
Fast large-scale optimization by unifying stochastic gradient and quasi-Newton methods

Jascha Sohl-Dickstein
Ben Poole
Surya Ganguli

JASCHA@{STANFORD.EDU,KHANACADEMY.ORG}
POOLE@CS.STANFORD.EDU
SGANGULI@STANFORD.EDU

Abstract

We present an algorithm for minimizing a sum of functions that combines the computational efficiency of stochastic gradient descent (SGD) with the second order curvature information leveraged by quasi-Newton methods. We unify these disparate approaches by maintaining an independent Hessian approximation for each contributing function in the sum. We maintain computational tractability and limit memory requirements even for high dimensional optimization problems by storing and manipulating these quadratic approximations in a shared, time evolving, low dimensional subspace. This algorithm contrasts with earlier stochastic second order techniques that treat the Hessian of each contributing function as a noisy approximation to the full Hessian, rather than as a target for direct estimation. Each update step requires only a single contributing function or minibatch evaluation (as in SGD), and each step is scaled using an approximate inverse Hessian and little to no adjustment of hyperparameters is required (as is typical for quasi-Newton methods). We experimentally demonstrate improved convergence on seven diverse optimization problems. The algorithm is released as open source Python and MATLAB packages.

1. Introduction

A common problem in optimization is to find a vector $\mathbf{x}^* \in \mathcal{R}^M$ which minimizes a function $F(\mathbf{x})$, where $F(\mathbf{x})$ is a sum of N computationally cheaper differentiable sub-

functions $f_i(\mathbf{x})$,

$$F(\mathbf{x}) = \sum_{i=1}^N f_i(\mathbf{x}), \quad (1)$$

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} F(\mathbf{x}). \quad (2)$$

Many optimization tasks fit this form (Boyd & Vandenberghe, 2004), including training of autoencoders, support vector machines, and logistic regression algorithms, as well as parameter estimation in probabilistic models. In these cases each subfunction corresponds to evaluating the objective on a separate data minibatch, thus the number of subfunctions N would be the datasize D divided by the minibatch size S . This scenario is commonly referred to in statistics as M-estimation (Huber, 1981).

There are two general approaches to efficiently optimizing a function of this form. The first is to use a quasi-Newton method (Dennis Jr & Moré, 1977), of which BFGS (Broyden, 1970; Fletcher, 1970; Goldfarb, 1970; Shanno, 1970) or LBFGS (Liu & Nocedal, 1989) are the most common choices. Quasi-Newton methods use the history of gradient evaluations to build up an approximation to the inverse Hessian of the objective function $F(\mathbf{x})$. By making descent steps which are scaled by the approximate inverse Hessian, and which are therefore longer in directions of shallow curvature and shorter in directions of steep curvature, quasi-Newton methods can be orders of magnitude faster than steepest descent. Additionally, quasi-Newton techniques typically require adjusting few or no hyperparameters, because they use the measured curvature of the objective function to set step lengths and directions. However, direct application of quasi-Newton methods requires calculating the gradient of the full objective function $F(\mathbf{x})$ at every proposed parameter setting \mathbf{x} , which can be very computationally expensive.

The second approach is to use a variant of Stochastic Gradient Descent (SGD) (Robbins & Monro, 1951; Bottou, 1991). In SGD, only one subfunction's gradient is evaluated per update step, and a small step is taken in the negative gradient direction. More recent descent techniques

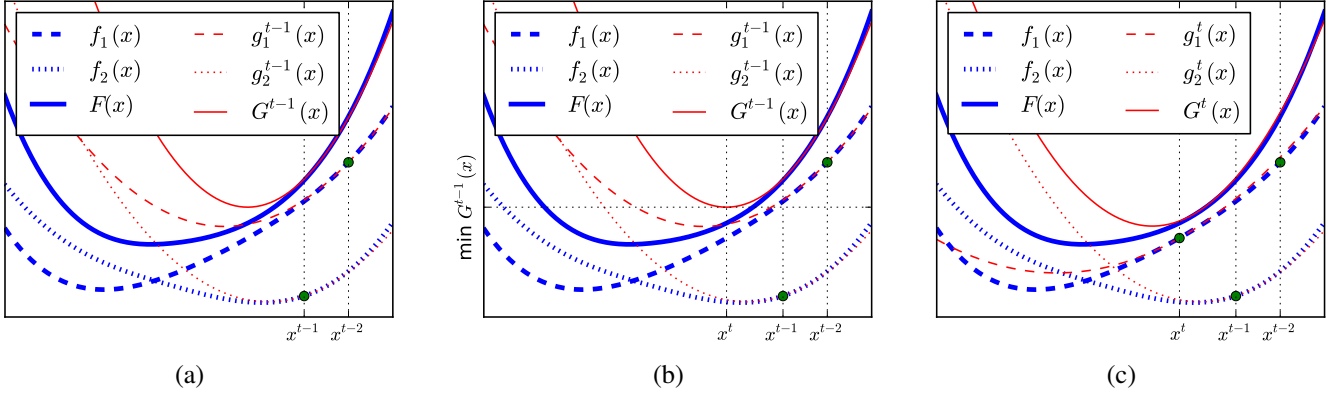


Figure 1. A cartoon illustrating the proposed optimization technique. (a) The objective function $F(x)$ (solid blue line) consists of a sum of two subfunctions (dashed blue lines), $F(x) = f_1(x) + f_2(x)$. At learning step $t-1$, $f_1(x)$ and $f_2(x)$ are approximated by quadratic functions $g_1^{t-1}(x)$ and $g_2^{t-1}(x)$ (red dashed lines). The sum of the approximating functions $G^{t-1}(x)$ (solid red line) approximates the full objective $F(x)$. The green dots indicate the parameter values at which each subfunction has been evaluated (b) The next parameter setting x^t is chosen by minimizing the approximating function $G^{t-1}(x)$ from the prior update step. See Equation 4. (c) After each parameter update, the quadratic approximation for one of the subfunctions is updated using a second order expansion around the new parameter vector x^t . See Equation 6. The constant and first order term in the expansion are evaluated exactly, and the second order term is estimated by performing BFGS on the subfunction’s history. In this case the approximating subfunction $g_1^t(x)$ is updated (long-dashed red line). This update is also reflected by a change in the full approximating function $G^t(x)$ (solid red line). Optimization proceeds by repeating these two illustrated update steps. In order to remain tractable in memory and computational overhead, optimization is performed in an adaptive low dimensional subspace determined by the history of gradients and positions.

