



Approximating Hessian matrices using Bayesian inference: a new approach for quasi-Newton methods in stochastic optimization

André Gustavo Carlon, Luis Espath & Raúl Tempone

To cite this article: André Gustavo Carlon, Luis Espath & Raúl Tempone (29 Apr 2024): Approximating Hessian matrices using Bayesian inference: a new approach for quasi-Newton methods in stochastic optimization, Optimization Methods and Software, DOI: [10.1080/10556788.2024.2339226](https://doi.org/10.1080/10556788.2024.2339226)

To link to this article: <https://doi.org/10.1080/10556788.2024.2339226>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 29 Apr 2024.



Submit your article to this journal [↗](#)



Article views: 159



View related articles [↗](#)



View Crossmark data [↗](#)

Approximating Hessian matrices using Bayesian inference: a new approach for quasi-Newton methods in stochastic optimization

André Gustavo Carlon^a, Luis Espath^b and Raúl Tempone^{a,c,d}

^aComputer, Electrical and Mathematical Sciences & Engineering Division (CEMSE), King Abdullah University of Science & Technology (KAUST), Thuwal, Saudi Arabia; ^bSchool of Mathematical Sciences, University of Nottingham, Nottingham, UK; ^cDepartment of Mathematics, RWTH Aachen University, Aachen, Germany; ^dAlexander von Humboldt Professor in Mathematics for Uncertainty Quantification, RWTH Aachen University, Aachen, Germany

ABSTRACT

Using quasi-Newton methods in stochastic optimization is not a trivial task given the difficulty of extracting curvature information from the noisy gradients. Moreover, pre-conditioning noisy gradient observations tend to amplify the noise. We propose a Bayesian approach to obtain a Hessian matrix approximation for stochastic optimization that minimizes the secant equations residue while retaining the extreme eigenvalues between a specified range. Thus, the proposed approach assists stochastic gradient descent to converge to local minima without augmenting gradient noise. We propose maximizing the log posterior using the Newton-CG method. Numerical results on a stochastic quadratic function and an ℓ_2 -regularized logistic regression problem are presented. In all the cases tested, our approach improves the convergence of stochastic gradient descent, compensating for the overhead of solving the log posterior maximization. In particular, pre-conditioning the stochastic gradient with the inverse of our Hessian approximation becomes more advantageous the larger the condition number of the problem is.

ARTICLE HISTORY

Received 8 September 2022
Accepted 13 March 2024

KEYWORDS

Stochastic optimization; quasi-Newton; Monte Carlo; variance reduction; control variates; machine learning

AMS SUBJECT CLASSIFICATIONS

65K10; 90C15; 62F15; 90C53

1. Introduction

First-order optimization methods are a common choice for performing local search when gradients are available. In the strongly convex case, the gradient descent method converges linearly, requiring a number of iterations proportional to the condition number [21, Theorem 2.1.15], that is, the ratio between the upper and lower bounds of the eigenvalues of the Hessian matrix in the search domain. However, first-order methods may not be a practical choice when the condition number for a problem is large, advancing little towards the optimum each iteration. One way to mitigate this dependency is to use quasi-Newton methods, which use the secant equation to obtain an approximation of the Hessian matrix

CONTACT André Gustavo Carlon  agcarlon@gmail.com  Computer, Electrical and Mathematical Sciences & Engineering Division (CEMSE), King Abdullah University of Science & Technology (KAUST), Thuwal 23955-6900, Saudi Arabia

© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

of the objective function from subsequent gradient evaluations. Then, one may precondition the gradient with the inverse of the Hessian approximation, furnishing a Newton-like method without the cost of computing the true Hessian or using a finite differences type of approximation [10].

In the case of stochastic optimization, the stochastic gradient descent method (SGD) is the canonical first-order method. The asymptotic convergence of SGD is dominated by noise in gradient evaluations (cf. [24, Theorem 3]), thus controlling the condition number for the problem does not improve convergence. However, controlling the relative statistical error of gradient estimates using a uniform bound recovers linear convergence per iteration for the minimization of strongly convex functions [8,24]. In this case, the convergence rate depends on the condition number of the problem; hence preconditioning the gradient with the Hessian inverse might be desirable. The natural approach is to use quasi-Newton methods to approximate the Hessian matrix, but devising quasi-Newton methods for stochastic optimization is challenging due to its noisy gradient observations. Naively fitting the secant equations on noisy gradient observations often leads to catastrophic errors in Hessian approximations [24, Section 4.3.3].

There are many works addressing the development of quasi-Newton methods in stochastic optimization. One of the first attempts is that of Schraudolph, Yu, and Günther [26] where they adapt the L-BFGS method to the stochastic setting by adding damping parameters that control the effect of noise, furnishing the oL-BFGS method. Their approach is empirical, and not much can be said about the quality of the Hessian approximation. Bordes and Pierre [3] propose a diagonal rescaling matrix for linear support vector machines based on oL-BFGS, providing an analysis of their approach. In their numerical tests, their method matched oL-BFGS in convergence per epoch but with a significant improvement in terms of runtimes. Hennig [15] uses a Bayesian approach to learn about an objective function from noisy gradient evaluations. Byrd et al. [6] use a sample average approximation to compute the search direction in Newton-CG and L-BFGS. As a consequence, they do not have a statistical error on their estimates but a bias. Hennig [15] models the prior distribution of the objective function as a Gaussian process with a squared exponential kernel and, using the linearity of Gaussian processes, derives prior distributions for the gradient and the Hessian of the objective function. Then, assuming Gaussian likelihoods for both the gradient and the Hessian of the objective function, Hennig [15] deduces closed-form equations for the mean of the (also Gaussian) posterior distribution. Sohl-Dickstein, Poole, and Ganguli [28] develop a quasi-Newton method for sums of functions where they keep a quadratic approximation for each of these functions, subsampling them at each iteration and updating the respective quadratic approximation with new gradient information. They use a shared, low-dimensional subspace to avoid the cost of storing a Hessian matrix for each function. Byrd et al. [7] propose a stochastic quasi-Newton method comprising a two-loop scheme where the outer loop updates the curvature pairs in an L-BFGS scheme (with large sample sizes), and the inner loop performs preconditioned stochastic gradient descent (with lower sample sizes). Thus, the method proposed by Byrd et al. [7] does not incorporate the information from the noisy gradients of the inner loops in the Hessian approximation. Moritz, Nishihara, and Jordan [20] improve the stochastic quasi-Newton method proposed by Byrd et al. [7] by combining it with a control variate method to estimate the gradient at each iteration, achieving linear convergence.

Gower, Goldfarb, and Richtárik [13] introduce a block BFGS method to update the Hessian approximation each iteration using a sample independent of the sample used to compute the gradient. The proposed block update fits the Hessian inverse to a sketched Hessian computed from a sub-sample of the data, but requires a Cholesky decomposition each iteration. Wang et al. [30] propose a damped quasi-Newton method for stochastic optimization that uses all noisy gradient evaluations to approximate the Hessian inverse. They provide a convergence proof for nonconvex problems, provided the Hessian eigenvalues have uniform lower bound. Li [18] preconditions the stochastic gradient estimates using a matrix computed to fit the secant equation while avoiding amplifying gradient estimate noise. Bollapragada et al. [2] use standard L-BFGS equations in stochastic optimization by increasing sample sizes to keep the relative error for the gradients uniformly bounded. Bollapragada et al. [2] recommend a stochastic line-search procedure to be used with their method. Goldfarb, Ren, and Bahamou [12] develop a quasi-Newton method for deep neural networks that approximates Hessian inverses as block matrices such that each block is decomposed as the kronecker products of smaller matrices. Lately, Shi et al. [27] build on top of their analysis in [31] to build a quasi-Newton method for stochastic optimization with a noise-resilient Wolfe–Armijo line search. Their method, however, requires the gradient estimate to be bounded above in norm by a known constant, which they use in their line search. We refer the interested reader to [4,19] for extensive reviews of quasi-Newton methods for stochastic optimization.

Most approaches to tackle the condition number issue in stochastic optimization use equations from deterministic quasi-Newton methods and somehow circumvent the randomness. Here, we propose a Bayesian framework to build a posterior distribution for the Hessian matrix from noisy gradient observations. In Bayesian inference, a posterior distribution is built from prior and likelihood distributions. We model the prior distribution for the Hessian matrix with a probability density function (pdf) that decreases exponentially as the Hessian moves farther away (in the Frobenius norm sense) from the initial Hessian estimate. This approach is motivated by classic deterministic quasi-Newton methods, where a Frobenius norm regularizer centred at the current Hessian is used to circumvent ill-posedness [10]. Moreover, we impose constraints on the Hessian matrix eigenvalues extremes, enforcing positive-definiteness and gradient smoothness. We model the likelihood using the secant equation, i.e. how likely it is to obtain the observed gradient differences for a given Hessian. The Hessian that maximizes the posterior distribution, i.e. the Maximum a Posteriori (MAP), is defined as the Hessian approximation.

The choice of prior and likelihood distributions guarantees that the negative log posterior is strongly-convex. Therefore, we opt to use the Newton-CG method to find the MAP. Although it is possible to compute the second-order derivatives of the negative log posterior of the Hessian directly, it could be expensive and memory intensive. Keep in mind that, for a d -dimensional problem, the Hessian matrix is a $d \times d$ matrix, thus, the second-order derivatives of the log posterior must be gathered in a fourth-order tensor with d^4 components. Instead of assembling this tensor, we compute its action over any symmetric matrix of adequate dimensions and then use the conjugate gradient method to find the Newton direction. Moreover, to enforce the stability of Newton-CG, we use the central-path interior point method to impose the prior conditions on the eigenvalues of the Hessian approximation, i.e. we start with a large penalization parameter and decrease it in steps, keeping the problem well-conditioned. If properly tuned, our approach not only guarantees that

Newton-CG converges to the MAP as it also results in a slow decay for the Hessian approximation's smallest eigenvalue, thus improving the stability for the SGD preconditioned with our Hessian approximation.

Other works address the problem of approximating a Hessian from noisy gradient observations. Hennig [15] also uses Bayesian inference; however, to build a posterior distribution of the objective function as a Gaussian process. The main difference between the proposed approach and that of Hennig [15] is in the way he models the prior and likelihood distributions. Moreover, Hennig [15] assumes independence between the components of the gradient noise in his analysis, and also that they have the same variance. In the same spirit of the present work, Li [18] also solves an optimization sub-problem each iteration to find a Hessian approximation from noisy gradients. His approach differs from ours in the way he uses the available information to estimate the preconditioning matrix; here, we use a Bayesian approach to build a posterior distribution of the Hessian matrix while [18] obtains a point estimate by minimizing a different objective function.

To illustrate the effectiveness of our approach, we solve two stochastic optimization problems. The first is minimizing a quadratic function, and the second is training an ℓ_2 regularized logistic regression model. In all cases, we start with SGD preconditioned with a diagonal matrix and update the preconditioning matrix in intervals using the proposed Bayesian approach. We study cases with up to 300 design variables and condition numbers of orders up to 10^9 . We present results combining the proposed Bayesian Hessian approach with different variations of SGD, namely, vanilla SGD, SGD-MICE [8], SVRG [16], and SARAH [22].

1.1. Problem definition

The stochastic optimization problem is stated as follows. Let ξ be the design variable in dimension d_ξ and θ be a vector-valued random variable in dimension d_θ equipped with a probability density function $\varrho: \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}^+$. Here, $\mathbb{E}[\cdot|\xi]$ and $\mathbb{V}[\cdot|\xi]$ are the expectation and variance operators conditioned on ξ , respectively. Aiming at minimizing the expectations conditioned on ξ , we state the problem as

$$\xi^* = \arg \min_{\xi \in \Xi} \mathbb{E}[f(\xi, \theta)|\xi], \quad (1)$$

where $\Xi \subset \mathbb{R}^{d_\xi}$ is a feasible set and $f: \Xi \times \mathbb{R}^{d_\theta} \rightarrow \mathbb{R}$. Through what follows, let $F(\xi) := \mathbb{E}[f(\xi, \theta)|\xi]$ where $F: \Xi \rightarrow \mathbb{R}$.

In minimizing (1) with respect to the design variable $\xi \in \Xi$, vanilla SGD is constructed with the update rule

$$\xi_{k+1} = \xi_k - \eta_k \mathbf{v}_k, \quad (2)$$

where η_k is the step size at iteration k , and \mathbf{v}_k is an unbiased estimator of the gradient of F at ξ_k . For instance, an unbiased estimator \mathbf{v}_k of the gradient of F at ξ_k may be constructed using a Monte Carlo estimator,

$$\nabla_\xi F(\xi_k) = \mathbb{E}[f(\xi_k, \theta)|\xi_k] \approx \mathbf{v}_k := \frac{1}{b} \sum_{i=1}^b \nabla_\xi f(\xi_k, \theta_i), \quad (3)$$

with b independent and identically distributed (iid) random variables $\theta_i \sim \varrho$. The estimator (3) is a random variable and its use in optimization algorithms gives rise to the so-called *Stochastic Optimizers*.

First-order methods are the most popular choice in stochastic optimization due to their robustness to noise. However, these methods might suffer from slow convergence when the condition number of the problem is large. Here, we are interested in finding a matrix $\hat{\mathbf{B}}_k$ that approximates $\nabla^2 F(\boldsymbol{\xi}_k)$ so that its inverse can be used as a preconditioner for the stochastic gradient,

$$\boldsymbol{\xi}_{k+1} = \boldsymbol{\xi}_k - \eta_k \hat{\mathbf{B}}_k^{-1} \mathbf{v}_k, \quad (4)$$

thus diminishing the effect of the condition number on the optimization convergence. In the deterministic setting, quasi-Newton methods circumvent the sensibility to the condition number by using the gradient observations to construct the Hessian matrix approximation at the current iterate. However, in the stochastic optimization setting, we first have to address the following questions:

- Q1:** How does one obtain a Hessian approximation from noisy gradient observations?
- Q2:** How do we avoid amplifying the noise in gradient estimates with our preconditioner matrix $\hat{\mathbf{B}}_k^{-1}$?
- Q3:** To mitigate the effect of the condition number, is it enough to have an accurate Hessian approximation, however, acting on noisy gradient estimates?

