
Active Probabilistic Inference on Matrices for Pre-Conditioning in Stochastic Optimization

Filip de Roos

Max Planck-Institute for Intelligent Systems and University of Tübingen, Germany
[filip.de.roos|ph]@tue.mpg.de

Philipp Hennig

Abstract

Pre-conditioning is a well-known concept that can significantly improve the convergence of optimization algorithms. For noise-free problems, where good pre-conditioners are not known a priori, iterative linear algebra methods offer one way to efficiently construct them. For the *stochastic* optimization problems that dominate contemporary machine learning, however, this approach is not readily available. We propose an iterative algorithm inspired by classic iterative linear solvers that uses a probabilistic model to *actively* infer a pre-conditioner in situations where Hessian-projections can only be constructed with strong Gaussian noise. The algorithm is empirically demonstrated to efficiently construct effective pre-conditioners for stochastic gradient descent and its variants. Experiments on problems of comparably low dimensionality show improved convergence. In very high-dimensional problems, such as those encountered in deep learning, the pre-conditioner effectively becomes an automatic learning-rate adaptation scheme, which we also empirically show to work well.

1 INTRODUCTION

Contemporary machine learning heavily features empirical risk minimization, on loss functions of the form

$$\mathcal{L}(w) = \frac{1}{|\mathcal{D}|} \sum_{i \in \mathcal{D}} l_i(w), \quad (1)$$

Proceedings of the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS) 2019, Naha, Okinawa, Japan. PMLR: Volume 89. Copyright 2019 by the author(s).

where \mathcal{D} is a data-set and w are parameters of the optimization problem (e.g. the weights of a neural network). Data sub-sampling in such problems gives rise to stochastic optimization problems: If gradients are computed over a randomly sampled batch $\mathcal{B} \subset \mathcal{D}$ of data-points, they give rise to a stochastic gradient

$$\nabla \tilde{\mathcal{L}}(w) = \frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \nabla l_i(w). \quad (2)$$

If \mathcal{D} is large and \mathcal{B} is sampled iid., then $\nabla \tilde{\mathcal{L}}$ is an unbiased estimator and, by the multivariate Central Limit Theorem, is approximately Gaussian distributed (assuming $|\mathcal{B}| \ll |\mathcal{D}|$) around the full-data gradient as (e.g. Balles et al., 2017; van der Vaart, 1998, §2.18)

$$p\left(\nabla \tilde{\mathcal{L}}(w) | \nabla \mathcal{L}\right) = \mathcal{N}\left(\nabla \tilde{\mathcal{L}}(w); \nabla \mathcal{L}(w), |\mathcal{B}|^{-1} \text{cov}(\nabla \mathcal{L})\right) \quad (3)$$

where $\text{cov}(\nabla \mathcal{L})$ is the empirical covariance over \mathcal{D} . For problems of large scale in both data and parameter-space, stochastic optimization using stochastic gradient descent (Robbins & Monro, 1951) and its by now many variants (e.g., momentum (Polyak, 1964), Adam (Kingma & Ba, 2014), etc.) are standards.

For (non-stochastic) gradient descent, it is a classic result that convergence depends on the condition-number of the objective's Hessian $B(w) := \nabla \nabla^\top \mathcal{L}(w)$. For example, Thm. 3.4 in Nocedal & Wright (Nocedal & Wright, 2006) states that the iterates of gradient descent with optimal local step sizes on a twice-differentiable objective $\mathcal{L}(w) : \mathbb{R}^N \rightarrow \mathbb{R}$ converge to a local optimum w_* such that, for sufficiently large i ,

$$\mathcal{L}(w_{i+1}) - \mathcal{L}(w_*) \leq r^2(\mathcal{L}(w_i) - \mathcal{L}(w_*)),$$
$$\text{with } r \in \left(\frac{\lambda_N - \lambda_1}{\lambda_N + \lambda_1}, 1\right), \quad (4)$$

where λ_N and λ_1 are the largest and smallest eigenvalue of the Hessian $B(w)$, respectively. In noise-free optimization, it is thus common practice to try and reduce the condition number $\kappa := \lambda_N/\lambda_1$ of the Hessian, by a linear re-scaling $\bar{w} = P^\top w$ of the input space using a *pre-conditioner* $P \in \mathbb{R}^{N \times N}$.

For ill-conditioned problems ($\kappa \gg 1$), effective pre-conditioning can drastically improve the convergence. Choosing a good pre-conditioner is a problem-dependent art. A good pre-conditioner should decrease the condition number of the problem and be cheap to apply, either through sparseness or low-rank structure.

Sometimes pre-conditioners can be constructed as a good ‘a-priori’ guess of the inverse Hessian B^{-1} . If no such good guess is available, then in the deterministic/non-stochastic setting, iterative linear-algebra methods can be used to build a low-rank pre-conditioner. For example, if P spans the leading eigenvectors of the Hessian, the corresponding eigen-values can be re-scaled to change the spectrum, and by extension the condition number. Iterative methods such as those based on the Lanczos process (Golub & C.F. Van Loan, 2012, §10.1) try to consecutively expand a helpful subspace by choosing B -conjugate vectors, which can fail in the presence of inexact computations. Due to these intricate instabilities (Trefethen & Bau III, 1997, p. 282) such algorithms tend not to work with the level of stochasticity encountered in practical machine learning applications.

