

Improving Levenberg-Marquardt Algorithm for Neural Networks

Omead Pooladzandi

Yiming Zhou

UCLA 420 Westwood Plaza Los Angeles, CA 90095

OPOOLADZ@UCLA.EDU

YIMINGZ0416@UCLA.EDU

Abstract

We explore the usage of the Levenberg-Marquardt (LM) algorithm for regression (non-linear least squares) and classification (generalized Gauss-Newton methods) tasks in neural networks. We compare the performance of the LM method with other popular first-order algorithms such as SGD and Adam[4], as well as other second-order algorithms such as L-BFGS [9], Hessian-Free [10] and KFAC[11]. We further speed up the LM method by using adaptive momentum, learning rate line search, and uphill step acceptance.

1. Introduction

First-order methods are easy to use and require moderate computational cost per iteration for optimizing neural networks. However, their hyper-parameters are cumbersome to tune, and the convergence rate can decrease as the neural network (NN) size becomes large. Second-order methods leverage the Hessian matrix (i.e. the curvature of the objective function) and enjoy faster convergence, robustness, and affine invariance. In a single step, instead of just advancing towards the minimum, it steps towards the global minimum under the assumption that the function is quadratic and its second-order expansion is a good approximation. This allows the optimizer to make big steps in low-curvature scenarios and small steps in high-curvature scenarios. However, computing and storing the Hessian $H = \nabla^2 f(x)$ and its inverse is expensive in computation (time complexity $O(n^3)$) and memory (space complexity $O(n^2)$). When the number of parameters in a NN becomes large it is prohibitively expensive to compute the full Hessian inverse. Several approaches, such as the Gauss-Newton method, the Limited-Memory Broyden–Fletcher–Goldfarb–Shanno (LBFGS) method, the Hessian-Free method, and the KFAC method [9–11], have been proposed to approximate curvature information making second-order algorithms more efficient.

In this work, we provide a version of the Levenberg-Marquardt (LM) Algorithm for Neural Networks which utilizes adaptive momentum, learning rate line search, and uphill step acceptance to increase convergence as well as performance. We compare the modified LM method to LBFGS, HF, KFAC to SGD and Adam for training a CNN on the MNIST dataset and an MLP on a noisy Sine regression task. For more on each method see Appendix A.

2. LM Algorithm for Neural Networks

There have been many updates to the standard Levenberg-Marquardt algorithm. In this section, we consider some of these extensions and combine the methods to get an updated version of LM.

2.1. Gauss-Newton Method

Gauss-Newton (GN) method was originally proposed to solve a nonlinear least square problem, where $r : \mathbb{R}^n \rightarrow \mathbb{R}^m$ denotes the residual $r(x) = [r_1(x), \dots, r_m(x)]^T$ of a n -vector x .

$$\min g(x) = \|r(x)\|_2^2 = \sum_{i=0}^m r_i(x)^2 \quad (1)$$

To optimize the function above, we can take a Newton step, which will do the following update

$$v_{\text{nt}} = -\nabla^2 g(x)^{-1} \nabla g(x) = -(J^T J + \sum_{i=1}^m r_i(x) \nabla^2 r_i(x))^{-1} J^T r \quad (2)$$

where $J = \frac{\partial r(x)}{\partial x}$ is the Jacobian matrix of $r(x)$. However in the Gauss-Newton step, we discard the second term in the inverse and simply do

$$v_{\text{gn}} = -(J^T J)^{-1} J^T r. \quad (3)$$

The Gauss-Newton step doesn't require second derivatives and always provides a descent direction when J has full column rank; because $H = 2J^T J$ is positive definite.

The Levenberg-Marquardt algorithm is a variant of the Gauss-Newton method which addresses the indefiniteness of the approximated Hessian by adding a regularization term λ

$$v_{\text{gn}} = -(J^T J + \lambda I)^{-1} J^T r. \quad (4)$$

Here λ defines a trust region where the quadratic approximation is a good fit. It blends linear and quadratic approximations. When $\lambda = 0$, the update rule is completely based on the second-order approximation, when λ is very large, the rule is completely based on the first-order approximation.

