Never look back - A modified EnKF method and its application to the training of neural networks without back propagation

Eldad Haber[*][†], Felix Lucka[‡] and Lars Ruthotto[§][†]

Abstract

In this work, we present a new derivative-free optimization method and investigate its use for training neural networks. Our method is motivated by the Ensemble Kalman Filter (EnKF), which has been used successfully for solving optimization problems that involve large-scale, highly nonlinear dynamical systems. A key benefit of the EnKF method is that it requires only the evaluation of the forward propagation but not its derivatives. Hence, in the context of neural networks it alleviates the need for back propagation and reduces the memory consumption dramatically. However, the method is not a pure "black-box" global optimization heuristic as it efficiently utilizes the structure of typical learning problems. Promising first results of the EnKF for training deep neural networks have been presented recently by Kovachki and Stuart. We propose an important modification of the EnKF that enables us to prove convergence of our method to the minimizer of a strongly convex function. Our method also bears similarity with implicit filtering and we demonstrate its potential for minimizing highly oscillatory functions using a simple example. Further, we provide numerical examples that demonstrate the potential of our method for training deep neural networks.

1 Introduction

Advances in numerical algorithms for training neural networks have been fueling the recent deep learning revolution. Virtually all training algorithms employ derivative-based optimization methods, and in particular, stochastic gradient descent (SGD) methods (see, e.g., [1, 2, 27, 26, 6, 13, 18, 11, 24] and reference within).

A key requirement of derivative-based optimization methods is an efficient way of computing the gradient of the objective function, which is used to update the parameters to be optimized. When training neural networks, this is commonly done using back propagation [22], i.e., propagating prediction errors from the output layer backwards through the network. Back propagation relies upon a means to store or quickly recompute hidden features and activations, which can be prohibitive when the neural network is deep or when the size of the data is large, e.g., when training a network to learn from videos or 3D datasets. While the storage problem is less pronounced when using stochastic gradient methods, the performance of these methods depends in a non-trivial way on parameters like batch size and learning rate [6], which often need to be tuned manually by solving the problem repeatedly. Also, SGD is very difficult to parallelize efficiently.

---

[*]Department of Earth and Ocean Science, The University of British Columbia, Vancouver, BC, Canada

[†]Xtract Inc., Vancouver, Canada

[‡]Computational Imaging, Centrum Wiskunde & Informatica (CWI), The Netherlands; Department of Computer Science, University College London, UK

[§]Department of Mathematics and Computer Science, Emory University, USA

Preprint. Work in progress.

To overcome some of the drawbacks associated with derivative-based learning algorithms, we propose a new derivative-free optimization method that can be seen as a slightly modified version of the Ensemble Kalman Filter (EnKF). Our method can either augment or replace the SGD for the training of neural networks. EnKF has a long track record in data assimilation and been applied to large-scale optimization problems, e.g., in weather and flow prediction [5, 19, 20, 8, 7, 4, 3]. In recent years, a version of the EnKF method that recovers an unknown parameter from noisy data has been applied to inverse problems [10, 9, 23]. Key advantages of the EnKF method for large-scale problems is that it can be easily parallelized and is derivative-free, i.e., it requires only evaluations of the forward operator. Another advantage of EnKF, that is particularly attractive for highly nonlinear forward problems is its potential to explore the objective function. However, EnKF is not a "black-box" metaheuristic for global optimization of arbitrary functions, like, e.g., genetic algorithms, which were also tested for deep learning recently [21]. It makes explicit and efficient use of the specific structure of learning problems.

In this paper, we continue the work of Kovachki and Stuart for training neural networks [15, 14], which is based on the EnKF presented in [23]. We modify this scheme and prove that our method converges to the global minimizer of a strongly convex function at a sublinear rate. This is an improvement over the method discussed in [23], which converges to the projection of the minimizer onto a low-dimensional subspace. Furthermore, our method can be parallelized easily, which can be attractive in a high-performance computing environment. Using an intuitive heuristic argument and a simple numerical experiment we also show how avoiding back propagation can be advantageous for problems that exhibit local high frequency oscillations. These properties motivate us to use our method for training deep neural networks even though we are not able to prove convergence to the global solution of such non-convex problems. Using derivative-free optimization methods for training deep neural networks has the advantage that no back propagation is needed and therefore, it is possible to work with arbitrary long networks without any memory limitations. We illustrate the potential of our method using a classical convolutional neural network for solving the MNIST problem of classifying hand-written digits [17] and show that our method converges faster than a SGD type method.

Our paper is organized as follows. In Sec. 2, we derive our method by modifying the version of the Ensemble Kalman Filter presented in [23]. In Sec. 3, we state and prove our convergence result for strongly convex functions and discuss the relation of our method to SGD. In Sec. 4, we present several practical improvements of our method. In Sec. 5, we present preliminary numerical results for strongly convex problems and the training of some common neural networks. In Sec. 6 we provide some discussion and concluding remarks.