Below, we propose a framework for the efficient construction of pre-conditioners in settings with noise-corrupted Hessian-vector products available. Our algorithm consists of three main components: We first build a probabilistic Gaussian inference model for matrices from noisy matrix-vector products (Section 2.1) by extending existing work on matrix-variate Gaussian inference. Then we construct an active algorithm that selects informative vectors aiming to explore the Hessian’s dominant eigen-directions (Section 2.2). The structure of this algorithm is inspired by that of the classic Arnoldi and Lanczos iterations designed for noise-free problems. Finally (Section 2.3), we provide some “plumbing” to empirically estimate hyper-parameters and efficiently construct a low-rank pre-conditioner and extend the algorithm to the case of high-dimensional models (Section 2.4). We evaluate the algorithm on some simple experiments to empirically study its properties as a way to construct pre-conditioners and test it on both a low-dimensional¹, and a high-dimensional deep learning problem. While we use pre-conditioning as the principal setting, the main contribution of our framework is the ability to construct matrix-valued estimates in the presence of noise, and to do so at complexity *linear* in the width and height of the matrix. It is thus applicable to problems of large scale, also in domains other than pre-conditioning and optimization in general. In contrast to a simple averaging of random observations, our algorithm actively chooses projections in an effort to improve the estimate.

¹Code repository.

1.1 Related Work

Our approach is related to optimization methods that try to emulate the behaviour of Newton’s method without incurring its cubic per-step cost. That is, iterative optimization updates $w_{i+1} = w_i - d_i$ that try to find a search direction d_i that is an approximate solution to the linear problem

$$B(w_i)d_i = \nabla\mathcal{L}(w_i). \quad (5)$$

This includes quasi-Newton methods like BFGS and its siblings (Dennis & Moré, 1977), and Hessian-free optimization (Pearlmutter, 1994; Martens, 2010). These methods try to keep track of the Hessian during the optimization. Pre-conditioning is a simpler approach that separates the estimation of the (inverse) Hessian from the on-line phase of optimization and moves it to an initialization phase. Our algorithm could in principle be run in every single step of the optimizer, such as in Hessian-free optimization. However, this would multiply the incurred cost, which is why we here only study its use for pre-conditioning.

There are stochastic variants of quasi-Newton methods and other algorithms originally constructed for noise-free optimization (e.g. Schraudolph et al., 2007; Byrd et al., 2016). These are generally based on collecting independent random (not actively designed) samples of quasi-Newton updates. Estimates can also be obtained by regularizing estimates (e.g. Wills & Schön, 2018) or partly update the stochastic gradient by reusing elements of a batch (e.g. Bollapragada et al., 2018). The conceptual difference between these methods and ours is that we actively try to design informative observations by explicitly taking evaluation uncertainty into account.

Our inference scheme is an extension of Gaussian models for inference on matrix elements, which started with early work by Dawid (1981) and was recently extended in the context of probabilistic numerical methods (Hennig & Kiefel, 2013; Hennig, 2015; Wills & Schön, 2018). Our primary addition to these works is the algebra required for dealing with structured observation noise.

2 THEORY

Our goal in this work is to construct an active inference algorithm for low-rank pre-conditioning matrices that can deal with data in the form of Hessian-vector multiplications corrupted by significant (not just infinitesimal) Gaussian noise. To this end, we will adopt a probabilistic viewpoint with a Gaussian observation likelihood, and design an active evaluation policy that aims to efficiently collect informative, non-redundant

Hessian-vector products. The algorithm will be designed so that it produces a Gaussian posterior measure over the Hessian of the objective function, such that the posterior mean is a low-rank matrix.

2.1 Matrix Inference

Bayesian inference on matrices $B \in \mathbb{R}^{N \times N}$ can be realised efficiently in a Gaussian framework by rearranging the matrix elements into a vector $\vec{B} \in \mathbb{R}^{N^2 \times 1}$, then performing standard Gaussian inference on this vector (e.g. Dawid, 1981). Although Hessian matrices are square and symmetric, the following derivations apply equally well to rectangular matrices. There are specializations for symmetric matrices (Hennig, 2015), but since they significantly complicate the derivations below, we use this weaker model.