2.2. Generalized Gauss-Newton Method

Consider a neural network (NN) with $L + 1$ layers, at the l -th layer, given the outputs from the preceding layer $x^{(l-1)}$ as input, $x^{(l)}$ is computed as $x^{(l)} = \phi(W^{(l)}x^{(l-1)} + b^{(l)})$, where $W^{(l)}$ and $b^{(l)}$ are the weights and bias, ϕ is the nonlinear activation function. \hat{y} is the output of the NN. Assume the dimension of the output of the last layer is c . The empirical average loss we want to minimize is:

$$f(\theta) = \frac{1}{N} \sum_{i=0}^N f_i(\theta) = \frac{1}{N} \sum_{i=1}^N \varepsilon(\hat{y}_i(\theta), y_i) \quad (5)$$

where θ denotes all the parameters in the NN, $\theta \in \mathbb{R}^n$ (n is the number of all the parameters) and $\varepsilon(\hat{y}_i(\theta), y_i)$ is a loss function based on the differences between the output \hat{y}_i and the ground truth y_i , ($\hat{y}_i, y_i \in \mathbb{R}^c$) for a given dataset $\{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$. The loss ε could be Mean Square Error for the regression task and Cross-Entropy for the classification task.

To derive the generalized Gauss-Newton method, we first examine the Hessian of $f_i(\theta)$ corresponding to a single data point.

$$\frac{\partial^2 f_i(\theta)}{\partial \theta} = J_i^T H_i J_i + \sum_{j=1}^c \left(\frac{\partial f_i(\theta)}{\partial \hat{y}_i} \right)_j \frac{\partial}{\partial \theta} (J_i)_j \quad (6)$$

where $J_i = \frac{\partial \hat{y}_i}{\partial \theta}$ and $H_i = \frac{\partial^2 f_i(\theta)}{\partial (\hat{y}_i)^2}$. In the Gauss-Newton method, we approximate the Hessian by ignoring the second term and only keeping $J_i^T H_i J_i$ where $J_i \in \mathbb{R}^{c \times n}$ and $H_i \in \mathbb{R}^{c \times c}$. Finally, $\frac{\partial^2 f_i(\theta)}{\partial \theta^2}$ is approximated by $H = \frac{1}{N} \sum_{i=1}^N J_i^T H_i J_i$ [15]. When $f_i(\theta) = \varepsilon(\hat{y}_i(\theta), y_i) = \frac{1}{2} \|\hat{y}_i(\theta) - y_i\|_2^2$, which is the Mean Square Error loss, $H_i = I$, then we have the same formula as stated in the nonlinear least square problem: $H = \frac{1}{N} \sum_{i=1}^N J_i^T J_i$.

When $f_i(\theta) = \varepsilon(\hat{y}_i(\theta), y_i)$ is the Cross Entropy loss, if we assume $\hat{y}_i \in \mathbb{R}^c$ is a vector which represents the probabilities of each class, and $y_i \in \mathbb{R}^c$ is a one-hot encoding vector with 1 for the groundtruth class k . Then $H_i = \frac{\partial^2 f_i(\theta)}{\partial(\hat{y}_i)^2}$ is a matrix in $\mathbb{R}^{n \times n}$ with only one non-zero entry at the position (k, k) , therefore the approximated H can be further simplified to

$$H = \frac{1}{N} \sum_{i=1}^N \nabla f_i(\theta) \nabla f_i(\theta)^T \quad (7)$$

Finally we add the regularization term as we did in the LM method and get $H = H + \lambda I$.

2.3. Dampening: Maximum Diagonal

To improve on the Gauss-Newton method, in the LM method we dampen the $J^T J$ matrix with a diagonal cut-off. The LM step therefore becomes: $d\theta = -(J^T J + \lambda D^T D)^{-1} \nabla f(\theta)^T$. Where $D^T D$ is a positive-definite matrix representing the relative scaling of the parameters and λ is the dampening parameter. We define $D^T D$ as the element-wise maximum of the diagonal of all the previous $J^T J$ and the current one. We initialize $D^T D = 0.01I$. This dampening matrix preserves rescaling invariance while being more robust at avoiding parameter evaporation [16].