2   The Ensemble Kalman Filter

Among the many variants of the Ensemble Kalman Filter (EnKF), we focus our attention on the recent version presented in [23]. This version simplifies much of the original EnKF discussion and allows for simple analysis and explanation. The method is an iterative method for the minimization of a function of the form

$$\phi(\boldsymbol{\theta}) = \mathcal{D}(F(\boldsymbol{\theta})), \qquad (2.1)$$

where $\boldsymbol{\theta} \in \mathbb{R}^n$ denotes the parameter to be recovered, $F : \mathbb{R}^n \to \mathbb{R}^m$ is the forward operator, and $\mathcal{D} : \mathbb{R}^m \to \mathbb{R}$ is a loss or misfit function. This formulation includes, e.g., linear least-squares problems that are obtained by choosing $\mathcal{D}(\mathbf{t}) = \frac{1}{2}\|\mathbf{t}\|^2$ and $F(\boldsymbol{\theta}) = \mathbf{A}\boldsymbol{\theta} - \mathbf{b}$ for fixed $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{b} \in \mathbb{R}^m$. For most classification problems, $F(\boldsymbol{\theta})$ is a neural network and $\mathcal{D}(\cdot)$ is the soft-max function. The key advantage of EnKF is that it does not rely on derivatives of F with respect to θ, which is usually computationally demanding, one only needs to evaluate the derivative of the loss function, which is computationally cheap. We will therefore call it "derivative-free" and use the short notation $\nabla\mathcal{D}(\theta)$ for $\nabla_t\mathcal{D}(t)$ evaluated at $t = F(\theta)$, but again stress that this is a rather different approach to other "derivative-free" optimization heuristics that try to minimize $\phi(\boldsymbol{\theta})$ without utilizing any of its structure, i.e., by evaluations of ϕ for given θ's only.

The first and simplest version of our algorithm consists of three steps that are listed in Algorithm 1. In step one of each iteration j we randomly choose k different particles $\boldsymbol{\theta}_i$, $i = 1, \ldots, k$ i.i.d. as $\boldsymbol{\theta} \sim \boldsymbol{\theta}_j + \boldsymbol{\omega}$. The distribution of $\boldsymbol{\omega}$ needs fulfill $\mathbb{E}\boldsymbol{\omega} = 0$, $\mathbb{C}\text{ov}(\boldsymbol{\theta}) = \sigma^2\mathbf{I}$ (where $\mathbf{I}$ denotes the identity matrix), and any higher moments up to forth order need to exist. Possible examples are given by

Algorithm 1 Derivative-free algorithm for solving $\min_{\boldsymbol{\theta}} \mathcal{D}(F(\boldsymbol{\theta}))$

Inputs: Starting guess $\bar{\boldsymbol{\theta}}_0 \in \mathbb{R}^n$, parameter $\sigma > 0$, number of particles k, covariance matrix $\boldsymbol{\Gamma}$, and maximum number of iterations N.

for $j = 0, 1, \ldots, N-1$ **do**
   [1] Draw k particles i.i.d. as $\boldsymbol{\theta}_i = \bar{\boldsymbol{\theta}}_j + \boldsymbol{\omega}_i$, $i = 1, \ldots, k$, with $\mathbb{E}(\boldsymbol{\omega}) = 0$ and $\mathrm{Cov}(\boldsymbol{\omega}) = \sigma^2 \mathbf{I}$. Set $\boldsymbol{\Omega}_j = [\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_k] \in \mathbb{R}^{n \times k}$.
   [2] Evaluate the forward operator F on each particle $\boldsymbol{\theta}_i$ and on $\boldsymbol{\theta}_j$ to construct

$$\mathbf{Q}_j = F(\bar{\boldsymbol{\theta}}_j + \boldsymbol{\Omega}_j) - F(\bar{\boldsymbol{\theta}}_j) := [F(\boldsymbol{\theta}_1) - F(\bar{\boldsymbol{\theta}}_j), \ldots, F(\boldsymbol{\theta}_k) - F(\bar{\boldsymbol{\theta}}_j)] \in \mathbb{R}^{m \times k} \qquad (2.2)$$

   [3] Set $\mathbf{H}_j = \mathbf{I}$ or $\mathbf{H}_j = \mathbf{Q}_j^\top \mathbf{Q}_j + \boldsymbol{\Gamma}$ and update

$$\bar{\boldsymbol{\theta}}_{j+1} = \bar{\boldsymbol{\theta}}_j - \mu_j \boldsymbol{\Omega}_j \mathbf{H}_j^{-1} \mathbf{Q}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j), \qquad (2.3)$$

   where μ_j is determined using a line search and $\nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j)$ denotes the gradient of the loss function, $\nabla_t \mathcal{D}(t)$ evaluated at $t = F(\bar{\boldsymbol{\theta}}_j)$.
end for
Output: Optimized parameters $\bar{\boldsymbol{\theta}}_N$.

$\boldsymbol{\omega} \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$, or $\boldsymbol{\omega} = \pm\sigma$ with equal probability (Rademacher distribution). The second step requires one to evaluate the forward operator (but not its derivatives) k times. In the third step of the algorithm, one needs to compute a step, by using a matrix, $\mathbf{H}_j$ and computing its inverse times a residual vector. The matrix can be chosen as an identity matrix or as $\mathbf{H}_j = \mathbf{Q}_j^\top \mathbf{Q}_j + \boldsymbol{\Gamma}$, where $\boldsymbol{\Gamma}$ is a data Covariance matrix. The latter leads to the usual Kalman Gain matrix [3, 10, 9, 23]. To update the solution, we choose the step size μ_j using a simple Armijo line search to ensure reduction of the objective function. We provide some justification for this in Sec. 3.2.