like IAG (Blatt et al., 2007), SAG (Roux et al., 2012), and MISO (Mairal, 2013; 2014) instead take update steps in the average gradient direction. For each update step, they evaluate the gradient of one subfunction, and update the average gradient using its new value. (Bach & Moulines, 2013) averages the iterates rather than the gradients. If the subfunctions are similar, then SGD can also be orders of magnitude faster than steepest descent on the full batch. However, because a different subfunction is evaluated for each update step, the gradients for each update step cannot be combined in a straightforward way to estimate the inverse Hessian of the full objective function. Additionally, efficient optimization with SGD typically involves tuning a number of hyperparameters, which can be a painstaking and frustrating process. (Le et al., 2011) compares the performance of stochastic gradient and quasi-Newton methods on neural network training, and finds both to be competitive.

Combining quasi-Newton and stochastic gradient methods could improve optimization time, and reduce the need to tweak optimization hyperparameters. This problem has been approached from a number of directions. In (Schraudolph et al., 2007; Suneag et al., 2009) a stochastic variant of LBFGS is proposed. In (Martens, 2010), (Byrd et al., 2011), and (Vinyals & Povey, 2011) stochastic versions of Hessian-free optimization are implemented and applied to optimization of deep networks. In (Lin et al., 2008) a trust region Newton method is used to train logistic regression

and linear SVMs using minibatches. In (Hennig, 2013) a nonparametric quasi-Newton algorithm is proposed based on noisy gradient observations and a Gaussian process prior. In (Byrd et al., 2014) LBFGS is performed, but with the contributing changes in gradient and position replaced by exactly computed Hessian vector products computed periodically on extra large minibatches. Stochastic meta-descent (Schraudolph, 1999), AdaGrad (Duchi et al., 2010), and SGD-QN (Bordes et al., 2009) rescale the gradient independently for each dimension, and can be viewed as accumulating something similar to a diagonal approximation to the Hessian. All of these techniques treat the Hessian on a subset of the data as a noisy approximation to the full Hessian. To reduce noise in the Hessian approximation, they rely on regularization and very large minibatches to descend $F(x)$. Thus, unfortunately each update step requires the evaluation of many subfunctions and/or yields a highly regularized (i.e. diagonal) approximation to the full Hessian.

We develop a novel second-order quasi-Newton technique that only requires the evaluation of *a single* subfunction per update step. In order to achieve this substantial simplification, we treat the full Hessian of each subfunction as a direct target for estimation, thereby maintaining a separate quadratic approximation of each subfunction. This approach differs from all previous work, which in contrast treats the Hessian of each subfunction as a noisy approximation to the full Hessian. Our approach allows us to com-

bine Hessian information from multiple subfunctions in a much more natural and efficient way than previous work, and avoids the requirement of large minibatches per update step to accurately estimate the full Hessian. Moreover, we develop a novel method to maintain computational tractability and limit the memory requirements of this quasi-Newton method in the face of high dimensional optimization problems (large M). We do this by storing and manipulating the subfunctions in a shared, adaptive low dimensional subspace, determined by the recent history of the gradients and positions.

Thus our optimization method can usefully estimate and utilize powerful second-order information while simultaneously combatting two potential sources of computational intractability: large numbers of subfunctions (large N) and a high-dimensional optimization domain (large M). Moreover, the use of a second order approximation means that minimal or no adjustment of hyperparameters is required. We refer to the resulting algorithm as Sum of Functions Optimizer (SFO). We demonstrate that the combination of techniques and new ideas inherent in SFO results in faster optimization on seven disparate example problems. Finally, we release the optimizer and the test suite as open source Python and MATLAB packages.

2. Algorithm

Our goal is to combine the benefits of stochastic and quasi-Newton optimization techniques. We first describe the general procedure by which we optimize the parameters \mathbf{x} . We then describe the construction of the shared low dimensional subspace which makes the algorithm tractable in terms of computational overhead and memory for large problems. This is followed by a description of the BFGS method by which an online Hessian approximation is maintained for each subfunction. Finally, we end this section with a review of implementation details.

2.1. Approximating Functions

We define a series of functions $G^t(\mathbf{x})$ intended to approximate $F(\mathbf{x})$,

$$G^t(\mathbf{x}) = \sum_{i=1}^N g_i^t(\mathbf{x}), \quad (3)$$

where the superscript t indicates the learning iteration. Each $g_i^t(\mathbf{x})$ serves as a quadratic approximation to the corresponding $f_i(\mathbf{x})$. The functions $g_i^t(\mathbf{x})$ will be stored, and one of them will be updated per learning step.

2.2. Update Steps

As is illustrated in Figure 1, optimization is performed by repeating the steps:

1. Choose a vector \mathbf{x}^t by minimizing the approximating objective function $G^{t-1}(\mathbf{x})$,

$$\mathbf{x}^t = \underset{\mathbf{x}}{\operatorname{argmin}} G^{t-1}(\mathbf{x}). \quad (4)$$

Since $G^{t-1}(\mathbf{x})$ is a sum of quadratic functions $g_i^{t-1}(\mathbf{x})$, it can be exactly minimized by a Newton step,

$$\mathbf{x}^t = \mathbf{x}^{t-1} - \eta^t (\mathbf{H}^{t-1})^{-1} \frac{\partial G^{t-1}(\mathbf{x}^{t-1})}{\partial \mathbf{x}}, \quad (5)$$

where \mathbf{H}^{t-1} is the Hessian of $G^{t-1}(\mathbf{x})$. The step length η^t is typically unity, and will be discussed in Section 3.5.

