $s$  approximated gradients during the optimization process and  $s$  is a fixed hyperparameter. Even though  $\mathcal{L}$  will dynamically change during the optimization, such an approach has several disadvantages in practice. Tuning parameter  $s$  is very difficult or almost impossible and the assumption that the dimensionality of  $\mathcal{L}$  should be constant during optimization is usually false. In our approach, dimensionality of  $\mathcal{L}$  varies and depends on the hardness of the optimization in different optimization stages.

We apply Principal Component Analysis (PCA, [18]) to create a subspace  $\mathcal{L}$  capturing particular variance  $\sigma > 0$  of the approximate gradients data. This data is either: the approximate gradients

computed in previous iterations from the antithetic formula or: the elements of the sum from that equation that are averaged over to obtain these gradients. For clarity of the exposition, from now on we will assume the former, but both variants are valid. Choosing  $\sigma$  is in practice much easier than  $s$  and leads to subspaces  $\mathcal{L}$  of varying dimensionalities throughout the optimization procedure, called by us from now on *ES-active subspaces*  $\mathcal{L}_{\text{active}}^{\text{ES}}$ .

---

**Algorithm 1** ASEBO Algorithm

---

**Hyperparameters:** number of iterations of full sampling  $l$ , smoothing parameter  $\sigma > 0$ , step size  $\eta$ , PCA threshold  $\epsilon$ , decay rate  $\gamma$ , total number of iterations  $T$ .

**Input:** blackbox function  $F$ , vector  $\theta_0 \in \mathbb{R}^d$  where optimization starts.  $\text{Cov}_0 \in \{0\}^{d \times d}$ ,  $p^0 = 0$ .

**Output:** vector  $\theta_T$ .

**for**  $t = 0, \dots, T - 1$  **do**

**if**  $t < l$  **then**

    Take  $n_t = d$ . Sample  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  from  $\mathcal{N}(0, \mathbf{I}_d)$  (independently).

**else**

1. Take top  $r$  eigenvalues  $\lambda_i$  of  $\text{Cov}_t$ , where  $r$  is smallest such that:  $\sum_{i=1}^r \lambda_i \geq \epsilon \sum_{i=1}^d \lambda_i$ , using its SVD as described in text and take  $n_t = r$ .
2. Take the corresponding eigenvectors  $\mathbf{u}_1, \dots, \mathbf{u}_r \in \mathbb{R}^d$  and let  $\mathbf{U} \in \mathbb{R}^{d \times r}$  be obtained by stacking them together. Let  $\mathbf{U}^{\text{act}} \in \mathbb{R}^{d \times r}$  be obtained from stacking together some orthonormal basis of  $\mathcal{L}_{\text{active}}^{\text{ES}} \stackrel{\text{def}}{=} \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_r\}$ . Let  $\mathbf{U}^\perp \in \mathbb{R}^{d \times (d-r)}$  be obtained from stacking together some orthonormal basis of the orthogonal complement  $\mathcal{L}_{\text{active}}^{\text{ES}, \perp}$  of  $\mathcal{L}_{\text{active}}^{\text{ES}}$ .
3. Sample  $n_t$  vectors  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  as follows: with probability  $1 - p^t$  from  $\mathcal{N}(0, \mathbf{U}^\perp (\mathbf{U}^\perp)^\top)$  and with probability  $p^t$  from  $\mathcal{N}(0, \mathbf{U}^{\text{act}} (\mathbf{U}^{\text{act}})^\top)$ .
4. Renormalize  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  such that marginal distributions  $\|\mathbf{g}_i\|_2$  are  $\chi(d)$ .

1. Compute  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t)$  as:  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t) = \frac{1}{2n_t\sigma} \sum_{j=1}^{n_t} (F(\theta_t + \mathbf{g}_j) - F(\theta_t - \mathbf{g}_j)) \mathbf{g}_j$ .
  2. Set  $\text{Cov}_{t+1} = \lambda \text{Cov}_t + (1 - \lambda) \Gamma$ , where  $\Gamma = \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t) (\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t))^\top$ .
  3. Set  $p^{t+1} = p_{\text{opt}}$  for  $p_{\text{opt}}$  output by Algorithm 2 and:  $\theta_{t+1} \leftarrow \theta_t + \eta \widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t)$ .
- 

These will be in turn applied to define covariance matrices encoding probabilistic distributions applied to construct sensing directions used for estimating  $\nabla F_\sigma(\theta)$ . The additional advantage of our approach is that PCA automatically filters out gradient noise.

We use our own online version of PCA with decaying weights (decay rate is defined by parameter  $\lambda > 0$ ). By tuning  $\lambda$  we can define the rate at which historical approximate gradient data is used to choose the right sensing directions, which will continuously decay. We consider a stream of approximate gradients  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_0), \dots, \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_i), \dots$  obtained during the optimization procedure. We maintain and update on-the-fly the covariance matrix  $\text{Cov}_t$ , where  $t$  stands for the number of completed iterations, in the form of the symmetric matrix SVD-decomposition  $\text{Cov}_t = \mathbf{Q}_t^\top \Sigma_t \mathbf{Q}_t \in \mathbb{R}^d$ . When the new approximate gradient  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t)$  arrives, the update of the covariance matrix is driven by the following equation, reflecting data decay process, where  $\mathbf{x}_t = \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t)$ :

$$\mathbf{Q}_{t+1}^\top \Sigma_{t+1} \mathbf{Q}_{t+1} = \lambda \mathbf{Q}_t^\top \Sigma_t \mathbf{Q}_t + (1 - \lambda) \mathbf{x}_t \mathbf{x}_t^\top, \quad (2)$$

To conduct the update cheaply, it suffices to observe that the RHS of Equation 2 can be rewritten as:  $\lambda \mathbf{Q}_t^\top \Sigma_t \mathbf{Q}_t + (1 - \lambda) \mathbf{x}_t \mathbf{x}_t^\top = \mathbf{Q}_t^\top (\lambda \Sigma_t + (1 - \lambda) \mathbf{Q}_t \mathbf{x}_t (\mathbf{Q}_t \mathbf{x}_t)^\top) \mathbf{Q}_t$ . Now, using the fact that for a matrix of the form  $\mathbf{D} + \mathbf{u} \mathbf{u}^\top$ , we can get its decomposition in time  $O(d^2)$  [13], we obtain an algorithm performing updates in quadratic time. That in practice suffices since the bottleneck of the computations is in querying  $F$  and additional overhead related to updating  $\mathcal{L}_{\text{active}}^{\text{ES}}$  is negligible.

**ES-active subspaces versus active subspaces:** Our mechanism for constructing  $\mathcal{L}_{\text{active}}^{\text{ES}}$  is inspired by the recent theory of active subspaces [8], developed to determine the most important directions in the space of input parameters of high-dimensional blackbox functions such as computer simulations.

The *active subspace* of a differentiable function  $F : \mathbb{R}^d \rightarrow \mathbb{R}$ , square-integrable with respect to the given probabilistic density function  $\rho : \mathbb{R}^d \rightarrow \mathbb{R}$ , is given as a linear subspace  $\mathcal{L}_{\text{active}}$  defined by the first  $r$  (for a fixed  $r < d$ ) eigenvectors of the following  $d \times d$  symmetric positive definite matrix:

$$\text{Cov} = \int_{\mathbf{x} \in \mathbb{R}^d} \nabla F(\mathbf{x}) \nabla F(\mathbf{x})^\top \rho(\mathbf{x}) d\mathbf{x} \quad (3)$$

Density function  $\rho$  determines where compact representation of  $F$  is needed. In our approach we do not assume that  $\nabla F$  exists, but the key difference between  $\mathcal{L}_{\text{active}}^{\text{ES}}$  and  $\mathcal{L}_{\text{active}}$  lies somewhere else.

The goal of ASEBO is to avoid approximating the exact gradient  $\nabla F(\mathbf{x}) \in \mathbb{R}^d$  which is what makes standard ES methods very expensive and which is done in [9] via gradient sketching techniques combined with finite difference approaches (standard methods of choice for ES baselines).

---

**Algorithm 2** Explore estimator via exponentiated sampling

---

**Hyperparameters:** smoothing parameter  $\sigma$ , horizon  $C$ , learning rate  $\alpha$ , probability regularizer  $\beta$ , initial probability parameter  $q_0^t \in (0, 1)$ .

**Input:** subspaces:  $\mathcal{L}_{\text{active}}^{\text{ES}}$ ,  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$ , function  $F$ , vector  $\theta_t$

**Output:**

**for**  $l = 1, \dots, C + 1$  **do**

1. Compute  $p_{l-1}^t = (1 - 2\beta)q_{l-1}^t + \beta$  and sample  $a_l^t \sim \text{Ber}(p_{l-1}^t)$ .
3. If  $a_l^t = 1$ , sample  $\mathbf{g}_l \sim \mathcal{N}(0, \sigma \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}})$ , otherwise sample  $\mathbf{g}_l \sim \mathcal{N}(0, \sigma \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES},\perp}})$ .
4. Compute  $v_l = \frac{1}{2\sigma} (F(\theta_t + \mathbf{g}_l) - F(\theta_t - \mathbf{g}_l))$ .
5. Set  $\mathbf{e}_l = (1 - 2\beta) \begin{bmatrix} \left( \frac{-a_l^t(\dim(\mathcal{L}_{\text{active}}^{\text{ES}})+2)}{(p_{l-1}^t)^3} \right) \\ \left( \frac{-(1-a_l^t)(\dim(\mathcal{L}_{\text{active}}^{\text{ES},\perp})+2)}{(1-p_{l-1}^t)^3} \right) \end{bmatrix} v_l^2$ .
6. Set  $q_l^t = \frac{q_{l-1}^t \exp(-\alpha \mathbf{e}_l(1))}{q_{l-1}^t \exp(-\alpha \mathbf{e}_l(1)) + (1 - q_{l-1}^t) \exp(-\alpha \mathbf{e}_l(2))}$ .

**Return:**  $p_C$ .

---

Instead, in ASEBO an ES-active subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$  is itself used to define sensing directions and the number of chosen samples  $k$  is given by the dimensionality of  $\mathcal{L}_{\text{active}}^{\text{ES}}$ . This drastically reduces sampling complexity, but comes at a price of risking the optimization to be trapped in the fixed lower-dimensional space that will not be representative for gradient data as optimization progresses. We propose a solution requiring only a constant number of extra queries to  $F$  in the next sections.

### 2.3 Exploration-exploitation trade-off: Adaptive Exploration Mechanism

The procedure described above needs to be accompanied with an exploration strategy that will determine how frequently to choose sensing directions outside the constructed on-the-fly lower-dimensional ES-subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$ . Our exploration strategies will be encoded by hybrid probabilistic distributions for sampling sensing directions. The frequency of sensing from the distributions with covariance matrices obtained from  $\mathcal{L}_{\text{active}}^{\text{ES}}$  (corresponding to exploitation) and from its orthogonal complement  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$  or entire space (corresponding to exploration) will be given by weights encoding the importance of exploitation versus exploration in any given iteration of the optimization. For a vector  $\mathbf{x} \in \mathbb{R}^d$  denote by  $\mathbf{x}_{\text{active}}$  its projection onto  $\mathcal{L}_{\text{active}}^{\text{ES}}$  and by  $\mathbf{x}_{\perp}$  its projection onto  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$ . The useful metric that can be used to update the above weights in an online manner in the  $t^{\text{th}}$  iteration of the algorithm is the ratio:  $r = \frac{\|(\nabla F_{\sigma}(\theta_t))_{\text{active}}\|_2}{\|(\nabla F_{\sigma}(\theta_t))_{\perp}\|_2}$ . Smaller values of  $r$  indicate that current active subspace is not representative enough for the gradient and more aggressive exploration needs to be conducted. In practice, we do not compute  $r$  explicitly, but rather its approximated version  $\hat{r}$ .

One can simply take:  $\hat{r} = \frac{\|(\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_{\sigma}(\theta_{t-1}))_{\text{active}}\|_2}{\|(\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_{\sigma}(\theta_{t-1}))_{\perp}\|_2}$ , where  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_{\sigma}(\theta_{t-1})$  is obtained in the previous iteration. But we can do better. It suffices to separately estimate  $\|(\nabla F_{\sigma}(\theta_t))_{\text{active}}\|_2$  and  $\|(\nabla F_{\sigma}(\theta_t))_{\perp}\|_2$ . However we do not aim to estimate  $(\nabla F_{\sigma}(\theta_t))_{\text{active}}$  and  $(\nabla F_{\sigma}(\theta_t))_{\perp}$ . That would be equivalent to computing exact estimate of  $\nabla F_{\sigma}(\theta_t)$ , defeating the purpose of ASEBO. Instead, we note that estimating the length of the unknown high-dimensional vector is much simpler than estimating the vector itself and can be done in the probabilistic manner with arbitrary precision via the set of dot-product queries of size independent from dimensionality  $d$  via compressed sensing methods. We refine this approach and propose more accurate contextual bandits method that also relies on dot-product queries applied in the ES-context, but aims to directly approximate optimal probabilities of sampling from  $\mathcal{L}_{\text{active}}^{\text{ES}}$  rather than approximating gradients components' lengths (see Algorithm 2 box, the compressed sensing baseline is presented in the Appendix). The related computational overhead is measured in constant number of extra function queries, negligible in practice.