2.4. LM with Adaptive Momentum

By combining the merits of the Levenberg-Marquardt (LM) method and the Conjugate Gradient (CG) method, one can incorporate an adaptive momentum term into the LM method. And can formulate training a NN as a constrained optimization problem [1].

We denote the NN parameters θ_t , the updates $d\theta_t$ at time step t , and H_t is the (approximated) Hessian. Two update θ_t and θ_{t-1} are non-interfering or mutually conjugate w.r.t. H when $d\theta_t^T H d\theta_{t-1} = 0$ Our goal is to maximize

$$\Phi_t = d\theta_t H_t d\theta_{t-1} \quad (8)$$

without compromising the need for a decrease of the loss function $f(\theta)$ so we add two constraints where $\delta Q_t < 0$.

$$d\theta_t^T H_t d\theta_t = (\delta P)^2 \quad df(\theta_t) = \delta Q_t \quad (9)$$

The constrained optimization problem can be solved analytically by introducing Lagrange multipliers z_1, z_2

$$\phi_t = \Phi_t + z_1(\delta Q_t - df(\theta_t)) + z_2[(\delta P)^2 - d\theta_t^T H_t d\theta_t] \quad (10)$$

To maximize the ϕ_t we get the following:

$$d\theta_t = -\frac{z_1}{2z_2} H_t^{-1} \nabla f(\theta_t) + \frac{1}{2z_2} d\theta_{t-1} \quad (11)$$

Define the following terms:

$$I_{GG} = \nabla f(\theta_t)^T H_t^{-1} \nabla f(\theta_t) \quad I_{GF} = \nabla f(\theta_t)^T d\theta_t \quad I_{FF} = d\theta_{t-1}^T H_t d\theta_{t-1} \quad (12)$$

Then z_1, z_2 can be evaluated in terms of known quantities

$$z_2 = \frac{1}{2} \left[\frac{I_{GG}(\delta P)^2 - (\delta Q_t)^2}{I_{FF} I_{GG} - I_{GF}^2} \right]^{-1/2} \quad z_1 = \frac{-2z_2 \delta Q_t + I_{GF}}{I_{GG}} \quad (13)$$

2.5. Learning Rate Line Search

Due to the trust region nature of the LM method, one will have successful iterations as well as failed ones. A successful iteration is defined as an iteration whose updated weight results in a lower loss than the previous iteration. In the case of a failed iteration, we perform a line search over the learning rate (LR) to minimize the loss function. We then update our network using those parameters. The algorithm is produced in Alg. 2 in the appendix. We do not count the line search as extra iterations since we are not recalculating the network's Jacobian or an inverse.

2.6. Uphill Steps

The LM algorithm requires a rule to accept or reject a proposed step. A natural and safe choice is to only accept steps that reduce the loss. But this may not be the most efficient acceptance rule. The Uphill step [16] rule accepts steps such that the algorithm is moving in the general direction of the previous iteration. Formally $\beta_i = \cos(d\theta^{new}, d\theta^{old})$. This denotes the cosine angle between the proposed direction $d\theta^{new}$ and the last accepted step $d\theta^{old}$. In words, we want to accept uphill moves if the angle β_i is acute, increasingly as $\beta_i \rightarrow 0$. Specifically the uphill acceptance rule becomes $(1 - \beta_i)^b f(\theta^{t+1}) \leq f(\theta^t)$ or more conservatively if $(1 - \beta_i)^b f(\theta^{t+1}) \leq \min(f(\theta^i), \dots, f(\theta^t))$. As b increases the algorithm will accept uphill steps more freely.

3. Evaluation and Comparison

3.1. Setup

We use a synthetic noisy Sine dataset and MNIST to evaluate the NN. We use the best hyper-parameters per optimizer. See appendix B for more details.