To address **Q1**, consider the curvature pair $\mathbf{y}_k := \nabla_{\boldsymbol{\xi}} F(\boldsymbol{\xi}_k) - \nabla_{\boldsymbol{\xi}} F(\boldsymbol{\xi}_{k-1})$ and $\mathbf{s}_k := \boldsymbol{\xi}_k - \boldsymbol{\xi}_{k-1}$. Thus, we first aim to find the Hessian approximation that minimizes the residue of the secant equation, that is,

$$\min_{\mathbf{B} \in \mathbb{R}^{d_{\boldsymbol{\xi}}} \times \mathbb{R}^{d_{\boldsymbol{\xi}}}} \sum_{\ell=1}^k \|\mathbf{B} \mathbf{s}_{\ell} - \mathbf{y}_{\ell}\|^2. \quad (5)$$

However, we do not have access to $\mathbf{y}_k := \mathbb{E}[\nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}_k, \boldsymbol{\theta}) | \boldsymbol{\xi}_k] - \mathbb{E}[\nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}_{k-1}, \boldsymbol{\theta}) | \boldsymbol{\xi}_{k-1}]$, instead, we may sample finitely many times $\hat{\mathbf{y}}_k(\boldsymbol{\theta}) := \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}_k, \boldsymbol{\theta}) - \nabla_{\boldsymbol{\xi}} f(\boldsymbol{\xi}_{k-1}, \boldsymbol{\theta})$. In Section 2, we discuss how to use Bayesian inference to find a matrix $\hat{\mathbf{B}}_k$ given the noisy curvature pairs $\hat{\mathbf{y}}$ and \mathbf{s} .

One of the main problems faced by preconditioning in stochastic optimization is the amplification of the noise in gradient estimates. This amplification is proportional to the smallest eigenvalue of matrix $\hat{\mathbf{B}}_k$,

$$\mathbb{E} \left[\left\| \hat{\mathbf{B}}_k^{-1} \mathbf{v}_k - \mathbb{E} \left[\hat{\mathbf{B}}_k^{-1} \mathbf{v}_k \mid \hat{\mathbf{B}}_k^{-1} \right] \right\|^2 \mid \hat{\mathbf{B}}_k^{-1} \right] \leq \left\| \hat{\mathbf{B}}_k^{-1} \right\|^2 \mathbb{E} \left[\|\mathbf{v}_k - \mathbb{E}[\mathbf{v}_k]\|^2 \mid \hat{\mathbf{B}}_k^{-1} \right], \quad (6)$$

by the matrix-norm compatibility. Thus, to tackle **Q2**, we impose a lower bound constraint on the eigenvalues of $\hat{\mathbf{B}}_k$, i.e. $\hat{\mathbf{B}}_k \succeq \tilde{\mu} \mathbf{I}$.

To address **Q3** in a numerical fashion, we investigate and compare stochastic gradients preconditioned with our Hessian inverse approximation with and without control in the relative statistical error.

Regarding notation, in this work, we use bold-face lower-case letters as vectors and bold-face upper-case letters as matrices. Moreover, we refer to first-order derivatives of functions with respect to matrices as gradients, even though they are also matrices.

2. Approximating the Hessian and its inverse using Bayesian inference

Important to what follows is the Monte Carlo estimator of \mathbf{y}_k

$$\bar{\mathbf{y}}_k := \frac{1}{b_k} \sum_{i=1}^{b_k} \hat{\mathbf{y}}_k(\boldsymbol{\theta}_i), \quad (7)$$

with b_k iid random variables $\boldsymbol{\theta}_i \sim \varrho$.

2.1. Bayesian formulation of the Hessian inference problem

For the pair of sets $Y := \{\bar{\mathbf{y}}_\ell\}_{\ell=1}^k$ and $S := \{\mathbf{s}_\ell\}_{\ell=1}^k$, we employ the Bayes' formula to build a posterior distribution of the Hessian \mathbf{B} given Y and S ,

$$\pi(\mathbf{B}|S, Y) \propto p(S, Y|\mathbf{B})\pi(\mathbf{B}), \quad (8)$$

where the likelihood distribution $p(S, Y|\mathbf{B})$ and the prior distribution $\pi(\mathbf{B})$ will be discussed in Sections 2.1.1 and 2.1.2, respectively. We opt to use as a Hessian approximation the matrix that minimizes the negative log posterior subject to the eigenvalues constraints

$$\mathbf{B}_k^* := \begin{cases} \arg \min_{\mathbf{B}} & (\mathcal{L}_k(\mathbf{B}) := -\log(\pi(\mathbf{B}|S, Y))) \\ \text{subject to} & \lambda_{\min}(\mathbf{B}) \geq \tilde{\mu} \\ & \lambda_{\max}(\mathbf{B}) \leq L \\ & \mathbf{B} = \mathbf{B}^\top, \end{cases} \quad (9)$$

where λ_{\min} and λ_{\max} are operators that return the minimum and maximum eigenvalues, respectively; L is the smoothness constant of F ; and $\tilde{\mu} > 0$ is an arbitrary parameter. If F is μ -convex, $\tilde{\mu}$ can be chosen as μ , however, a larger $\tilde{\mu}$ might be desirable to avoid amplifying the noise in the search direction. Since computing \mathbf{B}_k^* is not possible in general, we use an approximation $\hat{\mathbf{B}}_k$, obtained numerically, as a Hessian approximation.

2.1.1. Likelihood of observing curvature pairs given a Hessian

We use the secant equations to derive the likelihood of observing the curvature pairs (S, Y) given a Hessian \mathbf{B} . Let the difference in gradients be as in (7). From the observations $\hat{\mathbf{y}}_k$, we build a sample covariance matrix $\boldsymbol{\Sigma}_k$. Then, the covariance matrix of the estimator $\bar{\mathbf{y}}_k$ is given by

$$\bar{\boldsymbol{\Sigma}}_k := \text{Cov}[\bar{\mathbf{y}}_k] = \frac{\boldsymbol{\Sigma}_k}{b_k}. \quad (10)$$

We are interested in the inverse of $\bar{\boldsymbol{\Sigma}}_k$. However, to avoid singularity issues, we use a precision matrix approximation \mathbf{P}_ℓ instead of $\bar{\boldsymbol{\Sigma}}_\ell^{-1}$ for each $\bar{\mathbf{y}}_\ell \in Y$. We define \mathbf{P}_ℓ as

$$\mathbf{P}_\ell := (\bar{\boldsymbol{\Sigma}}_\ell + \sigma_P(\max(\bar{\boldsymbol{\Sigma}}_\ell))\mathbf{I}_{d_\xi})^{-1}, \quad (11)$$

where σ_P is a regularization parameter and \mathbf{I}_{d_ξ} is the identity matrix of size d_ξ . Then, we define the negative log likelihood of curvature pair \mathbf{s}_k and $\bar{\mathbf{y}}_k$ given a matrix \mathbf{B} as

$$-\log(p(\mathbf{s}_k, \bar{\mathbf{y}}_k | \mathbf{B})) = \frac{1}{2\nu} \|\mathbf{B}\mathbf{s}_k - \bar{\mathbf{y}}_k\|_{\mathbf{P}_k}^2 + C_k, \quad (12)$$

thus, for S and Y generated by a first-order stochastic gradient descent method,

$$-\log(p(S, Y | \mathbf{B})) = \frac{1}{2\nu} \sum_{\ell=1}^k \|\mathbf{B}\mathbf{s}_\ell - \bar{\mathbf{y}}_\ell\|_{\mathbf{P}_\ell}^2 + C. \quad (13)$$

In the above equations, $C = \sum_{i=\ell}^k C_\ell$ are constants and ν is used to normalize the negative log likelihood,

$$\nu = \sum_{\ell=1}^k \|\mathbf{P}_\ell(\hat{\mathbf{B}}_{k-1}\mathbf{s}_\ell - \bar{\mathbf{y}}_\ell)\| \|\mathbf{s}_\ell\|, \quad (14)$$

being $\hat{\mathbf{B}}_{k-1}$ a Hessian approximation at iteration $k-1$.

2.1.2. Prior distribution of the Hessian matrix

The prior distribution must encode the knowledge available about the Hessian before the observations of S and Y . In spirit of quasi-Newton methods, we want to discourage a Hessian that differs much from the current estimate. For this reason, we adopt a negative log prior of \mathbf{B} as the squared weighted Frobenius norm of the difference between \mathbf{B} and $\hat{\mathbf{B}}_{k-1}$,

$$\frac{1}{2} \|\mathbf{B} - \hat{\mathbf{B}}_{k-1}\|_{F, \mathbf{W}}^2 = \frac{1}{2} \|\mathbf{W}^{1/2}(\mathbf{B} - \hat{\mathbf{B}}_{k-1})\mathbf{W}^{1/2}\|_{F}^2, \quad (15)$$

for a nonsingular symmetric matrix of weights \mathbf{W} . Another information available about the matrix \mathbf{B} is that its eigenvalues must lay between $\tilde{\mu}$ and L , as described in the constraints of (9). To enforce these eigenvalue constraints on \mathbf{B} , we define the prior distribution from its negative logarithm as

$$-\log(\pi(\mathbf{B})) = \begin{cases} \frac{1}{2} \|\mathbf{B} - \hat{\mathbf{B}}_{k-1}\|_{F, \mathbf{W}}^2 + C & \text{if } \hat{\mu}\mathbf{I}_{d_\xi} \preceq \mathbf{B} \preceq \hat{L}\mathbf{I}_{d_\xi} \\ +\infty & \text{otherwise,} \end{cases} \quad (16)$$

where again C is a normalizing constant. Here, we opt to impose the eigenvalues constraints on the negative log prior as logarithmic barriers of form

$$-\log(\det(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})) - \log(\det(\hat{L}\mathbf{I}_{d_\xi} - \mathbf{B})), \quad (17)$$

where

$$\hat{\mu} := \frac{\tilde{\mu}}{\alpha}, \quad \hat{L} := L\alpha, \quad (18)$$

and $\alpha \geq 1$ is a relaxation parameter. Hence, we propose a family of priors with logarithmic barriers with hyperparameter β as

$$\begin{aligned} -\log(\pi(\mathbf{B}; \beta)) &= \frac{\rho}{2} \|\mathbf{B} - \hat{\mathbf{B}}_{k-1}\|_{F, \mathbf{W}}^2 - \beta \log(\det(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_\xi})) \\ &\quad - \beta \log\left(\det\left(\hat{\mathbf{L}} \mathbf{I}_{d_\xi} - \mathbf{B}\right)\right) + C, \end{aligned} \quad (19)$$

being supported on the set of symmetric positive-definite matrices with eigenvalues between $\hat{\mu}$ and $\hat{\mathbf{L}}$.

The parameter $\rho > 0$ weights the Frobenius regularization term and $\beta > 0$ weights the logarithmic barrier terms. Note that $\log(\pi(\mathbf{B}; \beta)) \rightarrow \log(\pi(\mathbf{B}))$ as $\beta \downarrow 0$ with $\rho = 1$.

Some options for \mathbf{W} are \mathbf{I}_{d_ξ} and $\hat{\mathbf{B}}_{k-1}^{-1/2}$. According to [10], the Frobenius norm as in (15) with weights $\hat{\mathbf{B}}_{k-1}^{-1/2}$ measures the relative error of \mathbf{B} with respect to $\hat{\mathbf{B}}_{k-1}$. In practice, estimating $\hat{\mathbf{B}}_{k-1}^{-1/2}$ can be expensive. In these cases, we recommend using $\mathbf{W} = \mathbf{I}_{d_\xi}$.

2.1.3. Posterior distribution of the Hessian matrix

Substituting (13) and (19) in (8), we obtain the posterior distribution of the Hessian matrix \mathbf{B} given curvature pairs \mathbf{S} and \mathbf{Y} . The negative log posterior with parameter β is

$$\begin{aligned} \mathcal{L}_k(\mathbf{B}; \beta) &:= \frac{1}{2\nu} \sum_{\ell=1}^k \|\mathbf{B} \mathbf{s}_\ell - \bar{\mathbf{y}}_\ell\|_{\mathbf{P}_\ell}^2 + \frac{\rho}{2} \|\mathbf{B} - \hat{\mathbf{B}}_{k-1}\|_{F, \mathbf{W}}^2 - \beta \log(\det(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_\xi})) \\ &\quad - \beta \log\left(\det\left(\hat{\mathbf{L}} \mathbf{I}_{d_\xi} - \mathbf{B}\right)\right) + C. \end{aligned} \quad (20)$$

The Equation (20) can be interpreted as the sum of the secant equation residues for each curvature pair (in the metrics induced by \mathbf{P}_ℓ), subject to the logarithmic barrier terms that impose the eigenvalues constraints, and a Frobenius norm regularization term.

