
Guided evolutionary strategies: Augmenting random search with surrogate gradients

Niru Maheswaranathan¹ Luke Metz¹ George Tucker¹ Dami Choi¹ Jascha Sohl-Dickstein¹

Abstract

Many applications in machine learning require optimizing a function whose true gradient is inaccessible, but where *surrogate* gradient information (directions that may be correlated with, but not necessarily identical to, the true gradient) is available instead. This arises when an approximate gradient is easier to compute than the full gradient (e.g. in meta-learning or unrolled optimization), or when a true gradient is intractable and is replaced with a surrogate (e.g. in certain reinforcement learning applications or training networks with discrete variables). We propose *Guided Evolutionary Strategies*, a method for optimally using surrogate gradient directions along with random search. We define a search distribution for evolutionary strategies that is elongated along a subspace spanned by the surrogate gradients. This allows us to estimate a descent direction which can then be passed to a first-order optimizer. We analytically and numerically characterize the trade-offs that result from tuning how strongly the search distribution is stretched along the guiding subspace, and use this to derive a setting of the hyperparameters that works well across problems. Finally, we apply our method to example problems, demonstrating an improvement over both standard evolutionary strategies and first-order methods that directly follow the surrogate gradient.

1. Introduction

Optimization in machine learning involves minimizing cost functions where gradient information may or may not be known. When gradient information is available, first-order

¹Google Research, Brain Team, Mountain View, CA, United States. Correspondence to: Niru Maheswaranathan <nirum@google.com>.

methods such as gradient descent are popular due to their ease of implementation, memory efficiency, and convergence guarantees (Sra et al., 2012). When gradient information is not available, however, we turn to zeroth-order optimization methods, including random search methods such as evolutionary strategies (Rechenberg, 1973; Nesterov & Spokoiny, 2011; Choromanski et al., 2018; Salimans et al., 2017).

However, what if only partial gradient information is available? That is, what if one has access to *surrogate* gradients that are correlated with the true gradient, but may be biased in some unknown fashion? Naively, there are two extremal approaches to optimization with surrogate gradients. On one hand, you could ignore the surrogate gradient information entirely and perform zeroth-order optimization, using methods such as evolutionary strategies to estimate a descent direction. These methods exhibit poor convergence properties when the parameter dimension is large (Duchi et al., 2015). On the other hand, you could directly feed the surrogate gradients to a first-order optimization algorithm. However, bias in the surrogate gradients will interfere with optimizing the target problem (Tucker et al., 2017). Ideally, we would like a method that combines the complementary strengths of these two approaches: we would like to combine the unbiased descent direction estimated with evolutionary strategies with the low-variance estimate given by the surrogate gradient. In this work, we propose a method for doing this called guided evolutionary strategies (Guided ES).

The critical assumption underlying Guided ES is that we have access to surrogate gradient information, but not the true gradient. This scenario arises in a wide variety of machine learning problems, which typically fall into two categories: cases where the true gradient is unknown or not defined, and cases where the true gradient is hard or expensive to compute. Examples of the former include: models with discrete stochastic variables (where straight through estimators (Bengio et al., 2013; van den Oord et al., 2017) or Concrete/Gumbel-Softmax methods (Maddison et al., 2016; Jang et al., 2016) are commonly used) and learned models in reinforcement learning (e.g. for Q functions (Watkins & Dayan, 1992; Mnih et al., 2013; 2015; Lillicrap et al., 2015)

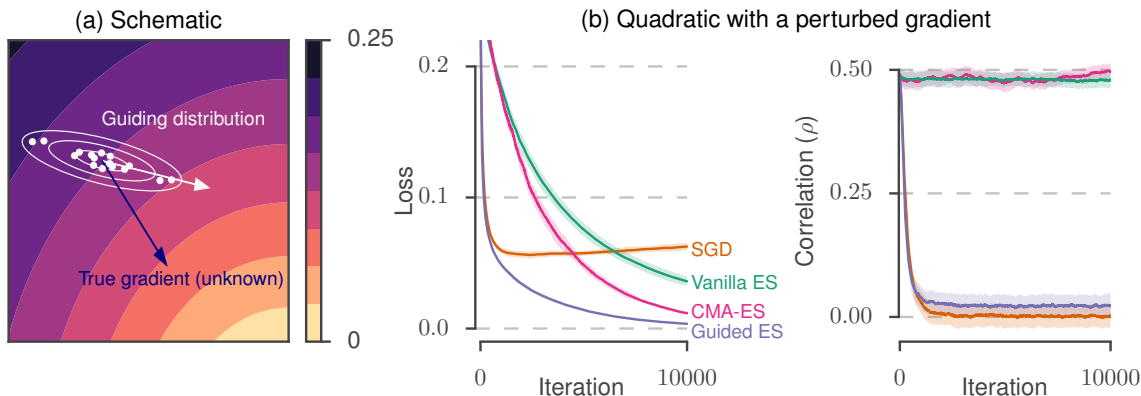


Figure 1. (a) Schematic of guided evolutionary strategies. We perform a random search using a distribution (white contours) elongated along a subspace (white arrow) which we are given instead of the true gradient (blue arrow). (b) Comparison of different algorithms on a quadratic loss, where a bias is explicitly added to the gradient to mimic situations where the true gradient is unknown. The loss (left) and correlation between surrogate and true gradient (right) are shown during optimization. See §4.1 for experimental details.

or value estimation (Mnih et al., 2016)). For the latter, examples include optimization using truncated backprop through time (Rumelhart et al., 1985; Williams & Peng, 1990; Wu et al., 2018). Surrogate gradients also arise in situations where the gradients are explicitly modified during training, as in feedback alignment (Lillicrap et al., 2014) and related methods (Nøkland, 2016).