At this stage, we point to an important modification to the EnKF as presented in in [3, 10, 9, 23], which allows us to better analyze its convergence properties. First, we re-sample the particles at each iteration. This allows us to use a simple convergence proof of the method to the global minimum of a strongly convex function. We suspect that this requirement can be relaxed in future work. A second change from the original algorithm is at step 2, where the original EnKF algorithm defines the matrix $\mathbf{Q}_j$ as the difference between $F(\bar{\boldsymbol{\theta}}_j + \Omega_j)$ and its mean over all the particles. Clearly, both versions are the same for a linear forward problem. Finally, the last difference between our method to the original EnKF method is that the original EnKF algorithm propagates all the particles while we only update the mean of their distribution. It is also possible to propagate all the particles but we avoid this here for simplicity. As we see next, these modifications allow us to prove convergence of our algorithm and the insight we obtain this way allows us to extend the method and propose ways to further accelerate it.

3 Convergence of our algorithm and comparison to existing methods

We first prove that when applied to strongly convex problems our method converges to the global minimizer. In Sec. 3.2 we compare our method to stochastic gradient methods and other methods commonly used in derivative-free optimization.

3.1 Convergence analysis for strongly convex problems

To prove convergence of the method we need some assumptions that are typically used also for the proof of convergence for stochastic gradient descent methods. We prove convergence for the case $\mathbf{H}_j = \mathbf{I}$ and note that it can be extended for any symmetric positive definite $\mathbf{H}_j$.

   A1: The function $\phi(\boldsymbol{\theta}) = \mathcal{D}(F(\boldsymbol{\theta}))$ is smooth and strongly convex, i.e., there exists a $L > 0$ such that for all $\boldsymbol{\theta}_1, \boldsymbol{\theta}_2 \in \mathbb{R}^n$

$$\mathcal{D}(F(\boldsymbol{\theta}_2)) - \mathcal{D}(F(\boldsymbol{\theta}_1)) \geq (\boldsymbol{\theta}_2 - \boldsymbol{\theta}_1)^\top \mathbf{J}(\boldsymbol{\theta}_1)^\top \nabla \mathcal{D}(\boldsymbol{\theta}_1) + \frac{L}{2} \|\boldsymbol{\theta}_2 - \boldsymbol{\theta}_1\|^2.$$

Here, $\mathbf{J}(\boldsymbol{\theta}) = \nabla_{\boldsymbol{\theta}} F(\boldsymbol{\theta})^\top \in \mathbb{R}^{m \times n}$ is the Jacobian of the mapping of F.

A2: The Taylor expansion error of F can be made sufficiently small, i.e., assuming that we sampled an $\boldsymbol{\omega}$ with $\|\boldsymbol{\omega}\| = \mathcal{O}(\sigma)$, we have that

$$F(\bar{\boldsymbol{\theta}} + \boldsymbol{\omega}) - F(\bar{\boldsymbol{\theta}}) = \mathbf{J}(\bar{\boldsymbol{\theta}})\boldsymbol{\omega} + \mathcal{O}(\sigma^2).$$

Using A2 we can write the matrix $\mathbf{Q}_j$ in (2.2) as

$$\mathbf{Q}_j = \mathbf{J}_j \boldsymbol{\Omega}_j + \mathcal{O}(\sigma^2).$$

Plugging this into (2.3) we obtain

$$\bar{\boldsymbol{\theta}}_{j+1} = \bar{\boldsymbol{\theta}}_j - \mu_j \left(\boldsymbol{\Omega}_j \boldsymbol{\Omega}_j^\top \mathbf{J}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) + \mathcal{O}(\sigma^2) \right). \tag{3.4}$$

From our assumptions on the distribution of the particles, it easily follows that $\mathbb{E}\, \boldsymbol{\Omega}_j \boldsymbol{\Omega}_j^\top = k\sigma^2 \mathbf{I}$. If we define $\mathbf{g}_j := \boldsymbol{\Omega}_j \boldsymbol{\Omega}_j^\top \mathbf{J}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j)$, we can compute that

$$\mathbb{E}\, \mathbf{g}_j = k\sigma^2 \mathbf{J}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) = k\sigma^2 \nabla \phi(\bar{\boldsymbol{\theta}}_j), \qquad \mathbb{E}\, \|\mathbf{g}_j\|^2 = \|\nabla \phi(\bar{\boldsymbol{\theta}}_j)\|^2 + \mathcal{O}(k^2 \sigma^4). \tag{3.5}$$

Now, let $\boldsymbol{\theta}^*$ be the minimizer of (2.1). Using A1 we have that

$$\phi(\boldsymbol{\theta}^*) - \phi(\bar{\boldsymbol{\theta}}_j) \geq (\boldsymbol{\theta}^* - \bar{\boldsymbol{\theta}}_j)^\top \mathbf{J}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) + \frac{L}{2}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2$$

$$\phi(\bar{\boldsymbol{\theta}}_j) - \phi(\boldsymbol{\theta}^*) \geq (\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*)^\top \mathbf{J}_*^\top \nabla \mathcal{D}(\boldsymbol{\theta}_*) + \frac{L}{2}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 = \frac{L}{2}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2,$$

where $\mathbf{J}_* := \mathbf{J}(\boldsymbol{\theta}^*)$, and we used that the gradient at $\boldsymbol{\theta}^*$ vanishes. Summing both inequalities we obtain

$$(\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*)^\top \mathbf{J}_j^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) \geq L\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2. \tag{3.6}$$