Assume that we have access to observations $Y \in \mathbb{R}^{N \times m}$ of matrix-vector products $Y = BS$ along the search directions $S \in \mathbb{R}^{N \times m}$. In the vectorized notation, this amounts to a linear projection of \vec{B} through the Kronecker product matrix $(I \otimes S^\top)$:

$$Y_{ab} = (I \otimes S^\top) \vec{B} = \sum_{ij} \delta_{ai} B_{ij} S_{jb} = [BS]_{ab}.$$

If the observations are exact (noise-free), the likelihood function is a Dirac distribution,

$$\begin{aligned} p(Y | B, S) &= \delta(\vec{Y} - (I \otimes S^\top)B) = \\ &= \lim_{\beta \rightarrow 0} \mathcal{N}(Y; (I \otimes S^\top)B, \beta \Lambda_0). \end{aligned} \quad (6)$$

For conjugate inference, we assign a Gaussian prior over B , with a prior mean matrix B_0 and a covariance consisting of a Kronecker product of 2 symmetric positive-definite matrices.

$$\begin{aligned} \mathcal{N}(B, B_0, V \otimes W) &= \frac{1}{((2\pi)^{n^2} / |V| |W|)^{1/2}} \\ &\cdot \exp\left(-\frac{1}{2}(\vec{B} - \vec{B}_0)^\top (V \otimes W)^{-1} (\vec{B} - \vec{B}_0)\right). \end{aligned} \quad (7)$$

This is an equivalent re-formulation of the *matrix-variate Gaussian* distribution (Dawid, 1981). For simplicity, and since we are inferring a Hessian matrix B (which is square and symmetric), we set $V = W$. This combination of prior (7) and likelihood (6) has previously been discussed (detailed derivation, e.g., in Hennig, 2015). It gives rise to a Gaussian posterior distribution, whose mean matrix is given by

$$\begin{aligned} B_m = B_0 + (W \otimes W)(I \otimes S) \vec{X} &= B_0 + W X S^\top W = \\ B_0 + (Y - B_0 S)(S^\top W S)^{-1} S^\top W & \end{aligned} \quad (8)$$

where X is found as the solution to the linear system

$$\underbrace{(W \otimes S^\top W S)}_{=:G} \vec{X} = \underbrace{(Y - B_0 S)}_{=: \vec{\Delta}}, \quad (9)$$

using the Kronecker product's property that $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$ (note that the matrix $(S^\top W S) \in \mathbb{R}^{m \times m}$ can be inverted in $\mathcal{O}(m^3)$).

2.1.1 Adding Noise

While noise-free observations of Gaussian-distributed matrices have been studied before, observation noise beyond the fully iid. case, (cf. Wills & Schön, 2018) is a challenging complication. A lightweight inference-scheme for noisy observations is one of the core contributions of this paper. In empirical risk minimization problems, mini-batching replaces a large or infinite sum (the population risk, Eq. (1)) into a small sum of individual loss functions (Eq. (2)). Analogous to Eq. (3), the Hessian \hat{B} of the batch risk $\hat{\mathcal{L}}$ is thus corrupted relative to the true Hessian by B a Gaussian likelihood:

$$\begin{aligned} \hat{B}(w) &= B(w) + \Gamma \quad \text{with} \quad \Gamma \sim \mathcal{N}(0, \Sigma) \\ &\quad \text{and} \quad \Sigma = \text{cov}(\nabla \nabla^\top \mathcal{L}) / |\mathcal{B}|. \end{aligned} \quad (10)$$

If we now compute matrix-vector products of this batch Hessian with vectors $S \in \mathbb{R}^{N \times m}$, even under the simplifying assumption of Kronecker structure $\Sigma = \Lambda \otimes \Lambda$ in the covariance, the observation likelihood becomes

$$p(Y | B, S) = \mathcal{N}(\vec{Y}; \vec{B} \vec{S}, \underbrace{\Lambda \otimes (S^\top \Lambda S)}_{=:R})_{ii}. \quad (11)$$

The subscript *ii* in Eq. (11) is there to highlight the diagonal structure of the right matrix due to the independent batches we use to calculate the Hessian-vector products. Relative to Eq (9), this changes the Gram matrix to be inverted from G to $(G + R)$.

To get the posterior of B in the noisy setting after m observations, instead of Eq. (9), we now have to solve the linear problem

$$\underbrace{(G + R)}_{\mathbb{R}^{nm \times nm}} \vec{X} = \vec{\Delta}. \quad (12)$$

This is a more challenging computation, since the sum of Kronecker products does not generally have an analytic inverse. However, Eq. (12) is a so-called matrix pencil problem, which can be efficiently addressed using a generalized eigen-decomposition (e.g. Golub & C.F Van Loan, 2012, §7.7). That is, by a matrix U and diagonal matrix $D = \text{diag}(\lambda)$ such that

$$GV = DRV \quad \text{with} \quad V^\top RV = I. \quad (13)$$

Eigen-decompositions distribute over a Kronecker product (C.F. Van Loan, 2000). And this property is inherited by the generalized eigen-decomposition, which offers a convenient way to rewrite a matrix in terms of the other. Eq. (12) can thus be solved with two generalized eigen-decompositions. The left and right parts of the Kronecker products of Eq. (12) are written with separate generalized eigen-decompositions as

$$\begin{aligned} WU &= \Lambda UD, & U^\top \Lambda U &= I \\ (S^\top W S)V &= (S^\top \Lambda S)_{ii} V \Omega, & V^\top (S^\top \Lambda S)_{ii} V &= I. \end{aligned}$$

U, D contain the generalized eigen-vectors and eigen-values from the left Kronecker term and V, Ω are analogous for the right Kronecker term.