3.2. Regression Problem

We begin our experiments with a vanilla LM method using the GN approximation of the Hessian matrix. We train a 2-layer MLP with ELU activations to do a regression of a noisy Sine function. We then add adaptive momentum, LR line search, and maximum diagonal dampening in LM, and compared. As we see, in Figure 1 (a), adding these methods to LM greatly improves the convergence.

In Figure 1 (b) compares modified LM to other first and Quasi/Gauss-Newton methods with tuned hyper-parameters. We see that our modified LM method was able to converge faster and to a lower loss compared to the rest. First order optimizers as well as KFAC converged to a loss of 0.5 whereas LBFGS and HF methods converged to around 0.25. In comparison, LM with max diagonal components of the GN matrix, adaptive momentum, and LR line search converged to 0.001.




Figure 1: Comparison of optimizers on Sine regression problem. The X-axis is iterations y-axis is loss. a) vanilla LM vs updated LM. b) Comparing 1st and 2nd order optimizers vs modified LM

As second-order methods are slower per iteration compared to first-order methods, we train a NN till the loss converges using the LM method and allow other optimizers the same time. In Figure 1 LM converges to a lower loss representing a better solution to the regression problem (see Figure 3).

3.3. MNIST Classification

Next, we train a Convolutional Neural Network on the MNIST dataset. In Figure 2 the LM method optimizes the NN in fewer epochs, converging to a lower loss with higher test and train accuracy

	LM	KFAC	HF	LBFGS	Adam	SGD
Time	0.9	3	4.7	5.6	0.1	0.1
Epochs to reach 90%	4	26	11	17	13	26
Converged	Yes	Yes	No	Yes	Yes	Yes
Accuracy (100 Epochs)	97.5%	96%	94%	91%	96.5%	96%

Table 1: Comparison of Timing, variance in appendix

compared to the rest of the optimizers. In our experiments, HF optimization tends to be unstable. We adjust the learning rate as well as increased the number of CG iterations for the HF method but we often found instabilities with both HF and LBFGS. We found LBFGS would often get stuck at an accuracy of 1%; we omit these runs from the results.

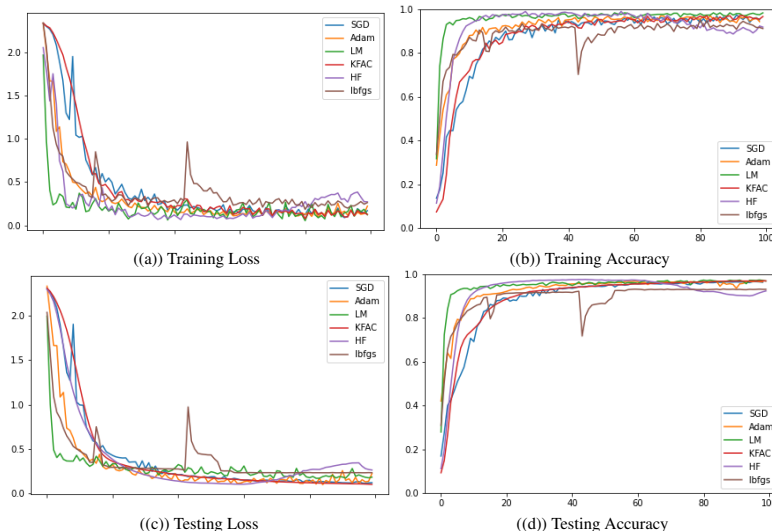


Figure 2: CNN on Full MNIST (100 Epochs)

Table 1 is the overall timing and results of the full MNIST experiment on a CNN. We ran each optimizer and timed them till they reached an accuracy of 90%. We see that the LM method reaches the best loss followed by Adam. Yet the time it took to run the first-order optimizers was less than other methods. LM converges to a lower loss and better accuracy than the other optimizers, even if it is slower. As we also observed from the noisy Sine regression before, the LM method can learn well from a small amount of data, as shown by the few numbers of iterations it took to reach the 90% accuracy mark. Figure 5 shows the ablation study on LM and demonstrates the effectiveness of learning rate line search, adaptive momentum, as well as uphill techniques. In Figure 5, we see that these adjustments significantly improve the stability of the loss and improve accuracy.