Let us assume that within Algorithm 1, we arrived at a concrete $\bar{\boldsymbol{\theta}}_j$. We now derive an inequality that describes how the distance $\|\bar{\boldsymbol{\theta}}_{j+1} - \boldsymbol{\theta}^*\|^2$ compares to $\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2$ *in expectation*. Technically, this means that all expected values here are conditioned on $\bar{\boldsymbol{\theta}}_j$ and only account for the randomness induced by drawing a new set of particles, which is manifest in $\boldsymbol{\Omega}_j$. We first use (3.4) and (3.5) to estimate:

$$\mathbb{E}\|\bar{\boldsymbol{\theta}}_{j+1} - \boldsymbol{\theta}^*\|^2 = \mathbb{E}\|\bar{\boldsymbol{\theta}}_j - \mu_j(\mathbf{g}_j + \mathcal{O}(\sigma^2)) - \boldsymbol{\theta}^*\|^2$$
$$= \mathbb{E}\, \|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 - 2\mu_j \mathbb{E}\, \mathbf{g}_j^\top (\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*) + \mathcal{O}(\mu_j \sigma^2) + \mathcal{O}(\mu_j^2 k^2 \sigma^4)$$
$$\leq \mathbb{E}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 - 2\mu_j k\sigma^2 \nabla \phi(\bar{\boldsymbol{\theta}}_j)^\top (\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*) + \mathcal{O}(\mu_j \sigma^2) + \mathcal{O}(\mu_j^2 k^2 \sigma^4)$$

Using (3.6) we obtain

$$\mathbb{E}\|\bar{\boldsymbol{\theta}}_{j+1} - \boldsymbol{\theta}^*\|^2 \leq = (1 - 2\mu_j L k\sigma^2)\mathbb{E}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 + \mathcal{O}(\mu_j \sigma^2) + \mathcal{O}(\mu_j^2 k^2 \sigma^4) \tag{3.7}$$

We can now prove convergence by induction when we choose the learning rate $\mu_j = \frac{1}{jLk\sigma^2}$.

Theorem 1 *At the j-th iteration, it holds that*

$$\mathbb{E}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 \leq \frac{C}{j}, \tag{3.8}$$

where the constant C only depends on the initial distance, $\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|$, the smoothness of ϕ, L, and k.

Proof 1 *For $j = 0$, it is clear that as we pick $\boldsymbol{\theta}_0$ manually, $\mathbb{E}\|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\| = \|\boldsymbol{\theta}_0 - \boldsymbol{\theta}^*\|$ can be bounded in the above way. The terms $\mathcal{O}(\mu_j \sigma^2) + \mathcal{O}(\mu_j^2 k^2 \sigma^4)$ in (3.7) all come from smoothness assumptions on ϕ and can thus be bounded accordingly. If we choose the learning rate $\mu_j = \frac{1}{jLk\sigma^2}$ and absorb remaining dependencies on k and L in C as well, we obtain*

$$\mathbb{E}\|\bar{\boldsymbol{\theta}}_{j+1} - \boldsymbol{\theta}^*\|^2 \leq = \left(1 - \frac{2}{j}\right) \mathbb{E}\|\bar{\boldsymbol{\theta}}_j - \boldsymbol{\theta}^*\|^2 + \frac{C}{j} + \frac{C}{j^2}. \tag{3.9}$$

If we then assume that the convergence holds at iteration j, we prove that it holds for $j + 1$:

$$\mathbb{E}\|\boldsymbol{\theta}_{j+1} - \boldsymbol{\theta}^*\|^2 \leq \left(1 - \frac{2}{j}\right)\frac{C}{j} + \frac{C}{j} + \frac{C}{j^2} = C\left(\frac{1}{j} - \frac{2}{j^2} + \frac{1}{j} + \frac{1}{j^2}\right) \leq \frac{C}{j+1}.$$

3.2 Comparison to other methods

It is interesting to note the difference between our derivative-free optimization algorithm to stochastic gradient descent (SGD) and its variants. To this end, note that the SGD iteration can be written as

$$\boldsymbol{\theta}_{j+1} = \boldsymbol{\theta}_j - \eta_j \mathbf{J}_j^\top \mathbf{X}_j \mathbf{X}_j^\top \nabla \mathcal{D}(\boldsymbol{\theta}_j),$$

where the rows of $\mathbf{X}_j \in \mathbb{R}^{s \times m}$ are randomly chosen rows of the $m \times m$ identity matrix and η_j is the learning rate. Although the choice of the learning rate in Theorem 1 is similar to the one commonly used to prove convergence of SGD, it is important to note one major difference between our update rule and the SGD step. In contrast to SGD step, the step size in our method can always be chosen so that the objective function is non-increasing. This is because the computed direction is a descent direction over the subspace spanned by $\boldsymbol{\Omega}_j$, since

$$-\nabla \phi^\top \boldsymbol{\Omega}_j \boldsymbol{\Omega}_j^\top \mathbf{J}_j^\top \nabla \mathcal{D} = -(\nabla \mathcal{D})^\top \mathbf{J}_j^\top \boldsymbol{\Omega}_j \boldsymbol{\Omega}_j^\top \mathbf{J}_j^\top \nabla \mathcal{D} \leq 0.$$

Hence, common line search strategies can be used to ensure that the objective is monotonically non-increasing. As we see in the numerical examples this makes the selection of hyper parameters easier than in the case of SGD.

Of course this property does not come without a price. This version of the method requires the propagation of all the data making it expensive for problems where the number of data is very large. Nonetheless, it is possible to combine the method with a stochastic selection of the data making it competitive with standard SGD iteration. In this case, the iteration has the form

$$\bar{\boldsymbol{\theta}}_{j+1} = \bar{\boldsymbol{\theta}}_j - \mu_j \boldsymbol{\Omega}_j \mathbf{H}_j^{-1} \mathbf{Q}_j^\top \mathbf{X}_j \mathbf{X}_j^\top \nabla \mathcal{D}(\boldsymbol{\theta}_j). \tag{3.10}$$

It is straight forward to extend the proof of convergence for this case when $\mathbf{X}_j$ and $\boldsymbol{\Omega}_j$ are uncorrelated.