2. Choose an index $j \in \{1 \dots N\}$, and update the corresponding approximating subfunction $g_j^t(\mathbf{x})$ using a second order power series around \mathbf{x}^t , while leaving all other subfunctions unchanged,

$$g_i^t(\mathbf{x}) = \begin{cases} g_i^{t-1}(\mathbf{x}) & i \neq j \\ \left[\begin{array}{l} f_i(\mathbf{x}^t) \\ + (\mathbf{x} - \mathbf{x}^t)^T f'_i(\mathbf{x}^t) \\ + \frac{1}{2} (\mathbf{x} - \mathbf{x}^t)^T \mathbf{H}_i^t (\mathbf{x} - \mathbf{x}^t) \end{array} \right] & i = j \end{cases}. \quad (6)$$

The constant and first order term in Equation 6 are set by evaluating the subfunction and gradient, $f_j(\mathbf{x}^t)$ and $f'_j(\mathbf{x}^t)$. The quadratic term \mathbf{H}_j^t is set by using the BFGS algorithm to generate an online approximation to the true Hessian of subfunction j based on its history of gradient evaluations (see Section 2.4). The Hessian of the summed approximating function $G^t(\mathbf{x})$ in Equation 5 is the sum of the Hessians for each $g_j^t(\mathbf{x})$, $\mathbf{H}^t = \sum_j \mathbf{H}_j^t$.

2.3. A Shared, Adaptive, Low-Dimensional Representation

The dimensionality M of $\mathbf{x} \in \mathcal{R}^M$ is typically large. As a result, the memory and computational cost of working directly with the matrices $\mathbf{H}_i^t \in \mathcal{R}^{M \times M}$ is typically prohibitive, as is the cost of storing the history terms $\Delta f'$ and $\Delta \mathbf{x}$ required by BFGS (see Section 2.4). To reduce the dimensionality from M to a tractable value, all history is instead stored and all updates computed in a lower dimensional subspace, with dimensionality between K_{min} and K_{max} . This subspace is constructed such that it includes the most recent gradient and position for every subfunction, and thus $K_{min} \geq 2N$. This guarantees that the subspace includes both the steepest gradient descent direction over the full batch, and the update directions from the most recent Newton steps (Equation 5).

For the results in this paper, $K_{min} = 2N$ and $K_{max} = 3N$. The subspace is represented by the orthonormal columns of a matrix $\mathbf{P}^t \in \mathcal{R}^{M \times K^t}$, $(\mathbf{P}^t)^T \mathbf{P}^t = \mathbf{I}$. K^t is the subspace dimensionality at optimization step t .

2.3.1. EXPANDING THE SUBSPACE WITH A NEW OBSERVATION

At each optimization step, an additional column is added to the subspace, expanding it to include the most recent gradient direction. This is done by first finding the component in the gradient vector which lies outside the existing subspace, and then appending that component to the current subspace,

$$\mathbf{q}_{\text{orth}} = f'_j(\mathbf{x}^t) - \mathbf{P}^{t-1} (\mathbf{P}^{t-1})^T f'_j(\mathbf{x}^t), \quad (7)$$

$$\mathbf{P}^t = \left[\mathbf{P}^{t-1} \quad \frac{\mathbf{q}_{\text{orth}}}{\|\mathbf{q}_{\text{orth}}\|} \right], \quad (8)$$

where j is the subfunction updated at time t . The new position \mathbf{x}^t is included automatically, since the position update was computed within the subspace \mathbf{P}^{t-1} . Vectors embedded in the subspace \mathbf{P}^{t-1} can be updated to lie in \mathbf{P}^t simply by appending a 0, since the first K^{t-1} dimensions of \mathbf{P}^t consist of \mathbf{P}^{t-1} .

2.3.2. RESTRICTING THE SIZE OF THE SUBSPACE

In order to prevent the dimensionality K^t of the subspace from growing too large, whenever $K^t > K_{max}$, the subspace is collapsed to only include the most recent gradient and position measurements from each subfunction. The orthonormal matrix representing this collapsed subspace is computed by a QR decomposition on the most recent gradients and positions. A new collapsed subspace is thus computed as,

$$\mathbf{P}' = \text{orth} \left(\begin{bmatrix} f'_1(\mathbf{x}^{\tau_1}) & \cdots & f'_N(\mathbf{x}^{\tau_N}) & \mathbf{x}^{\tau_1} & \cdots & \mathbf{x}^{\tau_N} \end{bmatrix} \right), \quad (9)$$

where τ_i^t indicates the learning step at which the i th subfunction was most recently evaluated, prior to the current learning step t . Vectors embedded in the prior subspace \mathbf{P} are projected into the new subspace \mathbf{P}' by multiplication with a projection matrix $\mathbf{T} = (\mathbf{P}')^T \mathbf{P}$. Vector components which point outside the subspace defined by the most recent positions and gradients are lost in this projection.

Note that the subspace \mathbf{P}' lies within the subspace \mathbf{P} . The QR decomposition and the projection matrix \mathbf{T} are thus both computed within \mathbf{P} , reducing the computational and memory cost (see Section 4.1).

2.4. Online Hessian Approximation

An independent online Hessian approximation \mathbf{H}_j^t is maintained for each subfunction j . It is computed by performing

BFGS on the history of gradient evaluations and positions for that single subfunction¹.

2.4.1. HISTORY MATRICES

For each subfunction j , we construct two matrices, $\Delta f'$ and $\Delta \mathbf{x}$. Each column of $\Delta f'$ holds the change in the gradient of subfunction j between successive evaluations of that subfunction, including all evaluations up until the present time. Each column of $\Delta \mathbf{x}$ holds the corresponding change in the position \mathbf{x} between successive evaluations. Both matrices are truncated after a number of columns L , meaning that they include information from only the prior $L + 1$ gradient evaluations for each subfunction. For all results in this paper, $L = 10$ (identical to the default history length for the LBFGS implementation used in Section 5).