The key idea in Guided ES is to keep track of a low-dimensional subspace, defined by the recent history of surrogate gradients during optimization, which we call the *guiding subspace*. We then perform a finite difference random search (as in evolutionary strategies) preferentially within this subspace. By concentrating our search samples in a low-dimensional subspace where the true gradient has non-negative support, we dramatically reduce the variance of the search direction.

Our contributions in this work are:

- a method for combining surrogate gradient information with random search,
- an analysis of the bias-variance tradeoff underlying the technique (§3.3),
- a scheme for choosing optimal hyperparameters for the method (§3.4), and
- applications to example problems (§4).

For a demo of the method, please see:

<https://github.com/brain-research/guided-evolutionary-strategies>

2. Related Work

This work builds upon a random search method known as evolutionary strategies (Rechenberg, 1973; Nesterov & Spokoiny, 2011), or ES for short, which generates a descent direction via finite differences over random perturbations of parameters. ES has seen a resurgence in popularity in recent years (Salimans et al., 2017; Mania et al., 2018; Ha, 2018; Ha & Schmidhuber, 2018; Houthoofd et al., 2018; Cui et al., 2018). Our method can primarily be thought of as a modification to ES where we augment the search distribution using surrogate gradients.

Extensions of ES that modify the search distribution use natural gradient updates in the search distribution (Wierstra et al., 2008) or construct non-Gaussian search distributions (Glasmachers et al., 2010). The idea of using gradients in concert with evolutionary algorithms was proposed by (Lehman et al., 2017b), who use gradients of a network with respect to its *inputs* (as opposed to parameters) to augment ES. Other methods for adapting the search distribution include covariance matrix adaptation ES (CMA-ES) (Hansen, 2016), which uses the recent history of descent steps to adapt the distribution over parameters, or variational optimization (Staines & Barber, 2012), which optimizes the parameters of a probability distribution over model weights. Guided ES, by contrast, adapts the search distribution using surrogate gradient information. In addition, we never need to work with or compute a full $n \times n$ covariance matrix. Finally, Hansen (2011) proposes a modification to CMA-ES that accommodates external information by directly injecting candidate solutions into the algorithm. By contrast, the algorithm proposed here augments the search distribution itself. More recently, Pourchot & Sigaud (2018) propose augmenting evolutionary strategies with policy gradient information

for reinforcement learning applications.

3. Guided evolutionary strategies

3.1. Vanilla ES

We wish to minimize a function $f(x)$ over a parameter space in n -dimensions ($x \in \mathbb{R}^n$), where ∇f is either unavailable or uninformative. A popular approach is to estimate a descent direction with stochastic finite differences (commonly referred to as evolutionary strategies (Rechenberg, 1973) or random search (Rastrigin, 1963)). Here, we use antithetic sampling (Owen, 2014) (using a pair of function evaluations at $x + \epsilon$ and $x - \epsilon$) to reduce variance. This estimator is defined as:

$$g = \frac{\beta}{2\sigma^2 P} \sum_{i=1}^P \epsilon_i (f(x + \epsilon_i) - f(x - \epsilon_i)), \quad (1)$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma^2 I)$, and P is the number of sample pairs. We will set P to one for all experiments, and when analyzing optimal hyperparameters. The overall scale of the estimate (β) and variance of the perturbations (σ^2) are constants, to be chosen as hyperparameters. This estimate solely relies on computing $2P$ function evaluations. However, it tends to have high variance, thus requiring a large number of samples to be practical, and scales poorly with the dimension n . We refer to this estimator as vanilla evolutionary strategies (or vanilla ES) in subsequent sections.

3.2. Guided search

Even when we do not have access to ∇f , we frequently have additional information about f , either from prior knowledge or gleaned from previous iterates during optimization. To formalize this, we assume we are given a set of vectors which may correspond to biased or corrupted gradients. That is, these vectors are correlated (but need not be perfectly aligned) with the true gradient. If we are given a single vector or surrogate gradient for a given parameter iterate, we can generate a subspace by keeping track of the previous k surrogate gradients encountered during optimization. We use U to denote an $n \times k$ orthonormal basis for the subspace spanned by these vectors (i.e., $U^T U = I_k$).

We leverage this information by changing the distribution of ϵ_i in eq. (1) to $\mathcal{N}(0, \sigma^2 \Sigma)$ with

$$\Sigma = \frac{\alpha}{n} I_n + \frac{1 - \alpha}{k} U U^T,$$

where k and n are the subspace and parameter dimensions, respectively, and α is a hyperparameter that trades off variance between the full parameter space and the subspace. Setting $\alpha = 1$ recovers the vanilla ES estimator (and ignores the guiding subspace), but as we show choosing $\alpha < 1$ can

result in significantly improved performance. The other hyperparameter is the scale β in (1), which controls the size of the estimated descent direction. The parameter σ^2 controls the overall scale of the variance, and will drop out of the analysis of the bias and variance below, due to the $\frac{1}{\sigma^2}$ factor in (1). In practice, if $f(x)$ is stochastic, then increasing σ^2 will dampen noise in the gradient estimate, while decreasing σ^2 reduces the error induced by third and higher-order terms in the Taylor expansion of f below. For an exploration of the effects of σ^2 in ES, see (Lehman et al., 2017a).