## 2.4 The Algorithm

ASEBO is given in the Algorithm 1 box. The algorithm we apply to score relative importance of sampling from the ES-active subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$  versus from outside  $\mathcal{L}_{\text{active}}^{\text{ES}}$  is in the Algorithm 2 box.

As we have already mentioned, it uses bandits method do determine optimal probability of sampling from  $\mathcal{L}_{\text{active}}^{\text{ES}}$ . In the next section we show that by using these techniques we can substantially reduce the variance of ES blackbox gradient estimators if ES-active subspaces approximate the gradient data well (which is the case for RL blackbox functions as presented in Fig. 1). Horizon lengths  $C$  in Algorithm 2 which determines the number of extra function queries should be in practice chosen as small constants. In each iteration of Algorithm 1 the number of function queries is proportional to the dimensionality of the ES-active subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$  rather than the original space.

## 3 Theoretical Results

We provide here theoretical guarantees for the ASEBO sampling mechanism (in Algorithm 1), where sensing directions  $\{\mathbf{g}_i\}$  at time  $t$  are sampled from the hybrid distribution  $\hat{P}$ : with probability  $p^t$  from  $\mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}})$  and with probability  $1 - p^t$  from  $\mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\perp}})$ .

Following notation in Algorithm 1, let  $\mathbf{U}^{\text{act}} \in \mathbb{R}^{d \times r}$  be obtained by stacking together vectors of some orthonormal basis of  $\mathcal{L}_{\text{active}}^{\text{ES}}$ , where  $\dim(\mathcal{L}_{\text{active}}^{\text{ES}}) = r$  and let  $\mathbf{U}^{\perp} \in \mathbb{R}^{d \times (d-r)}$  be obtained by stacking together vectors of some orthonormal basis of its orthogonal complement  $\mathcal{L}_{\text{active}}^{\text{ES}, \perp}$ . Denote by  $\sigma$  a smoothing parameter. We make the following regularity assumptions on  $F$ :

**Assumption 1.**  $F$  is  $L$ -Lipschitz, i.e. for all  $\theta, \theta' \in \mathbb{R}^d$ ,  $|F(\theta) - F(\theta')| \leq L\|\theta - \theta'\|_2$ .

**Assumption 2.**  $F$  has a  $\tau$ -smooth third order derivative tensor with respect to  $\sigma > 0$ , so that  $F(\theta + \sigma \mathbf{g}) = F(\theta) + \sigma \nabla F(\theta)^\top \mathbf{g} + \frac{\sigma^2}{2} \mathbf{g}^\top H(\theta) \mathbf{g} + \frac{1}{6} \sigma^3 F'''(\theta)[\mathbf{v}, \mathbf{v}, \mathbf{v}]$  for some  $\mathbf{v} \in \mathbb{R}^d$  ( $\|\mathbf{v}\|_2 \leq \|\mathbf{g}\|_2$ ) satisfying  $|F'''(\theta)[\mathbf{v}, \mathbf{v}, \mathbf{v}]| \leq \tau \|\mathbf{v}\|_2^3 \leq \tau \|\mathbf{g}\|_2^3$ .

Observe that:  $\mathbb{E}_{\mathbf{g} \sim \hat{P}}[\mathbf{g} \mathbf{g}^\top] = (p^t \mathbf{U}^{\text{act}} (\mathbf{U}^{\text{act}})^\top + (1 - p^t) \mathbf{U}^{\perp} (\mathbf{U}^{\perp})^\top)$ . Define  $C_1 = (p^t \mathbf{U}^{\text{act}} (\mathbf{U}^{\text{act}})^\top + (1 - p^t) \mathbf{U}^{\perp} (\mathbf{U}^{\perp})^\top)$ . Let  $\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{asebo}} F_\sigma(\theta) = C_1^{-1} \frac{F(\theta + \sigma \mathbf{g}) \mathbf{g} + F(\theta + \sigma \mathbf{g})(-\mathbf{g})}{2\sigma}$  be the gradient estimator corresponding to  $\hat{P}$ . We will assume that  $\sigma$  is small enough, i.e.  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}$  for some precision parameter  $\epsilon > 0$ . Our first result shows that under these assumptions, baseline and ASEBO estimators of  $\nabla_\sigma F(\theta)$  are also good estimators of  $\nabla F(\theta)$ :

**Lemma 3.1.** *If  $F$  satisfies Assumptions 1 and 2, the estimators  $\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{base}} F_\sigma(\theta)$  and  $\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{asebo}} F_\sigma(\theta)$  are close to the true gradient  $\nabla F(\theta)$ , i.e.:  $\left\| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_d)} \left[ \hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{base}} F_\sigma(\theta) \right] - \nabla F(\theta) \right\| \leq \epsilon$  and  $\left\| \mathbb{E}_{\mathbf{g} \sim \hat{P}} \left[ \hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{asebo}} F_\sigma(\theta) \right] - \nabla F(\theta) \right\| \leq \epsilon$ .*

### 3.1 Variance reduction via non isotropic sampling

We show now that under sampling strategy given by distribution  $\hat{P}$ , the variance of the gradient estimator can be made smaller by choosing the probability parameter  $p^t$  appropriately. Denote:  $d_{\text{active}} = \dim(\mathcal{L}_{\text{active}}^{\text{ES}})$  and  $d_{\perp} = \dim(\mathcal{L}_{\text{active}}^{\text{ES}, \perp})$ . Let  $\Gamma := \frac{d_{\text{active}} + 2}{p^t} s_{\mathbf{U}^{\text{act}}} + \frac{d_{\perp} + 2}{1 - p^t} s_{\mathbf{U}^{\perp}} - \|\nabla F(\theta)\|^2$ .

**Theorem 3.2.** *The following holds for  $s_{\mathbf{U}^{\text{act}}} = \|(\mathbf{U}^{\text{act}})^\top \nabla F(\theta)\|_2^2$  and  $s_{\mathbf{U}^{\perp}} = \|(\mathbf{U}^{\perp})^\top \nabla F(\theta)\|_2^2$ :*

1. *The variance of  $\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{asebo}} F_\sigma(\theta)$  is close to  $\Gamma$ , i.e.  $|\text{Var}[\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{asebo}} F_\sigma(\theta)] - \Gamma| \leq \epsilon$ .*
2. *The choice of  $p^t$  that minimizes  $\Gamma$  satisfies  $p_*^t := \frac{\sqrt{(s_{\mathbf{U}^{\text{act}}})(d_{\text{active}} + 2)}}{\sqrt{(s_{\mathbf{U}^{\text{act}}})(d_{\text{active}} + 2)} + \sqrt{(s_{\mathbf{U}^{\perp}})(d_{\perp} + 2)}}$  and the optimal variance  $\text{Var}_{\text{opt}}$  corresponding to  $p_*^t$  satisfies:  $|\text{Var}_{\text{opt}} - \Delta| \leq \epsilon$  for  $\Delta = \left[ \sqrt{(s_{\mathbf{U}^{\text{act}}})(d_{\text{active}} + 2)} + \sqrt{(s_{\mathbf{U}^{\perp}})(d_{\perp} + 2)} \right]^2 - \|\nabla F(\theta)\|^2$ .*
3.  $\text{Var}_{\text{opt}} \leq \text{Var}[\hat{\nabla}_{\text{MC}, k=1}^{\text{AT}, \text{base}} F_\sigma(\theta)] + \epsilon - \underbrace{\left| \sqrt{(s_{\mathbf{U}^{\perp}})(d_{\text{active}} + 2)} - \sqrt{(s_{\mathbf{U}^{\text{act}}})(d_{\perp} + 2)} \right|^2}_{\lambda} - 2\|\nabla F(\theta)\|^2$ .

Furthermore, slack variable  $\lambda$  is always nonnegative.

Theorem implies that when  $s_{\text{Uact}} = (1 - \alpha)\|\nabla F(\theta)\|_2^2$  and  $s_{\text{U}\perp} = \alpha\|\nabla F(\theta)\|_2^2$ , for some  $\alpha \in (0, 1)$ , we have:  $\text{Var}[\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, base}} F_\sigma(\theta)] \approx (d + 1)\|\nabla F(\theta)\|^2$  whereas  $\text{Var}_{\text{opt}} = \mathcal{O}((1 - \alpha)(d_{\text{active}} + 1) + \alpha(d_{\perp} + 1))$ . When  $d_{\text{active}} \ll d$  and  $\alpha \ll 1$ :  $\text{Var}_{\text{opt}} \ll \text{Var}[\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, base}} F_\sigma(\theta)]$ .

### 3.2 Adaptive Mirror Descent

In Theorem 3.2 we showed that for appropriate choices of  $\mathcal{L}_{\text{active}}^{\text{ES}}$  and  $p_t$ , the gradient estimator  $\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, asebo}} F_\sigma(\theta)$  will have significantly smaller variance than  $\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, base}} F_\sigma(\theta)$ . In this section we show that Algorithm 2 provides an adaptive way to choose  $p^t$ . Using tools from online learning theory, we provide regret guarantees that quantify the rate at which this Algorithm 2 minimizes the variance of  $\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, asebo}} F_\sigma(\theta)$  and converges to the optimal  $p_*^t$ .

Let  $\mathbf{p}_i^t = \begin{pmatrix} p_i^t \\ 1 - p_i^t \end{pmatrix}$ . The main component  $\Gamma$  of the variance of  $\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, asebo}} F_\sigma(\theta)$  as a function of  $\mathbf{p}_i^t$  equals  $\Gamma = \ell(\mathbf{p}_i^t) = \frac{d_{\text{active}} + 2}{\mathbf{p}_i^t(1)} s_{\text{Uact}} + \frac{d_{\perp} + 2}{\mathbf{p}_i^t(2)} s_{\text{U}\perp} - \|\nabla F(\theta)\|^2$  (Theorem 3.2). We have:

**Theorem 3.3.** Let  $\Delta_2$  be the a 2-d simplex. Under assumptions: 1 and 2, if  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1 - p^t)}{\tau d^3 \max(L, 1)}}$ ,  $\alpha = \frac{2\beta^2}{\sqrt{C[(d_{\text{active}} + 2)^2 s_{\text{Uact}}^2 + (d_{\perp} + 2)s_{\text{U}\perp}^2]}}$  and  $\epsilon = \frac{\beta^3}{2C(d+1)}$ , Algorithm 2 satisfies:

$$\frac{1}{C} \mathbb{E} \left[ \sum_{l=1}^C \ell(\mathbf{p}_l^t) \right] - \min_{\mathbf{p} \in \beta + (1 - 2\beta)\Delta_2} \ell(\mathbf{p}) \leq \frac{\text{Var}_{\text{opt}}}{\beta^2 \sqrt{C}} + \frac{1}{C}$$

## 4 Experiments

In our experiments we use different classes of high-dimensional blackbox functions: RL blackbox functions (where the input is a high-dimensional vector encoding a neural network policy  $\pi : \mathcal{S} \rightarrow \mathcal{A}$  mapping states  $s$  to actions  $a$  and the output is the cumulative reward obtained by an agent applying this policy in a particular environment) and functions from the recently open-sourced Nevergrad library [34]. In practice one can setup the hyperparameters used by Algorithm 2 as follows:  $\sigma = 0.01$ ,  $C = 10$ ,  $\alpha = 0.01$ ,  $\beta = 0.1$ ,  $q_0^t = 0.1$ . For each algorithm we used  $k = 5$  seeds and obtained curves are median-curves with inter-quartile ranges presented as shadowed regions.




Figure 2: Comparison of different blackbox optimization algorithms on OpenAI Gym tasks. All curves are median-curves obtained from  $k = 5$  seeds and with inter-quartile ranges presented as shadowed regions. For Reacher we present only 3 curves since LM-MA-ES and TRPO did not learn.

### 4.1 RL blackbox functions

We used the following environments from the OpenAI Gym library: Swimmer-v2, HalfCheetah-v2, Walker2d-v2, Reacher-v2, Pusher-v2 and Thrower-v2. In all experiments we used policies encoded by neural network architectures of two hidden layers and with tanh nonlinearities, with  $> 100$  parameters. For gradient-based optimization we use Adam. For this class of blackbox functions we compared ASEBO with other generic blackbox methods as well as those specializing in optimizing RL blackbox functions  $F$ , namely: (1) CMA-ES variants; we compare against two recently introduced algorithms designed for high-dimensional settings (we use the implementation

of VxD-CMA-ES in the pycma open-source implementation from <https://github.com/CMA-ES/pycma>), and that of LM-MA-ES from [26]), (2) Augmented Random Search (ARS) [25] (we use implementation released by the authors at <http://github.com/modestyachts/ARS>), (3) Proximal Policy Optimization (PPO) [32] and Trust Region Policy Optimization (TRPO) [31] (we use OpenAI baseline implementation [12]). The results for four environments are on Fig. 2.