Since we are interested in minimizing the negative log posterior in (20), it is our interest to know if it is strongly convex. The negative log-likelihood is convex but not strictly-convex until enough linearly independent curvature pairs $(\mathbf{s}_\ell, \mathbf{y}_\ell)$ are observed. Regarding the logarithmic barrier terms of the prior, the negative logarithm of the determinant defined on symmetric positive-definite matrices is a convex function [5, Section 3.1.5]; however, not a strictly-convex one. The Frobenius norm regularization term with $\mathbf{W} = \mathbf{I}_{d_\xi}$ is ρ -convex, therefore the negative log posterior is strongly convex with parameter ρ . Here, the regularization term works as a Tikhonov regularization term ($\|\mathbf{A}\|_F \equiv \|\text{vec}(\mathbf{A})\|_2$), guaranteeing uniqueness of solution. Properly tuning β and ρ guarantees that the problem of minimizing (20) is well-conditioned. A discussion on how to choose β is presented in Section 2.2.

2.2. Finding the posterior maximizer

Finding the \mathbf{B}_k^* that solves the problem in (9) is not a trivial task. Since the negative log posterior is ρ -convex, local search methods like the Newton method can be used to find its minimizer. However, the Newton method requires assembling the fourth-order tensor of second-order derivatives of the negative log posterior, which can be expensive in terms of

processing and memory allocation. Instead, we opt to derive the action of this fourth-order tensor of second-order derivatives over any given matrix and use the conjugate gradient method to find the Newton direction, a procedure known as the Newton-CG method. Since we have a constraint in (9) that \mathbf{B} should be symmetric, the first-order optimality condition is

$$\text{sym}(\nabla_{\mathbf{B}}\mathcal{L}_k(\mathbf{B};\beta)) = 0, \quad (21)$$

where the $\text{sym}(\cdot)$ operator returns the symmetric part of its argument, e.g. $\text{sym}(\mathbf{A}) = \frac{1}{2}(\mathbf{A} + \mathbf{A}^T)$. Thus, to use the Newton-CG method to solve this problem, every iteration we need to find the direction \mathbf{D} that satisfies

$$\text{sym}(\delta_{\mathbf{D}}\nabla_{\mathbf{B}}\mathcal{L}_k(\mathbf{B};\beta)) = -\text{sym}(\nabla_{\mathbf{B}}\mathcal{L}_k(\mathbf{B};\beta)), \quad (22)$$

where $\delta_{\mathbf{D}}$ is the directional derivative with respect to \mathbf{B} in the direction of \mathbf{D} ,

$$\delta_{\mathbf{D}}g(\mathbf{B}) = \lim_{\vartheta \downarrow 0} \frac{g(\mathbf{B} + \vartheta\mathbf{D}) - g(\mathbf{B})}{\vartheta}. \quad (23)$$

To avoid ill-conditioning, we need a large enough regularization parameter ρ , but not too large since it introduces bias in the posterior maximizer. Also, the bias given by the log-barriers grows with β , the logarithmic barrier parameter; but a small β negatively impacts the smoothness of the posterior distribution. We opt to keep ρ fixed and use an interior-point central-path method [5, Section 11.2.2] to decrease the parameter β ; the solution $\hat{\mathbf{B}}_k(\beta)$ of the approximated problem

$$\hat{\mathbf{B}}_k(\beta) = \arg \min_{\mathbf{B}} (\mathcal{L}_k(\mathbf{B};\beta)) \quad (24)$$

converges to the solution \mathbf{B}_k^* of (9) as $\beta \downarrow 0$. Thus, we start with a large β and a large tolerance tol and perform optimization until the norm of the gradient of the log posterior is below tol . Then, we decrease both β and tol by a factor γ and repeat the process. This procedure is repeated until the desired tolerance is achieved with the desired β . To minimize the negative log posterior for a given β , we use the Newton-CG method with backtracking-Armijo line search; we use the conjugate gradient method to find the Newton direction and then the backtracking-Armijo to find a step size. Therefore, to use the Newton-CG, we need the gradient of \mathcal{L}_k with respect to \mathbf{B} and the directional derivative of the gradient of \mathcal{L}_k with respect to \mathbf{B} in the direction of any arbitrary symmetric matrix $\mathbf{V} \in \mathbb{R}^{d_{\xi}} \times \mathbb{R}^{d_{\xi}}$. The pseudo-code for the minimization of the negative log posterior is presented in Algorithm 1. We use the central-path method to decrease β and tol every step outer loop, indexed as i in Algorithm 1. In lines 6 to 10, we have the Newton-CG method, where, each iteration, the conjugate gradient method is used to find the approximate Newton direction \mathbf{D} and a backtracking-Armijo line-search is used to find the step size.

Next, we will present how to compute $\nabla_{\mathbf{B}}\mathcal{L}_k(\hat{\mathbf{B}}_k; \beta)$ and $\delta_{\mathbf{V}}\nabla_{\mathbf{B}}\mathcal{L}_k(\hat{\mathbf{B}}_k; \beta)$ for any symmetric $\mathbf{V} \in \mathbb{R}^{d_{\xi}} \times \mathbb{R}^{d_{\xi}}$. Note that, since \mathbf{B} is a matrix, $\nabla_{\mathbf{B}}\mathcal{L}_k$ and $\delta_{\mathbf{V}}\nabla_{\mathbf{B}}\mathcal{L}_k$ are also symmetric matrices of the same dimensions of \mathbf{B} .

Algorithm 1 Pseudocode of the procedure to find the maximum a posteriori of the Hessian matrix of the objective function given noisy curvature pair observations.

```

1: procedure FINDMAP( $S, Y, \hat{\mathbf{B}}_{k-1}, \beta_0, tol_0, \gamma = 2, c_A = 10^{-4}$ )
2:    $\hat{\mathbf{B}}_k \leftarrow \hat{\mathbf{B}}_{k-1}$ 
3:   for  $i = 1, 2, \dots$  do
4:      $\beta_i = \beta_{i-1}/\gamma$ 
5:      $tol_i = tol_{i-1}/\gamma$ 
6:     while  $\|\nabla_{\mathbf{B}} \mathcal{L}_k(\hat{\mathbf{B}}_k; \beta_i)\|_F > tol_i$  do
7:       use the conjugate gradient method to find a symmetric  $\mathbf{D}$  such that
8:          $\text{sym}(\delta_{\mathbf{D}} \nabla_{\mathbf{B}} \mathcal{L}_k(\hat{\mathbf{B}}_k; \beta_i)) = -\text{sym}(\nabla_{\mathbf{B}} \mathcal{L}_k(\hat{\mathbf{B}}_k; \beta_i))$ 
9:       find  $\varsigma$  such that  $\mathcal{L}_k(\hat{\mathbf{B}}_k + \varsigma \mathbf{D}; \beta_i) \leq \mathcal{L}_k(\hat{\mathbf{B}}_k; \beta_i) + c_A \varsigma \nabla_{\mathbf{B}} \mathcal{L}_k(\hat{\mathbf{B}}_k; \beta_i)$ :  $\mathbf{D}$ 
       using backtracking-Armijo
10:       $\hat{\mathbf{B}}_k \leftarrow \hat{\mathbf{B}}_k + \varsigma \mathbf{D}$ 
11:    end while
12:  end for
13:  return  $\hat{\mathbf{B}}_k$ 
14: end procedure

```

2.3. Gradient of the negative log posterior

To compute $\nabla_{\mathbf{B}} \mathcal{L}_k$, we need to compute the gradients of the log prior and of the log likelihood,

$$\nabla_{\mathbf{B}} \mathcal{L}_k(\mathbf{B}; \beta) = -\nabla_{\mathbf{B}} \log(p(S, Y|\mathbf{B})) - \nabla_{\mathbf{B}} \log(\pi(\mathbf{B}; \beta)). \quad (25)$$

The gradient of the negative log likelihood can be calculated from (13) as

$$-\nabla_{\mathbf{B}} \log(p(S, Y|\mathbf{B})) = \frac{1}{2\nu} \sum_{\ell=1}^k \nabla_{\mathbf{B}} \|\mathbf{B} \mathbf{s}_{\ell} - \bar{\mathbf{y}}_{\ell}\|_{\mathbf{P}_{\ell}}^2 \quad (26)$$

$$= \frac{1}{\nu} \sum_{\ell=1}^k \mathbf{P}_{\ell} (\mathbf{B} \mathbf{s}_{\ell} - \bar{\mathbf{y}}_{\ell}) \otimes \mathbf{s}_{\ell}. \quad (27)$$

The gradient of the negative log prior can be derived from (19) as

$$\begin{aligned} -\nabla_{\mathbf{B}} \log(\pi(\mathbf{B}; \beta)) &= \frac{\rho}{2} \nabla_{\mathbf{B}} \|\mathbf{W}^{\frac{1}{2}} (\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W}^{\frac{1}{2}}\|_F^2 - \beta \nabla_{\mathbf{B}} \log(\det(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_z})) \\ &\quad - \beta \nabla_{\mathbf{B}} \log(\det(\hat{\mathbf{L}} \mathbf{I}_{d_z} - \mathbf{B})). \end{aligned} \quad (28)$$

The gradient of the Frobenius norm regularizer is

$$\frac{1}{2} \nabla_{\mathbf{B}} \|\mathbf{W}^{\frac{1}{2}} (\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W}^{\frac{1}{2}}\|_F^2 = \mathbf{W} (\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W}, \quad (29)$$

whereas the gradients of the logarithmic barrier terms are derived using the Jacobi formula,

$$-\nabla_{\mathbf{B}} \log(\det(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_z})) = -(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_z})^{-1} \quad (30)$$

$$-\nabla_{\mathbf{B}} \log \left(\det \left(\hat{\mathbf{L}}_{d_\xi} - \mathbf{B} \right) \right) = \left(\hat{\mathbf{L}}_{d_\xi} - \mathbf{B} \right)^{-1}. \quad (31)$$

Substituting (29), (30), and (31) in (28),

$$-\nabla_{\mathbf{B}} \log(\pi(\mathbf{B}; \beta)) = \rho \mathbf{W}(\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W} - \beta(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_\xi})^{-1} + \beta(\hat{\mathbf{L}}_{d_\xi} - \mathbf{B})^{-1}. \quad (32)$$

Finally, substituting the gradient of the negative log likelihood (27) and the gradient of the negative log prior (32) in (25), we obtain the gradient of the negative log posterior

$$\begin{aligned} \nabla_{\mathbf{B}} \mathcal{L}_k(\mathbf{B}; \beta) &= \frac{1}{v} \sum_{\ell=1}^k \mathbf{P}_\ell(\mathbf{B} \mathbf{s}_\ell - \bar{\mathbf{y}}_\ell) \otimes \mathbf{s}_\ell + \rho \mathbf{W}(\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W} \\ &\quad - \beta(\mathbf{B} - \hat{\mu} \mathbf{I}_{d_\xi})^{-1} + \beta(\hat{\mathbf{L}}_{d_\xi} - \mathbf{B})^{-1}. \end{aligned} \quad (33)$$

2.4. Directional derivative of the gradient of the negative log posterior

As discussed in Section 2.2, to find the posterior distribution maximizer using Newton-CG, we need to compute the directional derivative of the gradient of the negative log posterior in a given direction. Thus, for any symmetric matrix $\mathbf{V} \in \mathbb{R}^{d_\xi} \times \mathbb{R}^{d_\xi}$, we derive $\delta_{\mathbf{V}} \nabla_{\mathbf{B}} \mathcal{L}_k$ as

$$\delta_{\mathbf{V}} \nabla_{\mathbf{B}} \mathcal{L}_k(\mathbf{B}; \beta) = -\delta_{\mathbf{V}} \nabla_{\mathbf{B}} \log(p(\mathbf{S}, \mathbf{Y}|\mathbf{B})) - \delta_{\mathbf{V}} \nabla_{\mathbf{B}} \log(\pi(\mathbf{B}; \beta)). \quad (34)$$

Following from (27),

$$\begin{aligned} &-\delta_{\mathbf{V}} \nabla_{\mathbf{B}} \log(p(\mathbf{S}, \mathbf{Y}|\mathbf{B})) \\ &= -\lim_{\vartheta \downarrow 0} \frac{\nabla_{\mathbf{B}} \log p(\mathbf{S}, \mathbf{Y}|\mathbf{B} + \vartheta \mathbf{V}) - \nabla_{\mathbf{B}} \log p(\mathbf{S}, \mathbf{Y}|\mathbf{B})}{\vartheta} \end{aligned} \quad (35)$$

$$= \frac{1}{v} \lim_{\vartheta \downarrow 0} \frac{\sum_{\ell=1}^k \mathbf{P}_\ell((\mathbf{B} + \vartheta \mathbf{V}) \mathbf{s}_\ell - \bar{\mathbf{y}}_\ell) \otimes \mathbf{s}_\ell - \sum_{\ell=1}^k \mathbf{P}_\ell(\mathbf{B} \mathbf{s}_\ell - \bar{\mathbf{y}}_\ell) \otimes \mathbf{s}_\ell}{\vartheta} \quad (36)$$

$$= \frac{1}{v} \sum_{\ell=1}^k \mathbf{P}_\ell \mathbf{V} \mathbf{s}_\ell \otimes \mathbf{s}_\ell. \quad (37)$$

The directional derivative of the gradient of the negative log prior is derived from (32),

$$\begin{aligned} -\delta_{\mathbf{V}} \nabla_{\mathbf{B}} \log(\pi(\mathbf{B}; \beta)) &= \rho \delta_{\mathbf{V}} \left(\mathbf{W}(\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W} \right) \\ &\quad - \beta \delta_{\mathbf{V}} (\mathbf{B} - \hat{\mu} \mathbf{I}_{d_\xi})^{-1} + \beta \delta_{\mathbf{V}} (\hat{\mathbf{L}}_{d_\xi} - \mathbf{B})^{-1}. \end{aligned} \quad (38)$$