Samples of ϵ_i can be generated efficiently as $\epsilon_i = \sigma \sqrt{\frac{\alpha}{n}} \epsilon + \sigma \sqrt{\frac{1 - \alpha}{k}} U \epsilon'$ where $\epsilon \sim N(0, I_n)$ and $\epsilon' \sim N(0, I_k)$. Our estimator requires $2P$ function evaluations in addition to the cost of computing the surrogate gradient. Note that it is often possible to compute these function evaluations in parallel (Salimans et al., 2017).

Figure 1a depicts the geometry underlying our method. Instead of the true gradient (blue arrow), we are given a surrogate gradient (white arrow) which is correlated with the true gradient. We use this to form a guiding distribution (denoted with white contours) and use this to draw samples (white dots) which we use as part of a random search procedure. Figure 1b demonstrates the performance of the method on a toy problem, and is discussed in §4.1.

For the purposes of analysis, suppose ∇f exists. We can approximate the function in the local neighborhood of x using a second order Taylor approximation: $f(x + \epsilon) \approx f(x) + \epsilon^T \nabla f(x) + \frac{1}{2} \epsilon^T \nabla^2 f(x) \epsilon$. For the remainder of §3, we take this second order Taylor expansion to be exact. By substituting this expression into (1), we see that our estimate g is equal to

$$g = \frac{\beta}{\sigma^2 P} \sum_{i=1}^P (\epsilon_i \epsilon_i^T) \nabla f(x). \quad (2)$$

Note that even terms in the expansion cancel out in the expression for g due to antithetic sampling.

Importantly, regardless of the choice of α and β , the Guided ES estimator always provides a descent direction in expectation. The mean of the estimator in eq. (2) is $\mathbb{E}[g] = \beta \Sigma \nabla f(x)$ which corresponds to the gradient multiplied by a positive semi-definite (PSD) matrix, thus the update $(-\mathbb{E}[g])$ remains a descent direction. This desirable property ensures that α trades off variance for “safe” bias. That is, the bias will never produce an *ascent* direction when we are trying to *minimize* f .

Algorithm 1 provides pseudocode for guided evolutionary strategies. Note that computation of the loss at the perturbed parameters (the inner loop) may be done in parallel. The computational and memory costs of using Guided ES to compute parameter updates, compared to standard (vanilla)

ES and gradient descent, are outlined in §3.5.

3.3. Tradeoff between variance and safe bias

As we have alluded to, there is a bias-variance tradeoff lurking within our estimate g . In particular, by emphasizing the search in the full space (i.e., choosing α close to 1), we reduce the bias in our estimate at the cost of increased variance. Emphasizing the search along the guiding subspace (i.e., choosing α close to 0) will induce a bias in exchange for a potentially large reduction in variance, especially if the subspace dimension k is small relative to the parameter dimension n . Below, we analytically and numerically characterize this tradeoff.

The alignment between the k -dimensional orthonormal guiding subspace (U) and the true gradient ($\nabla f(x)$) is a key quantity for understanding the bias-variance tradeoff. We characterize this alignment using a k -dimensional vector of uncentered correlation coefficients ρ , whose elements are the correlation between the gradient and every column of U . That is, $\rho_i = \frac{\nabla f(x)^T U_i}{\|\nabla f(x)\|}$. This correlation $\|\rho\|_2$ varies between zero (if the gradient is orthogonal to the subspace) and one (if the gradient is fully contained in the subspace).

We define the normalized bias of the Guided ES gradient estimate as (see Supplementary Material A.1):

$$\tilde{b} = \left(\beta \frac{\alpha}{n} - 1 \right)^2 + \left(\beta^2 \frac{(1-\alpha)^2}{k^2} + 2\beta \frac{(1-\alpha)}{k} \left(\beta \frac{\alpha}{n} - 1 \right) \right) \|\rho\|_2^2 \quad (3)$$

where again β is a scale factor and α is part of the parameterization of the covariance matrix that trades off variance in the full parameter space for variance in the guiding subspace ($\Sigma = \frac{\alpha}{n}I + \frac{(1-\alpha)}{k}UU^T$). We see that the normalized squared bias consists of two terms: the first is a contribution from the search in the full space and is thus independent of ρ , whereas the second depends on the squared norm of the uncentered correlation, $\|\rho\|_2^2$.

Similarly, we can compute the normalized variance as (see Supplementary Material A.2):

$$\tilde{v} = \beta^2 \left(\frac{\alpha^2}{n^2} + \frac{\alpha}{n} \right) + \beta^2 \left(\frac{(1-\alpha)^2}{k^2} + 2\frac{\alpha(1-\alpha)}{kn} + \frac{(1-\alpha)}{k} \right) \|\rho\|_2^2 \quad (4)$$

Equations (3) and (4) quantify the bias and variance of our estimate as a function of the subspace and parameter dimensions (k and n), the parameters of the distribution (α and β), and the correlation $\|\rho\|_2$. Note that for simplicity we have set the number of pairs of function evaluations, P , to one. As P increases, the variance will decrease linearly, at the cost of extra function evaluations.