Although the convergence proof covers only strongly convex problems, one motivation to use our derivative-free method for highly nonlinear problems can be made by looking at the behavior of the method for problems with small, yet, high-frequency oscillations.

Example 1 *Let $l \in \mathbb{R}$, $\epsilon > 0$, $\mathbf{A} \in \mathbb{R}^{m \times n}$, and $\mathbf{B} \in \mathbb{R}^{m \times n}$ be given and*

$$F(\boldsymbol{\theta}) = S(\boldsymbol{\theta}) + \epsilon N(\boldsymbol{\theta}) = \mathbf{A}\boldsymbol{\theta} + \epsilon \sin(l\mathbf{B}\boldsymbol{\theta})$$

where S is a smooth function and N is an oscillatory one, and assume the least squares misfit function, $\mathcal{D}(\mathbf{t}) = \frac{1}{2}\|\mathbf{t}\|^2$. In this case the necessary conditions for a minimum are

$$\mathbf{J}^\top \nabla \mathcal{D}(\boldsymbol{\theta}) = (\mathbf{A}^\top + \epsilon l \mathbf{B}^\top \mathrm{diag}(\cos(l\mathbf{B}\boldsymbol{\theta})))(\mathbf{A}\boldsymbol{\theta} + \epsilon \sin(l\mathbf{B}\boldsymbol{\theta}) - \mathbf{d}^{\mathrm{obs}}) = 0$$

where $\mathbf{J} = \mathbf{A} + \epsilon l \mathrm{diag}(\cos(l\mathbf{B}\boldsymbol{\theta}))\mathbf{B}$. Clearly, the misfit function can have many local minima. For large values of l, the function presents some local highly oscillatory behavior as well as some "slow" low-frequency modes. As a result, if we are to compute the solution using a derivative-based technique, we would probably fail.

It is important to understand the cause of the oscillatory gradients. Note that the gradient of our function is simply

$$\nabla_{\boldsymbol{\theta}} \mathcal{D} = \mathbf{J}^\top \nabla \mathcal{D} = (\nabla S)^\top \mathcal{D} + \epsilon (\nabla N)^\top \nabla \mathcal{D}.$$

While $\nabla \mathcal{D}$ is smooth or oscillates mildly, the second part, $\mathbf{J}$ contains much higher oscillations. For the problem at hand, ∇S is linear, however, the oscillations in ∇N are amplified by a factor of l. Indeed, even if the function has small oscillations, its derivative has much larger oscillations and therefore, the solution can be difficult to obtain. The source of the oscillations can be tracked to the gradient of the oscillatory part N.

It is important to note that the approximate Jacobian $\mathbf{Q}_j = F(\bar{\boldsymbol{\theta}}_j + \boldsymbol{\Omega}_j) - F(\bar{\boldsymbol{\theta}})$, computed in our method does not suffer from this problem as long as we choose the elements of $\boldsymbol{\Omega}$ sufficiently large. This is because we only evaluate the forward operator and determining the step is similar to numerical differentiation. Thus, our method may present better properties for problems with highly oscillatory local behavior compared with derivative-based (i.e. back propagation) descent methods.

The above discussion shows that there is a strong link between the EnKF method and implicit filtering discussed in [12]. In fact, the EnKF can be seen as a stochastic version of implicit filtering, that is known to work well for noisy functions. Finally, the method can be seen as a particular implementation of a stochastic coordinate descent, where the coordinates are picked by the matrix $\boldsymbol{\Omega}_j$ and the gradient is evaluated numerically.

Algorithm 2 Derivative-free algorithm for solving $\min_{\boldsymbol{\theta}} \mathcal{D}(F(\boldsymbol{\theta}))$ with memory

---

Inputs: Starting guess $\bar{\boldsymbol{\theta}}_0 \in \mathbb{R}^n$, parameter $\sigma > 0$, number of particles k, covariance matrix $\boldsymbol{\Gamma}$, and maximum number of iterations N. Allocate $\boldsymbol{\Omega} = []$ and $\mathbf{Q} = []$

for $j = 0, 1, \ldots, N-1$ **do**

[1] Draw k particles i.i.d. as $\boldsymbol{\theta}_i = \bar{\boldsymbol{\theta}}_j + \boldsymbol{\omega}_i$, $i = 1, \ldots, k$, with $\mathbb{E}(\boldsymbol{\omega}) = 0$ and $\mathrm{Cov}(\boldsymbol{\omega}) = \sigma^2 \mathbf{I}$. Set $\boldsymbol{\Omega}_j = [\boldsymbol{\omega}_1, \ldots, \boldsymbol{\omega}_k] \in \mathbb{R}^{n \times k}$.

[2] Evaluate the forward operator F on each particle $\boldsymbol{\theta}_i$ and on $\boldsymbol{\theta}_j$ to construct

$$\mathbf{Q}_j = F(\bar{\boldsymbol{\theta}}_j + \boldsymbol{\Omega}_j) - F(\bar{\boldsymbol{\theta}}_j) := [F(\boldsymbol{\theta}_1) - F(\bar{\boldsymbol{\theta}}_j), \ldots, F(\boldsymbol{\theta}_k) - F(\bar{\boldsymbol{\theta}}_j)] \in \mathbb{R}^{m \times k} \quad (4.11)$$

[3] Update matrices $\boldsymbol{\Omega} \leftarrow [\boldsymbol{\Omega}, \boldsymbol{\Omega}_j]$ and $\mathbf{Q} \leftarrow [\mathbf{Q}, \mathbf{Q}_j]$ by adding columns.