The Frobenius regularizer term from (38) is derived as

$$\begin{aligned} \delta_{\mathbf{V}} \left(\mathbf{W}(\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W} \right) &= \lim_{\vartheta \downarrow 0} \frac{\mathbf{W}(\mathbf{B} + \vartheta \mathbf{V} - \hat{\mathbf{B}}_{k-1}) \mathbf{W} - \mathbf{W}(\mathbf{B} - \hat{\mathbf{B}}_{k-1}) \mathbf{W}}{\vartheta} \\ &= \mathbf{W} \mathbf{V} \mathbf{W}. \end{aligned} \quad (39)$$

Following the same rationale for the log-barrier terms,

$$\begin{aligned}
& \delta_V(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} \\
&= \lim_{\vartheta \downarrow 0} \frac{((\mathbf{B} + \vartheta\mathbf{V}) - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} - (\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{((\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})(\mathbf{I}_{d_\xi} + \vartheta(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V}))^{-1} - (\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{(\mathbf{I}_{d_\xi} + \vartheta(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V})^{-1}(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} - (\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{(\mathbf{I}_{d_\xi} - \vartheta(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V} + \mathcal{O}(\vartheta^2))(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} - (\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}}{\vartheta} \\
&= -(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V}(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}, \tag{40}
\end{aligned}$$

and,

$$\begin{aligned}
\delta_V(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1} &= \lim_{\vartheta \downarrow 0} \frac{((\hat{\mathbf{L}}\mathbf{I}_d - (\mathbf{B} + \vartheta\mathbf{V}))^{-1} - (\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{((\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})(\mathbf{I}_{d_\xi} - \vartheta(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}\mathbf{V}))^{-1} - (\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{(\mathbf{I}_{d_\xi} - \vartheta(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}\mathbf{V})^{-1}(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1} - (\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}}{\vartheta} \\
&= \lim_{\vartheta \downarrow 0} \frac{(\mathbf{I}_{d_\xi} + \vartheta(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}\mathbf{V} + \mathcal{O}(\vartheta^2))(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1} - (\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}}{\vartheta} \\
&= (\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}\mathbf{V}(\hat{\mathbf{L}}\mathbf{I}_d - \mathbf{B})^{-1}. \tag{41}
\end{aligned}$$

Substituting (39), (40), and (41) back into (38), we get the directional derivative of the gradient of the negative log prior,

$$\begin{aligned}
& -\delta_V\nabla_{\mathbf{B}}\log(\pi(\mathbf{B};\beta)) = \rho\mathbf{W}\mathbf{V}\mathbf{W} + \beta(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V}(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} \\
& \quad + \beta(\hat{\mathbf{L}}\mathbf{I}_{d_\xi} - \mathbf{B})^{-1}\mathbf{V}(\hat{\mathbf{L}}\mathbf{I}_{d_\xi} - \mathbf{B})^{-1}. \tag{42}
\end{aligned}$$

Finally, substituting (37) and (42) into (34), we obtain the directional derivative of the gradient of the negative log posterior,

$$\begin{aligned}
\delta_V\nabla_{\mathbf{B}}\mathcal{L}_k(\mathbf{B};\beta) &= \frac{1}{\nu} \sum_{\ell=1}^k \mathbf{P}_\ell \mathbf{V} \mathbf{s}_\ell \otimes \mathbf{s}_\ell + \rho\mathbf{W}\mathbf{V}\mathbf{W} \\
& \quad + \beta(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1}\mathbf{V}(\mathbf{B} - \hat{\mu}\mathbf{I}_{d_\xi})^{-1} + \beta(\hat{\mathbf{L}}\mathbf{I}_{d_\xi} - \mathbf{B})^{-1}\mathbf{V}(\hat{\mathbf{L}}\mathbf{I}_{d_\xi} - \mathbf{B})^{-1}. \tag{43}
\end{aligned}$$

2.5. Hessian approximation inverse

To use the Newton's method to solve (9), one needs to either solve the linear system $\mathbf{B}_k \mathbf{d}_k = \nabla_{\xi} F(\xi_k)$, or find the inverse of the Hessian approximation $\mathbf{H}_k = \mathbf{B}_k^{-1}$, thus $\mathbf{d}_k = \mathbf{H}_k \nabla_{\xi} F(\xi_k)$.

To avoid inverting $\hat{\mathbf{B}}_k$, we propose to use the Newton–Raphson method to find an approximation $\hat{\mathbf{H}}_k$ such that $\|\hat{\mathbf{B}}_k \hat{\mathbf{H}}_k - \mathbf{I}_{d_\xi}\|_F \leq \text{tol}_H$; we can start the Newton–Raphson from the previous Hessian inverse approximation, thus saving a significant computational time. The Newton–Raphson method for this problem is given by

$$\hat{\mathbf{H}}_k^{i+1} = \hat{\mathbf{H}}_k^i - (R'(\hat{\mathbf{H}}_k^i))^{-1} R(\hat{\mathbf{H}}_k^i), \quad (44)$$

where $R(\mathbf{H}) = \mathbf{B}_k \mathbf{H} - \mathbf{I}_{d_\xi}$ and $R'(\mathbf{H}) = \mathbf{B}_k$, thus,

$$\hat{\mathbf{H}}_k^{i+1} = (2\mathbf{I}_{d_\xi} - \hat{\mathbf{H}}_k^i \mathbf{B}_k) \hat{\mathbf{H}}_k^i. \quad (45)$$

The recursive procedure is repeated until $\|R(\hat{\mathbf{H}}_k^i)\|_F < \text{tol}_H$. Then, we set $\hat{\mathbf{H}}_k := \hat{\mathbf{H}}_k^i$ as an approximation of \mathbf{H}_k in optimization. The convergence of the Newton–Raphson method depends on the initial guess of the inverse, $\hat{\mathbf{H}}_k^0$; the method converges if the spectral norm of $R(\hat{\mathbf{H}}_k^0)$ is less than one [23]. We start from the inverse estimate in the previous iteration, $\hat{\mathbf{H}}_k^0 = \hat{\mathbf{H}}_{k-1}$. If the method does not converge, that is, if the residue increases, we reset $\hat{\mathbf{H}}_k^i = 2/(\hat{L} + \hat{\mu}) \mathbf{I}_{d_\xi}$. This starting point is similar to the one suggested by [23] for Hermitian positive definite $\hat{\mathbf{B}}_k$, $\hat{\mathbf{H}}_k^i = 2/(\lambda_{\max} + \lambda_{\min}) \mathbf{I}_{d_\xi}$, where λ_{\max} and λ_{\min} are, respectively, the largest and the smallest eigenvalues of $\hat{\mathbf{B}}_k$.

3. Preconditioned stochastic gradient descent methods

Instead of the common update of SGD as presented in (2), we investigate the case where the stochastic gradient estimate is preconditioned with our Hessian inverse approximation,

$$\xi_{k+1} = \xi_k - \eta_k \hat{\mathbf{H}}_k \mathbf{v}_k. \quad (46)$$

If the function to be minimized is strongly convex, SGD with fixed step size converges linearly to a neighbourhood of the optimum [14].

Here, we expect the preconditioned SGD to converge to a neighbourhood of the optimum faster than vanilla SGD. One possible side effect of the preconditioning of the stochastic gradient estimate is that the noise in the gradient might be amplified [18]. Therefore, it is fundamental to control the error in the gradient estimates using variance control techniques, e.g. MICE [8], SVRG [16], SARAH [22]. Note that the variance control methods mentioned before are specializations of (2), only differing in the way that they compute \mathbf{v}_k . Moreover, acknowledging the possibly prohibitive cost of updating the Hessian estimate every iteration, we update the Hessian after a given number of iterations or gradient evaluations.

Some variance control methods are based on computing differences between gradients at known points and are thus well-suited to be combined with our approach. The MICE algorithm [8] builds a hierarchy of iterates, then progressively computes gradient differences between subsequent iterates in this hierarchy to control the gradient estimator error. Since both the differences between the iterates and between the gradients are available, curvature pairs \mathbf{s} and $\bar{\mathbf{y}}$ can be obtained without any extra cost in terms of memory or gradient evaluations. If we refer to the algorithm in (2) using MICE to obtain \mathbf{v}_k as SGD-MICE, preconditioning the gradient estimates by our inverse Hessian approximation as in

(46) furnishes the SGD-MICE-Bay method. For the finite sum case, other control variates methods are available, e.g. SVRG [16], SARAH [22], SAGA [9], SAG [25]. However, SVRG and SARAH have an advantage over SAGA and SAG here in that curvature pairs can be obtained directly from the algorithms without extra memory or processing costs. We refer to SVRG and SARAH preconditioned by our inverse Hessian approximation as SVRG-Bay and SARAH-Bay, respectively.

3.1. Convergence of preconditioned stochastic gradient descent

Assume that the relative statistical error of the gradient estimator is uniformly bounded above by a constant ϵ and that an arbitrary lower-bound $\tilde{\mu}$ is imposed on the eigenvalues of the Hessian approximation, i.e. $\tilde{\mu} = \epsilon\mu$ with $\epsilon \geq 1$. On the one hand, a large value of ϵ decreases the noise amplification, as previously discussed, and at the same time improves the stability of the preconditioned stochastic gradient descent. On the other hand, if μ is indeed the smallest eigenvalue of the true Hessian matrix of F , imposing the constraint that the Hessian approximation must have a smallest eigenvalue $\tilde{\mu} \gg \mu$ causes the Hessian approximation to differ from the true Hessian. In our analysis, we seek to consider the effect of the choice of parameters ϵ , ϵ , and the step size η .

Assumption 3.1 (Regularity and convexity of the objective function): The objective function F has Lipschitz continuous second-order derivatives and is μ -convex. Namely, for all $\mathbf{x}, \mathbf{y} \in \Xi$,

$$\|\nabla^2 F(\mathbf{x}) - \nabla^2 F(\mathbf{y})\| \leq M \|\mathbf{x} - \mathbf{y}\| \quad (47)$$

$$F(\mathbf{y}) \geq F(\mathbf{x}) + \nabla F(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x}) + \frac{\mu}{2} \|\mathbf{y} - \mathbf{x}\|^2. \quad (48)$$

Assumption 3.2 (Unbiasedness and bounded statistical error of gradient estimator): The gradient estimator \mathbf{v}_k is an unbiased estimator of $\nabla F(\boldsymbol{\xi}_k)$ with an upper bound ϵ on its relative statistical error,

$$\mathbb{E}[\mathbf{v}_k | \boldsymbol{\xi}_k] = \nabla F(\boldsymbol{\xi}_k) \quad (49)$$

$$\mathbb{E}\left[\|\mathbf{v}_k - \mathbb{E}[\mathbf{v}_k | \boldsymbol{\xi}_k]\|^2 \mid \boldsymbol{\xi}_k\right] \leq \epsilon^2 \|\nabla F(\boldsymbol{\xi}_k)\|^2. \quad (50)$$

Assumption 3.3 (Bounded error of Hessian approximation): The approximation $\hat{\mathbf{B}}_k$ of $\nabla^2 F(\boldsymbol{\xi}_k)$ has a lower-bound $\tilde{\mu} := \epsilon\mu$ on its eigenvalues with $\epsilon \geq 1$ and the product of the true Hessian by the approximation inverse $\hat{\mathbf{H}}_k$ results in the identity matrix added by an error that is bounded above in norm by a linear function of ϵ ,

$$\hat{\mathbf{B}}_k \succeq \tilde{\mu} \mathbf{I}_{d_{\boldsymbol{\xi}}} := \epsilon\mu, \quad \epsilon \geq 1, \quad (51)$$

$$\nabla^2 F(\boldsymbol{\xi}_k) \hat{\mathbf{H}}_k = \mathbf{I} + \mathcal{E}_k, \quad (52)$$

$$\|\mathcal{E}_k\| \leq C_{\epsilon}(\epsilon - 1), \quad (53)$$

where $C_{\epsilon} > 0$ is a constant.

Theorem 3.1 (Convergence of preconditioned stochastic gradient descent): *If Assumptions 3.1, 3.2, and 3.3 are satisfied and*

$$\|\nabla F(\xi_k)\| < \frac{2\epsilon^2\mu^2(1-\epsilon-C_\epsilon(1+\epsilon)(\epsilon-1)-\frac{r}{\eta})}{M\eta(1+\epsilon^2)} \quad (54)$$

holds for all $k > 0$, then the gradients of the iterates generated by the preconditioned stochastic gradient descent as described in (46) converge linearly with rate $1-r$,

$$\|\nabla F(\xi_{k+1})\| \leq (1-r)^{k+1} \|\nabla F(\xi_0)\|. \quad (55)$$

Proof: From Assumption 3.1, we obtain an upper bound for the norm of the gradient at ξ_{k+1} ,

$$\|\nabla F(\xi_{k+1})\| \leq \|\nabla F(\xi_k) + \nabla^2 F(\xi_k)(\xi_{k+1} - \xi_k)\| + \frac{M}{2} \|\xi_{k+1} - \xi_k\|^2. \quad (56)$$

Using the preconditioned SGD update in (46), one can further develop the equation above as

$$\begin{aligned} \|\nabla F(\xi_{k+1})\| &\leq \left\| \nabla F(\xi_k) - \eta \nabla^2 F(\xi_k)(\hat{\mathbf{H}}_k \mathbf{v}_k) \right\| + \frac{M\eta^2}{2} \left\| \hat{\mathbf{H}}_k \mathbf{v}_k \right\|^2 \\ &\leq \left\| \nabla F(\xi_k) - \eta \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k \nabla F(\xi_k) \right\| \\ &\quad + \left\| \eta \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k (\mathbf{v}_k - \nabla F(\xi_k)) \right\| + \frac{M\eta^2}{2} \left\| \hat{\mathbf{H}}_k \mathbf{v}_k \right\|^2 \end{aligned} \quad (57)$$

$$\begin{aligned} &\leq \left\| \mathbf{I}_{d_\xi} - \eta \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k \right\| \|\nabla F(\xi_k)\| \\ &\quad + \eta \left\| \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k \right\| \|\mathbf{v}_k - \nabla F(\xi_k)\| + \frac{M\eta^2}{2} \left\| \hat{\mathbf{H}}_k \right\|^2 \|\mathbf{v}_k\|^2. \end{aligned} \quad (58)$$