4. Conclusion and Future Work

We implement the Levenberg-Marquardt method for neural networks and observe favorable convergence performance on both regression and classification tasks. Furthermore, learning rate line search, adaptive momentum, and uphill steps are all useful techniques to improve the performance of the LM method. With this in mind, we found this method to not be scalable due to memory and computational constraints. We believe that some of the methods explored in this paper can be applied to very recently proposed, cheap curvature-aware optimizers such as PSGD [5–8].

References

- [1] Nikolaos Ampazis and Stavros J Perantonis. Levenberg-marquardt algorithm with adaptive momentum for the efficient training of feedforward networks. In *Proceedings of the IEEE-INNS-ENNS International Joint Conference on Neural Networks. IJCNN 2000. Neural Computing: New Challenges and Perspectives for the New Millennium*, volume 1, pages 126–131. IEEE, 2000.
- [2] P.W.D. Charles. Project title. https://github.com/yaroslavvb/kfac_pytorch/blob/master/derivation.pdf, 2013.
- [3] Andrew Gibiansky. Hessian free optimization, 2014. URL <https://andrew.gibiansky.com/blog/machine-learning/hessian-free-optimization>.
- [4] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [5] Xi-Lin Li. Online second order methods for non-convex stochastic optimizations, 2018. URL <https://arxiv.org/abs/1803.09383>.
- [6] Xi-Lin Li. Preconditioned stochastic gradient descent. *IEEE Transactions on Neural Networks and Learning Systems*, 29(5):1454–1466, 2018. doi: 10.1109/TNNLS.2017.2672978.
- [7] Xi-Lin Li. Preconditioner on matrix lie group for sgd, 2018. URL <https://arxiv.org/abs/1809.10232>.
- [8] Xilin Li. Black box lie group preconditioners for sgd.
- [9] Dong C. Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *MATHEMATICAL PROGRAMMING*, 45:503–528, 1989.
- [10] James Martens. Deep learning via hessian-free optimization. In *ICML*, volume 27, pages 735–742, 2010.
- [11] James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pages 2408–2417, 2015.
- [12] Barak A Pearlmutter. Fast exact multiplication by the hessian. *Neural computation*, 6(1): 147–160, 1994.
- [13] Omead Pooladzandi, Jeffrey Jiang, Sunay Bhat, and Gregory Pottie. Causal structural hypothesis testing and data generation models. NeurIPS 2022 Workshop on Synthetic Data for Empowering ML Research.
- [14] Omead Pooladzandi, David Davini, and Baharan Mirzasoleiman. Adaptive second order coresets for data-efficient machine learning. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 17848–17869. PMLR, 17–23 Jul 2022. URL <https://proceedings.mlr.press/v162/pooladzandi22a.html>.

- [15] Yi Ren and Donald Goldfarb. Efficient subsampled gauss-newton and natural gradient methods for training neural networks. *arXiv preprint arXiv:1906.02353*, 2019.
- [16] Mark K. Transtrum and James P. Sethna. Improvements to the levenberg-marquardt algorithm for nonlinear least-squares minimization, 2012.

Appendix A. Theoretical Background

A.1. Newton Method

Newton's method is a second-order method which enjoys fast convergence, robustness and affine invariance. In a single step, instead of just advancing towards the minimum, it steps towards the global minimum under the assumption that the function is actually quadratic and its second order expansion is a good approximation. This allows the optimizer to make big steps in low-curvature scenarios and small steps in high-curvature scenarios. The Newton step v_{nt} is calculated by minimizing the second-order approximation in (14). However, computing and storing the Hessian $H = \nabla^2 f(x)$ and its inverse is expensive in computation (time complexity $O(n^3)$) and memory (space complexity $O(n^2)$). Several variations are proposed to approximate the Hessian H more efficiently.