2.4.2. BFGS UPDATES

The BFGS algorithm functions by iterating through the columns in $\Delta f'$ and $\Delta \mathbf{x}$, from oldest to most recent. Let s be a column index, and \mathbf{B}_s be the approximate Hessian for subfunction j after processing column s . For each s , the approximate Hessian matrix \mathbf{B}_s is set so that it obeys the secant equation $\Delta f'_s = \mathbf{B}_s \Delta \mathbf{x}_s$, where $\Delta f'_s$ and $\Delta \mathbf{x}_s$ are taken to refer to the s th columns of the gradient difference and position difference matrix respectively.

In addition to satisfying the secant equation, \mathbf{B}_s is chosen such that the difference between it and the prior estimate \mathbf{B}_{s-1} has the smallest weighted Frobenius norm². This produces the standard BFGS update equation

$$\mathbf{B}_s = \mathbf{B}_{s-1} + \frac{\Delta f'_s \Delta f'^T_s}{\Delta f'^T_s \Delta \mathbf{x}_s} - \frac{\mathbf{B}_{s-1} \Delta \mathbf{x}_s \Delta \mathbf{x}_s^T \mathbf{B}_{s-1}}{\Delta \mathbf{x}_s^T \mathbf{B}_{s-1} \Delta \mathbf{x}_s}. \quad (10)$$

The final update is used as the approximate Hessian for subfunction j , $\mathbf{H}_j^t = \mathbf{B}_{\max(s)}$.

3. Implementation Details

Here we briefly review additional design choices that were made when implementing this algorithm. Each of these choices is presented more thoroughly in Appendix C. Supplemental Figure C.1 demonstrates that the optimizer performance is robust to changes in several of these design

¹We additionally experimented with Symmetric Rank 1 (Dennis Jr & Moré, 1977) updates to the approximate Hessian, but found they performed worse than BFGS. See Supplemental Figure C.1.

²The weighted Frobenius norm is defined as $\|\mathbf{E}\|_{F, \mathbf{W}} = \|\mathbf{W} \mathbf{E} \mathbf{W}\|_F$. For BFGS, $\mathbf{W} = \mathbf{B}_s^{-\frac{1}{2}}$ (Papakonstantinou, 2009). Equivalently, in BFGS the unweighted Frobenius norm is minimized after performing a linear change of variables to map the new approximate Hessian to the identity matrix.

Optimizer	Computation per pass	Memory use
SFO	$\mathcal{O}(QN + MN^2)$	$\mathcal{O}(MN)$
SFO, ‘sweet spot’	$\mathcal{O}(QN)$	$\mathcal{O}(MN)$
LBFGS	$\mathcal{O}(QN + ML)$	$\mathcal{O}(ML)$
SGD	$\mathcal{O}(QN)$	$\mathcal{O}(M)$
AdaGrad	$\mathcal{O}(QN)$	$\mathcal{O}(M)$
SAG	$\mathcal{O}(QN)$	$\mathcal{O}(MN)$

Table 1. Leading terms in the computational cost and memory requirements for SFO and several competing algorithms. Q is the cost of computing the value and gradient for a single subfunction, M is the number of data dimensions, N is the number of subfunctions, and L is the number of history terms retained. “SFO, ‘sweet spot’” refers to the case discussed in Section 4.1.1 where the minibatch size is adjusted to match computational overhead to subfunction evaluation cost. For this table, it is assumed that $M \gg N \gg L$.

choices.

3.1. BFGS Initialization

The first time a subfunction is evaluated (before there is sufficient history to run BFGS), the approximate Hessian \mathbf{H}_j^t is set to the identity times the median eigenvalue of the average Hessian of the other active subfunctions. For later evaluations, the initial BFGS matrix is set to a scaled identity matrix, $\mathbf{B}_0 = \beta \mathbf{I}$, where β is the minimum eigenvalue found by solving the squared secant equation for the full history. See Appendix C.1 for details and motivation.

3.2. Enforcing Positive Definiteness

It is typical in quasi-Newton techniques to enforce that the Hessian approximation remain positive definite. In SFO each \mathbf{H}_i^t is constrained to be positive definite by an explicit eigendecomposition and hard thresholding. This is computationally cheap due to the shared low dimensional subspace (Section 2.3). This is described in detail in Appendix C.2.

3.3. Choosing a Target Subfunction

The subfunction j to update in Equation 6 is chosen to be the one farthest from the current location \mathbf{x}^t , using the current Hessian approximation \mathbf{H}^t as the metric. This is described more formally in Appendix C.3. As illustrated in Supplemental Figure C.1, this distance based choice outperforms the commonly used random choice of subfunction.

3.4. Growing the Number of Active Subfunctions

For many problems of the form in Equation 1, the gradient information is nearly identical between the different subfunctions early in learning. We therefore begin with only