[4] Set $\mathbf{H}_j = \mathbf{I}$ or $\mathbf{H}_j = \mathbf{Q}^\top \mathbf{Q} + \boldsymbol{\Gamma}$ and update

$$\bar{\boldsymbol{\theta}}_{j+1} = \bar{\boldsymbol{\theta}}_j - \mu_j \boldsymbol{\Omega} \mathbf{H}_j^{-1} \mathbf{Q}^\top \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) \quad (4.12)$$

where μ_j is determined using a line search and $\nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j)$ denotes the gradient of the loss function, $\nabla_t \mathcal{D}(t)$ evaluated at $t = F(\bar{\boldsymbol{\theta}}_j)$.

Output: Optimized parameters $\bar{\boldsymbol{\theta}}_N$.

end for

---

4   Practical improvements and extensions

While Algorithm 1 can be used directly for the solution of the training problem, some simple modifications can boost its performance considerably.

4.1   Reducing the number of particles per iteration

The computational cost at each iteration scales linearly with the number of particles used and therefore, it is desirable to reduce the number of particles used at each iteration. To this end, we propose to re-use the matrices computed from previous particles at earlier iterations at the computation. The modified algorithm is described in Algorithm 2. This version of the algorithm is similar to the variance reduction SGD method proposed in [16]. In practice, one allocates memory for, say ℓ vectors for $\mathbf{Q}$ and $\boldsymbol{\Omega}$ and saves only the last ℓ vectors. It is important to see that this version of the method requires much less forward calculations at each step. We will show how this version of the code can be advantageous compared to the "vanilla" version presented in Algorithm 1.

4.2   A Gauss-Newton like update

The update of $\bar{\boldsymbol{\theta}}$ proposed in (2.3) or (4.12) reduces to a steepest descent algorithm when we use a full basis to sample the problem. However, the original EnKF algorithm uses a Kalman gain matrix to update the solution. It is well known that Kalman filtering and the Kalman gain can be interpreted as a Gauss-Newton iteration [25]. We can thus use the matrix $\mathbf{Q}$ to assess an approximation of the Hessian. Noting that

$$\mathbb{E}\, \mathbf{Q}\mathbf{Q}^\top \approx \mathbb{E}\, \mathbf{J}\boldsymbol{\Omega}\boldsymbol{\Omega}^\top \mathbf{J}^\top = \sigma^2 k \mathbf{J}\mathbf{J}^\top$$

, one can use the update

$$\bar{\boldsymbol{\theta}}_{j+1} = \bar{\boldsymbol{\theta}}_j - \frac{\mu_j}{\sigma^2 k} \boldsymbol{\Omega} \mathbf{Q}^\top (\mathbf{Q}\mathbf{Q}^\top + \boldsymbol{\Gamma})^{-1} \nabla \mathcal{D}(\bar{\boldsymbol{\theta}}_j) \quad (4.13)$$

where $\boldsymbol{\Gamma}$ is a data covariance matrix. This update is similar to the common EnKF presented in the original work [3]. It is important to realize that the system (4.13) needs not be formed and it is possible to use k steps of the conjugate gradient method, preconditioned by $\boldsymbol{\Gamma}$ to solve the system exactly.
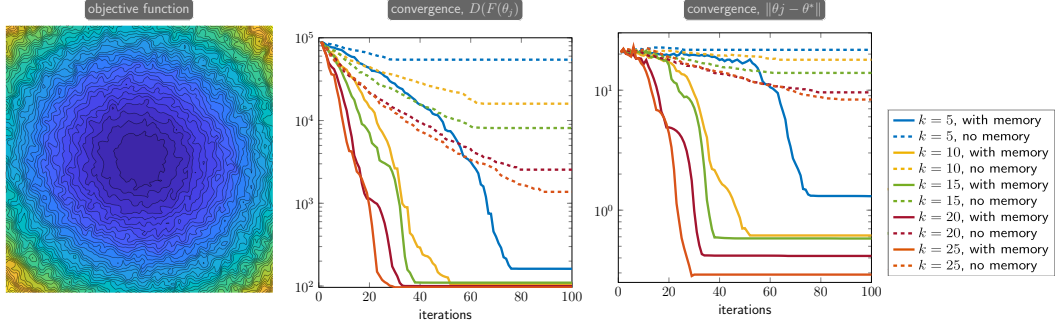
Figure 1: Convergence of our methods for a regression problem with highly oscillatory forward operator as described in Example 1. Left: Plot of the objective function for $m = 300, n = 200, l = 20$, and $\epsilon = 1$ along two randomly chosen orthogonal directions. Center: Reduction of the objective function value for different values of k (indicated by colors) and different algorithms (dashed lines indicate Algorithm 1 and solid line represent Algorithm 2). Right indicates the distance to the true solution. Note the improvements achieved by increasing k and storing previous forward evaluations.

5 Numerical examples

In this section, we present results from two numerical experiments. First, we demonstrate our algorithm's potential to converge on non-smooth objective functions. Second, we compare our method to a SGD scheme for the MNIST problem of classifying hand-written digits.