$$\hat{f}(x + v) = f(x) + \nabla f(x)^T v + \frac{1}{2} v^T \nabla^2 f(x) v \quad (14)$$

$$v_{\text{nt}} = -\nabla^2 f(x)^{-1} \nabla f(x) \quad (15)$$

A.2. Quasi-Newton Method

Quasi-Newton method uses inverse Hessian estimation to compute each update iteration. The most popular one is the Broyden–Fletcher–Goldfarb–Shanno (BFGS) update.

$$H_{k+1}^{-1} = \left(I - \frac{sy^T}{y^T s}\right) H_k^{-1} \left(I - \frac{ys^T}{y^T s}\right) + \frac{ss^T}{y^T s} \quad (16)$$

$$s = x_{k+1} - x_k \quad y = \nabla f(x_{k+1}) - \nabla f(x_k) \quad (17)$$

where H_k^{-1} is the approximation of inverse Hessian at iteration k . Then to further tackle the memory issue, limited-memory BFGS (L-BFGS) was proposed to avoid storing H_k and H_k^{-1} by evaluating H_k recursively for m iterations ($j = k - 1, \dots, k - m$).

$$H_{j+1}^{-1} = \left(I - \frac{s_j y_j^T}{y_j^T s_j}\right) H_j^{-1} \left(I - \frac{y_j s_j^T}{y_j^T s_j}\right) + \frac{s_j s_j^T}{y_j^T s_j} \quad (18)$$

$$s_j = x_{j+1} - x_j \quad y_j = \nabla f(x_{j+1}) - \nabla f(x_j) \quad (19)$$

A.3. Gauss-Newton Method

Gauss-Newton method was originally proposed to solve the nonlinear least squares problems, where $r : \mathbb{R}^n \rightarrow \mathbb{R}^m$ denotes the residual $r(x) = [r_1(x), \dots, r_m(x)]^T$ of a n -vector x .

$$\min \quad g(x) = \|r(x)\|_2^2 = \sum_{i=0}^m r_i(x)^2 \quad (20)$$

To optimize the function above, we can take a Newton step, which will do the following update

$$v_{\text{nt}} = -\nabla^2 g(x)^{-1} \nabla g(x) = -\left(J^T J + \sum_{i=1}^m r_i(x) \nabla^2 r_i(x)\right)^{-1} J^T r \quad (21)$$

where $J = \frac{\partial r(x)}{\partial x}$ is the Jacobian matrix of $r(x)$. However in the Gauss-Newton step, we discard the second term in the inverse and simply do

$$v_{\text{gn}} = -(J^T J)^{-1} J^T r \quad (22)$$

So the Gauss-Newton step doesn't require second derivatives and always provides a descent direction when J has full column rank; because $H = 2J^T J$ is positive definite.

Levenberg-Marquardt algorithm is a variant of Gauss-Newton method which addresses the indefiniteness of the approximated Hessian by adding a regularization term λ

$$v_{\text{gn}} = -(J^T J + \lambda I)^{-1} J^T r \quad (23)$$

Here λ essentially defines a trust region where the quadratic approximation is a good fit. It blends the linear and quadratic approximations. When $\lambda = 0$, the update rule is completely based on the second-order approximation, when λ is very large, the rule is completely based on the first-order approximation.

A.4. Generalized Gauss-Newton Method

Consider a neural network (NN) with $L + 1$ layers, at the l -th layer, given the outputs from the preceding layer $x^{(l-1)}$ as input, $x^{(l)}$ is computed as $x^{(l)} = \phi(W^{(l)} x^{(l-1)} + b^{(l)})$, where $W^{(l)}$ and $b^{(l)}$ are the weight and bias, ϕ is the nonlinear activation function. \hat{y} is the output of the NN. Assume the dimension of output of the last layer is c . The empirical average loss we want to minimize is:

$$f(\theta) = \frac{1}{N} \sum_{i=0}^N f_i(\theta) = \frac{1}{N} \sum_{i=1}^N \varepsilon(\hat{y}_i(\theta), y_i) \quad (24)$$

where θ denotes all the parameters in the NN, $\theta \in \mathbb{R}^n$ (n is the number of all the parameters) and $\varepsilon(\hat{y}_i(\theta), y_i)$ is a loss function based on the differences between the output \hat{y}_i and the groundtruth y_i , ($\hat{y}_i, y_i \in \mathbb{R}^c$) for a given dataset $\{(x_1, y_1), (x_2, y_2), \dots, (x_N, y_N)\}$. The loss ε could be Mean Square Error for the regression task and Cross Entropy for the classification task.