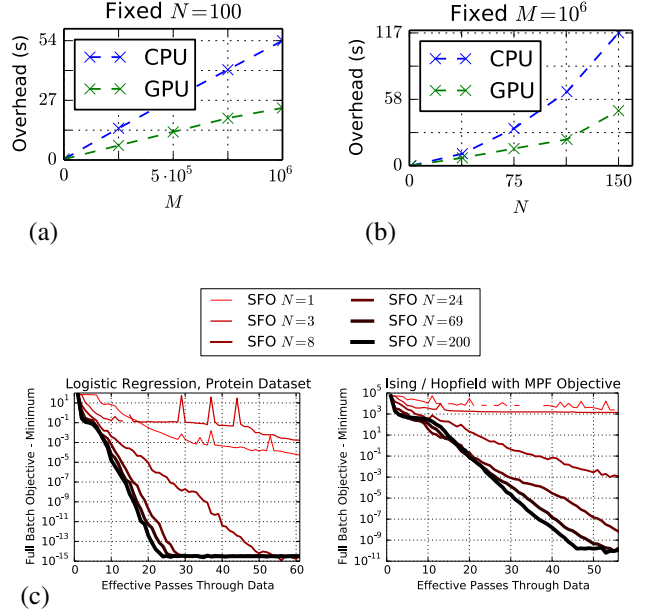


Figure 2. An exploration of computational overhead and optimizer performance, especially as the number of minibatches or subfunctions N is adjusted. (a) Computational overhead required by SFO for a full pass through all the subfunctions as a function of dimensionality M for fixed $N = 100$. (b) Computational overhead of SFO as a function of N for fixed $M = 10^6$. Both plots show the computational time required for a full pass of the optimizer, excluding time spent computing the target objective and gradient. This time is dominated by the $\mathcal{O}(MN^2)$ cost per pass of N iterations of subspace projection. CPU indicates that all computations were performed on a 2012 Intel i7-3970X CPU (6 cores, 3.5 GHz). GPU indicates that subspace projection was performed on a GeForce GTX 660 Ti GPU. (c) Optimization performance on the two convex problems in Section 5 as a function of the number of minibatches N . Note that near maximal performance is achieved after breaking the target problem into only a small number of minibatches.

two active subfunctions, and expand the active set whenever the length of the standard error in the gradient across subfunctions exceeds the length of the gradient. This process is described in detail in Appendix C.4. As illustrated in Supplemental Figure C.1, performance only differs from the case where all subfunctions are initially active for the first several optimization passes.

3.5. Detecting Bad Updates

Small eigenvalues in the Hessian can cause update steps to overshoot severely (ie, if higher than second order terms come to dominate within a distance which is shorter than the suggested update step). It is therefore typical in quasi-Newton methods such as BFGS, LBFGS, and Hessian-free optimization to detect and reject bad proposed update steps, for instance by a line search. In SFO, bad update steps

are detected by comparing the measured subfunction value $f_j(\mathbf{x}^t)$ to its quadratic approximation $g_j^{t-1}(\mathbf{x}^t)$. This is discussed in detail in Section C.5.

4. Properties

4.1. Computational Overhead and Storage Cost

Table 1 compares the cost of SFO to competing algorithms. The dominant computational costs are the $\mathcal{O}(MN)$ cost of projecting the M dimensional gradient and current parameter values into and out of the $\mathcal{O}(N)$ dimensional active subspace for each learning iteration, and the $\mathcal{O}(Q)$ cost of evaluating a single subfunction. The dominant memory cost is $\mathcal{O}(MN)$, and stems from storing the active subspace \mathbf{P}^t . Table A.1 in the Supplemental Material provides the contribution to the computational cost of each component of SFO. Figure 2 plots the computational overhead per a full pass through all the subfunctions associated with SFO as a function of M and N . If each of the N subfunctions corresponds to a minibatch, then the computational overhead can be shrunk as described in Section 4.1.1.

Without the low dimensional subspace, the leading term in the computational cost of SFO would be the far larger $\mathcal{O}(M^{2.4})$ cost per iteration of inverting the approximate Hessian matrix in the full M dimensional parameter space, and the leading memory cost would be the far larger $\mathcal{O}(M^2N)$ from storing an $M \times M$ dimensional Hessian for all N subfunctions.

4.1.1. IDEAL MINIBATCH SIZE

Many objective functions consist of a sum over a number of data points D , where D is often larger than M . For example, D could be the number of training samples in a supervised learning problem, or data points in maximum likelihood estimation. To control the computational overhead of SFO in such a regime, it is useful to choose each subfunction in Equation 3 to itself be a sum over a minibatch of data points of size S , yielding $N = \frac{D}{S}$. This leads to a computational cost of evaluating a single subfunction and gradient of $\mathcal{O}(Q) = \mathcal{O}(MS)$. The computational cost of projecting this gradient from the full space to the shared N dimensional adaptive subspace, on the other hand, is $\mathcal{O}(MN) = \mathcal{O}(M\frac{D}{S})$. Therefore, in order for the costs of function evaluation and projection to be the same order, the minibatch size S should be proportional to \sqrt{D} , yielding

$$N \propto \sqrt{D}. \quad (11)$$

The constant of proportionality should be chosen small enough that the majority of time is spent evaluating the subfunction. In most problems of interest, $\sqrt{D} \ll M$, justifying the relevance of the regime in which the number of

subfunctions N is much less than the number of parameters M . Finally, the computational and memory costs of SFO are the same for sparse and non-sparse objective functions, while Q is often much smaller for a sparse objective. Thus the ideal $S(N)$ is likely to be larger (smaller) for sparse objective functions.

Note that as illustrated in Figure 2c and Figure 3 performance is very good even for small N .

4.2. Convergence

Concurrent work by (Mairal, 2013) considers a similar algorithm to that described in Section 2.2, but with \mathbf{H}_i^t a scalar constant rather than a matrix. Proposition 6.1 in (Mairal, 2013) shows that in the case that each g_i majorizes its respective f_i , and subject to some additional smoothness constraints, $G^t(\mathbf{x})$ monotonically decreases, and \mathbf{x}^* is an asymptotic stationary point. Proposition 6.2 in (Mairal, 2013) further shows that for strongly convex f_i , the algorithm exhibits a linear convergence rate to \mathbf{x}^* . A near identical proof should hold for a simplified version of SFO, with random subfunction update order, and with \mathbf{H}_i^t regularized in order to guarantee satisfaction of the majorization condition.

5. Experimental Results

We compared our optimization technique to several competing optimization techniques for seven objective functions. The results are illustrated in Figures 3 and 4, and the optimization techniques and objectives are described below. For all problems our method outperformed all other techniques in the comparison.

Open source code which implements the proposed technique and all competing optimizers, and which directly generates the plots in Figures 1, 2, and 3, is provided at <https://github.com/Sohl-Dickstein/Sum-of-Functions-Optimizer>.