$$\begin{aligned} \vec{\Delta} &= (W \otimes S^\top W S + \Lambda \otimes (S^\top \Lambda S)_{ii}) \vec{X} \\ &= (\Lambda U \otimes (S^\top \Lambda S)_{ii} V) (D \otimes \Omega + I \otimes I) (U^{-1} \otimes V^{-1}) \vec{X} \\ &= (U^{-T} \otimes V^{-T}) (D \otimes \Omega + I \otimes I) (U^{-1} \otimes V^{-1}) \vec{X} \end{aligned}$$

In the first step above, the left matrix is expressed in terms of the right matrix in both terms by means of the generalized eigen-decomposition. In remaining step, the conjugacy property in Eq. (13) is used to simplify the expression to the inversion of a diagonal matrix and the Kronecker product of the generalized eigen-vectors. The solution now becomes

$$\begin{aligned} \vec{X} &= (U \otimes V) \underbrace{(D \otimes \Omega + I \otimes I)^{-1} \overline{U^\top \Delta V}}_{\substack{\frac{1}{(D_{jj} \Omega_{ii} + 1)} \odot U^\top \Delta V = \Psi_{ji}}} \vec{V} \\ &= (U \otimes V) \vec{\Psi} = U \Psi V^\top \end{aligned} \quad (14)$$

where \odot refers to the Hadamard product of the two matrices. Using this form, we can represent the posterior mean estimate for B with Eq. (8) where X is replaced with the solution from Eq. (14).

$$B_m = B_0 + \underbrace{W X}_{\mathbb{R}^{N \times m}} \underbrace{S^\top W}_{\mathbb{R}^{m \times N}}$$

This matrix can generally not be written in a simpler analytic form. If the prior mean B_0 is chosen as a simple matrix (e.g. a scaled identity), then B_m would admit fast $\mathcal{O}(Nm)$ multiplication and inversion (using the matrix inversion lemma) due to its low-rank outer-product structure.

2.2 Active Inference

The preceding section constructed an inference algorithm that turns noisy projections of a matrix into a low-rank estimate for the latent matrix. The second ingredient of our proposed algorithm, outlined in this section, is an active policy that chooses non-redundant

projection directions to efficiently improve the posterior mean. Algorithm 1 summarizes as pseudo-code.

The structure of this algorithm is motivated, albeit not exactly matched to, that of stationary iterative linear solvers and eigen-solvers such as GMRES (Saad & Schultz, 1986) and Conjugate gradients (Hestenes & Stiefel, 1952) (and the corresponding eigen-value-iterations, the Arnoldi process and the Lanczos process. Cf. (Golub & C.F Van Loan, 2012, §10) and (Trefethen & Bau III, 1997, §VI)). These algorithms can be interpreted as optimization methods that iteratively expand a low-rank approximation to (e.g. in the case of Conjugate gradients / Lanczos) the Hessian of a quadratic problem, then solve the quadratic problem within the span of this approximation. In our algorithm, the exact low-rank approximation is replaced by the posterior mean *estimate* arising from the Bayesian inference routine described in Section 2.1. This leads the algorithm to suppress search directions that are co-linear with those collected in previous iterations, focussing instead on the efficient collection of new information.

Readers familiar with linear solvers like conjugate gradients will be able to recognise the structural similarity of Algorithm 1 to linear solvers, with two differing aspects. Each iteration constructs a projection direction s_i , collects one matrix-vector multiplication, $y_i = \tilde{B} s_i$ and rescales them by a step size β_i (here set to 1 and omitted). A linear solver would update the solution x_i and residual r_i using s_i , y_i and β_i but we let the algorithm stay at x_0 and sample new search directions and projections. The core difference to a solver is in line 7: Where the classic solvers would perform a Gram-Schmidt step, we instead explicitly perform Gaussian inference on the Hessian B . In the noise-free limit the proposed method would choose the same search directions as projection methods, a superclass of iterative solvers containing algorithms such as GMRES and CG (Hennig, 2015).

2.3 Algorithmic Details

Like all numerical methods, our algorithm requires the tuning of some internal parameters and some implementational engineering. We set all free parameters in an empirical fashion, via cheap pre-computations, and use standard linear-algebra tools for the internals.

2.3.1 Estimating Parameters

The Gaussian prior (7) and likelihood (11) (also used as parameters of Alg. 1) have several parameters. For simplicity and to limit computational cost, we set all these to scaled identity matrices: prior mean $B_0 = b_0 I$,

variance $W = w_0 I$ and noise covariance $\Lambda = \lambda_0 I$. We set these parameters empirically: Before the method tries to estimate the Hessian, it gathers gradients and Hessian-gradient products locally, from a small number of initial batches. Then we set the parameters of the prior over B as the empirical estimates

$$1/b_0 = \sqrt{\frac{s^\top B s}{s^\top B B s}} \quad \text{and} \quad w_0 = \frac{s^\top B s}{s^\top s}. \quad (15)$$

The noise variance for the likelihood set by an empirical estimate

$$\lambda_0 = (\mathbb{E}[g^2] - \underbrace{(\mathbb{E}[g])^2}_{\bar{g}}) / \sqrt{s^\top s}.$$