Table 1: Median rewards obtained across  $k = 5$  seeds for seven RL environments. For each environment the top two performing algorithms are bolded, while the bottom two are shown in red.

Environment	Timesteps	Median reward after # timesteps						
		ASEBO	ES	ARS	VxD-CMA	LM-MA	TRPO	PPO
HalfCheetah	$5 \cdot 10^7$	<b>3821</b>	1530	<b>2420</b>	-144	1632	-512	1514
Swimmer	$10^7$	<b>358</b>	36	348	<b>367</b>	297	110	52
Walker2d	$5 \cdot 10^7$	<b>9941</b>	347	1112	1	<b>18065</b>	3011	2377
Hopper	$10^7$	<b>99949</b>	626	1091	42	<b>100199</b>	1663	1310
Reacher	$10^5$	-11	-10	-12	-1391	-173	-112	-196
Pusher	$10^5$	-46	-48	-45	-1001	-467	-120	-316
Thrower	$10^5$	-89	-96	-90	-796	-737	-85	-175

Sampling complexity is measured in the number of timesteps (environment transitions) used by the algorithms. ASEBO is the only algorithm that performs consistently across all seven environments (see: Table 1), outperforming CMA-ES variants on all tasks aside from VxD-CMA-ES on Swimmer and LM-MA-ES on Walker2d. For environments such as Reacher, Thrower and Pusher, these methods perform poorly, drastically underperforming even Vanilla ES. On Fig. 3, we demonstrate the common problem of state-of-the-art CMA-ES methods: if the number of samples  $n$  is not carefully tuned, the algorithm does not learn. ASEBO does not have this problem since  $n$  is learned on-the-fly.




Figure 3: Sensitivity analysis for CMA-ES variants on the HalfCheetah (HC) and Walker2d (WA) tasks. In each setting, we run  $k = 5$  seeds, solely changing the number of samples per iteration (or population size)  $n$ .

## 4.2 Nevergrad blackbox functions

We tested functions: sphere, rastrigin, rosenbrock and lunacek (from the class of Bi-Rastrigin/Lunacek’s No.02 functions). All tested functions are 1000-dimensional. The results are presented on Fig. 4. ASEBO is the most reliable method across different functions.




Figure 4: Comparison of median-curves obtained from  $k = 5$  seeds for different algorithms on Nevergrad functions [34]. Inter-quartile ranges are presented as shadowed regions.



## 5 Conclusion

We proposed a new algorithm for optimizing high-dimensional blackbox functions. ASEBO adjusts on-the-fly the strategy of choosing gradient sensing directions to the hardness of the problem at the current stage of optimization and can be applied for both RL and non-RL problems. We provided theoretical guarantees for our method and exhaustive empirical validation.

## References

- [1] S. Agrawal, N. R. Devanur, and L. Li. Contextual bandits with global constraints and objective. *CoRR*, abs/1506.03374, 2015.
- [2] S. Ahmad and K. B. Thomas. Flight optimization system ( flops ) hybrid electric aircraft design capability. 2013.
- [3] Y. Akimoto, A. Auger, and N. Hansen. Comparison-based natural gradient optimization in high dimension. *GECCO*, 2014.
- [4] Y. Akimoto and N. Hansen. Projection-Based Restricted Covariance Matrix Adaptation for High Dimension. *GECCO*, 2016.
- [5] R. A. Bridges, A. D. Gruber, C. Felder, M. E. Verma, and C. Hoff. Active manifolds: A non-linear analogue to active subspaces. *CoRR*, abs/1904.13386, 2019.
- [6] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba. OpenAI Gym, 2016.
- [7] K. Choromanski, M. Rowland, V. Sindhvani, R. E. Turner, and A. Weller. Structured evolution with compact architectures for scalable policy optimization. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pages 969–977, 2018.
- [8] P. G. Constantine. *Active Subspaces - Emerging Ideas for Dimension Reduction in Parameter Studies*, volume 2 of *SIAM spotlights*. SIAM, 2015.
- [9] P. G. Constantine, A. Eftekhari, and M. B. Wakin. Computing active subspaces efficiently with gradient sketching. In *6th IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing, CAMSAP 2015, Cancun, Mexico, December 13-16, 2015*, pages 353–356, 2015.
- [10] P. G. Constantine, C. Kent, and T. Bui-Thanh. Accelerating markov chain monte carlo with active subspaces. *SIAM J. Scientific Computing*, 38(5), 2016.
- [11] E. Conti, V. Madhavan, F. P. Such, J. Lehman, K. O. Stanley, and J. Clune. Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada.*, pages 5032–5043, 2018.
- [12] P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, Y. Wu, and P. Zhokhov. Openai baselines. <https://github.com/openai/baselines>, 2017.
- [13] G. H. Golub. Some modified matrix eigenvalue problems. *SIAM*, 15, 1973.
- [14] D. Ha and J. Schmidhuber. Recurrent world models facilitate policy evolution. *NeurIPS*, 2018.
- [15] P. Hämäläinen, A. Babadi, X. Ma, and J. Lehtinen. Ppo-cma: Proximal policy optimization with covariance matrix adaptation. *CoRR*, abs/1810.02541, 2018.
- [16] N. Hansen and A. Ostermeier. Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. In *Evolutionary Computation, 1996., Proceedings of IEEE International Conference on*, pages 312–317. IEEE, 1996.

- [17] R. Houthoofd, Y. Chen, P. Isola, B. Stadie, F. Wolski, O. Jonathan Ho, and P. Abbeel. Evolved policy gradients. *NeurIPS*, 2018.
- [18] I. Jolliffe. Principal component analysis. *Series: Springer Series in Statistics*, XXIX, 2002.
- [19] J. Lehman, J. Chen, J. Clune, and K. O. Stanley. ES is more than just a traditional finite-difference approximator. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018, Kyoto, Japan, July 15-19, 2018*, pages 450–457, 2018.
- [20] C. Li, H. Farkhoor, R. Liu, and J. Yosinski. Measuring the intrinsic dimension of objective landscapes. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*, 2018.
- [21] G. Liu, L. Zhao, F. Yang, J. Bian, T. Qin, N. Yu, and T.-Y. Liu. Trust region evolution strategies. In *AAAI*, 2019.
- [22] I. Loshchilov. A computationally efficient limited memory cma-es for large scale optimization. *GECCO*, 2014.
- [23] I. Loshchilov, T. Glasmachers, and H. Beyer. Large scale black-box optimization by limited-memory matrix adaptation. *IEEE Transactions on Evolutionary Computation*, 2019.
- [24] N. Maheswaranathan, L. Metz, G. Tucker, and J. Sohl-Dickstein. Guided evolutionary strategies: escaping the curse of dimensionality in random search. *CoRR*, abs/1806.10230, 2018.
- [25] H. Mania, A. Guy, and B. Recht. Simple random search provides a competitive approach to reinforcement learning. *CoRR*, abs/1803.07055, 2018.
- [26] N. Müller and T. Glasmachers. Challenges in high-dimensional reinforcement learning with evolution strategies. *Parallel Problem Solving from Nature – PPSN XV*, 2018.
- [27] Y. Nesterov and V. Spokoiny. Random gradient-free minimization of convex functions. *Found. Comput. Math.*, 17(2):527–566, Apr. 2017.
- [28] R. Ros and N. Hansen. A simple modification in cma-es achieving linear time and space complexity. In G. Rudolph, T. Jansen, N. Beume, S. Lucas, and C. Poloni, editors, *Parallel Problem Solving from Nature – PPSN X*, pages 296–305, 2008.
- [29] M. Rowland, K. Choromanski, F. Chalus, A. Pacchiano, T. Sarlos, R. E. Turner, and A. Weller. Geometrically coupled monte carlo sampling. In *NeurIPS*, 2018.
- [30] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever. Evolution strategies as a scalable alternative to reinforcement learning. 2017.
- [31] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, pages 1889–1897, 2015.
- [32] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [33] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune. Deep neuroevolution: Genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. *CoRR*, abs/1712.06567, 2017.
- [34] O. Teytaud and J. Rapin. Nevergrad: An open source tool for derivative-free optimization. <https://code.fb.com/ai-research/nevergrad/>, 2018.
- [35] K. Varelas, A. Auger, D. Brockhoff, N. Hansen, O. A. Elhara, Y. Semet, R. Kassab, and F. Barbaresco. A comparative study of large-scale variants of cma-es. *PPSN XV 2018 - 15th International Conference on Parallel Problem Solving from Nature*, 2018.
- [36] G. Zhang and J. Hinkle. Resnet-based isosurface learning for dimensionality reduction in high-dimensional function approximation with limited data. *CoRR*, 2019.
- [37] Z. Zhou, X. Li, and R. N. Zare. Optimizing chemical reactions with deep reinforcement learning. *ACS Central Science*, 3(12):1337–1344, 2017. PMID: 29296675.

# APPENDIX: From Complexity to Simplicity: Adaptive ES-Active Subspaces for Blackbox Optimization

## 6 Theoretical Results

Throughout this section we will assume the sensings directions  $\{\mathbf{g}_i\}$  at time  $t$  are sampled from one of the following families of distributions:

$$\widehat{P} = \begin{cases} \mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}}) & \text{with probability } p^t \\ \mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}, \perp}}) & \text{with probability } 1 - p^t \end{cases}$$

Where  $p^t$  is a probability parameter with values in  $[0, 1]$ .

Denote an by  $\mathbf{U}^{\text{act}} \in \mathbb{R}^{d \times d_{\text{active}}}$  an orthonormal basis of the active subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$  and  $\mathbf{U}^{\perp} \in \mathbb{R}^{d \times (d - d_{\text{active}})}$  an orthonormal basis of  $\mathcal{L}_{\text{active}}^{\text{ES}, \perp}$ .

Let's start by computing the covariance matrix of  $\widehat{P}$ :

$$\mathbb{E}_{\mathbf{g} \sim P_i} [\mathbf{g}\mathbf{g}^{\top}] = \underbrace{(p^t \mathbf{U}^{\text{act}} (\mathbf{U}^{\text{act}})^{\top} + (1 - p^t) \mathbf{U}^{\perp} (\mathbf{U}^{\perp})^{\top})}_{\mathbf{C}_1}$$

In order to simplify the notation of the proofs in this section we use the following conventions:

$$\begin{aligned} z_{ES} &= \widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, base}} F_{\sigma}(\theta) \\ z_1 &= \widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, asebo}} F_{\sigma}(\theta) \end{aligned}$$

Where  $z_1$  is the ASEBO gradient estimator resulting from using sampling mechanism  $\widehat{P}$ .

**Notational simplification** To simplify notation we also use  $\mathbf{U}$  instead of  $\mathbf{U}^{\text{act}}$ ,  $\mathbf{I}_{\mathbf{U}}$  instead of  $\mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}}$  and  $\mathbf{I}_{\mathbf{U}^{\perp}}$  instead of  $\mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}, \perp}}$ .

Let  $\epsilon > 0$  be the precision parameter. We choose  $\sigma$  with the goal of making the bias between the expectation of our gradient estimators and the true gradient of  $F$  smaller than  $\epsilon$ . Throughout this section we assume  $\sigma$  is small enough:

$$0 < \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1 - p^t)}{\tau d^3 \max(L, 1)}}$$

### 6.1 Gradient Estimators, their bias and their variance.

In this section we aim to produce theoretical guarantees regarding the bias and variance of our proposed gradient estimators. We show that under the right assumptions, the isotropic and non isotropic versions of the evolution Strategies estimators have small bias, and

We make the following assumptions on  $F$ :

**Assumption 1.**  $F$  is  $L$ -Lipschitz. For all  $\theta, \theta' \in \mathbb{R}^d$ ,  $|F(\theta) - F(\theta')| \leq L \|\theta - \theta'\|$ .

**Assumption 2.**  $F$  has a  $\tau$ -smooth third order derivative tensor, so that  $F(\theta + \sigma \mathbf{g}) = F(\theta) + \sigma \nabla F(\theta)^{\top} \mathbf{g} + \frac{\sigma^2}{2} \mathbf{g}^{\top} H(\theta) \mathbf{g} + \frac{1}{6} \sigma^3 F'''(\theta)[v, v, v]$  with  $v \in [0, \mathbf{g}]$  satisfying  $|F'''(\theta)[v, v, v]| \leq \tau \|v\|^3 \leq \tau \|\mathbf{g}\|^3$ .

Let  $d_{\text{active}}$  and  $d_{\perp}$  denote the dimensionality of  $\mathcal{L}_{(\text{active})}$  and  $\mathcal{L}_{\perp}$  respectively.