5.1. Optimizers

SFO refers to Sum of Functions Optimizer, and is the new algorithm presented in this paper. *SAG* refers to Stochastic Average Gradient method, with the trailing number providing the Lipschitz constant. *SGD* refers to Stochastic Gradient Descent, with the trailing number indicating the step size. *ADAGrad* indicates the AdaGrad algorithm, with the trailing number indicating the initial step size. *LBFGS* refers to the limited memory BFGS algorithm. *LBFGS minibatch* repeatedly chooses one tenth of the subfunctions, and runs LBFGS for ten iterations on them. *Hessian-free* refers to Hessian-free optimization.

For *SAG*, *SGD*, and *ADAGrad* the hyperparameter was cho-

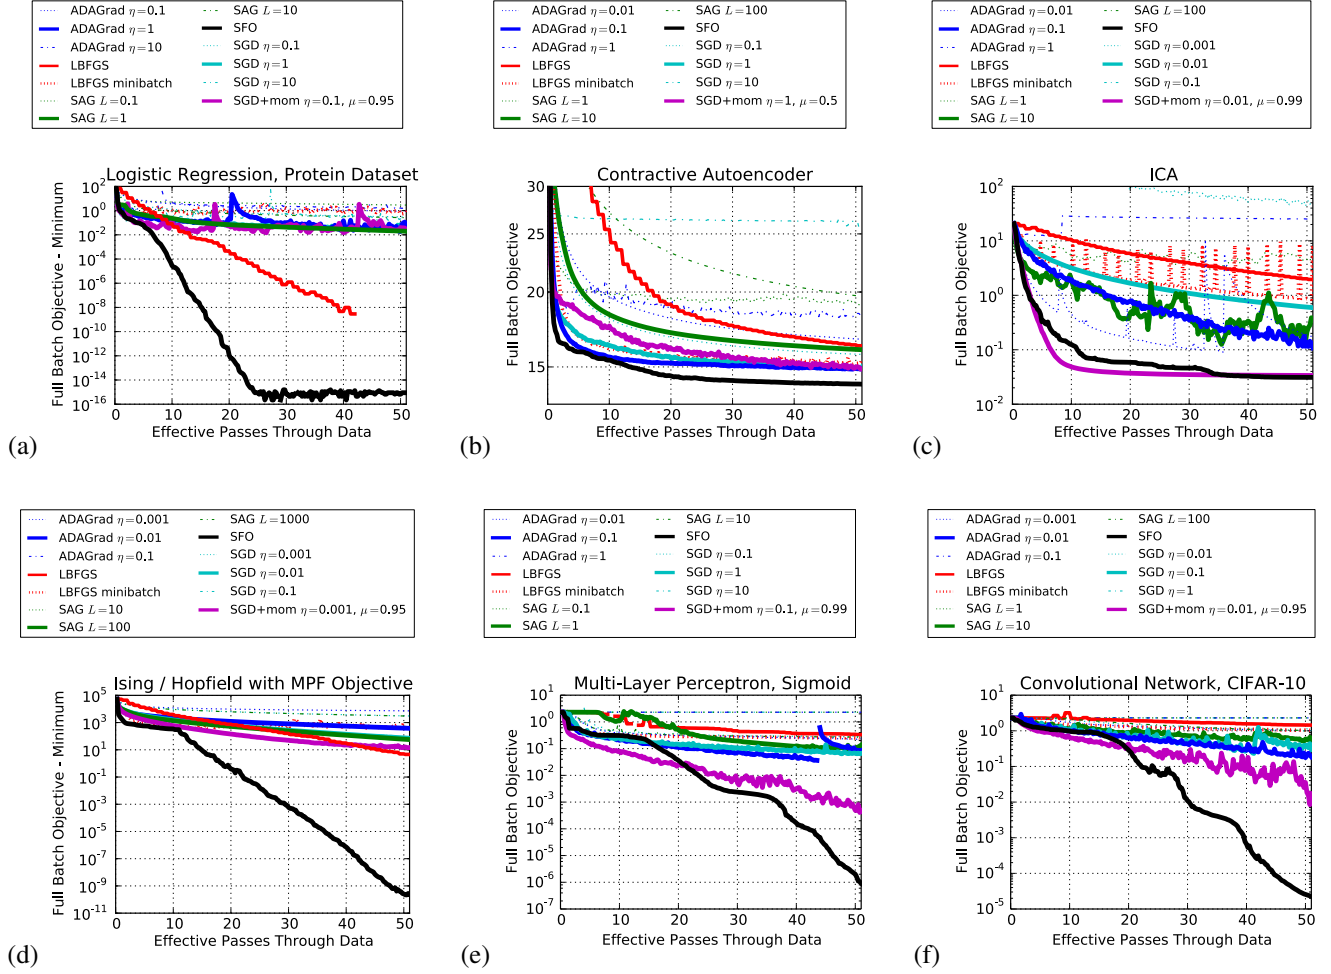


Figure 3. A comparison of SFO to competing optimization techniques for six objective functions. The bold lines indicate the best performing hyperparameter for each optimizer. Note that unlike all other techniques besides LBFGS, SFO does not require tuning hyperparameters (for instance, the displayed SGD+momentum traces are the best out of 32 hyperparameter configurations). The objective functions shown are (a) a logistic regression problem, (b) a contractive autoencoder trained on MNIST digits, (c) an Independent Component Analysis (ICA) model trained on MNIST digits, (d) an Ising model / Hopfield associative memory trained using Minimum Probability Flow, (e) a multi-layer perceptron with sigmoidal units trained on MNIST digits, and (f) a multilayer convolutional network with rectified linear units trained on CIFAR-10. The logistic regression and MPF Ising objectives are convex, and their objective values are plotted relative to the global minimum.

sen by a grid search. The best hyperparameter value, and the hyperparameter values immediately larger and smaller in the grid search, are shown in the plots and legends for each model in Figure 3. In *SGD+momentum* the two hyperparameters for both step size and momentum coefficient were chosen by a grid search, but only the best parameter values are shown. The grid-searched momenta were 0.5, 0.9, 0.95, and 0.99, and the grid-searched step lengths were all integer powers of ten between 10^{-5} and 10^2 . For *Hessian-free*, the hyperparameters, source code, and objective function are identical to those used in (Martens, 2010), and the training data was divided into four “chunks.” For all other experiments and optimizers the training data was divided into $N = 100$ minibatches (or subfunctions).