Taking the expectation conditioned on ξ_k and using Assumption 3.2,

$$\begin{aligned} \mathbb{E} \left[\|\nabla F(\xi_{k+1})\| \mid \xi_k \right] &\leq \left\| \mathbf{I}_{d_\xi} - \eta \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k \right\| \|\nabla F(\xi_k)\| \\ &\quad + \eta \epsilon \left\| \nabla^2 F(\xi_k) \hat{\mathbf{H}}_k \right\| \|\nabla F(\xi_k)\| + \frac{M\eta^2}{2} (1+\epsilon^2) \left\| \hat{\mathbf{H}}_k \right\|^2 \|\nabla F(\xi_k)\|^2. \end{aligned} \quad (59)$$

Using Assumption 3.3,

$$\begin{aligned} \mathbb{E} \left[\|\nabla F(\xi_{k+1})\| \mid \xi_k \right] &\leq \left(\left\| \mathbf{I}_{d_\xi} - \eta (\mathbf{I}_{d_\xi} + \mathcal{E}_k) \right\| + \eta \epsilon \left\| \mathbf{I}_{d_\xi} + \mathcal{E}_k \right\| \right) \|\nabla F(\xi_k)\| \\ &\quad + \frac{M\eta^2}{2\bar{\mu}^2} (1+\epsilon^2) \|\nabla F(\xi_k)\|^2 \\ &\leq ((1-\eta) + \eta \|\mathcal{E}_k\| + \eta \epsilon (1 + \|\mathcal{E}_k\|)) \|\nabla F(\xi_k)\| \\ &\quad + \frac{M\eta^2}{2\bar{\mu}^2} (1+\epsilon^2) \|\nabla F(\xi_k)\|^2 \\ &\leq (1 - \eta(1-\epsilon) + \eta(1+\epsilon)C_\epsilon(\epsilon-1)) \|\nabla F(\xi_k)\| \end{aligned}$$

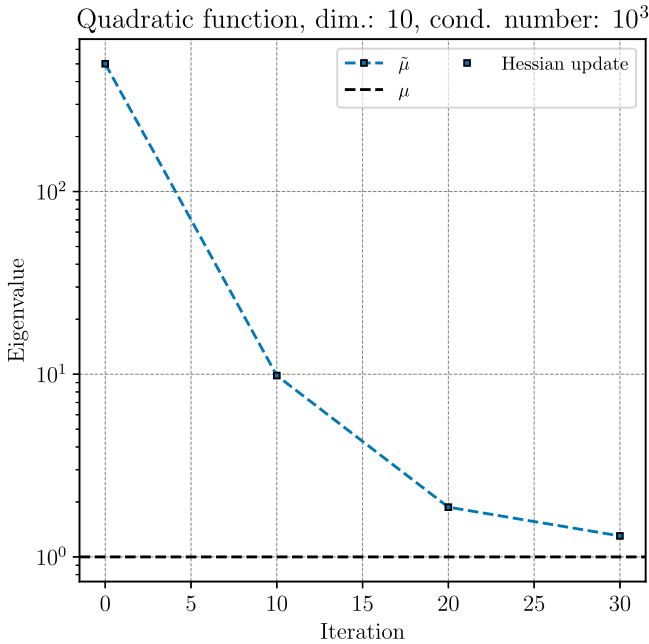


Figure 1. Value of $\tilde{\mu}$ versus iteration for SGD-MICE preconditioned with our Hessian inverse approximation applied to the minimization of the quadratic problem in example in Section 4.1.

$$+ \frac{M\eta^2}{2\varepsilon^2\mu^2}(1 + \varepsilon^2) \|\nabla F(\xi_k)\|^2. \quad (60)$$

Letting the right hand side of the above inequality be strictly less than $(1 - r) \|\nabla F(\xi_k)\|$ results in (54). Taking expectation and unrolling the recursion complete the proof. ■

Theorem 3.1 shows that increasing the parameter ε that controls the smallest eigenvalue of the Hessian approximation improves the stability of preconditioned stochastic gradient descent. Moreover, controlling ε and the step size η to satisfy (54) for a given r guarantees global convergence. By letting $\varepsilon = 0$, thus resulting in deterministic gradient observations; $\varepsilon = 1$, meaning that $\tilde{\mu} = \mu$; and unitary step size $\eta = 1$, one recovers the quadratic convergence of the Newton method with the condition that $\|\nabla F(\xi_k)\| < 2\mu^2/M$ (see (60)).

Since our analysis depends on quantities that are hard to characterize in practice, we do not use it to control ε and η . Instead, we start optimization with $\varepsilon = (L + \mu)/(2\mu)$ and control the parameter ρ of the Frobenius regularization of the prior; the regularization term penalizes large changes in the Hessian approximation. Thus, at every Hessian approximation update, ε progressively decays. Figure 1 illustrates this behaviour for the Example in Section 4.1.

3.2. Time complexity

Given that the interval between Hessian updates is set to be $\mathcal{O}(d_\xi)$, then the extra cost of finding the Hessian approximation using our Bayesian approach is $\mathcal{O}(d_\xi^2)$. The number

of Newton iterations and conjugate gradient iterations do not depend on the problem's dimensionality but on the tolerance set for each algorithm. The cost of each conjugate gradient iteration is the cost of evaluating $\delta_V \nabla_{\mathbf{B}} \mathcal{L}_k(\mathbf{B}, \beta)$ in (43), which is $\mathcal{O}(d_{\xi}^3)$ due to matrix multiplications. If, instead of using the conjugate gradient method to compute the Newton direction, one builds the fourth-order tensor of second-order derivatives of the log posterior with respect to the Hessian, the time complexity increases to $\mathcal{O}(d_{\xi}^4)$.

In each iteration, we have to compute P_k , the precision matrix for each curvature pair. If a full matrix is computed, we have a complexity of $\mathcal{O}(d_{\xi}^2)$ each iteration. If, however, we opt to ignore the off-diagonal terms, this complexity is improved to $\mathcal{O}(d_{\xi})$. We can further decrease the memory overhead of our method by storing only a scalar, the trace of the covariance matrix, for each curvature pair. The reasoning behind this approach is that the trace of the covariance matrix works as a weight for the curvature pairs; in the curvature pairs where the gradient difference is better estimated, the respective curvature pair has a larger weight when computing the likelihood. When our Bayesian Hessian approximation is used with the MICE gradient estimator, the trace of the covariance matrix is available without any extra cost. In numerical tests, we observed that keeping the trace of the covariance matrix resulted in the best option regarding the balance between accuracy and computational time.

We do not discuss here the cost to compute \mathbf{v}_k satisfying condition (50), as it depends on the method used to estimate it. For example, if \mathbf{v}_k is a Monte Carlo estimator with adaptive batch sizes, the cost of evaluating \mathbf{v}_k is $\mathcal{O}(\epsilon^{-2} \|\nabla F(\xi_k)\|^{-2})$. If the conditions of Theorem 1 are satisfied, the gradient sampling cost is $\mathcal{O}(\epsilon^{-2}(1-r)^{-2k})$.

4. Numerical examples

In this section, we solve some numerical examples with and without our Bayesian approach to validate its performance. In all cases, we present the convergence of the optimality gap, $F(\xi_k) - F(\xi^*)$, and in some cases, the convergence of the central-path Newton-CG used to find the Hessian approximations. The parameters of the Bayesian approach are kept fixed for all the numerical examples, namely, the logarithmic barrier parameter, $\beta = 10^{-2}$; the Frobenius norm regularizer parameter, $\rho = 10^{-2}$; the tolerance on the norm of the gradient of the negative log posterior, $tol = 10^{-6}$; the constraint on the lower bound of the eigenvalues of \mathbf{B} , $\tilde{\mu}$, is set as the strong-convexity parameter μ ; the relaxation parameter, $\alpha = 1.05$; the number of central-path steps is six; and the factor of decrease of both tol_i and β_i for each i th central-path step is $\gamma = 2$. Moreover, to keep our Bayesian approach competitive in terms of computational time and memory allocation, we do not compute the full matrix P_{ℓ} for each curvature pair, using the inverse of the trace of the covariance matrix instead. Note that the trace of the covariance matrix of the difference between gradients is available from MICE since it is used to control the relative error of the gradient [8].

The results presented in this section are obtained using a python implementation of the proposed method whose source code can be found at GitHub.¹ Moreover, the source code of the MICE gradient estimator, also used in this section, can be found at Bitbucket.² A repository with all the code necessary to reproduce the results presented in this section is available at GitHub.³

Table 1. Step sizes and batch sizes used for Example 4.1.

Method	Step size	Batch size
SGD	$\frac{1}{L\sqrt{k}}$	1
SGD-MICE	$\frac{1}{(L+\mu)(1+\epsilon^2)}$	$\epsilon = 1$
SGD-Bay (decr.)	$\frac{1}{\sqrt{k}}$	10
SGD-Bay (fixed)	1	1000
SGD-MICE-Bay	$1/(1 + \epsilon^2)$	$\epsilon = 0.5$

Notes: Here, L is the Lipschitz constant of ∇f , μ is the strong-convexity constant, k is the iteration, and ϵ is the tolerance on the relative statistical error.

4.1. Quadratic function

The first numerical test is a quadratic function with noise added to its gradient. We want to minimize $\mathbb{E}_\theta[f(\xi, \theta)|\xi]$ with

$$f(\xi, \theta) = \frac{1}{2}\xi \cdot A \xi - \mathbf{b} \cdot \xi + \xi \cdot \sqrt{\Sigma} \theta, \quad (61)$$

where A is a square symmetric positive-definite matrix with $1 \preceq A \preceq \kappa$, Σ is a symmetric and positive definite covariance matrix, θ is a vector of independent standard normal distributed random variables, and the vector \mathbf{b} is defined as $\mathbf{b} = \mathbf{1}_{d_\xi}$, where $\mathbf{1}_{d_\xi}$ is a vector of ones of dimension d_ξ . The eigenvalues of A are sampled uniformly between 1 and κ , with necessarily one eigenvalue being 1 and another being κ . The gradient of f with respect to ξ is

$$\nabla f(\xi, \theta) = A \xi - \mathbf{b} + \sqrt{\Sigma} \theta, \quad (62)$$

thus, the covariance matrix of ∇f is Σ .

We solve this problem for two different values of the condition number κ , namely 10^3 and 10^6 , and for a fixed number of dimensions $d_\xi = 10$. We compare the performances of different methods in the solution of this problem: SGD, SGD-MICE, SGD-Bay with decreasing step size, SGD-Bay with fixed step size, and SGD-MICE-Bay. In all cases, we run the optimization until 10^7 gradients are evaluated. For SGD-Bay, since the number of gradient evaluations per iteration is fixed, we know the number of iterations a priori and thus choose the interval between Hessian updates to have 15 equally distributed updates. For SGD-MICE-Bay, we opt to have an update every d_ξ iterations. The step sizes and batch sizes used for each method are presented in Table 1, where, for the cases when MICE is used, we show the tolerance ϵ on the relative statistical error of the gradient.

In Figure 2, we present the optimality gap versus iteration for both $\kappa = 10^3$ (left) and $\kappa = 10^6$ (right) for the previously mentioned methods. In the cases where our Bayesian approach is used, the Hessian updates are illustrated as coloured squares with black edges. For the case where $\kappa = 1000$, SGD-MICE-Bay achieved, in 39 iterations, a lower optimality gap than SGD in 10^7 iterations. Moreover, SGD presents oscillations around the optimum, making it difficult to have a reliable optimum estimate. In the case of SGD-Bay, with and without the decreasing step size, we observe oscillations of larger amplitude due to the preconditioning with our Hessian approach. This behaviour is due to the lack of variance control of SGD; not even the decreasing step size is enough to avoid the oscillations. In the case of an even larger condition number, $\kappa = 10^6$, SGD-MICE-Bay far outperforms

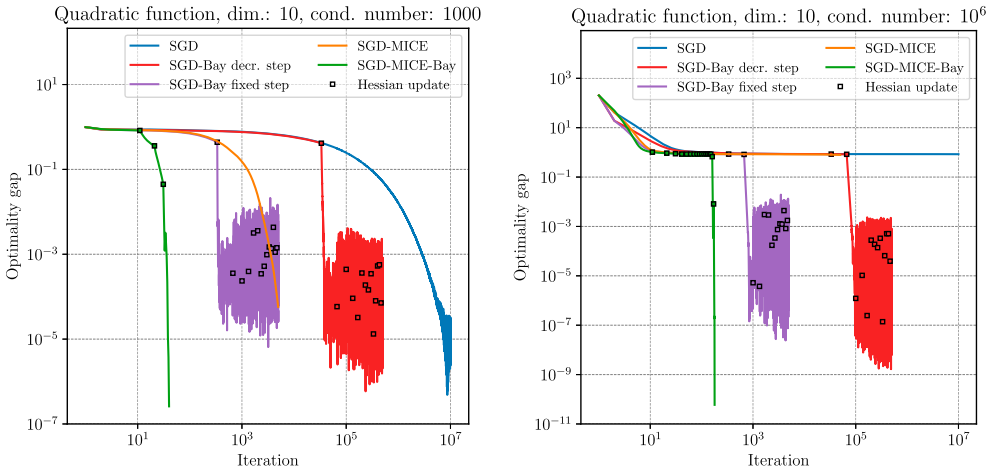


Figure 2. Convergence of the optimality gap versus iteration with $\kappa = 1000$ (left) and $\kappa = 10^6$ (right) for Example 4.1. The Hessian updates are marked with coloured squares with black edges. In both cases, SGD-MICE-Bay reduced the optimality gap more than the other methods and in a much smaller number of iterations. Moreover, SGD-MICE-Bay did not oscillate around the optimum like both SGD-Bay cases.

the other methods while requiring 178 iterations. Both SGD and SGD-MICE, being first-order methods, cannot improve the optimality gap significantly once close to the optimum. As for SGD-Bay, the same oscillating behaviour is observed for both the fixed and decreasing step cases. From these results, it is clear that, in large condition numbers cases, using second-order information is necessary for an efficient convergence. Moreover, we observe that controlling the variance of the gradient estimator is of central importance to maintaining convergence in the stochastic setting, avoiding the amplification of noise due to the preconditioning.