Under these assumptions,  $\frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} = (\mathbf{g}^{\top} \nabla F(\theta_t)) + \xi_{\mathbf{g}}(\theta_t)$  such that  $\xi_{\mathbf{g}}(\theta_t) \leq \frac{\tau}{6} \sigma^2 \|\mathbf{g}\|^3$ , uniformly over all  $\theta_t$ . We relax the constants slightly. If  $F$ 's third order derivative tensor is smooth with constant  $\tau$ :

$$\left| \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} - \mathbf{g}^\top \nabla F(\theta_t) \right| \leq \tau \sigma^2 \|\mathbf{g}\|^3.$$

Recall the following definitions:

- **Evolution Strategies Gradient.** Let  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})$ . The ES gradient is defined as  $z_{ES} = \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g}$ .
- **$\hat{P}$  Nonisotropic Gradient.** Let  $\mathbf{g} \sim \hat{P}$ . The  $\hat{P}$  gradient is defined as  $z_1 = \mathbf{C}_1^{-1} \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g}$ .

The following inequalities hold:

$$\begin{aligned} \|\xi_{\mathbf{g}}(\theta_t) \mathbf{g}\|^2 &\leq \frac{\tau}{6} \sigma^2 \|\mathbf{g}\|^4 \\ \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}]\|^2 &\leq \frac{\sigma^4 \tau^2}{36} \left( \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^4] \right)^2 \leq \frac{\sigma^4 \tau^2 d^4}{4} \\ \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}]\|^2 &\leq \frac{\sigma^4 \tau^2}{36} \left( \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})} [\|\mathbf{g}\|^4] \right)^2 \leq \frac{\sigma^4 \tau^2 d_{\text{active}}^4}{4} \\ \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}]\|^2 &\leq \frac{\sigma^4 \tau^2}{36} \left( \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})} [\|\mathbf{g}\|^4] \right)^2 \leq \frac{\sigma^4 \tau^2 d_{\perp}^4}{4} \end{aligned}$$

**Bounding the Bias** The first result in this section is to show that under the right conditions the ES gradient estimators in both the isotropic and non isotropic cases can be close to the true gradient provided the function satisfies Assumptions 1 and 2. Theorem 6.1 deals with the isotropic case and Theorem 6.2 with the non isotropic case. The combination of these results yields the proof of Lemma 3.1 in the main text.

**Theorem 6.1.** *The evolution strategies gradient estimator  $z_{ES}$  satisfies:*

$$\|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}] - \nabla F(\theta_t)\| \leq 3\tau \sigma^2 d^2 \quad (4)$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}:$$

$$\|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}] - \nabla F(\theta_t)\| \leq \epsilon \quad (5)$$

*Proof.* Notice that  $\|\mathbf{g}\|^4 = (\sum_{i=1}^d \mathbf{g}(i)^2)^2 \leq d \sum_{i=1}^d \mathbf{g}(i)^4$ . Where we denote  $\mathbf{g}(i)$  as the  $i$ -th entry of the  $d$ -dimensional vector  $\mathbf{g} \in \mathbb{R}^d$ . Since  $\mathbb{E}[\mathbf{g}(i)^4] = 3$  for all  $i$ :

$$\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^4] \leq 3d^2$$

And therefore:

$$\left\| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g} \right] - \nabla F(\theta_t) \right\| \leq \tau \sigma^2 \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^4] \leq 3\tau \sigma^2 d^2$$

□

A similar result holds for the  $z_1$  gradient.

**Theorem 6.2.** *The non isotropic  $\hat{P}$  gradient estimator satisfies:*

$$\|\mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1] - \nabla F(\theta_t)\| \leq \frac{3\sigma^2 \tau}{p^t} d_{\text{active}}^2 + \frac{3\sigma^2 \tau}{1-p^t} d_{\perp}^2$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}:$$

$$\|\mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1] - \nabla F(\theta_t)\| \leq \epsilon$$

*Proof.* Expanding  $\mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1]$  yields:

$$\begin{aligned} \mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1] &= \mathbf{C}_1^{-1} \mathbb{E}_{\mathbf{g} \sim \hat{P}} \left[ \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g} \right] \\ &= \mathbf{C}_1^{-1} \mathbb{E}_{\mathbf{g} \sim \hat{P}} [\mathbf{g} \mathbf{g}^\top \nabla F(\theta_t) + \xi_{\mathbf{g}}(\theta_t) \mathbf{g}] \\ &= \nabla F(\theta_t) + \frac{1}{p^t} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}] + \frac{1}{1-p^t} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}] \end{aligned}$$

By a similar argument as in the proof of Theorem 6.1:

$$\begin{aligned} \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}]\| &\leq 3\tau\sigma^2 d_{\text{active}}^2 \\ \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})} [\xi_{\mathbf{g}}(\theta_t) \mathbf{g}]\| &\leq 3\tau\sigma^2 d_{\perp}^2 \end{aligned}$$

The result follows.  $\square$

**Towards bounding the variance** We start by showing how under the right assumptions the expected squared norm of the ES gradients are also bounded away from the squared norms of the true gradients. The distance between the square norms of the expectation of the ES gradient and the true gradient of  $F$  are also bounded. Theorem 6.3 deals with the isotropic ES estimator and Theorem 6.4 with its non isotropic counterpart:

**Theorem 6.3.** *If  $F$  satisfies Assumption 1 and 2:*

$$\left| \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}]\|^2 - \|\nabla F(\theta_t)\|^2 \right| \leq 105\tau^2\sigma^4 d^4 + 6\tau\sigma^2 L d^2 \quad (6)$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}.$$

$$\left| \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}]\|^2 - \|\nabla F(\theta_t)\|^2 \right| \leq \epsilon \quad (7)$$

*Proof.*

$$\begin{aligned} \left| \left\| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g} \right] \right\|^2 - \|\nabla F(\theta_t)\|^2 \right| &\leq \tau^2 (\sigma^2 \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^4])^2 + \\ &\quad 2\tau\sigma^2 \|\nabla F(\theta_t)\| \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^4]\| \\ &\leq 105\tau^2\sigma^4 d^4 + 6\tau\sigma^2 L d^2 \end{aligned}$$

$\square$

**Theorem 6.4.** *If  $F$  satisfies Assumption 1 and 2:*

$$\begin{aligned} \left| \left\| \mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1] \right\|^2 - \|\nabla F(\theta_t)\|^2 \right| &\leq \frac{1}{(p^t)^2} \frac{\sigma^4 \tau^2 d_{\text{active}}^4}{4} + \frac{1}{(1-p^t)^2} \frac{\sigma^4 \tau^2 d_{\perp}^4}{4} + \frac{2}{p^t} L \frac{\sigma^2 \tau d_{\text{active}}^2}{4} + \\ &\quad \frac{2}{1-p^t} L \frac{\sigma^2 \tau d_{\perp}^2}{4} + \frac{2}{p^t(1-p^t)} \frac{\sigma^4 \tau^2 d_{\text{active}}^2 d_{\perp}^2}{16} \end{aligned}$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}.$$

$$\left| \left\| \mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1] \right\|^2 - \|\nabla F(\theta_t)\|^2 \right| \leq \epsilon$$

*Proof.* Consider the following expansion of  $\mathbb{E} [z_1]$ .

$$\begin{aligned} \|\mathbb{E}_{\mathbf{g} \sim \hat{P}}[z_1]\|^2 &= \|\nabla F(\theta_t)\|^2 + \frac{1}{(p^t)^2} \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_U)}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}]\|^2 + \left(\frac{1}{1-p^t}\right)^2 \|\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{U^\perp})}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}]\|^2 \\ &\quad + \frac{2}{p^t} \langle \nabla F(\theta_t), \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_U)}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}] \rangle + \frac{2}{1-p^t} \langle \nabla F(\theta_t), \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{U^\perp})}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}] \rangle \\ &\quad + \frac{2}{p^t(1-p^t)} \langle \nabla \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_U)}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}], \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{U^\perp})}[\xi_{\mathbf{g}}(\theta_t)\mathbf{g}] \rangle \end{aligned}$$

And therefore by Cauchy Schwartz:

$$\begin{aligned} \left| \|\mathbb{E}_{\mathbf{g} \sim \hat{P}}[z_1]\|^2 - \|\nabla F(\theta_t)\|^2 \right| &\leq \frac{1}{(p^t)^2} \frac{\sigma^4 \tau^2 d_{\text{active}}^4}{4} + \frac{1}{(1-p^t)^2} \frac{\sigma^4 \tau^2 d_{\perp}^4}{4} \\ &\quad + \frac{2}{p^t} L \frac{\sigma^2 \tau d_{\text{active}}^2}{4} + \frac{2}{1-p^t} L \frac{\sigma^2 \tau d_{\perp}^2}{4} + \frac{2}{p^t(1-p^t)} \frac{\sigma^4 \tau^2 d_{\text{active}}^2 d_{\perp}^2}{16} \end{aligned}$$

As desired.  $\square$

**Bounding the variance of  $z_{ES}$  and  $z_1$ .** We have now the necessary ingredients for bounding the variance of the ES isotropic and non isotropic estimators. We start by showing in theorem 6.5 that the variance of the isotropic estimator is roughly of the order of  $(d+1)\|\nabla F(\theta_t)\|^2$ . In contrast, Theorem 6.6 characterizes the variance of  $z_1$  the non isotropic ES gradient estimator in terms of the  $\nabla F(\theta_t)$  decomposition along the subspaces spanned by  $\mathbf{U}$  and  $\mathbf{U}^\perp$ . In the following section 6.2 we show that with an appropriate choice of the probabilities  $p^t$ ,  $1-p^t$ , and provided the subspace decomposition is adequate, the variance of the non isotropic gradient estimator can be much smaller than the variance of the  $z_{ES}$ .

**Theorem 6.5.** *If  $F$  satisfies Assumption 1 and 2, the variance of the ES estimator satisfies:*

$$|\text{Var}_{ES} - (d+1)\|\nabla F(\theta_t)\|^2| \leq 105\tau^2\sigma^4 d^4 + 6\tau\sigma^2 L d^2 + 15d^3\sigma^2 L\tau + 105\tau^2\sigma^4 d^4$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}.$$

$$|\text{Var}_{ES} - (d+1)\|\nabla F(\theta_t)\|^2| \leq \epsilon$$

*Proof.* The second moment of the ES estimator satisfies:

$$\begin{aligned} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}^\top z_{ES}] &= \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \frac{(F(\theta_t + \sigma\mathbf{g}) - F(\theta_t - \sigma\mathbf{g}))^2}{2^2\sigma^2} \mathbf{g}^\top \mathbf{g} \right] \\ &= \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ (\mathbf{g}^\top \nabla F(\theta_t) + \xi_{\mathbf{g}}(\theta_t))^2 \mathbf{g}^\top \mathbf{g} \right] \\ &= \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \nabla F(\theta_t)^\top \mathbf{g} \mathbf{g}^\top \mathbf{g} \mathbf{g}^\top \nabla F(\theta_t) + 2\nabla F(\theta_t)^\top \mathbf{g} \mathbf{g}^\top \mathbf{g} \xi_{\mathbf{g}}(\theta_t) + \xi_{\mathbf{g}}(\theta_t)^2 \mathbf{g}^\top \mathbf{g} \right] \\ &= (d+2)\|\nabla F(\theta_t)\|^2 + 2\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\nabla F(\theta_t)^\top \mathbf{g} \mathbf{g}^\top \mathbf{g} \xi_{\mathbf{g}}(\theta_t)] + \\ &\quad \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\xi_{\mathbf{g}}(\theta_t)^2 \mathbf{g}^\top \mathbf{g}] \end{aligned}$$

Under Assumption 1 and 2, the following bound for the second and third terms of the last equality holds:

$$\left| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\nabla F(\theta_t)^\top \mathbf{g} \mathbf{g}^\top \mathbf{g} \xi_{\mathbf{g}}(\theta_t)] \right| \leq \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\nabla F(\theta_t)\| \|\mathbf{g}\|^6] \leq 15d^3\sigma^2 L\tau$$

And:

$$\left| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\xi_{\mathbf{g}}(\theta_t)^2 \mathbf{g}^\top \mathbf{g}] \right| \leq \tau^2 \sigma^4 \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [\|\mathbf{g}\|^8] \leq 105\tau^2 \sigma^4 d^4$$

Therefore:

$$\text{Var}_{\text{ES}} = \underbrace{\mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \frac{(F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g}))^2}{2^2 \sigma^2} \mathbf{g}^\top \mathbf{g} \right]}_{\diamond} - \underbrace{\left\| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} \left[ \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{g} \right] \right\|^2}_{\spadesuit}$$

After coalescing the bounds derived in the preceding section, we can obtain the following bound on the term  $\diamond$ :

$$|\diamond - (d+2)\|\nabla F(\theta_t)\|^2| \leq 15d^3\sigma^2 L\tau + 105\tau^2\sigma^4 d^4$$

Notice that by virtue of 6, the following bound on term  $\spadesuit$  of the previous equation holds:

$$|\spadesuit - \|\nabla F(\theta_t)\|^2| \leq 105\tau^2\sigma^4 d^4 + 6\tau\sigma^2 Ld^2$$

Combining these two inequalities the result follows.  $\square$

A similar theorem holds for  $z_1$ .