Figure 2 explores the tradeoff between normalized bias and variance for different settings of the relevant hyperparameters (α and β) for example values of $\|\rho\|_2 = 0.23$, $k = 3$, and $n = 100$. Figure 2c shows the sum of the normalized bias plus variance, the global minimum of which (blue star) can be used to choose optimal values for the hyperparameters, discussed in the next section.

3.4. Choosing optimal hyperparameters by minimizing error in the estimate

The expressions for the normalized bias and variance depend on the subspace and parameter dimensions (k and n , respectively), the hyperparameters of the guiding distribution (α and β) and the uncentered correlation between the true gradient and the subspace ($\|\rho\|_2$). All of these quantities except for the correlation $\|\rho\|_2$ are known or defined in advance.

To choose optimal hyperparameters, we minimize the sum of the normalized bias and variance, (equivalent to the expected normalized squared error in the gradient estimate, $\tilde{b} + \tilde{v} = \frac{\mathbb{E}[\|g - \nabla f(x)\|_2^2]}{\|\nabla f(x)\|_2^2}$). This objective becomes:

$$\tilde{b} + \tilde{v} = \left[2\beta^2 \frac{\alpha^2}{n^2} + (\beta^2 - 2\beta) \frac{\alpha}{n} + 1 \right] + \left[2\beta^2 \frac{(1-\alpha)^2}{k^2} + 4\beta^2 \frac{\alpha(1-\alpha)}{kn} + (\beta^2 - 2\beta) \frac{(1-\alpha)}{k} \right] \|\rho\|_2^2, \quad (5)$$

subject to the feasibility constraints $\beta \geq 0$ and $0 \leq \alpha \leq 1$.

As further motivation for this hyperparameter objective, in the simple case that $f(x) = \frac{1}{2}\|x\|_2^2$ then minimizing eq. (5) also results in the hyperparameters that cause SGD to most rapidly descend $f(x)$. See Supplementary Material C for a derivation of this relationship.

We can solve for the optimal tradeoff (α^*) and scale (β^*) hyperparameters as a function of $\|\rho\|_2$, k , and n . Figure 3a shows the optimal value for the tradeoff hyperparameter (α^*) in the 2D plane spanned by the correlation ($\|\rho\|_2$) and ratio of the subspace dimension to the parameter dimension $\frac{k}{n}$. Remarkably, we see that for large regions of the ($\|\rho\|_2, \frac{k}{n}$) plane, the optimal value for α is either 0 or 1. In the upper left (blue) region, the subspace is of high quality (highly correlated with the true gradient) and small relative to the full space, so the optimal solution is to place all of the weight in the subspace, setting α to zero (therefore $\Sigma \propto UU^T$). In the bottom right (orange) region, we have the opposite scenario, where the subspace is large and low-quality, thus the optimal solution is to place all of the weight in the full space, setting α to one (equivalent to vanilla ES, $\Sigma \propto I$). The strip in the middle is an intermediate regime where the

Algorithm 1 Guided Evolutionary Strategies

Require: Initial parameters $x^{(0)}$, Loss function $f(x)$, Learning rate η , Hyperparameters $(\alpha, \beta, \sigma^2, P)$

- 1: **for** $t = 1$ to T **do**
- 2: Get surrogate gradient $\nabla f(x^{(t)})$
- 3: Update low-dimensional guiding subspace U with the surrogate gradient
- 4: Define search covariance $\Sigma = \frac{\alpha}{n}I + \frac{1-\alpha}{k}UU^T$
- 5: **for** $i = 1$ to P **do**
- 6: Sample perturbation $\epsilon_i \sim \mathcal{N}(0, \sigma^2\Sigma)$
- 7: Compute antithetic pair of losses $f(x + \epsilon_i)$ and $f(x - \epsilon_i)$
- 8: **end for**
- 9: Compute Guided ES gradient estimate $g = \frac{\beta}{2\sigma^2P} \sum_{i=1}^P \epsilon_i [f(x + \epsilon_i) - f(x - \epsilon_i)]$
- 10: Update parameters using gradient descent $x \leftarrow x - \eta g$
- 11: **end for**

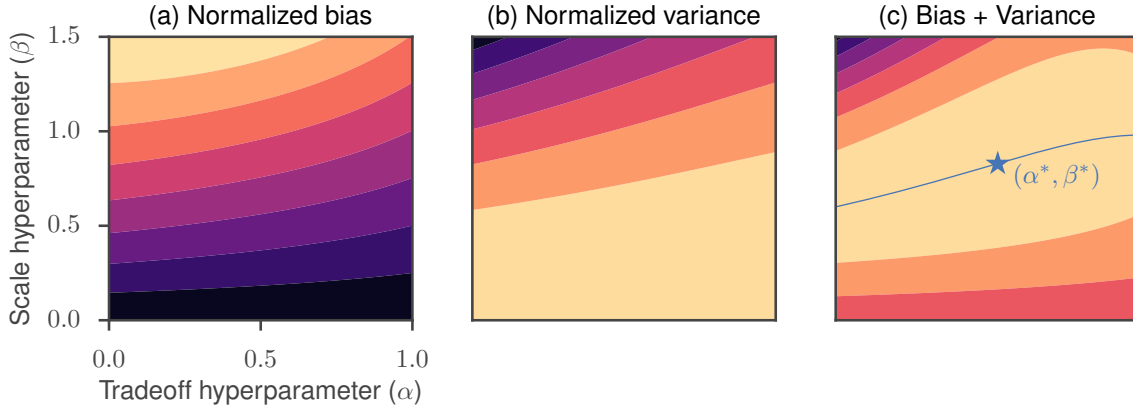


Figure 2. Exploring the tradeoff between variance and safe bias in Guided ES. Contour plots of normalized bias \tilde{b} (a), normalized variance \tilde{v} (b), and the sum of both (c) are shown as a function of the tradeoff (α) and scale (β) hyperparameters, for a fixed $\|\rho\|_2 = 0.23$. For these plots, the subspace dimension was set to $k = 3$ and the parameter dimension was set to $n = 100$. The blue line in (c) denotes the optimal β for every value of α , and the star denotes the global optimum.

optimal α is between 0 and 1.