To derive the generalized Gauss-Newton method, we first examine the Hessian of $f_i(\theta)$ corresponding to a single data point.

$$\frac{\partial^2 f_i(\theta)}{\partial \theta} = J_i^T H_i J_i + \sum_{j=1}^c \left(\frac{\partial f_i(\theta)}{\partial \hat{y}_i} \right)_j \frac{\partial}{\partial \theta} (J_i)_j \quad (25)$$

where $J_i = \frac{\partial \hat{y}_i}{\partial \theta}$ and $H_i = \frac{\partial^2 f_i(\theta)}{\partial (\hat{y}_i)^2}$. In the Gauss-Newton method, we approximate the Hessian by ignoring the second term and only keeping $J_i^T H_i J_i$ where $J_i \in \mathbb{R}^{c \times n}$ and $H_i \in \mathbb{R}^{c \times c}$. Finally, we approximate $\frac{\partial^2 f_i(\theta)}{\partial \theta^2}$ is approximated by [15]

$$H = \frac{1}{N} \sum_{i=1}^N J_i^T H_i J_i \quad (26)$$

When $f_i(\theta) = \varepsilon(\hat{y}_i(\theta), y_i) = \frac{1}{2} \|\hat{y}_i(\theta) - y_i\|_2^2$, which is the Mean Square Error loss, $H_i = I$, then we have the same formula as stated in the nonlinear least square problem.

$$H = \frac{1}{N} \sum_{i=1}^N J_i^T J_i \quad (27)$$

When $f_i(\theta) = \varepsilon(\hat{y}_i(\theta), y_i)$ is the Cross Entropy loss, if we assume $\hat{y}_i \in \mathbb{R}^c$ is a vector which represents the probabilities of each class, and $y_i \in \mathbb{R}^c$ is a one-hot encoding vector with 1 for the groundtruth class k . Then $H_i = \frac{\partial^2 f_i(\theta)}{\partial(\hat{y}_i)^2}$ is a matrix in $\mathbb{R}^{n \times n}$ with only one non-zero entry at the position (k, k) , therefore the approximated H can be further simplified to

$$H = \frac{1}{N} \sum_{i=1}^N \nabla f_i(\theta) \nabla f_i(\theta)^T \quad (28)$$

Finally we add the regularization term as we did in the LM method and get $H = H + \lambda I$

A.5. Hessian-Free

Hessian-free optimization uses the insights from Newton’s method but comes up with a better way to minimize the quadratic function. Just as in Newton’s method, given a function f , we find the second-order Taylor expansion:

$$f(x + v) \approx f(x) + \nabla f(x)^T v + v^T H(f)v \quad (29)$$

We then find the minimum of this approximation (the best Δx) using conjugate gradient, take an iterative step to $x = x + v$, and continue this iteration until we have convergence.

Algorithm 1 Hessian-Free

Let f be any function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ which we wish to minimize.

1. Initialize:
Let $i = 0$ and $x_i = x_0$ be some initial guess.
2. Set up quadratic expansion:
At the current x_n , compute the gradient $\nabla f(x_n)$ and Hessian $H(f)(x_n)$, and consider the following Taylor expansion of f :

$$f(x + v) \approx f(x) + \nabla f(x)^T v + v^T H(f)v$$

3. Conjugate gradient:
Compute x_{n+1} using the Conjugate Gradient algorithm for quadratic functions on the current Taylor expansion.
 4. Iterate:
Repeat steps 2 and 3 until x_n has converged.
-

Fortunately in the Hessian free algorithm, it is possible to circumvent the full Hessian expansion in step 2, and instead only requires us to calculate the much cheaper Hessian vector product Hv .