**Theorem 6.6.** Let  $\Gamma = \left( \frac{d_{\text{active}}+2}{p^t} \|\mathbf{U}^\top \nabla F(\theta_t)\|^2 + \frac{d_{\perp}+2}{1-p^t} \|(\mathbf{U}^\perp)^\top F(\theta_t)\|^2 - \|\nabla F(\theta_t)\|^2 \right)$ .

$$\begin{aligned} |Var_{\hat{P}} - \Gamma| &\leq \frac{1}{p^t} (15d_{\text{active}}^3\sigma^2 L\tau + 105\tau^2\sigma^4 d_{\text{active}}^4) + \\ &\frac{1}{1-p^t} (15d_{\perp}^3\sigma^2 L\tau + 105\tau^2\sigma^4 d_{\perp}^4) + \\ &\frac{1}{(p^t)^2} \frac{\sigma^4\tau^2 d_{\text{active}}^4}{4} + \frac{1}{(1-p^t)^2} \frac{\sigma^4\tau^2 d_{\perp}^4}{4} + \\ &\frac{2}{p^t} L \frac{\sigma^2\tau d_{\text{active}}^2}{4} + \frac{2}{1-p^t} L \frac{\sigma^2\tau d_{\perp}^2}{4} + \\ &\frac{2}{p^t(1-p^t)} \frac{\sigma^4\tau^2 d_{\text{active}}^2 d_{\perp}^2}{16} \end{aligned}$$

$$\text{If } \sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}.$$

$$\left| Var_{\hat{P}} - \left( \frac{d_{\text{active}}+2}{p^t} \|\mathbf{U}^\top \nabla F(\theta_t)\|^2 + \frac{d_{\perp}+2}{1-p^t} \|(\mathbf{U}^\perp)^\top F(\theta_t)\|^2 - \|\nabla F(\theta_t)\|^2 \right) \right| \leq \epsilon$$

*Proof.* The second moment of  $z_1$  satisfies:

$$\begin{aligned} \mathbb{E}_{\mathbf{g} \sim \hat{P}} [z_1^\top z_1] &= \frac{1}{p^t} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})} \left[ \frac{(F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g}))^2}{2^2 \sigma^2} \mathbf{g}^\top \mathbf{g} \right] + \\ &\frac{1}{1-p^t} \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})} \left[ \frac{(F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g}))^2}{2^2 \sigma^2} \mathbf{g}^\top \mathbf{g} \right] \end{aligned}$$

Notice that:

$$\text{Var}_{\hat{P}} = \underbrace{\mathbb{E}_{\mathbf{g} \sim \hat{P}} \left[ \frac{(F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g}))^2}{2^2 \sigma^2} \mathbf{g}^\top \mathbf{C}_1^{-2} \mathbf{g} \right]}_{\diamond} - \underbrace{\left\| \mathbb{E}_{\mathbf{g} \sim \hat{P}} \left[ \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \mathbf{C}_1^{-1} \mathbf{g} \right] \right\|^2}_{\spadesuit}$$

By a similar argument as that in the previous theorem, we conclude:

$$\left| \diamond - \frac{d_{\text{active}} + 2}{p^t} \|\mathbf{U}^\top \nabla F(\theta_t)\|^2 - \frac{(d_{\perp} + 2)}{1 - p^t} \|(\mathbf{U}^\perp)^\top \nabla F(\theta_t)\|^2 \right| \leq \frac{1}{p^t} (15d_{\text{active}}^3 \sigma^2 L \tau + 105\tau^2 \sigma^4 d_{\text{active}}^4) + \frac{1}{1 - p^t} (15d_{\perp}^3 \sigma^2 L \tau + 105\tau^2 \sigma^4 d_{\perp}^4)$$

By Theorem 6.4:

$$\begin{aligned} |\spadesuit - \|\nabla F(\theta_t)\|^2| &\leq \frac{1}{(p^t)^2} \frac{\sigma^4 \tau^2 d_{\text{active}}^4}{4} + \frac{1}{(1 - p^t)^2} \frac{\sigma^4 \tau^2 d_{\perp}^4}{4} + \frac{2}{p^t} L \frac{\sigma^2 \tau d_{\text{active}}^2}{4} + \frac{2}{1 - p^t} L \frac{\sigma^2 \tau d_{\perp}^2}{4} + \\ &\quad \frac{2}{p^t(1 - p^t)} \frac{\sigma^4 \tau^2 d_{\text{active}}^2 d_{\perp}^2}{16} \end{aligned}$$

□

## 6.2 Variance reduction via non isotropic sampling

The first result of this section is to condense the theorems in the previous sections into a single result (see Theorem 6.7). Lemma 6.8 then shows what the variance corresponding to the optimal choice of parameter  $p^t$  is. Theorem 6.9 then provides conditions under which the approximate variance (without considering the bias terms) corresponding to the optimal non isotropic estimator is smaller than the variance of the isotropic one. Finally Theorem 6.10 takes into account the bias and states the final result of this section. The combination of these results yield the proof of Theorem 3.2 in the main section of the paper.

**Theorem 6.7.** *Let  $\epsilon > 0$ . If  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1 - p^t)}{\tau d^3 \max(L, 1)}}$  then:*

$$\left\| \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I})} [z_{ES}] - \nabla F(\theta_t) \right\| \leq \epsilon \quad (8)$$

$$\left\| \mathbb{E}_{\mathbf{g} \sim \hat{\mathcal{P}}} [z_1] - \nabla F(\theta_t) \right\| \leq \epsilon \quad (9)$$

and

$$|\text{Var}_{ES} - (d + 1) \|\nabla F(\theta_t)\|^2| \leq \epsilon \quad (10)$$

$$\left| \text{Var}_{\hat{\mathcal{P}}} - \left( \frac{d_{\text{active}} + 2}{p^t} \|\mathbf{U}^\top \nabla F(\theta_t)\|^2 + \frac{d_{\perp} + 2}{1 - p^t} \|(\mathbf{U}^\perp)^\top \nabla F(\theta_t)\|^2 - \|\nabla F(\theta_t)\|^2 \right) \right| \leq \epsilon \quad (11)$$

We say that in this case:

$$\text{Var}_{\hat{\mathcal{P}}} \approx \underbrace{\left( \frac{d_{\text{active}} + 2}{p^t} \|\mathbf{U}^\top \nabla F(\theta_t)\|^2 + \frac{d_{\perp} + 2}{1 - p^t} \|(\mathbf{U}^\perp)^\top \nabla F(\theta_t)\|^2 - \|\nabla F(\theta_t)\|^2 \right)}_{\text{Var}_{\hat{\mathcal{P}}}^M}$$

and  $\text{Var}_{ES} \approx (d + 1) \|\nabla F(\theta_t)\|^2$ . We refer to  $\text{Var}_{\hat{\mathcal{P}}}^M$  as the "main component" of the variance  $\text{Var}_{\hat{\mathcal{P}}}$ . Similarly we define  $\text{Var}_{ES}^M = (d + 1) \|\nabla F(\theta_t)\|^2$  and use the same name, "main component" of the variance  $\text{Var}_{ES}^M$ .

The optimal  $p^t$ , that which minimizes  $\text{Var}_{\hat{\mathcal{P}}}^M$  equals:

$$(p^t)^* = \frac{\|(\nabla F(\theta_t))_{\text{active}}\| \sqrt{d_{\text{active}} + 2}}{\|(\nabla F(\theta_t))_{\text{active}}\| \sqrt{d_{\text{active}} + 2} + \|(\nabla F(\theta_t))_{\perp}\| \sqrt{d_{\perp} + 2}}$$



*Proof.* Roughly the same argument as above yields the desired result.  $\square$

**Lemma 6.8.** *The optimal variance  $\text{Var}_{\hat{p}^*}^M$  corresponding to  $(p^t)^*$  equals:*

$$\left[ \left\| (\nabla F(\theta_t))_{\text{active}} \right\| \sqrt{d_{\text{active}} + 2} + \left\| (\nabla F(\theta_t))_{\perp} \right\| \sqrt{d_{\perp} + 2} \right]^2 - \left\| \nabla F(\theta_t) \right\|^2 \quad (12)$$

*Proof.* The statement follows directly from substituting the expression for  $(p^t)^*$  into the variance formula.  $\square$

**Theorem 6.9.**  *$\text{Var}_{\hat{p}^*}^M \leq \text{Var}_{ES}^M$  if*

$$\left| \sqrt{d_{\text{active}} + 2} \left\| (\nabla F(\theta_t))_{\perp} \right\| - \sqrt{d_{\perp} + 2} \left\| \nabla F(\theta_t)_{\text{active}} \right\| \right| \geq \sqrt{2} \left\| \nabla F(\theta_t) \right\|$$

*Proof.* By definition,  $\text{Var}_{\hat{p}^*}^M < \text{Var}_{ES}^M$  if:

$$\left( \left\| (\nabla F(\theta_t))_{\text{active}} \right\| \sqrt{d_{\text{active}} + 2} + \left\| (\nabla F(\theta_t))_{\perp} \right\| \sqrt{d_{\perp} + 2} \right)^2 < \left\| \nabla F(\theta_t) \right\|^2 (d + 2) \quad (13)$$

Let  $a_1 = \sqrt{d_{\text{active}} + 2}$ ,  $a_2 = \sqrt{d_{\perp} + 2}$ ,  $b_1 = \left\| (\nabla F(\theta_t))_{\text{active}} \right\|$ ,  $b_2 = \left\| (\nabla F(\theta_t))_{\perp} \right\|$ ,  $a = \sqrt{d + 2}$  and  $b = \left\| \nabla F(x_t) \right\|$ .

The following relationships hold:  $b_1^2 + b_2^2 = b^2$  and  $a_1^2 + a_2^2 - 2 = a^2$ . The bound we want to prove in Equation 13 reduces to finding conditions under which:

$$(a_1 b_1 + a_2 b_2)^2 \leq (b_1^2 + b_2^2)(a_1^2 + a_2^2 - 2)$$

Which holds iff:

$$2b_1^2 + 2b_2^2 \leq a_2^2 b_1^2 + a_1^2 b_2^2 - 2a_1 a_2 b_1 b_2$$

The later holds iff:

$$|a_1 b_2 - a_2 b_1| \geq \sqrt{2} b$$

Which holds iff:

$$\left| \sqrt{d_{\text{active}} + 2} \left\| (\nabla F(x_t))_{\perp} \right\| - \sqrt{d_{\perp} + 2} \left\| (\nabla F(x_t))_{\text{active}} \right\| \right| \geq \sqrt{2} \left\| \nabla F(x_t) \right\|$$

$\square$

The inequality is strict for example when  $\left\| (\nabla F(x_t))_{\perp} \right\| = 0$  and  $d_{\perp} \geq 1$ .

This in turn implies that, after taking into account the bias terms:

**Theorem 6.10.** *If  $\epsilon > 0$ . If  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min((p^t)^*, 1 - (p^t)^*)}{\tau d^3 \max(L, 1)}}$ , we denote by  $\text{Var}_{(\hat{p}^*)}$  as the variance of the gradient estimator  $z_1$  corresponding to the optimal (for  $\text{Var}_{\hat{p}^*}^M$ ) probability  $(p^t)^*$  and*

$$\left| \sqrt{d_{\text{active}} + 2} \left\| (\nabla F(\theta_t))_{\perp} \right\| - \sqrt{d_{\perp} + 2} \left\| \nabla F(\theta_t)_{\text{active}} \right\| \right| \geq \sqrt{2} \left\| \nabla F(\theta_t) \right\|$$

*Then:*

$$\text{Var}_{\hat{p}^*} \leq \text{Var}_{ES} + \epsilon$$

### 6.3 Adaptive Mirror Descent for variance reduction.

In this section we propose an adaptive procedure to learn the optimal probability parameter  $(p^t)^*$  (as introduced in the previous section) this is necessary since as it can be inferred from the discussion in section 6.2, the optimal variance depends of unknown parameters such as the projection of the true gradient onto the subspaces spanned by  $\mathbf{U}$  and  $\mathbf{U}^\perp$ . The final result of this section 6.15 corresponds to Theorem 3.3 in the main section of the text.

Let  $\mathbf{p}^l = \begin{pmatrix} p^l \\ 1-p^l \end{pmatrix}$ . The main component  $\Gamma$  of the variance of  $\widehat{\nabla}_{\text{MC}, k=1}^{\text{AT, asebo}} F_\sigma(\theta)$  as a function of  $\mathbf{p}^l$  equals (Lemma 3.2) :

$$\Gamma = \ell(\mathbf{p}^l) = \frac{d_{\text{active}} + 2}{\mathbf{p}^l(1)} s_{\mathbf{U}^{\text{act}}} + \frac{d_\perp + 2}{\mathbf{p}^l(2)} s_{\mathbf{U}^\perp} - \|\nabla F(\theta)\|^2. \quad (14)$$

In order to avoid the gradients to explote in norm, we parametrise  $l$  as follows:

$$\mathbf{p}^l = (1 - 2\beta)\mathbf{q}^l + \begin{pmatrix} \beta \\ \beta \end{pmatrix}$$

For  $\mathbf{q}^l \in \Delta_2$  and  $\beta \in (0, 1)$ , the boundary probability bias.