We can also derive an analytic expression for *when* this transition in optimal hyperparameters occurs. To do this, we use the reparameterization $\theta = \begin{pmatrix} \alpha\beta \\ (1-\alpha)\beta \end{pmatrix}$. This allows us to express the objective in (5) as a least squares problem $\frac{1}{2}\|A\theta - b\|_2^2$, subject to a non-negativity constraint ($\theta \succeq 0$), where A and b depend solely on the problem data: k , n , and $\|\rho\|_2$ (see Supplementary Material B.1 for details). In addition, A is always a positive semi-definite matrix, so the reparameterized problem is convex. We are particularly interested in the point where the non-negativity constraint becomes tight, as this is when there is a change in the optimal hyperparameters. As we have formulated the problem of finding optimal hyperparameters as a convex problem with a convex constraint, we can borrow tools from convex analysis to write down an equation for when the constraint becomes tight. In particular, formulating the Lagrange dual of the problem and solving for the KKT conditions allows

us to identify this point using the complementary slackness conditions (Boyd & Vandenberghe, 2004). This yields the equations $\|\rho\|_2 = \sqrt{\frac{k+4}{n+4}}$ and $\|\rho\|_2 = \sqrt{\frac{k}{n}}$ (see Supplementary Material B.2), which are shown in Figure 3a, and line up with the numerical solution. Figure 3b further demonstrates this tradeoff, specifically for the intermediate regime $\left(\sqrt{\frac{k}{n}} \leq \|\rho\|_2 \leq \sqrt{\frac{k+4}{n+4}}\right)$. For fixed $n = 100$, we plot four curves for k ranging from 1 to 30. As $\|\rho\|_2$ increases, the optimal hyperparameters sweep out a curve from $(\alpha^* = 1, \beta^* = \frac{n}{n+2})$ to $(\alpha^* = 0, \beta^* = \frac{k}{k+2})$.

In practice, the correlation between the gradient and the guiding subspace is typically unknown. However, we find that ignoring $\|\rho\|_2$ and setting $\beta = 2$ and $\alpha = \frac{1}{2}$ works well (these are the values used for all experiments in this paper). A direction for future work would be to estimate the correlation $\|\rho\|_2$ online, and to use this to choose hyperparameters by minimizing equation (5).

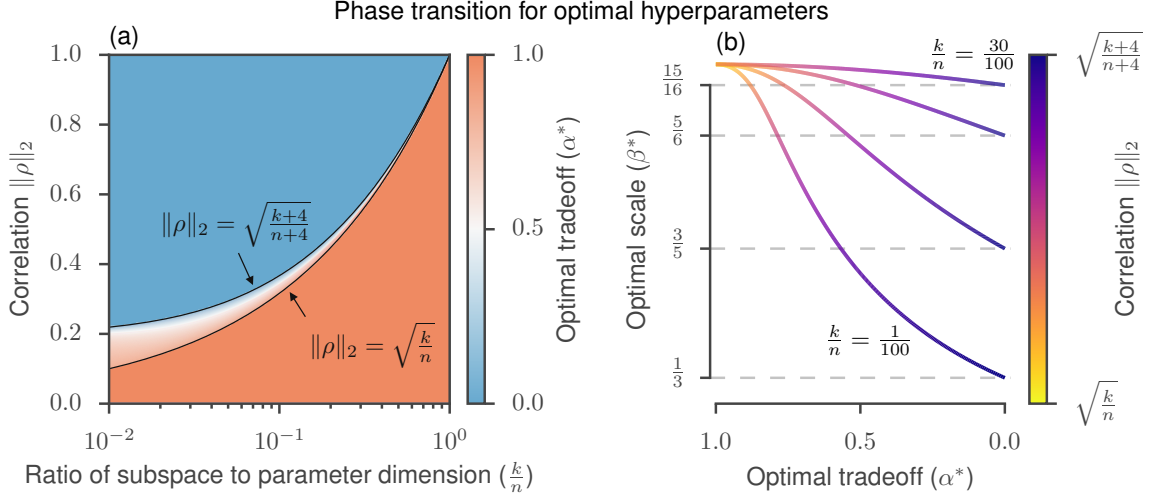


Figure 3. Choosing optimal hyperparameters. (a) Different regimes of optimal hyperparameters in the $(\frac{k}{n}, \|\rho\|_2)$ plane are shown as shaded regions. See §3.4 for details. (b) As $\|\rho\|_2$ increases, the optimal hyperparameters sweep out a curve from $(\alpha^* = 1, \beta^* = \frac{n}{n+2})$ to $(\alpha^* = 0, \beta^* = \frac{k}{k+2})$.