Since batch elements can vary quite extremely, however, we use a more robust choice, by setting it to the median of the variance estimates. The mean gradient (\bar{g}) from the initial sampling is used for the first iteration of the Hessian-inference scheme, line 2. A new search direction along which a Hessian-vector product is to be computed, is obtained by applying the inverse of the current estimated Hessian to a stochastic gradient:

$$s_{i+1} = -\hat{B}_i^{-1} \nabla \tilde{\mathcal{L}}(w). \quad (16)$$

The estimated Hessian is updated by first solving Eq. (12) and using the result in Eq. (8) to get the posterior mean: It is a sum of the prior mean and a low-rank outer product of two $N \times m$ matrices, with N the number of parameters and m the number of observations/iterations. A diagonal prior mean offers efficient inversion of the estimated Hessian by the matrix inversion lemma, and subsequent multiplication in Eq. (16).

In the experiments below, the algorithm is used to construct a pre-conditioner for stochastic optimizers. In this application, it runs at the beginning of an optimization process, collecting several batches “in-place”, each time computing a noisy matrix-Hessian product. To find the dominant k eigen-directions, especially if noise is significant, the solver usually requires a number $m > k$ of iterations, producing a posterior estimate B_m of rank m . To reduce the computational cost and memory requirements in the subsequent actual optimization run, we reduce the rank of this approximation down to k using a standard fast singular value decomposition and obtain the singular values Σ and the left singular vectors U and estimate the Hessian $B \approx U \Sigma U^\top$. This is equivalent to taking the symmetric part of a polar decomposition which yields the closest symmetric approximation to a matrix in Frobenius norm (Higham, 1988). For an N -dimensional problem and a pre-conditioner of rank k , this method requires the storage of $\mathcal{O}(Nk)$ numbers and has a computational overhead of $\mathcal{O}(Nk)$ additional FLOPS compared to a standard sgd update. Such overhead is negligible

for batch-sizes $|\mathcal{B}| > k$, the typical case in machine learning.

Algorithm 1 Active probabilistic solver for linear problems of the form $Bx = b$, where multiplications with the matrix B can only be performed corrupted by noise. When used in optimization, the target vector b is set to the (noisy) gradient $\nabla \tilde{\mathcal{L}}$ (which amounts to searching for the Newton direction).

```

1 procedure INFHESS( $x_0, B(\cdot), b, p(B), p(Y | S, B)$ )
2    $r_0 = B_{\text{true}} \cdot x_0 - b$  // initial noisy gradient
3   for  $i = 1, \dots$  do
4      $s_i = -B_{i-1}^{-1} r_{i-1}$  // step direction
5      $y_i = B_{\text{true}} \cdot s_i$  // observe
6      $r_i = B_{\text{true}} \cdot x_0 - b$  // new noisy gradient at  $x_0$ 
7      $B_i = \text{INFER}(B | Y_i, S_i, p(B), p(Y | S, B))$ 
7     // estimate  $B$ , using results from Section 2.1
8   end for
9 end procedure

```

2.3.2 Pre-conditioning

A pre-conditioner P is constructed to rescale the stochastic gradients in the direction of the singular vectors.

$$P = \alpha(I + U[\beta I_k / \sqrt{\Sigma} - I_k]U^\top) \quad (17)$$

By rescaling the gradients with P^2 , the linear system in Eq. (5) is transformed into $P^\top \tilde{B}(w_i) P P^{-1} d_i = P^\top \nabla \tilde{\mathcal{L}}(w_i)$. The goal, as outlined above, is to reduce the condition number of the transformed Hessian $P^\top \tilde{B}(w_i) P$. If the estimated vectors U, Σ are indeed the real eigen-vectors and eigen-values, this approach would rescale these directions to have length β . Theoretically $\beta = 1$ would be ideal if the real eigen-pairs are used. When instead an approximation of the subspace is used with poor approximations of the eigen-values $\tilde{\lambda}_i$, it is possible to scale a direction too much so the eigen-vectors corresponding to the largest eigen-values become the new smallest eigen-values. In our experiments this rarely happened because the Hessian contained many eigen-values $\lambda_i \ll 1$ and so $\beta = 1$ could be used. The improved condition number allows usage of a larger step-size. We scale up by the empirical improvement of the condition number, i.e. the fixed step-size of sgd η is multiplied with $\alpha^2 = \Sigma_1 / \Sigma_k$ (in Eq. (17)), the ratio between largest and the smallest estimated eigen-value. Using the current notation we can write the pre-conditioned sgd update as

$$w_{i+1} = w_i - \eta P^2 \nabla \tilde{\mathcal{L}}(w_i).$$




Figure 1: Comparison of sgd and pre-conditioned sgd on the linear test problem, along with other baselines. Details in text. The plots show data for four different choices of batch size $|\mathcal{B}|$ (and thus varying observation noise). For fairness to sgd, the abscissa is scaled in the number of data points loaded from disk (as opposed to the number of steps). Due to the noisy setting, vanilla CG is unstable and diverges in the first few steps.