5.2. Objective Functions

A detailed description of all target objective functions in Figure 3 is included in Section B of the Supplemental Material. In brief, they consisted of:

- A logistic regression objective, chosen to be the same as one used in (Roux et al., 2012).
- A contractive autoencoder with 784 visible units, and 256 hidden units, similar to the one in (Rifai et al., 2011).
- An Independent Components Analysis (ICA) (Bell & Sejnowski, 1995) model with Student’s t-distribution prior.

- An Ising model / Hopfield network trained using code from (Hillar et al., 2012) implementing MPF (Sohl-Dickstein et al., 2011b;a).
- A multilayer perceptron with a similar architecture to (Hinton et al., 2012), with layer sizes of 784, 1200, 1200, and 10. Training used Theano (Bergstra & Breuleux, 2010).
- A deep convolutional network with max pooling and rectified linear units, similar to (Goodfellow & Warde-Farley, 2013a), with two convolutional layers with 48 and 128 units, and one fully connected layer with 240 units. Training used Theano and Pylearn2 (Goodfellow & Warde-Farley, 2013b).

The logistic regression and Ising model / Hopfield objectives are convex, and are plotted relative to their global minimum. The global minimum was taken to be the smallest value achieved on the objective by any optimizer.

In Figure 4, a twelve layer neural network was trained on cross entropy reconstruction error for the CURVES dataset. This objective, and the parameter initialization, was chosen to be identical to an experiment in (Martens, 2010).

6. Future Directions

We perform optimization in an $\mathcal{O}(N)$ dimensional subspace. It may be possible, however, to drastically reduce the dimensionality of the active subspace without significantly reducing optimization performance. For instance, the subspace could be determined by accumulating, in an online fashion, the leading eigenvectors of the covariance matrix of the gradients of the subfunctions, as well as the leading eigenvectors of the covariance matrix of the update steps. This would reduce memory requirements and computational overhead even for large numbers of subfunctions (large N).

Most portions of the presented algorithm are naively parallelizable. The $g_i^t(\mathbf{x})$ functions can be updated asynchronously, and can even be updated using function and gradient evaluations from old positions \mathbf{x}^τ , where $\tau < t$. Developing a parallelized version of this algorithm could make it a useful tool for massive scale optimization problems. Similarly, it may be possible to adapt this algorithm to an online / infinite data context by replacing subfunctions in a rolling fashion.

Quadratic functions are often a poor match to the geometry of the objective function (Pascanu et al., 2012). Neither the dynamically updated subspace nor the use of independent approximating subfunctions $g_i^t(\mathbf{x})$ which are fit to the true subfunctions $f_i(\mathbf{x})$ depend on the functional form of $g_i^t(\mathbf{x})$. Exploring non-quadratic approximating subfunctions has the potential to greatly improve performance.




Figure 4. A comparison of SFO to Hessian-free optimization for a twelve layer neural network trained on the CURVES dataset. This problem is identical to an experiment in (Martens, 2010), and the Hessian-free convergence trace was generated using source code from the same paper. SFO converges in approximately one tenth the number of effective passes through the data as Hessian-free optimization.

Section 3.1 initializes the approximate Hessian using a diagonal matrix. Instead, it might be effective to initialize the approximate Hessian for each subfunction using the average approximate Hessian from all other subfunctions. Where individual subfunctions diverged they would overwrite this initialization. This would take advantage of the fact that the Hessians for different subfunctions are very similar for many objective functions.

Recent work has explored the non-asymptotic convergence properties of stochastic optimization algorithms (Bach & Moulines, 2011). It may be fruitful to pursue a similar analysis in the context of SFO.

Finally, the natural gradient (Amari, 1998) can greatly accelerate optimization by removing the effect of dependencies and relative scalings between parameters. The natural gradient can be simply combined with other optimization methods by performing a change of variables, such that in the new parameter space the natural gradient and the ordinary gradient are identical (Sohl-Dickstein, 2012). It should be straightforward to incorporate this change-of-variables technique into SFO.

7. Conclusion

We have presented an optimization technique which combines the benefits of LBFGS-style quasi-Newton optimization and stochastic gradient descent. It does this by using BFGS to maintain an independent quadratic approximation for each contributing subfunction (or minibatch) in an objective function. Each optimization step then alternates between descending the quadratic approximation of the full objective, and evaluating a single subfunction and updating the quadratic approximation for that single subfunction. This procedure is made tractable in memory and computational time by working in a shared low dimensional subspace defined by the history of gradient evaluations.