3.5. Computational and memory cost

Here, we outline the computational and memory costs of Guided ES and compare them to standard (vanilla) evolutionary strategies and gradient descent. As elsewhere in the paper, we define the parameter dimension as n , the Guided ES subspace dimension as k , and the number of pairs of function evaluations (for evolutionary strategies) as P . We denote the cost of computing the full loss as F_0 , and we assume that at every iteration we compute a surrogate gradient (for Guided ES and gradient descent) which has cost F_1 . Note that for standard training of neural networks with backpropagation, these quantities have similar cost ($F_1 \approx 2F_0$), however for some applications (such as unrolled optimization discussed in §4.2) these can be very different.

4. Applications

4.1. Quadratic function with a biased gradient

We first test our method on a toy problem where we control the bias of the surrogate gradient explicitly. We generated random quadratic problems of the form $f(x) = \frac{1}{2}\|Ax - b\|_2^2$ where the entries of A and b were drawn independently from a standard normal distribution, but rather than allow the optimizers to use the true gradient, we (for illustrative purposes) added a random bias to generate surrogate gradients. Figure 1b compares the performance of stochastic gradient descent (SGD) with standard (vanilla) evolutionary strategies (ES), CMA-ES, and Guided ES. For this, and all of the results in this paper, we set the hyperparameters as $\beta = 2$ and $\alpha = \frac{1}{2}$, as described above.

We see that Guided ES proceeds in two phases: it initially quickly descends the loss as it follows the biased gradient, and then transitions into random search. Vanilla ES and CMA-ES, however, do not get to take advantage of the information available in the surrogate gradient, and converge more slowly. We see this also in the plot of the uncentered correlation (ρ) between the true gradient and the surrogate gradient in Figure 1c.

This is exactly the transition discussed in §3.4, where as ρ varies we move from the regime where we want to only search in the subspace to the regime where we want to ignore the subspace (diagrammed in Figure 3a). Although we know ρ in this case, we find the practical choice of fixing the hyperparameters (discussed in §3.4) independently of ρ works well. Further experimental details are provided in Supplementary Material D.1.

4.2. Unrolled optimization

Another application where surrogate gradients are available is in *unrolled* optimization. Unrolled optimization refers to taking derivatives through an optimization process. For example, this approach has been used to optimize hyperparameters (Domke, 2012; Maclaurin et al., 2015; Baydin et al., 2017), to stabilize training (Metz et al., 2016), and even to train neural networks to act as optimizers (Andrychowicz et al., 2016; Wichrowska et al., 2017; Li & Malik, 2017; Lv et al., 2017). Taking derivatives through optimization with a large number of steps is costly, so a common approach is to instead choose a small number of unrolled steps, and

Table 1. Per-iteration compute and memory costs for gradient descent, standard (vanilla) evolutionary strategies, and the method proposed in this paper, guided evolutionary strategies. Here, F_0 is the cost of a function evaluation, F_1 is the cost of computing a surrogate gradient, n is the parameter dimension, k is the subspace dimension used for the guiding subspace, and P is the number of pairs of function evaluations used for the evolutionary strategies algorithms.

Algorithm	Computational cost	Memory cost
Gradient descent	F_1	n
Vanilla evolutionary strategies	$2PF_0$	n
Guided evolutionary strategies	$F_1 + 2PF_0$	$(k + 1)n$

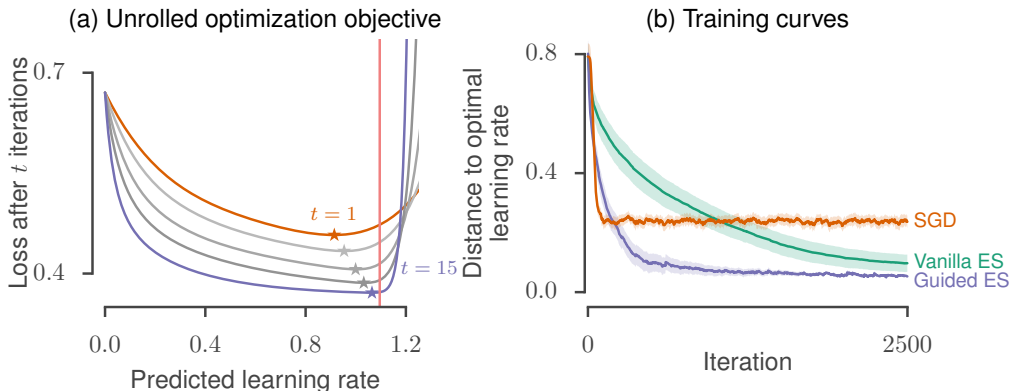


Figure 4. Unrolled optimization. (a) Bias in the loss landscape of unrolled optimization for small numbers of unrolled optimization steps (t). (b) Training curves (shown as distance from the optimum) for training a multi-layer perceptron to predict the optimal learning rate as a function of the eigenvalues of the function to optimize. See §4.2 for details.

use that as a target for training. However, (Wu et al., 2018) recently showed that this approach yields biased gradients.

To demonstrate the utility of Guided ES here, we trained multi-layer perceptrons (MLP) to predict the learning rate for a target problem, using as input the eigenvalues of the Hessian at the current iterate. Figure 4a shows the bias induced by unrolled optimization, as the number of optimization steps ranges from one iteration (orange) to 15 (blue). We compute the surrogate gradient of the parameters in the MLP using the loss after one SGD step. Figure 4b, we show the absolute value of the difference between the optimal learning rate and the MLP prediction for different optimization algorithms. Further experimental details are provided in Supplementary Material D.2.