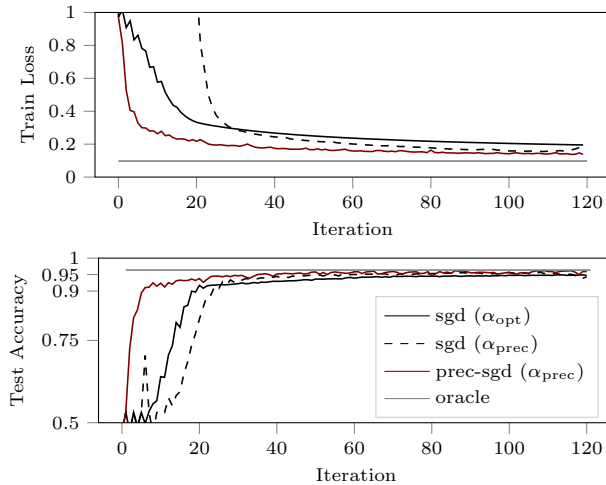


Figure 2: Progress of pre-conditioned and standard sgd on the MNIST logistic regression problem (details in text). Top plot shows progress on the raw optimization objective (train loss), the bottom plot shows generalization in terms of test accuracy. Two lines are shown for sgd. The solid one uses a step-size α_{opt} optimized for good performance. This directly induces a step-size α_{prec} for the pre-conditioner. For comparison, we also show progress of sgd when directly using this step-size—it makes sgd unstable.

2.4 High-Dimensional Modification

Deep learning has become the benchmark for stochastic optimizers and it imposes many new constraints on the optimizers. Due to the large number of parameters, even if we would have access to the fraction of eigen-directions which make the Hessian ill-conditioned it would likely be inefficient to use because each vector has the same size as the network, and the approach would only work if a few of the eigen-values of the

Hessian have a clear separation in magnitude. This changes the functionality we can expect from a pre-conditioner to keeping the parameter updates relevant with respect to the current condition number rather than finding all directions with problematic curvature.

Some simplifications are required in order to adapt the algorithm for deep learning. The most important change is that we approximate the Hessian as a layer-wise block matrix, effectively treating each layer of a deep net as an independent task. Hessian-vector products are calculated using automatic differentiation (Pearlmutter, 1994). Good estimates of the eigen-values for the algorithm proved difficult to get because of large deviations, likely due to the noisy Hessian-vector products. To alleviate this problem we changed the multiplicative update of the step-size presented in section 2.3.2 to redefining the step-length. Each time the algorithm is used to build a pre-conditioner, a new step-length is set to the scalar prior mean in Eq. (15). The last modification is how the empirical parameters of the prior (section 2.3.1) are set. $1/b_0$ is used as the new step-length and gave better results when a smaller step-size of $s^\top B s / s^\top B B s$ was used and λ_0 is estimated to $\lambda_0 = \sqrt{(\sum [g^\top g] - \hat{g}^\top \hat{g}) / n}$, with $\hat{g} = \sum g_{(\mathcal{B})}$. All the parameters of the prior and likelihood (b_0 , w_0 and λ_0) are shared among the layers. No clear improvement was visible when treating the parameters separately for each layer.

3 RESULTS

We analyze and compare the performance of the proposed algorithm on a simple test problem along with two standard applications in machine learning.

3.1 Regression

Figure 1 shows results from a conceptual test setup designed to showcase the algorithm’s potential: An ill-conditioned linear problem of a scale chosen such that the analytical solution can still be found for comparison. We used linear parametric least-squares regression on the SARCOS data-set (Vijayakumar & Schaal, 2000) as a test setup. The data-set contains $|\mathcal{D}| = 44,484$ observations of a uni-variate² response function $y_i = f(x_i)$ in a 21-dimensional space $x_i \in \mathbb{R}^{21}$. We used the polynomial feature functions $\phi(x) = A[x, \text{vec}(xx^\top)] \in \mathbb{R}^{253}$, with a linear mapping, A , manually designed to make the problem ill-conditioned. The model $f(x) = \phi(x)^\top w$ with a quadratic loss function then yields a quadratic optimization problem,

$$w_* = \arg \min_w \alpha \|w\|^2 + \frac{1}{|\mathcal{D}|} \|\Phi^\top w - \mathbf{y}\|^2 = \alpha \|w\|^2 + \frac{1}{|\mathcal{D}|} \sum_{i=1}^{|\mathcal{D}|} (\phi(x_i)^\top w - y_i)^2, \quad (18)$$

where $\Phi \in \mathbb{R}^{253 \times 44,484}$ is the map from weights to data. The exact solution of this problem is given by the regularized least-squares estimate $w_* = (\Phi\Phi^\top/|\mathcal{D}| + \alpha\mathbf{I})^{-1}\Phi\mathbf{y}/|\mathcal{D}|$. For this problem size, this solution can be computed easily, providing an oracle baseline for our experiments. But if the number of features were higher (e.g. $\gtrsim 10^4$), then exact solutions would not be tractable. One could instead compute, as in deep learning, batch gradients from Eq. (18), and also produce a noisy approximation of the Hessian $B = (\Phi\Phi^\top/|\mathcal{D}| + \alpha\mathbf{I})$ as the low rank matrix