One crucial aspect of our method is the cost to update the Hessian approximation, i.e. the cost of the Newton-CG method in this setting. For SGD-MICE-Bay in the $\kappa = 1000$ case, we present in Figure 3 the Newton-CG convergence of the gradient norm $\|\nabla_{\mathbf{B}} \mathcal{L}\|$ per Newton iteration. The black numbers on top of each update denote the number of conjugate gradient steps needed to find the Newton direction. Each colour represents a different central-path step, i.e. at each colour change, both β and tol are decreased by a factor $\gamma = 2$, with each tol denoted by a dash-dotted line. For this problem, the runtime of each of the Hessian updates is at fractions of seconds. Given that $d_{\xi} = 10$, the Hessian matrix has 100 components. Still, in the worst case, 51 conjugate gradient iterations were needed to find the Newton direction, showing that the problem of finding $\hat{\mathbf{B}}$ is well-conditioned. Moreover, the total number of Newton iterations in the three cases did not exceed 25.

Controlling the eigenvalues maxima and minima is essential to keep the stability and efficiency of preconditioned SGD. To motivate why our approach is better suited to stochastic optimization than using the BFGS formulae to approximate the Hessian, we, for every Hessian update of SGD-MICE-Bay, also compute a Hessian approximation using BFGS. In Figure 4, we present, for both $\kappa = 1000$ (left) and $\kappa = 10^6$ (right), the extreme eigenvalues for our Bayesian approach, for BFGS, and the true extremes. The same curvature pairs are used for both our Bayesian approach and BFGS. For both condition number

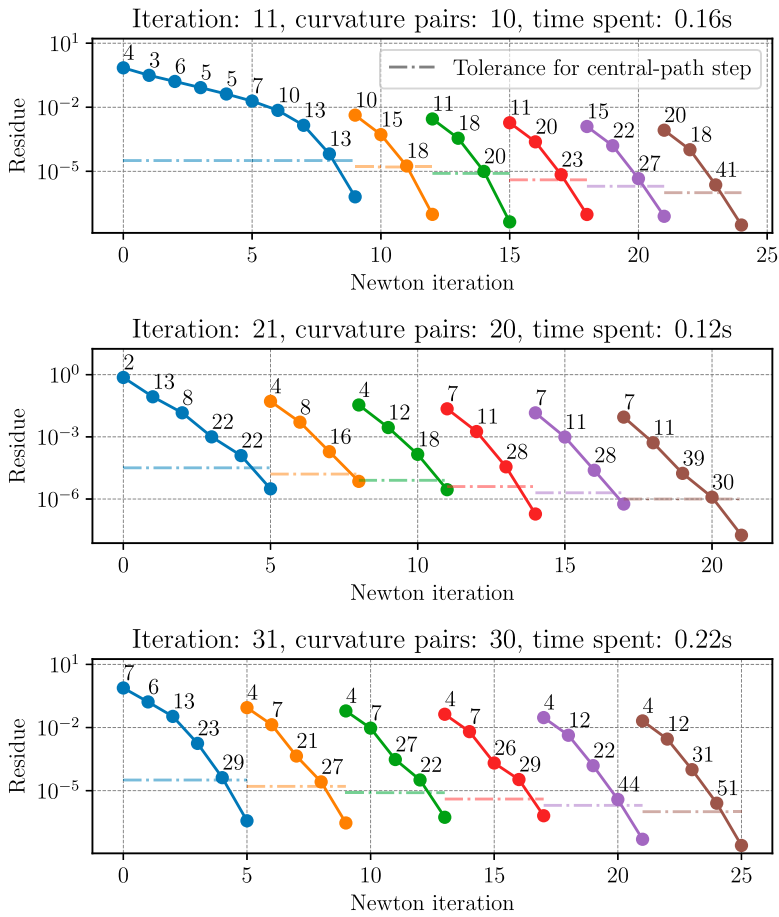


Figure 3. For each of the three Hessian updates of SGD-MICE-Bay with $\kappa = 1000$ in Example 4.1, we present the convergence of the gradient norm for the Newton-CG method. The number of curvature pairs available and the time taken are presented for each Hessian update. Each colour represents a different step of the central-path method with a different logarithmic barrier parameter β and a different residue tolerance. The tolerance for each central-path step is presented as a dash-dotted line. The number above each Newton iteration represents the number of conjugate gradient iterations needed to find the Newton direction.

cases, the BFGS method can keep the minimum eigenvalue above the actual value; however, the largest eigenvalue far exceeds the true largest eigenvalue. In contrast, our Bayesian approach keeps the eigenvalues of the Hessian approximation between the extreme values due to the logarithmic barrier constraints.

As motivated in (6) and Theorem 3.1, increasing $\hat{\mu}$ decreases the noise amplification of preconditioning the gradients by our Bayesian Hessian approximation. Figure 5 presents a comparison between SGD-Bay with a fixed step size for different values of $\hat{\mu}$. One thing to notice is the reduction of the amplitude of the oscillations as $\hat{\mu}$ increases. Moreover, we observe a clear linear convergence for the more conservative cases of $\hat{\mu} = 1000$ and $\hat{\mu} = 10,000$. Comparing these results with the ones presented in Figure 2 for the same condition number $\kappa = 10^6$, it is clear that increasing $\hat{\mu}$ improves the performance of SGD-Bay even

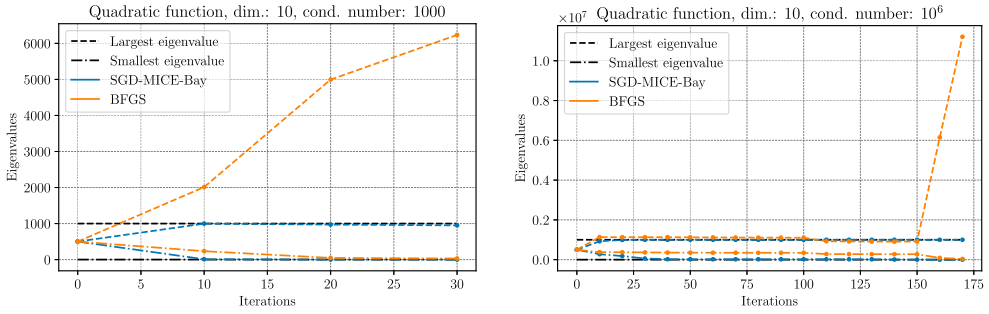


Figure 4. Maximum and minimum eigenvalues for both our Bayesian approach and BFGS for Example 4.1 with $\kappa = 1000$ (left) and $\kappa = 10^6$ (right). We use the same curvature pairs for both cases, obtained from SGD-MICE-Bay. The true extremes of the eigenvalues are presented as black lines. Dashed lines represent the largest eigenvalues, and dash-dotted lines represent the smallest eigenvalues. It is clear that our approach gets closer to the actual extreme eigenvalues than BFGS, which is farther from the smallest eigenvalue than our method and exceeds the value of the largest eigenvalue.




Figure 5. Optimality gap versus iteration for SGD-Bay with fixed step size on Example 4.1 for different values of $\hat{\mu}$. The coloured squares with black edges represent the Hessian updates of each method. The different values of $\hat{\mu}$ result in different amplitudes for the oscillations around the optimum of SGD-Bay; larger values of $\hat{\mu}$ result in smaller amplitudes of oscillations (cf. (6)).

without controlling the gradient statistical error or decreasing the step size. However, in the general convex case, SGD-Bay with a fixed step size will converge to a neighbourhood of the optimum and oscillate around it with a fixed amplitude; increasing $\hat{\mu}$ or decreasing the step size may reduce the amplitude of the oscillations but not circumvent them.

Table 2. For each of the datasets used in Example 4.2, we present its data size, number of features, and condition number of the optimization problem.

dataset	size	num. of features	κ
<i>mushrooms</i>	8124	112	5.25×10^5
<i>ijcnn1</i>	49,990	22	3.76×10^4
<i>w8a</i>	49,749	300	2.91×10^5
<i>cod-rna</i>	59,535	8	6.48×10^9
<i>HIGGS</i>	11,000,000	28	2.50×10^4

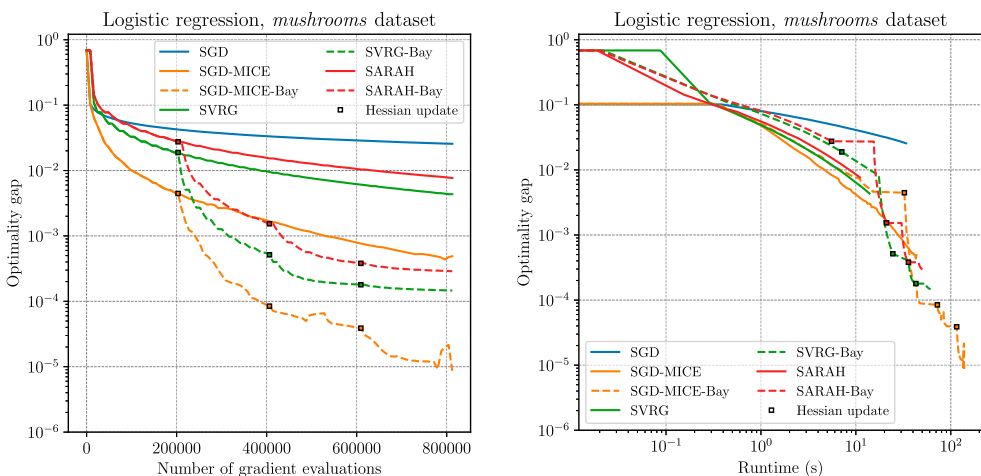
4.2. Logistic regression

Training logistic regression models for classification of finite data is an empirical risk minimization problem, thus a problem of minimizing a sum of functions. When training a model on a large dataset, it is common to use SGD to reduce the iteration cost by subsampling the data points used in each iteration. Since inexact information about the gradient is used in each iteration, deterministic quasi-Newton methods usually fail to recover the Hessian and its inverse. Here, we use our Bayesian framework to recover the Hessian from noisy gradient observations and precondition the gradient estimates with the Hessian inverse. We compare plain SGD, SGD-MICE, SGD-MICE-Bay, SVRG, SVRG-Bay, SARAH, and SARAH-Bay on the task of training a binary classification logistic regression model with ℓ_2 regularization on five different datasets, *mushrooms*, *ijcnn1*, *w8a* [17], *cod-rna* [29], and *HIGGS* [1].⁴

The ℓ_2 -regularized log loss function we seek to minimize is

$$F(\xi) = \frac{1}{N} \sum_{i=1}^N f(\xi, \theta_i = (\mathbf{x}_i, y_i)) = \frac{1}{N} \sum_{i=1}^N \log(1 + \exp(-y_i \xi \cdot \mathbf{x}_i)) + \frac{\lambda}{2} \|\xi\|^2, \quad (63)$$

where each data point (\mathbf{x}_i, y_i) is such that $\mathbf{x}_i \in \mathbb{R}^{d_\xi}$ and $y_i \in \{-1, 1\}$. For all datasets, we use a regularization parameter $\lambda = 10^{-5}$.

**Figure 6.** Convergence of the optimality gap versus iteration (left) and runtime (right) for Example 4.2, *mushrooms* dataset. The Hessian updates are marked with coloured squares with black edges.

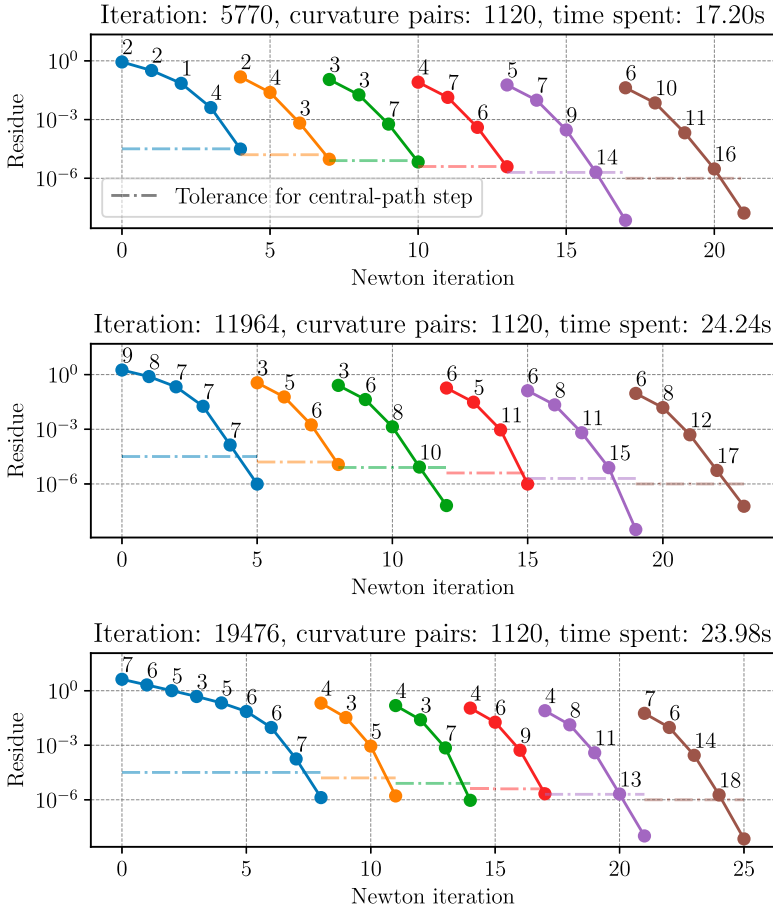


Figure 7. Convergence of the Newton-CG method for each of the Hessian updates of SGD-MICE-Bay in Example 4.2, *mushrooms* dataset. The number of curvature pairs available and the time taken are presented for each Hessian update. Each colour represents a different step of the central-path method with a different logarithmic barrier parameter β and a different residue tolerance. The tolerance for each central-path step is presented as a dash-dotted line. The number above each Newton iteration represents the number of conjugate gradient iterations needed to find the Newton direction.