4.3. Synthesizing gradients for a guiding subspace

Next, we explore using Guided ES in the scenario where the surrogate gradient is not provided, but instead we train a model to generate surrogate gradients (we call these synthetic gradients). In real-world applications, training a model to produce synthetic gradients is the basis of model-based and actor-critic methods in RL (Lillicrap et al., 2015; Heess et al., 2015) and has been applied to decouple training across neural network layers (Jaderberg et al., 2016) and to generate policy gradients (Houthoofd et al., 2018). A key

challenge with such an approach is that early in training, the model generating the synthetic gradients is untrained, and thus will produce biased gradients. In general, it is unclear during training when following these synthetic gradients will be beneficial.

We define a parametric model, $M(x; \theta)$ (an MLP), which provides synthetic gradients for the target problem f . The target model $M(\cdot)$ is trained online to minimize mean squared error against evaluations of $f(x)$. Figure 5 compares vanilla ES, Guided ES, and the Adam optimizer (Kingma & Ba, 2014). We show training curves for these methods in Figure 5a, and the correlation between the synthetic gradient and true gradients for Guided ES in Figure 5b. Despite the fact that the quality of the synthetic gradients varies wildly during optimization, Guided ES consistently makes progress on the target problem. Further experimental details are provided in Supplementary Material D.3.

5. Discussion

We have introduced guided evolutionary strategies (Guided ES), an optimization algorithm which combines the benefits of first-order methods and random search, when we have access to surrogate gradients that are correlated with the true gradient. We analyzed the bias-variance tradeoff inherent in

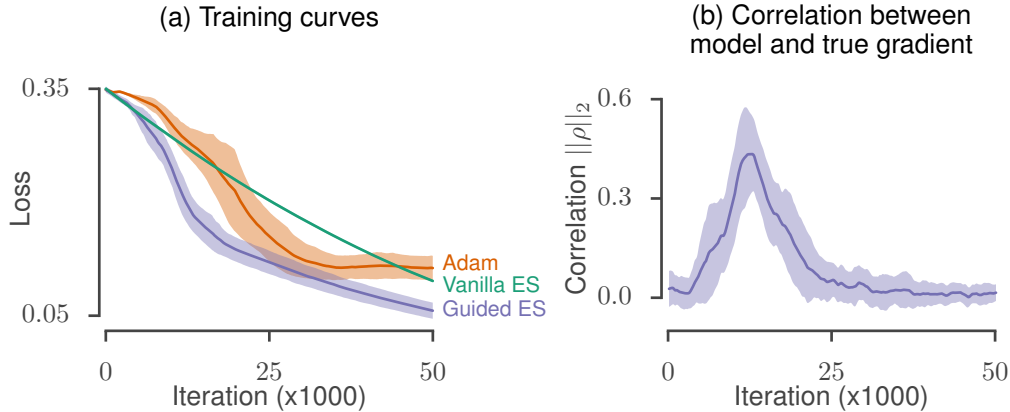


Figure 5. Synthetic gradients serving as the guiding subspace for Guided ES. (a) Loss curves when using synthetic gradients to minimize a target quadratic problem. (b) Correlation between the synthetic update direction and the true gradient during optimization for Guided ES.

our method analytically, and demonstrated the generality of the technique by applying it to unrolled optimization, synthetic gradients, and training neural networks with discrete variables.

One caveat of the method is that up to second order, the expected Guided ES estimator is equivalent to multiplying the true gradient by a positive semi-definite (PSD) matrix. The product of PSD matrices is not in general PSD, so care must be taken when passing the Guided ES update vector to an optimization method that also multiplies its input by a PSD matrix (such as RMSProp or Adam). Future work will determine if this is an issue in practice.

The nearly binary nature of the tradeoff in optimal hyperparameters (Figure 3) suggests that methods that switch between biased gradient descent and vanilla evolutionary strategies would also likely work, if the switch between methods occurs close to the appropriate value of $\|\rho\|_2$. An interesting direction for future work is to estimate the correlation $\|\rho\|_2$ during optimization, and use this to adapt the hyperparameters online.

In conclusion, Guided ES is a general technique which has widespread applications in machine learning due to the ubiquity of problems where the gradient is either uninformative or only noisy gradient estimates are tractable.

ACKNOWLEDGEMENTS

The authors would like to thank Alex Alemi and Ben Poole for discussions, and Roy Frostig, Katherine Lee, Benjamin Eysenbach, and Brian Cheung for comments on the manuscript.

References

- Andrychowicz, M., Denil, M., Gomez, S., Hoffman, M. W., Pfau, D., Schaul, T., and de Freitas, N. Learning to learn by gradient descent by gradient descent. In *Advances in Neural Information Processing Systems*, pp. 3981–3989, 2016.
- Baydin, A. G., Cornish, R., Rubio, D. M., Schmidt, M., and Wood, F. Online learning rate adaptation with hypergradient descent. *arXiv preprint arXiv:1703.04782*, 2017.
- Bengio, Y., Léonard, N., and Courville, A. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Boyd, S. and Vandenberghe, L. *Convex optimization*. Cambridge university press, 2004.
- Choromanski, K., Rowland, M., Sindhvani, V., Turner, R. E., and Weller, A. Structured evolution with compact architectures for scalable policy optimization. *arXiv preprint arXiv:1804.02395*, 2018.
- Cui, X., Zhang, W., Tüske, Z., and Picheny, M. Evolutionary stochastic gradient descent for optimization of deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 6049–6059, 2018.
- Domke, J. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pp. 318–326, 2012.
- Duchi, J. C., Jordan, M. I., Wainwright, M. J., and Wibisono, A. Optimal rates for zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on Information Theory*, 61(5):2788–2806, 2015.