$$\tilde{B} = \alpha\mathbf{I} + \frac{1}{|\mathcal{B}|} \sum_{b \in \mathcal{B}} \phi(x_b)\phi(x_b)^\top \quad (19)$$

where \mathcal{B} is a batch. Clearly, multiplying an arbitrary vector with this low-rank matrix has cost $\mathcal{O}(|\mathcal{B}|)$, thus providing the functionality required for our noisy solver. Figure 1 compares the progress of vanilla sgd with that of pre-conditioned sgd if the pre-conditioner is constructed with our algorithm. In each case, the construction of the pre-conditioner was performed in 16 iterations of the inference algorithm. Even in the largest case of $\mathcal{B} = 256$, this amounts to 4096 data read, and thus only a minuscule fraction of the overall runtime.

An alternative approach would be to compute the inverse of \tilde{B} separately for each batch, then average over the batch-least-squares estimate $\tilde{w} = \tilde{B}^{-1}\Phi\mathbf{y}/|\mathcal{B}|$. In our toy setup, this can again be done directly in the

²The data-set contains 7 such univariate target variables. Following convention, we used the first one.

feature space. In an application with larger feature space, this is still feasible using the matrix-inversion lemma on Eq. (19), instead inverting a dense matrix of size $|\mathcal{B}| \times |\mathcal{B}|$. The Figure also shows the progression of this stochastic estimate (labelled as *avg-inv*, always using $|\mathcal{B}| = 256$ since smaller batch-sizes did not work at all). It performs much worse unless the batch-size is increased, which highlights the advantage of the active selection of projection directions for identifying appropriate eigen-vectors. A third option is to use an iterative solver with noise-corrupted observations to approach the optimum. In figure 1 a barely visible line labelled CG can be seen which used the method of conjugate gradients with a batch-size of 256. This method required a batch-size $|\mathcal{B}| > 10000$ to show reasonable convergence on the training objective but would still perform poorly on the test set.

3.2 Logistic Regression

Figure 2 shows an analogous experiment on a more realistic, and *non-linear* problem: Classic linear logistic regression on the digits 3 and 5 from the MNIST data-set (i.e. using linear features $\phi(x) = x$, and $p(y|x) = \sigma(\phi(x)^\top w)$). Here we used the model proposed in Rasmussen & Williams (2005, §3.4), which defines a convex, non-linear regularized empirical risk minimization problem that again allows the construction of stochastic gradients, and associated noisy Hessian-vector products. Analogous to Figure 1, Figure 2 shows progress of sgd and pre-conditioned sgd. As before, this problem is actually just small enough to compute an exact solution by Newton optimization (gray baseline in plot). And as before, computation of the pre-conditioner takes up a small fraction of the optimization runtime.

3.3 Deep Learning

For a high-dimensional test bed, we used a deep net consisting of convolutional and fully-connected layers (see appendix A for details) to classify the CIFAR-10 data-set (Krizhevsky, 2009). The proposed algorithm was implemented in PyTorch (Paszke et al., 2017), using the modifications listed in section 2.4.

To stabilize the algorithm, a predetermined fixed learning-rate was used for the first epoch of the pre-conditioned sgd. Figure 3 and 4 compare the convergence for the proposed algorithm against sgd for training loss and test loss respectively on CIFAR-10. In both figures we see that the pre-conditioned sgd has similar performance to a well-tuned sgd regardless of the initial learning rate. To keep the cost of computa-




Figure 3: Training loss for sgd (left) and pre-conditioned sgd (right) on the CIFAR-10 data-set for different learning rates and batch-size of 32 over 250 epochs. Both graphs share y-axis and colors to facilitate comparison between the optimizers. The solid lines represent the mean of several individual runs plotted as translucent.




Figure 4: Test loss for sgd (left) and pre-conditioned sgd (right) on the CIFAR-10 data-set for different learning rates and batch-size of 32 over 250 epochs. Both graphs share y-axis and colors to facilitate comparison between the optimizers. These graphs were collected at the same runs as the results in Fig. 3.




Figure 5: Evolution of the estimated learning rate α over 250 epochs for different initial values.

tions and storage low we used a rank 2 approximation of the Hessian that was recalculated at the beginning of every epoch. The cost of building the rank 2 approximation was 2–5% of the total computational cost per epoch.

The improved convergence of pre-conditioned sgd over normal sgd is mainly attributed the adaptive step-size, which seems to capture the general scale of the curvature to efficiently make progress. This approach offers a more rigorous way to update the step-length over popular empirical approaches using exponentially decaying learning-rates or division upon plateauing. By studying the scale of the found learning rate, see Fig. 5, we see that regardless of the initial value, all α follow the same trajectory although spanning values across four orders of magnitude.