For SGD, only a singleton sample is used, i.e. the minibatch size is set to 1. Also, as in the example in Section 4.1, a decreasing schedule is used for the step size. For SGD-MICE, the sample size is adaptively controlled to keep the relative error of the gradient estimates below $\epsilon = 0.5$ with a step size of $\eta = 2/((L + \mu)(1 + \epsilon^2))$ [8]. For the *HIGGS* dataset, the choice of $\epsilon = 0.5$ was too conservative, thus we used $\epsilon = 0.8$. As for SGD-MICE-Bay, we used the same values of ϵ , however, since we precondition the gradient estimates with the inverse of the Hessian, we use a step size of $\eta = (1 + \epsilon^2)^{-1}$. Therefore, with the initial Hessian guess with diagonal entries $B_{ii} = \frac{L + \mu}{2}$, SGD-MICE-Bay recovers exactly SGD-MICE with the optimal step size. Similarly, for SVRG and SARAH, we use a fixed step size of $\eta_k = 0.1/L$, whereas for SVRG-Bay and SARAH-Bay we use $\eta_k = 0.1$. For SVRG, SVRG-Bay, SARAH, and SARAH-Bay, we start with a full batch, iterate using a mini-batch size of five, and restart every two epochs.




Figure 8. Convergence of the optimality gap versus iteration (left) and runtime (right) for Example 4.2, *ijcnn1* dataset. The Hessian updates are marked with coloured squares with black edges.

In Table 2, we present the sizes, number of features, and condition numbers of the resulting optimization problems for the five datasets. The used datasets differ significantly in their characteristics.

In Figure 6, we present the convergence for the optimality gap versus iterations (left) and versus runtime in seconds (right) for the *mushrooms* dataset. For SGD, SGD-MICE, SVRG, and SARAH, we present their convergences as solid lines and their respective preconditioned cases as dashed lines. For all the methods, preconditioning the gradient estimates using our Bayesian approach improved the final optimality gap in at least one order of magnitude. In terms of convergence of the optimality gap versus runtime, we observe that computing the Hessian approximation adds an overhead to the optimization process; however, the improvement in convergence compensates for this overhead. In this example, SGD-MICE-Bay achieved an optimality gap ten times smaller than that found by SGD-MICE and more than two orders of magnitude smaller than the one found by SGD. Preconditioning the gradient estimates of SVRG and SARAH with our Bayesian Hessian inverse approximation reduced the final optimality gap found by more than one order of magnitude.

In Figure 7, the convergence of the Newton-CG method is presented for each of the Hessian updates of SGD-MICE-Bay for the *mushrooms* dataset. Note that a relatively small number of Newton steps are needed for each central-path step; 8 in the worst case. Moreover, the number of conjugate gradient steps does not exceed 18, which is a remarkably small number of iterations considering that we have a problem with 112×112 variables. This result indicates that the problem is well-conditioned.

Figure 8 presents the convergence for the *ijcnn1* dataset. Here, as in the case of the *mushrooms* dataset, our Bayesian approach improved the convergence in all cases. However, the most remarkable improvement is the one of SVRG-Bay, which achieved an optimality gap of more than two orders of magnitude smaller than the vanilla SVRG with just one Hessian update. Both SVRG-Bay and SGD-MICE-Bay achieved very similar values of optimality gap, with SVRG-Bay requiring less runtime than SGD-MICE-Bay. In Figure 9,

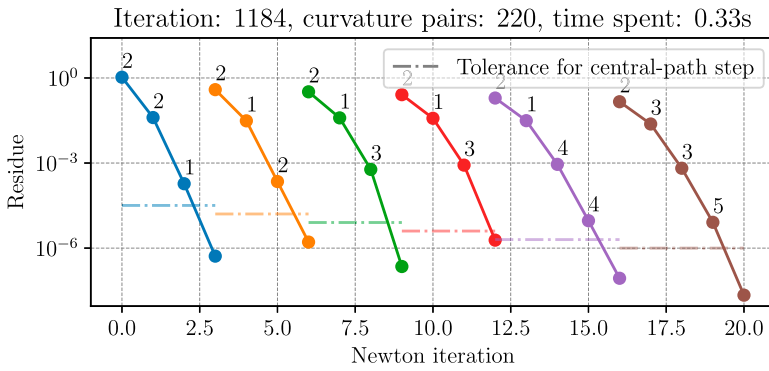


Figure 9. Convergence of the Newton-CG method for each of the Hessian updates of SGD-MICE-Bay in Example 4.2, *ijcnn1* dataset. The number of curvature pairs available and the time taken are presented for each Hessian update. Each colour represents a different step of the central-path method with a different logarithmic barrier parameter β and a different residue tolerance. The tolerance for each central-path step is presented as a dash-dotted line. The number above each Newton iteration represents the number of conjugate gradient iterations needed to find the Newton direction.

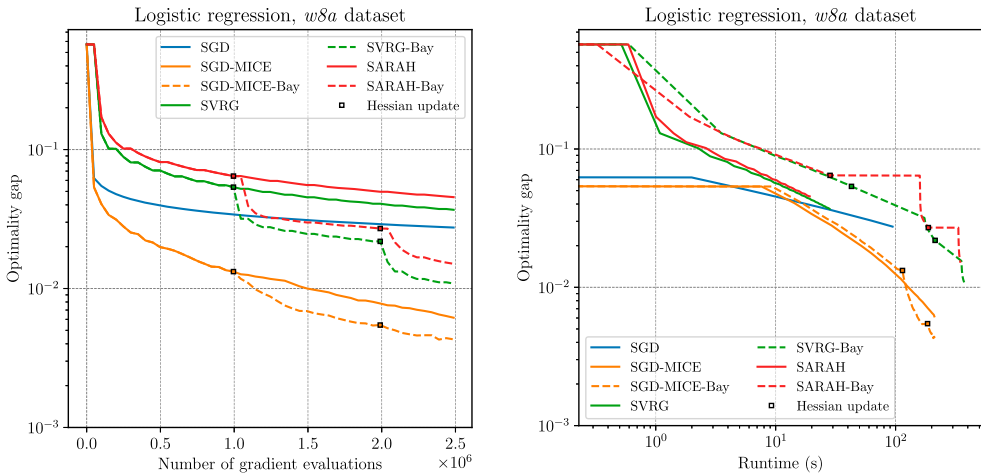


Figure 10. Convergence of the optimality gap versus iteration (left) and runtime (right) for Example 4.2, *w8a* dataset. The Hessian updates are marked with coloured squares with black edges.

we present the convergence of Newton-CG for the one Hessian update of SGD-MICE-Bay for the *ijcnn1* dataset. Note that the time spent computing the Hessian approximation is 0.33 seconds; for comparison, computing the true Hessian took 0.54 seconds in an average of 100 evaluations.

The *w8a* dataset is the case with the largest number of features, 300, meaning that the Hessian to be approximated by our Bayesian approach has 300×300 components. The convergence for the *w8a* dataset is presented in Figure 10. Coupling SVRG, SARAH, and SGD-MICE with our Bayesian approach improved the convergence in all cases, but it was more significant for SVRG and SARAH. However, in terms of runtime, SGD-MICE was able to perform better than its competitors, and SGD-MICE-Bay had a small edge over SGD-MICE. Figure 11 presents the convergence of Newton-CG for SGD-MICE-Bay on

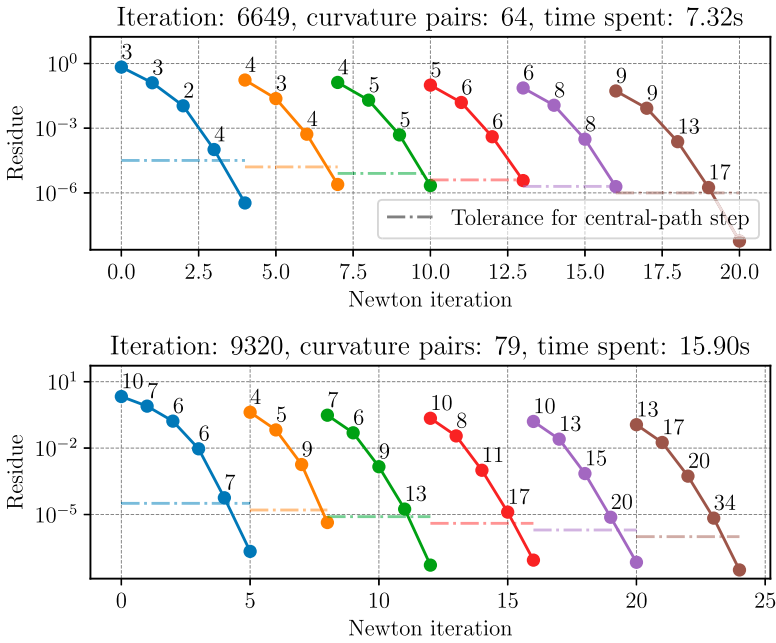


Figure 11. Convergence of the Newton-CG method for each of the Hessian updates of SGD-MICE-Bay in Example 4.2, *w8a* dataset. The number of curvature pairs available and the time taken are presented for each Hessian update. Each colour represents a different step of the central-path method with a different logarithmic barrier parameter β and a different residue tolerance. The tolerance for each central-path step is presented as a dash-dotted line. The number above each Newton iteration represents the number of conjugate gradient iterations needed to find the Newton direction.

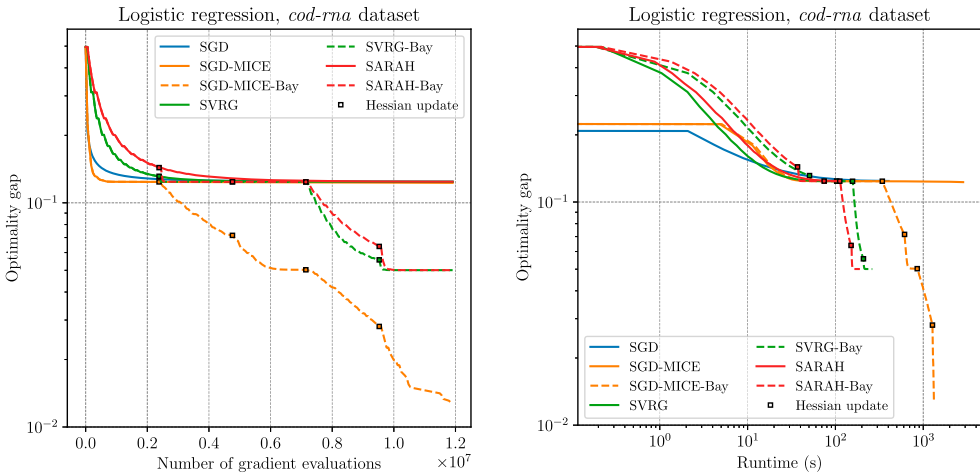


Figure 12. Convergence of the optimality gap versus iteration (left) and runtime (right) for Example 4.2, *cod-rna* dataset. The Hessian updates are marked with coloured squares with black edges.

the *w8a* dataset. As in the case of the *mushrooms* dataset, a small number of Newton steps (a maximum of 24) and conjugate gradient steps (a maximum of 34) were needed despite the large dimensionality of the problem at 90,000 variables. Even with the large Hessian,