Notice that  $\Gamma$  is a convex function of  $\mathbf{p}$  and also a convex function of  $\mathbf{q}$ . With a slight abuse of notation we denote  $\ell(\mathbf{q}^l)$  as the loss parametrized by  $\mathbf{q}^l$  (which satisfies  $\ell(\mathbf{q}^l) = \ell(\mathbf{p}^l)$ ).

The gradient  $\nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l)$  equals:

$$\nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l) = (1 - 2\beta) \begin{pmatrix} -\frac{d_{\text{active}} + 2}{((1-2\beta)\mathbf{q}^l(1) + \beta)^2} s_{\mathbf{U}^{\text{ort}}} \\ -\frac{d_\perp + 2}{((1-2\beta)\mathbf{q}^l(2) + \beta)^2} s_{\mathbf{U}^\perp} \end{pmatrix},$$

And can be approximated (at the cost of some bias) using function evaluations.

**Lemma 6.11.** *The gradient  $\nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l)$  satisfies:*

$$\left\| \nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l) - \mathbb{E} \left[ (1 - 2\beta) \begin{pmatrix} -\frac{a_l(d_{\text{active}} + 2)}{((1-2\beta)\mathbf{p}^l(1) + \beta)^3} v_l^2 \\ -\frac{(1-a_l)(d_\perp + 2)}{((1-2\beta)\mathbf{p}^l(2) + \beta)^3} \end{pmatrix} v_l^2 \right] \right\| \leq \frac{\epsilon(d+2)}{(\min(\mathbf{p}^l(1), \mathbf{p}^l(2)))^3} \leq \frac{\epsilon(d+2)}{\beta^3},$$

where  $v_l = \frac{1}{2\sigma} (F(\theta + \mathbf{g}_l) - F(\theta - \mathbf{g}_l))$ .

*Proof.* We start with some notation borrowed from the previous section:

$$\xi_{\mathbf{g}_l}^{(2)}(\theta) = \underbrace{\left( \frac{F(\theta + \sigma \mathbf{g}_l) - F(\theta - \sigma \mathbf{g}_l)}{2\sigma} \right)^2}_{v_l^2} - (\mathbf{g}_l^\top \nabla F(\theta_t))^2$$

Observe that:

$$\begin{aligned} |\xi_{\mathbf{g}_l}^{(2)}(\theta)| &= \left| \left( \frac{F(\theta + \sigma \mathbf{g}_l) - F(\theta - \sigma \mathbf{g}_l)}{2\sigma} \right)^2 - (\mathbf{g}_l^\top \nabla F(\theta))^2 \right| \\ &\leq \xi_{\mathbf{g}_l}(\theta)^2 + 2 |\mathbf{g}_l^\top \nabla F(\theta) \xi_{\mathbf{g}_l}(\theta)| \\ &\leq \sigma^4 \tau^2 \|\mathbf{g}_l\|^6 + 2\sigma^2 \tau L \|\mathbf{g}_l\|^4 \end{aligned}$$

Since  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}$ .

$$\mathbb{E} \left[ |\xi_{\mathbf{g}_l}^{(2)}(\theta)| \right] \leq \epsilon$$

The result follows. □

Let  $p^l = (1 - 2\beta)q^l + \beta$  be the probability that we choose to sample from the subspace  $\mathcal{L}_{\text{active}}^{\text{ES}}$  and  $1 - p^l$  the probability that we choose to sample from  $\mathcal{L}_{\text{active}}^{\text{ES}, \perp}$ . Let  $a_l$  be a Bernoulli random variable  $a_l \in \{0, 1\}$  with  $\mathbb{E} \left[ \begin{pmatrix} a_l \\ 1 - a_l \end{pmatrix} \right] = \mathbf{p}^l$ . Define the stochastic gradient (with respect to  $\mathbf{q}^l$ ):

$$\mathbf{e}_l = (1 - 2\beta) \left[ \begin{pmatrix} -\frac{a_l(d_{\text{active}} + 2)}{p_l^3} \\ -\frac{(1 - a_l)(d_{\perp} + 2)}{(1 - p_l)^3} \end{pmatrix} \right] v_l^2$$

By definition this random vector (conditioned on the choice of  $\mathbf{p}^l$ ) satisfies:

$$\|\mathbb{E} [e_l] - \nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l)\| \leq \frac{\epsilon(d + 2)}{(\min(\mathbf{p}^l(1), \mathbf{p}^l(2)))^3} \leq \frac{\epsilon(d + 2)}{\beta^3},$$

If  $\epsilon$  is chosen small enough, the bias can be driven to be arbitrarily small.

### 6.3.1 Mirror descent

We treat this problem as that of minimizing the loss  $\ell$  over the two dimensional simplex and resort to adapt a version of Mirror descent for it. As opposed the case of projected gradient descent, mirror descent performs updates that are adapted to the geometry of the simplex, ensuring the iterates always belong to the simplex and no projection step is necessary. The mirror descent updates are:

$$\begin{aligned} \mathbf{q}_l(1) &= \frac{\mathbf{q}_{l-1}(1) \exp(-\alpha \mathbf{e}_l(1))}{\mathbf{q}_{l-1}(1) \exp(-\alpha \mathbf{e}_l(1)) + (\mathbf{q}_{l-1}(2)) \exp(-\alpha \mathbf{e}_l(2))} \\ \mathbf{q}_l(2) &= \frac{\mathbf{q}_{l-1}(2) \exp(-\alpha \mathbf{e}_l(2))}{\mathbf{q}_{l-1}(1) \exp(-\alpha \mathbf{e}_l(1)) + (\mathbf{q}_{l-1}(2)) \exp(-\alpha \mathbf{e}_l(2))} \end{aligned}$$

For a step size parameter  $\alpha$ .

### 6.4 Regret guarantees

Using the notation in <https://www.stat.berkeley.edu/~bartlett/courses/2014fall-cs294stat260/lectures/mirror-descent-notes.pdf>, In this case let  $R(\mathbf{q}) = \mathbf{q}(1) \log(\mathbf{q}(1)) + \mathbf{q}(2) \log(\mathbf{q}(2)) - \mathbf{q}(1) - \mathbf{q}(2)$  and therefore:

$$\nabla R(\mathbf{q}) = \begin{pmatrix} \log(\mathbf{q}(1)) \\ \log(\mathbf{q}(2)) \end{pmatrix} \quad (15)$$

The Fenchel conjugate of  $R$  equals:

$$R^*(\mathbf{q}) = e^{\mathbf{q}(1)} + e^{\mathbf{q}(2)} \quad (16)$$

And therefore the gradient of the Fenchel conjugate equals:

$$\nabla R^*(\mathbf{q}) = \begin{pmatrix} \exp(\mathbf{q}(1)) \\ \exp(\mathbf{q}(2)) \end{pmatrix} \quad (17)$$

And:

$$D_R(\mathbf{q}_1, \mathbf{q}_2) = \mathbf{q}_1(1) \log \left( \frac{\mathbf{q}_1(1)}{\mathbf{q}_2(1)} \right) + \mathbf{q}_1(2) \log \left( \frac{\mathbf{q}_1(2)}{\mathbf{q}_2(2)} \right) + \mathbf{q}_2(1) - \mathbf{q}_1(1) + \mathbf{q}_2(2) - \mathbf{q}_1(2) \quad (18)$$

Recall the update behind Mirror descent takes the form (stepsize  $\alpha$ ):

1. Play  $\begin{pmatrix} a_l \\ 1 - a_l \end{pmatrix}$  such that  $\mathbb{E}[\begin{pmatrix} a_l \\ 1 - a_l \end{pmatrix}] = \mathbf{p}^l$ .
2. Let  $w_{l+1} = \nabla R^* (\nabla R(\mathbf{p}^l) - \alpha \mathbf{e}_l)$

3. Let  $\mathbf{p}^{l+1} = \arg \min_{\mathbf{p} \in \Delta_2} D_R(\mathbf{p}, w_{t+1})$

Recall the general definition of Bregman divergence:

$$D_\Psi(u, v) = \Psi(u) - \Psi(v) - \langle \nabla \Psi(v), u - v \rangle \quad (19)$$

The following regret guarantee holds for Mirror descent (see <https://www.stat.berkeley.edu/~bartlett/courses/2014fall-cs294stat260/lectures/mirror-descent-notes.pdf>):

**Theorem 6.12.** *If at time  $l$  a convex loss function  $f_l$  is revealed to the player and the player performs the mirror descent step using  $\nabla f_l$  as a proxy linear function, with actions (from the mirror descent step)  $a_l$  at time  $l$ , for any  $a$  in the intersection of all of  $f_l$ 's domains, the following regret bound holds:*

$$\sum_{l=1}^C (f_l(a_l) - f_l(a)) \leq \sum_{l=1}^C \nabla f_l(a_l)^\top (a_l - a) \quad (20)$$

$$\leq \frac{1}{\alpha} \left( R(a) - R(a_1) + \sum_{l=1}^C D_{R^*}(\nabla R(a_l) - \alpha \nabla f_l(a_l), \nabla R(a_l)) \right) \quad (21)$$

Also remember that if  $R^*$  is  $\theta$ -smooth with respect to some norm  $\|\cdot\|$ , we can upper bound  $D_{R^*}$ . The former ( $R^*$  being  $\theta$ -smooth) holds if  $R$  is  $\frac{1}{\theta}$ -strongly convex with respect to the dual norm  $\|\cdot\|_*$ . When  $R$  equals the entropy, this is 1-strongly convex with respect to the  $L_1$  norm and hence  $R^*$  is 1-strongly smooth with respect to the  $L_\infty$  norm:

$$D_{R^*}(a, b) \leq \frac{\|a - b\|_\infty^2}{2} \quad (22)$$

In our case, let  $f_l(\mathbf{q}) = \mathbf{e}_l^\top \mathbf{q}$ . Using the upper bound previously described for  $D_R$ . For any  $\mathbf{q} \in \Delta_2$ :

$$\sum_{l=1}^C f_l(\mathbf{q}^l) - f_l(\mathbf{q}) \leq \frac{1}{\alpha} \left( R(\mathbf{q}) - R(\mathbf{q}^1) + \alpha^2 \sum_{l=1}^C \frac{\|\nabla f_l(\mathbf{q}^l)\|_\infty^2}{2} \right)$$

Taking expectations, since  $|\mathbb{E}[f_l(\mathbf{q})|\mathbf{q}^l] - \nabla_{\mathbf{q}^l}^\top \ell(\mathbf{q}^l)\mathbf{q}| \leq \frac{\epsilon(d+2)}{\beta^3}$  we obtain the following result:

$$\left( \sum_{l=1}^C \nabla_{\mathbf{q}^l}^\top \ell(\mathbf{q}^l) (\mathbf{q}^l - \mathbf{q}) \right) - C \frac{2\epsilon(d+2)}{\beta^3} \leq \frac{1}{\alpha} \left( R(\mathbf{q}) - \mathbb{E}[R(\mathbf{q}^1)] + \alpha^2 \sum_{l=1}^C \frac{\mathbb{E}[\|\mathbf{e}_l\|_\infty^2]}{2} \right)$$

Now we bound the Right Hand side of the expression above. Notice that  $R(\mathbf{q}) \leq 2$  and  $R(\mathbf{q}_1) \geq 0$ . We can also bound the expectation  $\mathbb{E}[\|\mathbf{e}_l\|_\infty^2]$ .

**Lemma 6.13.**  $|\mathbb{E}[\|\mathbf{e}_l\|^2] - \|\nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l)\|^2| \leq \frac{\epsilon(d+1)}{\beta^3}$

*Proof.* A similar calculation as in Lemma 6.11 yields the desired result.  $\square$

Since:

$$\mathbb{E}[\|\mathbf{e}_l\|_\infty^2] \leq \mathbb{E}[\|\mathbf{e}_l\|^2] \quad (23)$$

And  $\|\nabla_{\mathbf{q}^l} \ell(\mathbf{q}^l)\|^2 \leq \frac{1}{\beta^4} ((d_{\text{active}} + 2)^2 s_{\mathbf{U}^{\text{ort}}}^2 + (d_\perp + 2)^2 s_{\mathbf{U}^\perp}^2)$ .