5.1 Application to nonlinear regression

We apply our method to regression problem described in Example 1, which involves the highly oscillatory forward problem

$$F(\boldsymbol{\theta}) = \mathbf{A}\boldsymbol{\theta} + \epsilon \sin(l\mathbf{B}\boldsymbol{\theta}). \tag{5.14}$$

Here, we choose $\mathbf{A}$ and $\mathbf{B}$ as random matrices of size 300×200, we set l to 20, and use $\epsilon = 1$. Figure 1 visualizes the resulting objective function and the results obtained using both versions of our method as listed in Algorithm 1 and Algorithm 2 and varying the number of particles $k \in \{5, 10, 15, 20, 25\}$. As can be seen from the plot of the objective function, this problem is challenging for gradient-based methods. Similar to implicit filtering methods [12], our method is better-equipped to deal with this non-smoothness because of the spatial sampling when approximating the gradient. We note that the convergence accelerates when more particles are used and the benefit of storing intermediate evaluations of the forward model is obvious.

5.2 Image classification using Convolutional Neural Networks

We use our method to train the weights of a simple neural network to classify the hand-written digits in the MNIST data set [17]. The data set consists of 60,000 gray scale images of resolution 28×28 that is divided into 50,000 training and 10,000 test images associated with their labels.

Here, we denote by $F(\theta, \mathbf{x})$ the forward propagation of an image $\mathbf{x}$ through the CNN consisting of two hidden layers. The first one uses a convolution with 5×5 stencils and 32 output channels and a ReLU activation function. This layer is followed by an average pooling operator and a second convolutional layer with 5×5 stencils, 64 output channels, and also ReLu activation. The last part of the forward model is a second average pooling operator, reducing the image size to 7×7. The network has $52,000$ weights.

As common in classification using deep neural networks, a fully-connected layer followed by a softmax hypothesis function, H, is applied to the output of the neural network and the result is compared to the given label using a cross entropy loss function. In the notation of our method, we interpret these parts as being associated with the misfit function $\mathcal{D}$ but note that this function depends on the weights of the fully connected layer, denoted by $\mathbf{W}$. Overall, denoting the image-label pairs
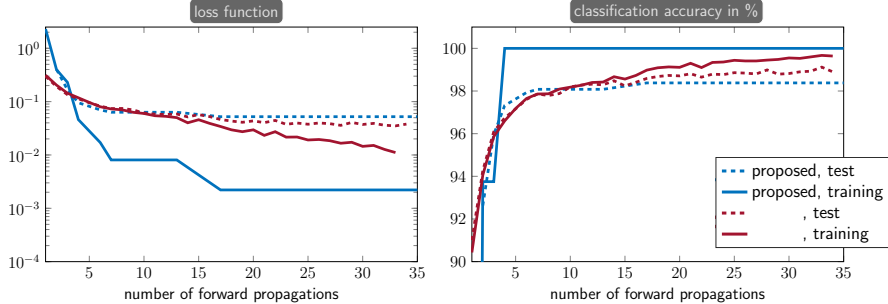
7

Figure 2: Convergence of our derivative-free method with variable projection (blue) and the stochastic gradient method ADAM (red) for training a convolutional neural network with 2 layers to classify the MNIST dataset. Left: Convergence of the loss function on the training (solid line) and test data (dashed line). Right: classification accuracy. The x-axis is scaled according to the number of forward propagations. It can be seen that ADAM generalizes slightly better while our proposed method converges considerably faster.

by $\{(\mathbf{x}_i, \mathbf{c}_i)\}_{i=1}^s$ we phrase the learning problem as an optimization problem

$$\min_{\mathbf{W}, \boldsymbol{\theta}} \frac{1}{s} \sum_{i=1}^s D(H(\mathbf{W}, F(\boldsymbol{\theta}, \mathbf{x}_i)) + R(\mathbf{W}, \boldsymbol{\theta}), \tag{5.15}$$

where R is a regularizer. Here, for simplicity we use a weight decay regularizer for $\mathbf{W}$ with a weight of 100 and no regularization on $\boldsymbol{\theta}$.

While the objective function in 5.15 is non-convex in the weights $\boldsymbol{\theta}$, it is convex with respect to $\mathbf{W}$ for the softmax cross entropy loss function and regularizer at hand. Hence, for a given choice of $\boldsymbol{\theta}$ the optimal $\mathbf{W}^*(\boldsymbol{\theta})$ can be computed efficiently using, e.g., Newton's method, which is feasible in our case since the number of elements in $\mathbf{W}$ is only $31,370$. We exploit this structure in our method in the following way. In each iteration, we evaluate the forward operator F for all the examples in the training data set. Then, we solve the classification problem using 10 iterations of an inexact Newton method using up to 20 iterations of a Conjugate Gradient (CG) scheme to determine the direction. Keeping $\mathbf{W}^*(\theta)$ we approximate the Jacobian of the neural network by using Algorithm 1 with one mini batch consisting of 16 examples and four particles. Line-search and updating of the current value of $\boldsymbol{\theta}$ is done as described above.

Figure 2 shows the convergence of our method in comparison with the stochastic gradient method ADAM with a mini batch size of 16 and a learning rate $\eta_j = 10^{-3}/\sqrt{j}$. The number of epochs is chosen so that an equal number of forward propagations is performed. Note that the x-axis scales with the number of forward propagations and we do not account for back-propagation in ADAM and the classification solves in our method. ADAM generalizes slightly better in this case, achieving a test accuracy of (99.12% vs. 98.38%) but our method converges considerably faster both in terms of forward propagations and in terms of runtime (238.5 vs. 1,751 secs on a TITAN X with CUDNN 7.1 in MATLAB2017b).