- Glasmachers, T., Schaul, T., Yi, S., Wierstra, D., and Schmidhuber, J. Exponential natural evolution strategies. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pp. 393–400. ACM, 2010.
- Ha, D. Neuroevolution for deep reinforcement learning problems. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 421–431. ACM, 2018.
- Ha, D. and Schmidhuber, J. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- Hansen, N. Injecting external solutions into CMA-ES. *arXiv preprint arXiv:1110.4181*, 2011.
- Hansen, N. The CMA evolution strategy: A tutorial. *arXiv preprint arXiv:1604.00772*, 2016.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Erez, T., and Tassa, Y. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, pp. 2944–2952, 2015.
- Houthoofd, R., Chen, R. Y., Isola, P., Stadie, B. C., Wolski, F., Ho, J., and Abbeel, P. Evolved policy gradients. *arXiv preprint arXiv:1802.04821*, 2018.
- Jaderberg, M., Czarnecki, W. M., Osindero, S., Vinyals, O., Graves, A., Silver, D., and Kavukcuoglu, K. Decoupled neural interfaces using synthetic gradients. *arXiv preprint arXiv:1608.05343*, 2016.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lehman, J., Chen, J., Clune, J., and Stanley, K. O. Es is more than just a traditional finite-difference approximator. *arXiv preprint arXiv:1712.06568*, 2017a.
- Lehman, J., Chen, J., Clune, J., and Stanley, K. O. Safe mutations for deep and recurrent neural networks through output gradients. *arXiv preprint arXiv:1712.06563*, 2017b.
- Li, K. and Malik, J. Learning to optimize. *International Conference on Learning Representations*, 2017.
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. Random feedback weights support learning in deep neural networks. *arXiv preprint arXiv:1411.0247*, 2014.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Lv, K., Jiang, S., and Li, J. Learning gradient descent: Better generalization and longer horizons. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2247–2255. JMLR. org, 2017.
- Maclaurin, D., Duvenaud, D., and Adams, R. Gradient-based hyperparameter optimization through reversible learning. In *International Conference on Machine Learning*, pp. 2113–2122, 2015.
- Maddison, C. J., Mnih, A., and Teh, Y. W. The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*, 2016.
- Mania, H., Guy, A., and Recht, B. Simple random search provides a competitive approach to reinforcement learning. *arXiv preprint arXiv:1803.07055*, 2018.
- Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163*, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540): 529, 2015.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning. In *International Conference on Machine Learning*, pp. 1928–1937, 2016.
- Nesterov, Y. and Spokoiny, V. Random gradient-free minimization of convex functions. Technical report, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 2011.
- Nøkland, A. Direct feedback alignment provides learning in deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 1037–1045, 2016.
- Owen, A. B. Monte carlo theory, methods and examples (book draft), 2014.
- Pourchot, A. and Sigaud, O. CEM-RL: Combining evolutionary and gradient-based methods for policy search. *arXiv preprint arXiv:1810.01222*, 2018.
- Rastrigin, L. About convergence of random search method in extremal control of multi-parameter systems. *Avtomat. i Telemekh.*, 24(11):1467–1473, 1963.

- Rechenberg, I. *Evolutionsstrategie—optimierung technischer systeme nach prinzipien der biologischen evolution*. 1973.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science, 1985.
- Salimans, T., Ho, J., Chen, X., Sidor, S., and Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864*, 2017.
- Sra, S., Nowozin, S., and Wright, S. J. *Optimization for machine learning*. MIT Press, 2012.
- Staines, J. and Barber, D. Variational optimization. *arXiv preprint arXiv:1212.4507*, 2012.
- Tucker, G., Mnih, A., Maddison, C. J., Lawson, J., and Sohl-Dickstein, J. Rebar: Low-variance, unbiased gradient estimates for discrete latent variable models. In *Advances in Neural Information Processing Systems*, pp. 2624–2633, 2017.
- van den Oord, A., Vinyals, O., et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.
- Watkins, C. J. and Dayan, P. Q-learning. *Machine learning*, 8(3-4):279–292, 1992.
- Wichrowska, O., Maheswaranathan, N., Hoffman, M. W., Colmenarejo, S. G., Denil, M., de Freitas, N., and Sohl-Dickstein, J. Learned optimizers that scale and generalize. *International Conference on Machine Learning*, 2017.
- Wierstra, D., Schaul, T., Peters, J., and Schmidhuber, J. Natural evolution strategies. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, pp. 3381–3387. IEEE, 2008.
- Williams, R. J. and Peng, J. An efficient gradient-based algorithm for on-line training of recurrent network trajectories. *Neural computation*, 2(4):490–501, 1990.
- Wu, Y., Ren, M., Liao, R., and Grosse, R. Understanding short-horizon bias in stochastic meta-optimization. *arXiv preprint arXiv:1803.02021*, 2018.