4 CONCLUSION

We have proposed an active probabilistic inference algorithm to efficiently construct pre-conditioners in stochastic optimization problems. It consists of three conceptual ingredients: First, a matrix-valued Gaussian inference scheme that can deal with structured Gaussian noise in observed matrix-vector products. Second, an active evaluation scheme aiming to collect informative, non-redundant projections. Third, additional statistical and linear algebra machinery to empirically estimate hyper-parameters and arrive at a low-rank pre-conditioner. The resulting algorithm was shown to significantly improve the behaviour of sgd in imperfectly conditioned problems, even in the case of severe observation noise typical for contemporary machine learning problems. It scales from low- to medium- and high-dimensional problems, where its behaviour qualitatively adapts from full and stable pre-conditioning to low-rank pre-conditioning and, eventually, scalar adaptation of the learning rate of sgd-type optimization.

Acknowledgements

Filip de Roos acknowledges support by the International Max Planck Research School for Intelligent Systems. Philipp Hennig gratefully acknowledges financial support by the European Research Council through ERC StG Action 757275 / PANAMA.

References

- Balles, L., J. Romero, & P. Hennig (Aug. 2017). “Coupling Adaptive Batch Sizes with Learning Rates”. In: *Proceedings Conference on Uncertainty in Artificial Intelligence (UAI) 2017*. Association for Uncertainty in Artificial Intelligence (AUAI), pp. 410–419.
- Bollapragada, R. et al. (2018). “A Progressive Batching L-BFGS Method for Machine Learning”. In: *ArXiv e-prints*. arXiv: 1802.05374.
- Byrd, R.H., S.L Hansen, J. Nocedal, & Y. Singer (2016). “A stochastic quasi-Newton method for large-scale optimization”. In: *SIAM Journal on Optimization* 26.2, pp. 1008–1031.
- Dawid, A.P. (1981). “Some matrix-variate distribution theory: Notational considerations and a Bayesian application”. In: *Biometrika* 68.1, pp. 265–274.
- Dennis, J.E. & J.J. Moré (1977). “Quasi-Newton methods, motivation and theory”. In: *SIAM Review* 19.1, pp. 46–89.
- Golub, G.H & C.F Van Loan (2012). *Matrix computations*. Vol. 3. JHU Press.
- Hennig, P. (2015). “Probabilistic Interpretation of Linear Solvers”. In: *SIAM Journal on Optimization* 25.1, pp. 234–260.
- Hennig, P. & M. Kiefel (2013). “Quasi-Newton method: A new direction”. In: *Journal of Machine Learning Research* 14.Mar, pp. 843–865.
- Hestenes, M.R. & E. Stiefel (1952). “Methods of conjugate gradients for solving linear systems”. In: *Journal of Research of the National Bureau of Standards* 49.6, pp. 409–436.
- Higham, N.J. (1988). “Computing a nearest symmetric positive semidefinite matrix”. In: *Linear Algebra and its Applications* 103, pp. 103–118.
- Kingma, D.P. & J. Ba (2014). “Adam: A method for stochastic optimization”. In: *arXiv preprint 1412.6980*.
- Krizhevsky, A. (2009). *Learning multiple layers of features from tiny images*. Tech. rep. Citeseer.
- Martens, J. (2010). “Deep learning via Hessian-free optimization”. In: *International Conference on Machine Learning (ICML)*.
- Nocedal, J. & S.J. Wright (2006). *Numerical Optimization*. second. New York, NY, USA: Springer.
- Paszke, A. et al. (2017). “Automatic differentiation in PyTorch”. In: *NIPS-Workshops*.
- Pearlmutter, B.A. (1994). “Fast exact multiplication by the Hessian”. In: *Neural computation* 6.1, pp. 147–160.
- Polyak, B. T. (1964). “Some methods of speeding up the convergence of iteration methods”. In: *USSR Computational Mathematics and Mathematical Physics* 4.5, pp. 1–17.
- Rasmussen, C.E. & C.K.I. Williams (2005). *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. The MIT Press.
- Robbins, H. & S. Monro (1951). “A stochastic approximation method”. In: *The Annals of Mathematical Statistics*, pp. 400–407.
- Saad, Y. & M.H. Schultz (1986). “GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems”. In: *SIAM Journal on scientific and statistical computing* 7.3, pp. 856–869.
- Schraudolph, N.N, J. Yu, & S. Günter (2007). “A stochastic quasi-Newton method for online convex optimization”. In: *Artificial Intelligence and Statistics*, pp. 436–443.
- Trefethen, L.N. & D. Bau III (1997). *Numerical Linear Algebra*. SIAM.
- van der Vaart, A.W. (1998). *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Van Loan, C.F. (2000). “The ubiquitous Kronecker product”. In: *Journal of Computational and Applied Mathematics* 123, pp. 85–100.
- Vijayakumar, S. & S. Schaal (2000). “Locally Weighted Projection Regression: Incremental Real Time Learning in High Dimensional Space”. In: *Proceedings of the Seventeenth International Conference on Machine Learning*, pp. 1079–1086.
- Wills, A. & T. Schön (2018). “Stochastic quasi-Newton with adaptive step lengths for large-scale problems”. In: *ArXiv e-prints*. arXiv: 1802.04310.