We obtain the following bound:

$$\left( \sum_{l=1}^C \nabla_{\mathbf{q}^l}^\top \ell(\mathbf{q}^l) (\mathbf{q}^l - \mathbf{q}) \right) - C \frac{2\epsilon(d+2)}{\beta^3} \leq \frac{2}{\alpha} + \frac{\alpha C}{2\beta^4} ((d_{\text{active}} + 2)^2 s_{\mathbf{U}^{\text{ort}}}^2 + (d_\perp + 2)^2 s_{\mathbf{U}^\perp}^2) + \frac{\alpha C \epsilon(d+1)}{\beta^3}$$

The following theorem follows:

**Theorem 6.14.** If  $\alpha = \frac{2\beta^2}{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}}$  and  $\epsilon = \frac{\beta^3}{2C(d+1)}$ , for any  $\mathbf{q} \in \Delta_2$ :

$$\left( \sum_{l=1}^C \nabla_{\mathbf{q}^l}^{\top} \ell(\mathbf{q}^l) (\mathbf{q}^l - \mathbf{q}) \right) \leq \frac{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}}{\beta^2} + 1$$

*Proof.* Plugging in this value of  $\alpha$ :

$$\begin{aligned} \left( \sum_{l=1}^C \nabla_{\mathbf{q}^l}^{\top} \ell(\mathbf{q}^l) (\mathbf{q}^l - \mathbf{q}) \right) &\leq \frac{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}}{\beta^2} \\ &+ \left( 1 + \frac{2\beta^2}{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}} \right) \frac{C\epsilon(d+1)}{\beta^3} \end{aligned}$$

By setting  $\epsilon = \frac{\beta^3}{2C(d+1)}$  the result follows. Assuming  $C$  is large enough so that  $\alpha < 1$ .  $\square$

Since  $\ell(\mathbf{q})$  is a convex function of  $\mathbf{q}$  for all  $l$  and  $\mathbf{q} \in \Delta_2$ :

$$\ell(\mathbf{q}^l) - \ell(\mathbf{q}) \leq \nabla_{\mathbf{q}^l}^{\top} \ell(\mathbf{q}^l) (\mathbf{q}^l - \mathbf{q})$$

Which in turn implies the main result of this section:

**Theorem 6.15.** If  $\alpha = \frac{2\beta^2}{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}}$  and  $\epsilon = \frac{\beta^3}{2C(d+1)}$ , for any  $\mathbf{q} \in \Delta_2$ :

$$\mathbb{E} \left[ \sum_{l=1}^C \ell(\mathbf{q}^l) - \ell(\mathbf{q}) \right] \leq \frac{\sqrt{C}\sqrt{(d_{\text{active}}+2)^2s_{\mathbf{U}_{\text{ort}}}^2+(d_{\perp}+2)s_{\mathbf{U}_{\perp}}^2}}{\beta^2} + 1$$

*This is equivalent to the result stated in the main paper.*

## 7 Additional Implementation Details

In this section we present additional details on our experimental results, for both the RL tasks and Nevergrad functions.

### 7.1 Reinforcement Learning Experiment Details

We provide additional details regarding the RL experiments below.

**State Normalization.** State-of-the-art policy optimization baselines such as PPO/TRPO [12] and the original ARS [25] apply state normalization as part of the implementation. In particular, the algorithms maintain a component-wise running average of mean  $\bar{s}$  and standard deviation  $\sigma(s)$  of the state. When at given state  $s_t$ , the algorithm computes the normalized state  $\tilde{s}_t = \frac{s_t - \bar{s}}{\sigma(s)}$  before inputting to the policy network to compute actions  $a_t = \pi(\tilde{s}_t)$ . For PPO/TRPO, since the optimization is based on back-propagation of neural networks, properly scaling the inputs  $s_t \rightarrow \tilde{s}_t$  is critical for the performance. In all experiments, we remove state normalization mechanism from the implementation to test the robustness of various blackbox optimization algorithms. Notice that as reported by [25], state normalization was not needed in ARS to learn good policies for RL tasks under consideration in this paper. As a result, we observe that PPO/TRPO underperform other ES algorithms for most tasks.

**Benchmark Environments.** Benchmark environments are from OpenAI gym [6]. These environments have variable sizes of observation space and action space: Swimmer-v2  $|\mathcal{S}| = 8$ ,  $|\mathcal{A}| = 2$ ; Hopper-v2  $|\mathcal{S}| = 11$ ,  $|\mathcal{A}| = 3$ ; HalfCheetah-v2  $|\mathcal{S}| = 17$ ,  $|\mathcal{A}| = 6$ ; Thrower-v2  $|\mathcal{S}| = 23$ ,  $|\mathcal{A}| = 7$ ; Pusher-v2  $|\mathcal{S}| = 23$ ,  $|\mathcal{A}| = 7$ ; Walker2d-v2  $|\mathcal{S}| = 17$ ,  $|\mathcal{A}| = 6$ ; Reacher-v2  $|\mathcal{S}| = 11$ ,  $|\mathcal{A}| = 2$ . All environments have a natural termination condition specified in the simulation environment.

**Policy Architecture.** All baseline algorithms involve training a parameterized policy  $\pi_\theta(a|s)$  using sample gradient estimates generated from the environment. The policy architecture is shared across all algorithms: a 2-layer feed-forward neural network with tanh non-linearity and  $h$  hidden units per layer. The input to the network is the state  $s \in \mathcal{S}$ . For all ES-based algorithms (Vanilla ES, CMA-ES, ARS and ASEBO), the output of the network is the action  $a_\theta(s) \in \mathcal{A}$ . For policy optimization algorithms (PPO, TRPO), the output of the network is a mean of Gaussian  $\mu_\theta(s)$  and we draw actions from a factorized Gaussian distribution  $a \sim \mathcal{N}(\mu_\theta(s), \sigma^2 \mathbb{I})$  where we separately parameterize a standard deviation parameter  $\sigma$  shared across dimensions. The sizes of hidden layers where: 4 for LQR, 16 for Swimmer-v2, Hopper-v2 and Reacher-v2, 32 for HalfCheetah-v2 and Walker2d-v2, reflecting the difficulty of each task.

**Optimization.** Our method (ASEBO) and most baselines (Vanilla ES, ARS, CMA-ES and PPO) apply SGD based methods and we apply the Adam optimizer to stabilize the gradients.

### 7.1.1 Baseline Algorithms

**Vanilla ES.** Vanilla ES is the simplest evolutionary algorithm applied in RL tasks [7, 30]. We apply the antithetic sampling scheme as applied in [7]. Our implementation does not rank the rewards as in [30], and as previously discussed does not include observation normalization.

**CMA-ES variants.** Covariance Matrix Adaptation Evolution Strategy is a state-of-the-art and popular black box optimization algorithm [16]. V<sub>k</sub>D-CMA-ES and LM-MA-ES are recently proposed variant designed for high dimensional blackbox functions. For V<sub>k</sub>D-CMA-ES we use the open source implementation from *pycma* available at <http://github.com/CMA-ES/pycma>. We use the default hyper-parameters in the original code base with the standard deviation parameter  $\sigma = 1.0$ . For LM-MA-ES we use the implementation from [26].

**ARS.** Augmented Random Search [25] is based on the code released by the original paper. We use the standard deviation  $\sigma = 0.02$  and learning rate  $\eta = 0.01$ . The hyper-parameters are tuned on top of the default hyper-parameters in the original code base. We remove the observation normalization utility in the original code for fair comparison.

**ASEBO.** We propose Adaptive Sample Efficient Blackbox Optimization in this work. Our algorithms have the following hyper-parameters: the covariance decay parameter  $\lambda = 0.995$  (slow decay), proportion of variance of the active (PCA) space  $\epsilon = 0.995$ , standard deviation parameter  $\sigma = 0.02$ . We set the learning rate  $\eta = 0.02$ .

**Trust Region Policy Optimization.** Trust Region Policy Optimization (TRPO) is based on the implementation of OpenAI baseline [12]. We use the default training hyper-parameters in the code base: we collect  $N = 1024$  samples per batch to compute a policy gradient, with the trust region size parameter  $\epsilon = 0.01$ . We remove the observation normalization utility in the original code for fair comparison.

**Proximal Policy Optimization.** Proximal Policy Optimization (PPO) [32] is also based on the implementation of OpenAI baseline [12]. We use the default hyper-parameters in the code base: we collect  $N = 2048$  samples per batch to compute policy gradients and set the clipping coefficient  $\epsilon = 0.2$ . The learning rate is set to be  $\alpha = 3 \cdot 10^{-5}$  for all environments. We remove the observation normalization utility in the original code for fair comparison.

## 7.2 Nevergrad Experiment Details

**Function Settings.** We tested the following functions: cigar, ellipsoid, sphere, sphere2, rosenbrock, rastrigin and lunacek. In each case we used  $d = 1000$  to evaluate ASEBO in a high dimensional setting.

**Algorithm Hyper-Parameters.** We use the same hyper-parameters across all functions. For ASEBO and VanillaES, we use  $\eta = 0.02$ . For ASEBO we set  $\lambda = 0.99$ . For Vkd-CMA-ES we use the default parameters from the pycma package, and for LM-MA-ES we use the implementation from [26].

## 8 The Algorithm - Additional Details & Analysis

We provide here few variations of the ASEBO algorithm from the main body of the paper, namely:

---

### Algorithm 3 ASEBO Algorithm - extended version

---

**Hyperparameters:** number of iterations of full sampling  $l$ , smoothing parameter  $\sigma > 0$ , step size  $\eta$ , PCA threshold  $\epsilon$ , decay rate  $\gamma$ , total number of iterations  $T$ .

**Input:** blackbox function  $F$ , vector  $\theta_0 \in \mathbb{R}^d$  where optimization starts.  $\text{Cov}_0 \in \{0\}^{d \times d}$ ,  $p^0 = 0$ .

**Output:** vector  $\theta_T$ .

**for**  $t = 0, \dots, T - 1$  **do**

**if**  $t < l$  **then**

    Take  $n_t = d$ . Sample  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  from  $\mathcal{N}(0, \mathbf{I}_d)$  (independently).

**else**

    1. Take top  $r$  eigenvalues  $\lambda_i$  of  $\text{Cov}_t$ , where  $r$  is smallest such that:  $\sum_{i=1}^r \lambda_i \geq \epsilon \sum_{i=1}^d \lambda_i$ , using its SVD as described in text and take  $n_t = r$ .

    2. Take the corresponding eigenvectors  $\mathbf{u}_1, \dots, \mathbf{u}_r \in \mathbb{R}^d$  and let  $\mathbf{U} \in \mathbb{R}^{d \times r}$  be obtained by stacking them together. Let  $\mathbf{U}^{\text{act}} \in \mathbb{R}^{d \times r}$  be obtained from stacking together some orthonormal basis of  $\mathcal{L}_{\text{active}}^{\text{ES}} \stackrel{\text{def}}{=} \text{span}\{\mathbf{u}_1, \dots, \mathbf{u}_r\}$ . Let  $\mathbf{U}^\perp \in \mathbb{R}^{d \times (d-r)}$  be obtained from stacking together some orthonormal basis of the orthogonal complement  $\mathcal{L}_{\text{active}}^{\text{ES}, \perp}$  of  $\mathcal{L}_{\text{active}}^{\text{ES}}$ .

    3. Sample  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  from  $\mathcal{N}(0, \sigma \Sigma)$  (independently), where  $\Sigma = \frac{1-p^t}{d} \mathbf{I}_d + \frac{p^t}{r} \mathbf{U} \mathbf{U}^\top (\mathbf{V0})$  or sample  $n_t$  vectors  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  as follows: with probability  $1 - p^t$  from  $\mathcal{N}(0, \mathbf{U}^\perp (\mathbf{U}^\perp)^\top)$  and with probability  $p^t$  from  $\mathcal{N}(0, \mathbf{U}^{\text{act}} (\mathbf{U}^{\text{act}})^\top)$  (**V1**).

    4. Renormalize  $\mathbf{g}_1, \dots, \mathbf{g}_{n_t}$  such that marginal distributions  $\|\mathbf{g}_i\|_2$  are  $\chi(d)$ .

  1. Compute  $\widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t)$  as:

$$\widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t) = \frac{1}{2n_t \sigma} \sum_{j=1}^{n_t} (F(\theta_t + \mathbf{g}_j) - F(\theta_t - \mathbf{g}_j)) \mathbf{g}_j.$$

  2. Set  $\text{Cov}_{t+1} = \lambda \text{Cov}_t + (1 - \lambda) \Gamma$ , where  $\Gamma = \widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t) (\widehat{\nabla}_{\text{MC}}^{\text{AT}} F_\sigma(\theta_t))^\top$ .

  3. Set  $p^{t+1} = p_{\text{opt}}$  for  $p_{\text{opt}}$  output by Algorithm 2 (from the main body) or  $p^{t+1} = \frac{\widehat{r}}{\widehat{r}+1}$ , where:

$$\widehat{r} = \frac{\|(\widehat{\nabla} F_\sigma(\theta_t))_{\text{active}}\|_2}{\|(\widehat{\nabla} F_\sigma(\theta_t))_{\perp}\|_2},$$

is computed by Algorithm 4 (see: below) and scalars  $\|(\widehat{\nabla} F_\sigma(\theta_t))_{\text{active}}\|_2$ ,  $\|(\widehat{\nabla} F_\sigma(\theta_t))_{\perp}\|_2$  stand for the estimates of  $\|(\nabla F_\sigma(\theta_t))_{\text{active}}\|_2$  and  $\|(\nabla F_\sigma(\theta_t))_{\perp}\|_2$ .