6 Conclusions

In this paper, we present a new derivative-free method for minimizing functions that are a concatenation of a misfit or loss function and a forward operator. Our method is derivative-free in the sense that the Jacobian of the forward operator is estimated using its evaluations at randomly chosen points around the current iterate while the derivative of the misfit is computed exactly. Our method can be seen as a modified version of the EnKF with improved convergence properties for strongly convex problems. Borrowing ideas from variance reduction methods we present an improved version of our method that stores intermediate evaluations of the forward operator.

Our numerical experiments show two benefits of our method. First, we use a simple nonlinear least-squares problem with an oscillatory forward operator to demonstrate the smoothing property

of our method. In fact our method can be seen as a randomized version of the implicit filtering [12] applied to a randomly chosen low-dimensional subspace. Second, motivated by the works [15, 14], we show that our method achieves comparable classification results to the SGS variant ADAM for the MNIST data set. In particular, our method parallelizes better and achieves a lower loss value within only a few forward propagations. Notable only four particles are used to optimize around 50,000 weights which motives the use of our method also for more challenging learning tasks.

7   Acknowledgements

References

[1] D. P. Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. *Math Programming B*, 129:123–165, 2011.

[2] L Bottou. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade*. Springer, Berlin, Heidelberg, 2012.

[3] G. Evensen and P. J. van Leeuwen. An ensemble Kalman smoother for nonlinear dynamics. *Mon. Wea. Rev.*, 128:1852–1867, 2000.

[4] Geir Evensen. The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dynam.*, 53, 2003.

[5] A. Fahimuddin. *4D Seismic History Matching Using the Ensemble Kalman Filter (EnKF): Possibilities and Challenges*. Ph.d. dissertation, The University of Bergen, 2010. Norway.

[6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.

[7] P. L. Houtekamer and H. L. Mitchell. A sequential ensemble Kalman filter for atmospheric data assimilation. *Mon. Wea. Rev.*, 129:123–137, January 2001.

[8] Brian R. Hunt, Eric J. Kostelich, and Istvan Szunyogh. Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D*, 230:112–126, 2007.

[9] Marco A. Iglesias. Iterative regularization for ensemble data assimilation in reservoir models. *Computational Geosciences*, 19(1):177–212, Feb 2015.

[10] Marco A Iglesias, Kody J H Law, and Andrew M Stuart. Ensemble kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.

[11] A. Juditsky, G. Lan, A. Nemirovski, and A. Shapiro. Stochastic approximation approach to stochastic programming. *SIAM J. Sci. Optim.*, 19:1574–1609, 2009.

[12] C.T. Kelley. *Implicit Filtering*. SIAM, Philadelphia, 2011.

[13] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On Large-Batch Training for Deep Learning: Generalization Gap and Sharp Minima. *arXiv.org*, 2016.

[14] Nikola Kovachki and Andrew M Stuart. Derivative-free ensemble methods for machine learning tasks (slides for presentation at CM+X Workshop available at http://cmx.caltech.edu/ipml-program.html), February 2018.

[15] Nikola Kovachki and Andrew M Stuart. Ensemble Kalman inversion: a derivative-free technique for machine learning tasks. *in preparation*, 2018.

[16] N. Le Roux, M. Schmidt, and F. Bach. A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets. *ArXiv e-prints*, February 2012.

[17] Y LeCun, B E Boser, and J S Denker. Handwritten digit recognition with a back-propagation network. *Advances in Neural Networks*, 1990.

[18] J.J. Linderoth, A. Shapiro, and S. Wright. The empirical behavior of sampling methods for stochastic programming. *Annals of Operations Research*, 142:215–241, 2006.

[19] V. Nenna, A. Pidlisecky, and R. Knight. Application Of An Extended Kalman Filter Approach To Inversion Of Time-lapse Electrical Resistivity Imaging Data For Monitoring Infiltration. *Water Resources Research*, 78:120–132, 2011.

[20] Edward Ott, Brian R Hunt, Istvan Szunyogh, Aleksey V Zimin, Eric J Kostelich, Matteo Corazza, Eugenia Kalnay, DJ Patil, and James A Yorke. A local ensemble kalman filter for atmospheric data assimilation. *Tellus A*, 56(5):415–428, 2004.

[21] F. Petroski Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune. Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning. *ArXiv e-prints*, December 2017.

[22] D.E. Rumelhart, Geoffrey Hinton, and J. Williams, R. Learning representations by back-propagating errors. *Nature*, 323(6088):533–538, 1986.

[23] C. Schillings and A. M. Stuart. Analysis of the ensemble kalman filter for inverse problems. *SIAM J. Numerical Analysis*, 39:1264–1290, 2017.

[24] A. Shapiro, D. Dentcheva, and D. Ruszczynski. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia, 2009.

[25] Micheal Verhaegen and Paul Van Dooren. Numerical aspects of different kalman filter implementations. *IEEE Transactions on Automatic Control*, 31(10):907–917, 1986.

[26] J. Weston, F. Ratle, H. Mobahi, and R. Collobert. Deep learning via semi-supervised embedding. *Neural Networks: Tricks of the Trade Lecture Notes in Computer Science Volume 7700*, pages 639–655, 2012.

[27] Chiyuan Zhang, Qianli Liao, Alexander Rakhlin, Karthik Sridharan, Brando Miranda, Noah Golowich, and Tomaso Poggio. Theory of deep learning iii: Generalization properties of sgd. Technical report, Center for Brains, Minds and Machines (CBMM), 2017.