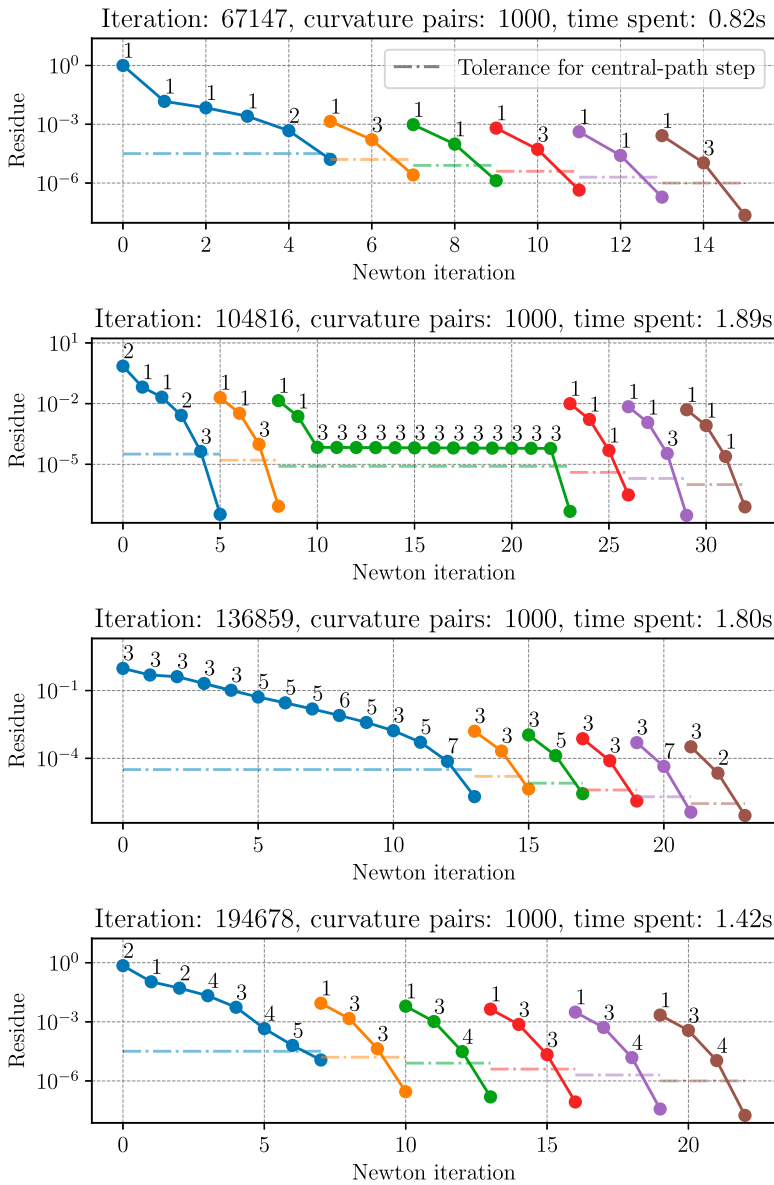


Figure 13. Convergence of the Newton-CG method for each of the Hessian updates of SGD-MICE-Bay in Example 4.2, *cod-rna* dataset. The number of curvature pairs available and the time taken are presented for each Hessian update. Each colour represents a different step of the central-path method with a different logarithmic barrier parameter β and a different residue tolerance. The tolerance for each central-path step is presented as a dash-dotted line. The number above each Newton iteration represents the number of conjugate gradient iterations needed to find the Newton direction.

the time spent to find an approximation (7.32 s in the first update and 15.9 s in the second) is competitive with the time of computing the true Hessian, of 8.52 s in an average of 100 evaluations. Remember that our approach has stability advantages over using the true Hessian inverse as a preconditioner due to the control of the noise amplification given by the logarithmic barrier on the smallest eigenvalue.




Figure 14. Convergence of the optimality gap versus iteration (left) and runtime (right) for Example 4.2, HIGGS dataset. The Hessian updates are marked with coloured squares with black edges.

Training the logistic regression model to the *cod-rna* dataset is the problem with the largest condition number of those studied here, at 6.48×10^9 . We expect our Bayesian approach to be able to assist the optimization methods in converging to the true solution despite the large condition number of the problem. The convergence of the optimality gap for the *cod-rna* dataset is presented in Figure 12. Note that SVRG, SARAH, and SGD-MICE get stuck after some iterations, improving very little for most of the optimization process. Their preconditioned counterparts, however, perform much better, being able to decrease the optimality gap continuously. When looking at the convergence of the optimality gap versus runtime for SGD-MICE-Bay, we observe a sharp decrease after the first Hessian update that follows until its stop. Figure 13 presents the convergence of Newton-CG for the four Hessian updates of SGD-MICE-Bay. In the second Hessian update, the third central-path step, the Newton-CG method got temporarily stuck for 12 iterations, leading us to believe that the parameters β and/or ρ were not properly set for this example. Still, the Newton-CG method could resume the Hessian approximation in a reasonable time, at 1.89 s, showing the robustness of our approach.

The HIGGS dataset contains 11×10^6 data points, being the largest studied here. Thus, since our Hessian approximation is independent of the data size, we expect to have a time advantage compared to computing the true Hessian. Figure 14 presents the convergence of the optimality gap for the HIGGS dataset. Here we observe that, in contrast to the previous results, our Bayesian approach worsened the performance of SARAH and SVRG. For SGD-MICE, however, our method improved the final optimality gap by more than one order of magnitude. Also, concerning runtime, SGD-MICE-Bay performed better than the other methods, despite doing 9 Hessian updates. The longest Hessian update was the last one, lasting 2.36 s. In comparison, the average time taken to compute a Hessian for this example is 130.58 s. One possible interpretation for the bad performance of SVRG-Bay and SARAH-Bay is that, contrary to SGD-MICE, SVRG and SARAH do not control the relative statistical error in the gradient estimates.

5. Conclusion

We presented a Bayesian approach to approximate Hessians of functions given noisy observations of their gradients. This problem is of great importance in stochastic optimization, where pre-conditioning gradient estimates with the inverse of the Hessian matrix can improve the convergence of stochastic gradient descent methods. The proposed Bayesian approach takes into consideration not only the secant equations as in deterministic quasi-Newton methods but also the noise in the gradients. To mitigate the known effect of amplification of the gradient noise, we control the smallest eigenvalue of the Hessian approximation. To maximize the log-posterior of the Hessian matrix, we use a Newton-CG method with a central-path approach to impose the eigenvalue constraints. The numerical results presented show that our approach is not only interesting from the theoretical perspective but also results in practical advantage. In both a stochastic quadratic equation and in the training of a logistic regression model with ℓ_2 regularization, using our Bayesian approximation of the Hessian matrix resulted in better convergence rates without a significant increase in runtime. For future research, we suggest using low-rank or diagonal approximations of the Hessians.

Notes

1. <https://github.com/agcarlon/bayhess>
2. <https://bitbucket.org/agcarlon/mice>
3. https://github.com/agcarlon/bayhess_numerics
4. Obtained from LibSVM datasets at <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>, except for the HIGGS dataset that was obtained from <https://archive.ics.uci.edu> [11].

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This work was partially supported by the KAUST Office of Sponsored Research (OSR) under Award number OSR-2019-CRG8-4033 in the KAUST Competitive Research Grants Program Round 8, the Alexander von Humboldt Foundation, and the Flexible Interdisciplinary Research Collaboration Fund at the University of Nottingham Project ID 7466664.

Notes on contributors

Dr *André Gustavo Carlon* completed his Bachelor's degree in Civil Engineering at the Universidade Federal de Santa Catarina (UFSC) in Brazil in 2013. He then pursued his Master's and PhD degrees at the same institution, defending his Master's thesis on structural optimization in 2015 and his Ph.D. dissertation on uncertainty quantification in 2019. Dr. Carlon has since been a postdoctoral fellow in the stochastic numerics group at King Abdullah University of Science and Technology (KAUST) in Saudi Arabia. His research focuses on stochastic optimization, structural optimization, reliability engineering, and optimal experimental design.

Dr *Luis Espath* obtained his Engineering Diploma from the Pontifical Catholic University of Rio Grande do Sul (PUCRS) in Brazil in 2007. He then completed both his Master's and Doctoral

degrees at the Federal University of Rio Grande do Sul (UFRGS) in Brazil, concluding his studies in 2013. Dr Espath served as a postdoctoral fellow at King Abdullah University of Science and Technology (KAUST) until 2019. He then joined RWTH Aachen University in Germany as a Research Scientist, where he defended his Habilitation thesis in Mathematics in the field of Theoretical Mechanics in 2022. His research interests include theoretical and computational mechanics, uncertainty quantification, stochastic optimization, and machine learning.

Dr *Raül Tempone* received his PhD from the Royal Institute of Technology in Stockholm, Sweden, in 2002. He subsequently held postdoctoral positions at the Institute for Computational Engineering and Sciences (ICES) at the University of Texas at Austin and served as an Assistant Professor at Florida State University in Tallahassee. Since 2009, Dr. Tempone has been a Full Professor at King Abdullah University of Science and Technology (KAUST) in Thuwal, Saudi Arabia. In 2018, he was awarded the Alexander von Humboldt Professorship at RWTH Aachen University in Germany. Dr. Tempone's research focuses on numerical analysis, specifically the development and analysis of efficient and robust numerical methods for problems involving stochastic models and differential equations in various fields, including computational mechanics, quantitative finance, biological and chemical modeling, and wireless communications.

References

- [1] P. Baldi, P. Sadowski, and D. Whiteson, *Searching for exotic particles in high-energy physics with deep learning*, Nat. Commun. 5 (2014), pp. 1–9.
- [2] R. Bollapragada, J. Nocedal, D. Mudigere, H.J. Shi, and P.T.P. Tang, *A progressive batching L-BFGS method for machine learning*, in *International Conference on Machine Learning*, PMLR, 2018, pp. 620–629.
- [3] A. Bordes, L. Bottou, and P. Gallinari, *Sgd-qn: Careful quasi-Newton stochastic gradient descent*, J. Mach. Learn. Res. 10 (2009), pp. 1737–1754.
- [4] L. Bottou, F.E. Curtis, and J. Nocedal, *Optimization methods for large-scale machine learning*, SIAM Rev. 60 (2018), pp. 223–311.
- [5] S. Boyd, S.P. Boyd, and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, 2004.
- [6] R.H. Byrd, G.M. Chin, W. Neveitt, and J. Nocedal, *On the use of stochastic hessian information in optimization methods for machine learning*, SIAM. J. Optim. 21 (2011), pp. 977–995.
- [7] R.H. Byrd, S.L. Hansen, J. Nocedal, and Y. Singer, *A stochastic quasi-Newton method for large-scale optimization*, SIAM. J. Optim. 26 (2016), pp. 1008–1031.
- [8] A. Carlon, L. Espath, R. Lopez, and R. Tempone, *Multi-iteration stochastic optimizers*, preprint (2020). Available at arXiv.
- [9] A. Defazio, F. Bach, and S. Lacoste-Julien, *SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives*, in *Advances in Neural Information Processing Systems*, Montreal, QC, Canada, 2014, pp. 1646–1654.
- [10] J.E. Dennis Jr and J.J. Moré, *Quasi-Newton methods, motivation and theory*, SIAM Rev. 19 (1977), pp. 46–89.
- [11] D. Dua and C. Graff, *UCI machine learning repository* (2022). Available at <http://archive.ics.uci.edu/ml>.
- [12] D. Goldfarb, Y. Ren, and A. Bahamou, *Practical quasi-Newton methods for training deep neural networks*, Adv. Neural. Inf. Process. Syst. 33 (2020), pp. 2386–2396.
- [13] R. Gower, D. Goldfarb, and P. Richtárik, *Stochastic block BFGS: Squeezing more curvature out of data*, in *International Conference on Machine Learning*, PMLR, 2016, pp. 1869–1878.
- [14] R.M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin, and P. Richtárik, *SGD: General analysis and improved rates*, in *International Conference on Machine Learning*, PMLR, 2019, pp. 5200–5209.
- [15] P. Hennig, *Fast probabilistic optimization from noisy gradients*, in *International Conference on Machine Learning*, PMLR, 2013, pp. 62–70.

- [16] R. Johnson and T. Zhang, *Accelerating stochastic gradient descent using predictive variance reduction*, Adv. Neural. Inf. Process. Syst. 26 (2013), pp. 315–323.
- [17] R. Kohavi, *Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid*, in KDD, Portland, OR, USA, Vol. 96, 1996, pp. 202–207.
- [18] X.L. Li, *Preconditioned stochastic gradient descent*, IEEE. Trans. Neural. Netw. Learn. Syst. 29 (2017), pp. 1454–1466.
- [19] A. Mokhtari and A. Ribeiro, *Stochastic quasi-Newton methods*, Proc. IEEE 108 (2020), pp. 1906–1922.
- [20] P. Moritz, R. Nishihara, and M. Jordan, *A linearly-convergent stochastic L-BFGS algorithm*, in *Artificial Intelligence and Statistics*, Cadiz, Spain, 2016, pp. 249–258.
- [21] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic Course*, Vol. 87, Springer Science & Business Media, Cham, Switzerland, 2013.
- [22] L.M. Nguyen, J. Liu, K. Scheinberg, and M. Takáč, *SARAH: A novel method for machine learning problems using stochastic recursive gradient*, in *International Conference on Machine Learning*, PMLR, 2017, pp. 2613–2621.
- [23] V. Pan and J. Reif, *Efficient parallel solution of linear systems*, in *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, Providence, RI, USA, 1985, pp. 143–152.
- [24] B.T. Polyak, *Introduction to Optimization. Optimization Software*, Vol. 1, Inc., Publications Division, New York, 1987.
- [25] M. Schmidt, N. Le Roux, and F. Bach, *Minimizing finite sums with the stochastic average gradient*, Math. Program. 162 (2017), pp. 83–112.
- [26] N.N. Schraudolph, J. Yu, and S. Günter, *A stochastic quasi-Newton method for online convex optimization*, in *Artificial Intelligence and Statistics*, PMLR, 2007, pp. 436–443.
- [27] H.J.M. Shi, Y. Xie, R. Byrd, and J. Nocedal, *A noise-tolerant quasi-Newton algorithm for unconstrained optimization*, SIAM. J. Optim. 32 (2022), pp. 29–55.
- [28] J. Sohl-Dickstein, B. Poole, and S. Ganguli, *Fast large-scale optimization by unifying stochastic gradient and quasi-Newton methods*, in *International Conference on Machine Learning*, Beijing, China, 2014, pp. 604–612.
- [29] A.V. Uzilov, J.M. Keegan, and D.H. Mathews, *Detection of non-coding rnas on the basis of predicted secondary structure formation free energy change*, BMC. Bioinformatics. 7 (2006), pp. 1–30.
- [30] X. Wang, S. Ma, D. Goldfarb, and W. Liu, *Stochastic quasi-Newton methods for nonconvex stochastic optimization*, SIAM. J. Optim. 27 (2017), pp. 927–956.
- [31] Y. Xie, R. Byrd, and J. Nocedal, *Analysis of the BFGS method with errors*, preprint (2019). Available at arXiv, arXiv:1901.09063.