  4. Set  $\theta_{t+1} \leftarrow \theta_t + \eta \widehat{\nabla}_{\text{MC}}^{\text{AT}} F(\theta_t)$ .

---

- we propose one more method for sampling from heterogeneous distributions (see: version **V0** in Algorithm 3; the default one that we present in the main body is called **V1** here),
- we propose to use compressed sensing techniques (Algorithm 4) as an alternative to the contextual bandits method from the main body (Algorithm 2); the bandits method can be seen as an extension of the compressed sensing techniques.

---

**Algorithm 4** Explore estimator via compressed sensing

---

**Hyperparameters:** smoothing parameter  $\sigma$ , horizon  $C$ .

**Input:** subspaces:  $\mathcal{L}_{\text{active}}^{\text{ES}}$ ,  $\mathcal{L}_{\text{active}}^{\text{ES},\perp}$ , function  $F$ , vector  $\theta_t$ .

**Output:** ratio  $\hat{r}$ .

1. Initialize square norm averages  $s_0^{\text{active}} = s_0^\perp = 0$ .

**for**  $l = 1, \dots, C$  **do**

1. Sample  $\mathbf{g}_l^{\text{active}} \sim \mathcal{N}(0, \sigma \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}})$ .
2. Sample  $\mathbf{g}_l^\perp \sim \mathcal{N}(0, \sigma \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES},\perp}})$ .
3. Ask for  $F(\theta_t \pm \mathbf{g}_l^{\text{type}})$  for type  $\in \{\text{active}, \perp\}$ .
4. Compute  $v_l^{\text{type}} = \frac{1}{2\sigma} (F(\theta_t + \mathbf{g}_l^{\text{type}}) - F(\theta_t - \mathbf{g}_l^{\text{type}}))$ .
5. Compute  $s_l^{\text{active}} = \frac{l-1}{l} * s_{l-1}^{\text{active}} + \frac{(v_l^{\text{active}})^2}{l}$ .
6. Compute  $s_l^\perp = \frac{l-1}{l} * s_{l-1}^\perp + \frac{(v_l^\perp)^2}{l}$ .

**Return:**  $\hat{r} = \sqrt{\frac{s_C^{\text{active}}}{s_C^\perp}}$ .

---

### 8.1 Estimating the sensing ratio $r$ .

In this section we provide guarantees for the estimation of the ratio  $r$  as specified in Section 2.3 for Algorithm 4. Recall the definitions  $s_{\mathbf{U}^{\text{act}}} = \|\mathbf{U}^\top \nabla F(\theta_t)\|^2$  and  $s_{\mathbf{U}^\perp} = \|(\mathbf{U}^\perp)^\top \nabla F(\theta_t)\|^2$ .

Since  $\left| \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} - \mathbf{g}^\top \nabla F(\theta_t) \right| \leq \xi_{\mathbf{g}}(\theta_t)$ , when  $\mathbf{g} \sim \hat{P}$ , we recognize two cases. If  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}})$  the distribution of  $\frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \approx N(0, \|\mathbf{U}^\top \nabla F(\theta_t)\|^2)$ . Analogously when  $\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_{\mathbf{U}^\perp})$  the distribution of  $\frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \approx N(0, \|(\mathbf{U}^\perp)^\top \nabla F(\theta_t)\|^2)$ .

**Theorem 8.1.** Let  $0 < s < C$  and let  $\mathbf{g}_i \sim \mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES}}})$  for  $i = 1, \dots, s$  and  $\mathbf{g}_i \sim \mathcal{N}(0, \mathbf{I}_{\mathcal{L}_{\text{active}}^{\text{ES},\perp}})$  for  $i = s + 1, \dots, C$ . Let  $\hat{s}_{\mathbf{U}^{\text{ort}}} := \frac{1}{s} \sum_{j=1}^s \left( \frac{F(\theta + \sigma \mathbf{g}_j) - F(\theta - \sigma \mathbf{g}_j)}{2\sigma} \right)^2$ ,  $\hat{s}_{\mathbf{U}^\perp} := \frac{1}{C-s} \sum_{j=1}^{C-s} \left( \frac{F(\theta + \sigma \mathbf{g}_j) - F(\theta - \sigma \mathbf{g}_j)}{2\sigma} \right)^2$  and let  $\hat{r} = \sqrt{\frac{\hat{s}_{\mathbf{U}^{\text{ort}}}}{\hat{s}_{\mathbf{U}^\perp}}}$ . Given  $u, \epsilon > 0$  and  $\delta \in (\epsilon, 1)$ , the following holds.

1. If  $C = 2s$  for  $s \geq \frac{16}{u^2} \log\left(\frac{8}{\delta}\right)$  and under the mechanism from Algorithm 4 or
2. If  $\{\mathbf{g}_i\}_{i=1}^C$  are samples generated under  $\hat{P}$ ,  $\min(p^t, 1 - p^t) > u$  and  $C \geq \max\left(\frac{8}{(p^t - u)u^2}, \frac{8}{(1 - p^t - u)u^2}, \frac{2p^t + 2u/3}{u^2}\right) \log\left(\frac{12}{\delta}\right)$ ,

then with probability at least  $1 - \delta$ :

$$\sqrt{\frac{s_{\mathbf{U}^{\text{ort}}}(1 - u) - \frac{2\epsilon}{\delta}}{s_{\mathbf{U}^\perp}(1 + u) + \frac{2\epsilon}{\delta}}} \leq \hat{r} \leq \sqrt{\frac{s_{\mathbf{U}^{\text{act}}}(1 + u) + \frac{2\epsilon}{\delta}}{s_{\mathbf{U}^\perp}(1 - u) - \frac{2\epsilon}{\delta}}}.$$

*Proof.* First observe we introduce some notation.

$$\xi_{\mathbf{g}}^{(2)}(\theta_t) = \left( \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \right)^2 - (\mathbf{g}^\top \nabla F(\theta_t))^2$$

Observe that:

$$\begin{aligned} \xi_{\mathbf{g}}^{(2)}(\theta_t) &= \left| \left( \frac{F(\theta_t + \sigma \mathbf{g}) - F(\theta_t - \sigma \mathbf{g})}{2\sigma} \right)^2 - (\mathbf{g}^\top \nabla F(\theta_t))^2 \right| \\ &\leq \xi_{\mathbf{g}}(\theta_t)^2 + 2 |\mathbf{g}^\top \nabla F(\theta_t) \xi_{\mathbf{g}}(\theta_t)| \\ &\leq \sigma^4 \tau^2 \|\mathbf{g}\|^6 + 2\sigma^2 \tau L \|\mathbf{g}\|^4 \end{aligned}$$



Let  $\hat{s}_V = \hat{s}_V^0 + \frac{1}{s} \sum_{j=1}^s \xi_{\mathbf{g}_j}(\theta_t)$  and  $\hat{s}_{V^\perp} = \hat{s}_{V^\perp}^0 + \frac{1}{C-s} \sum_{j=1}^{C-s} \xi_{\mathbf{g}_j}(\theta_t)$ . Where  $s_{\mathbf{U}^{\text{act}}}^0 = \frac{1}{s} \sum_{j=1}^s (\nabla F(\theta_t)^\top \mathbf{g}_j)^2$  and  $s_{\mathbf{U}^\perp}^0 = \frac{1}{C-s} \sum_{j=s}^C (\nabla F(\theta_t)^\top \mathbf{g}_j)^2$ .

Notice that  $\nabla F(\theta_t)^\top \mathbf{g}$  is distributed as a Gaussian Random variable (with variance depending on the support of the covariance of  $\mathbf{g}$ ).

By concentration of squared gaussian random variables:

$$\begin{aligned} \mathbb{P} [|\hat{s}_V^0 - s_{\mathbf{U}^{\text{act}}}| \geq us_{\mathbf{U}^{\text{act}}}] &\leq 2 \exp\left(-\frac{su^2}{8}\right) \\ \mathbb{P} [|\hat{s}_{V^\perp}^0 - s_{\mathbf{U}^\perp}| \geq us_{\mathbf{U}^\perp}] &\leq 2 \exp\left(-\frac{(k-s)u^2}{8}\right) \end{aligned}$$

Consequently, with probability  $1 - 2 \exp\left(-\frac{su^2}{8}\right) - 2 \exp\left(-\frac{(k-s)u^2}{8}\right)$ , it holds that:

$$\frac{s_{\mathbf{U}^{\text{act}}}}{s_{\mathbf{U}^\perp}} \left(\frac{1+u}{1-u}\right) \geq \frac{\hat{s}_V^0}{\hat{s}_{V^\perp}^0} \geq \frac{s_{\mathbf{U}^{\text{act}}}}{s_{\mathbf{U}^\perp}} \left(\frac{1-u}{1+u}\right)$$

Notice that by Markov's inequality:

$$\mathbb{P} \left( \xi_{\mathbf{g}}(\theta_t) \geq \frac{2\epsilon}{\delta} \right) \leq \delta \frac{\mathbb{E} [\sigma^4 \tau^2 \|\mathbf{g}\|^6 + 2\sigma^2 \tau L \|\mathbf{g}\|^4]}{2\epsilon} \leq \frac{\delta}{2} \quad (24)$$

Since  $\sigma < \frac{1}{35} \sqrt{\frac{\epsilon \min(p^t, 1-p^t)}{\tau d^3 \max(L, 1)}}$ ,  $\mathbb{E} [\sigma^4 \tau^2 \|\mathbf{g}\|^6 + 2\sigma^2 \tau L \|\mathbf{g}\|^4] \leq \epsilon$ .

Regardless if  $\mathbf{g}$  was sampled from  $\mathcal{N}(0, \mathbf{I})$ ,  $\mathcal{N}(0, \mathbf{I}_U)$  or  $\mathcal{N}(0, \mathbf{I}_{U^\perp})$ .

### Case 1.

By definition,  $C = 2s$ , and therefore  $C - s = C/2$  and therefore:

$$\begin{aligned} \mathbb{P} [|\hat{s}_V^0 - s_{\mathbf{U}^{\text{act}}}| \geq us_{\mathbf{U}^{\text{act}}}] &\leq 2 \exp\left(-\frac{Cu^2}{16}\right) \\ \mathbb{P} [|\hat{s}_{V^\perp}^0 - s_{\mathbf{U}^\perp}| \geq us_{\mathbf{U}^\perp}] &\leq 2 \exp\left(-\frac{Cu^2}{16}\right) \end{aligned}$$

We require that:

$$\begin{aligned} 2 \exp\left(-\frac{Cu^2}{16}\right) &\leq \delta/4 \\ 2 \exp\left(-\frac{Cu^2}{16}\right) &\leq \delta/4 \end{aligned}$$

### Case 2.

In fact, by concentration results on Bernoulli variables, given  $\alpha > 0$ ,  $|s - kp^t| \leq k\alpha$  and  $|(k-s) - (1-p^t)k| \leq k\alpha$  with probability at least  $1 - \exp\left(-\frac{k\alpha^2}{2p^t}\right) - \exp\left(-\frac{k\alpha^2}{2p^t + 2\alpha/3}\right)$ .

Let  $\alpha = u$ . Conditioning on the events that  $|s - kp^t| \leq uk$  and  $|(k-s) - (1-p^t)k| \leq uk$ . We seek to ensure that:

$$\begin{aligned} 2 \exp\left(-\frac{(p^t - u)ku^2}{8}\right) &\leq \delta/6 \\ 2 \exp\left(-\frac{(1-p^t - u)ku^2}{8}\right) &\leq \delta/6 \\ \exp\left(-\frac{ku^2}{2p^t}\right) &\leq \delta/6 \\ \exp\left(-\frac{ku^2}{2p^t + 2u/3}\right) &\leq \delta/6 \end{aligned}$$

**Case 1 and 2** The following inequalities hold:

$$\frac{s_{\mathbf{U}^{\text{act}}}(1-u) - \frac{2\epsilon}{\delta}}{s_{\mathbf{U}^{\perp}}(1+u) + \frac{2\epsilon}{\delta}} \leq \frac{\hat{s}_V^0 - \frac{2\epsilon}{\delta}}{\hat{s}_{V^{\perp}}^0 + \frac{2\epsilon}{\delta}} \leq \frac{\hat{s}_V}{\hat{s}_{V^{\perp}}} \leq \frac{\hat{s}_V^0 + \frac{2\epsilon}{\delta}}{\hat{s}_{V^{\perp}}^0 - \frac{2\epsilon}{\delta}} \leq \frac{s_{\mathbf{U}^{\text{act}}}(1+u) + \frac{2\epsilon}{\delta}}{s_{\mathbf{U}^{\perp}}(1-u) - \frac{2\epsilon}{\delta}}$$

The union bound yields the desired result. And therefore the result follows.  $\square$