References

- Amari, Shun-Ichi. Natural Gradient Works Efficiently in Learning. *Neural Computation*, 10(2):251–276, 1998. ISSN 08997667. doi: 10.1162/089976698300017746.
- Bach, F and Moulines, E. Non-strongly-convex smooth stochastic approximation with convergence rate $O(1/n)$. *Neural Information Processing Systems*, 2013.
- Bach, FR and Moulines, E. Non-Asymptotic Analysis of Stochastic Approximation Algorithms for Machine Learning. *Neural Information Processing Systems*, 2011.
- Bell, AJ and Sejnowski, TJ. An information-maximization approach to blind separation and blind deconvolution. *Neural computation*, 1995.
- Bergstra, J and Breuleux, O. Theano: a CPU and GPU math expression compiler. *Proceedings of the Python for Scientific Computing Conference (SciPy)*, 2010.
- Blatt, Doron, Hero, Alfred O, and Gauchman, Hillel. A convergent incremental gradient method with a constant step size. *SIAM Journal on Optimization*, 18(1): 29–51, 2007.
- Bordes, Antoine, Bottou, Léon, and Gallinari, Patrick. SGD-QN: Careful quasi-Newton stochastic gradient descent. *The Journal of Machine Learning Research*, 10:1737–1754, 2009.
- Bottou, Léon. Stochastic gradient learning in neural networks. *Proceedings of Neuro-Nimes*, 91:8, 1991.
- Boyd, S P and Vandenberghe, L. *Convex optimization*. Cambridge Univ Press, 2004. ISBN 0521833787.
- Broyden, CG. The convergence of a class of double-rank minimization algorithms 2. The new algorithm. *IMA Journal of Applied Mathematics*, 1970.
- Byrd, RH, Hansen, SL, Nocedal, J, and Singer, Y. A Stochastic Quasi-Newton Method for Large-Scale Optimization. *arXiv preprint arXiv:1401.7020*, 2014.
- Byrd, RH Richard H, Chin, GM Gillian M, Neveitt, Will, and Nocedal, Jorge. On the use of stochastic hessian information in optimization methods for machine learning. *SIAM Journal on Optimization*, 21(3):977–995, 2011.
- Dennis Jr, John E and Moré, Jorge J. Quasi-Newton methods, motivation and theory. *SIAM review*, 19(1):46–89, 1977.
- Duchi, John, Hazan, Elad, and Singer, Yoram. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2010.
- Fletcher, R. A new approach to variable metric algorithms. *The computer journal*, 1970.
- Goldfarb, D. A family of variable-metric methods derived by variational means. *Mathematics of computation*, 1970.
- Goodfellow, IJ and Warde-Farley, D. Maxout networks. *arXiv:1302.4389*, 2013a.
- Goodfellow, IJ and Warde-Farley, D. Pylearn2: a machine learning research library. *arXiv:1308.4214*, 2013b.
- Hennig, P. Fast probabilistic optimization from noisy gradients. *International Conference on Machine Learning*, 2013.
- Hillar, Christopher, Sohl-Dickstein, Jascha, and Koepsell, Kilian. Efficient and optimal binary Hopfield associative memory storage using minimum probability flow. *arXiv*, 1204.2916, April 2012.
- Hinton, Geoffrey E., Srivastava, Nitish, Krizhevsky, Alex, Sutskever, Ilya, and Salakhutdinov, Ruslan R. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv:1207.0580*, 2012.
- Huber, PJ. Robust statistics. *Wiley, New York*, 1981.
- Le, Quoc V, Ngiam, Jiquan, Coates, Adam, Lahiri, Abhik, Prochnow, Bobby, and Ng, Andrew Y. On optimization methods for deep learning. *International Conference on Machine Learning*, 2011.
- Lin, Chih-Jen, Weng, Ruby C, and Keerthi, S Sathiya. Trust region newton method for logistic regression. *The Journal of Machine Learning Research*, 9:627–650, 2008.
- Liu, Dong C DC and Nocedal, Jorge. On the limited memory BFGS method for large scale optimization. *Mathematical programming*, 45(1-3):503–528, 1989.
- Mairal, J. Optimization with First-Order Surrogate Functions. *International Conference on Machine Learning*, 2013.
- Mairal, Julien. Incremental Majorization-Minimization Optimization with Application to Large-Scale Machine Learning. *arXiv:1402.4419*, February 2014.
- Martens, James. Deep learning via Hessian-free optimization. In *Proceedings of the 27th International Conference on Machine Learning (ICML)*, volume 951, pp. 2010, 2010.
- Papakonstantinou, JM. *Historical Development of the BFGS Secant Method and Its Characterization Properties*. 2009.
- Pascanu, Razvan, Mikolov, Tomas, and Bengio, Yoshua. On the difficulty of training Recurrent Neural Networks. *arXiv preprint arXiv:1211.5063*, November 2012.
- Rifai, Salah, Vincent, Pascal, Muller, Xavier, Glorot, Xavier, and Bengio, Yoshua. Contractive auto-encoders: Explicit invariance during feature extraction. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 833–840, 2011.
- Robbins, Herbert and Monro, Sutton. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407, 1951.
- Roux, N Le, Schmidt, M, and Bach, F. A Stochastic Gradient Method with an Exponential Convergence Rate for Finite Training Sets. *NIPS*, 2012.
- Schraudolph, Nicol, Yu, Jin, and Günter, Simon. A stochastic quasi-Newton method for online convex optimization. *AIstats*, 2007.
- Schraudolph, Nicol N. Local gain adaptation in stochastic gradient descent. In *Artificial Neural Networks, 1999. ICANN 99. Ninth International Conference on (Conf. Publ. No. 470)*, volume 2, pp. 569–574. IET, 1999.
- Shanno, DF. Conditioning of quasi-Newton methods for function minimization. *Mathematics of computation*, 1970.
- Sohl-Dickstein, Jascha. The Natural Gradient by Analogy to Signal Whitening, and Recipes and Tricks for its Use. *arXiv:1205.1828v1*, May 2012.
- Sohl-Dickstein, Jascha, Battaglino, Peter, and DeWeese, Michael. New Method for Parameter Estimation in Probabilistic Models: Minimum Probability Flow. *Physical Review Letters*, 107(22):11–14, November 2011a. ISSN 0031-9007. doi: 10.1103/PhysRevLett.107.220601.
- Sohl-Dickstein, Jascha, Battaglino, Peter B., and DeWeese, Michael R. Minimum Probability Flow Learning. *International Conference on Machine Learning*, 107(22):11–14, November 2011b. ISSN 0031-9007. doi: 10.1103/PhysRevLett.107.220601.
- Sunehag, Peter, Trunpf, Jochen, Vishwanathan, S V N, and Schraudolph, Nicol. Variable metric stochastic approximation theory. *arXiv preprint arXiv:0908.3529*, August 2009.
- Vinyals, Oriol and Povey, Daniel. Krylov subspace descent for deep learning. *arXiv preprint arXiv:1111.4259*, 2011.