To see this, consider the first component of Hv . The i -th row of the Hessian contains partial derivatives of the form $\frac{\partial^2}{\partial x_i \partial x_j} f$ where j is the column index. As per normal matrix-vector multiplication, the i -th row of Hv is the dot product of v and i th row of H . The i -th row of H can be expressed as the gradient of the derivative of f with respect to x_i , yielding [3]:

$$(Hv)_i = \sum_{j=1}^N \frac{\partial^2 f}{\partial x_i \partial x_j}(x) \cdot v_j = \nabla \frac{\partial f}{\partial x_i}(x) \cdot v$$

One can utilize auto-differentiation packages to calculate the exact Hessian vector product as detailed by [12] The operator is defined

$$\mathcal{R}_v\{f(\mathbf{x})\} \equiv \left. \frac{\partial}{\partial t} f(\mathbf{x} + t\mathbf{v}) \right|_{t=0}$$

so $H\mathbf{v} = \mathcal{R}_v\{\nabla_{\mathbf{x}}\}$. (To avoid clutter we will usually write $\mathcal{R}\{\cdot\}$ instead of $\mathcal{R}_v\{\cdot\}$) We can then take all the equations of a procedure that calculates a gradient, *e.g.* the backpropagation procedure, and we can apply the $\mathcal{R}_v\{\cdot\}$ operator to each equation. Because $\mathcal{R}\{\cdot\}$ is a differential operator, it obeys the usual rules for differential operators, such as:

$$\begin{aligned} \mathcal{R}\{cf(\mathbf{x})\} &= c\mathcal{R}\{f(\mathbf{x})\} \\ \mathcal{R}\{f(\mathbf{x}) + g(\mathbf{x})\} &= \mathcal{R}\{f(\mathbf{x})\} + \mathcal{R}\{g(\mathbf{x})\} \\ \mathcal{R}\{f(\mathbf{x})g(\mathbf{x})\} &= \mathcal{R}\{f(\mathbf{x})\}g(\mathbf{x}) + f(\mathbf{x})\mathcal{R}\{g(\mathbf{x})\} \\ \mathcal{R}\{f(g(\mathbf{x}))\} &= f'(g(\mathbf{x}))\mathcal{R}\{g(\mathbf{x})\} \\ \mathcal{R}\left\{\frac{df(\mathbf{x})}{dt}\right\} &= \frac{d\mathcal{R}\{f(\mathbf{x})\}}{dt} \end{aligned}$$

Also note that

$$\mathcal{R}\{\mathbf{x}\} = \mathbf{v}$$

These rules allow one to use existing auto-differentiation packages to calculate the vector $\mathcal{R}\{\nabla_{\mathbf{x}}\}$ which is precisely the vector which we desire Hv .

In practice, HF sees fewer applications than SGD because its updates are much more expensive to compute, as they involve running linear conjugate gradient (CG) for potentially hundreds of iterations. Each of these iterations requires a matrix-vector product with the curvature matrix (which are as expensive to compute as the stochastic gradient on the current mini-batch). Next, HF's estimate of the curvature matrix must remain fixed while CG iterates. This means that the method is able to go through much less data than SGD can in a comparable amount of time, making it less well suited to stochastic optimizations.

A.6. KFAC

To get the full understand behind the KFAC method one needs to dive into ample theory covered in the paper. Instead of reproducing this, I will instead provide a sketch for the KFAC method in the Newton Step, we well as talk about some differences between KFAC and other methods. The following notes on KFAC are reproduced from [2] for completion.

A.6.1. LINEAR NETWORK

Given a matrix parameter W and a goal of modeling Y as follows

$$Y = B^T W A \quad (30)$$

We define a prediction error as

$$r = \hat{Y} - B^T W A. \quad (31)$$

We minimize the prediction error by minimizing over W in the trace loss

$$f = \frac{1}{2} \text{tr}(r^T r) \quad (32)$$

To minimize we take the first derivative:

$$df = \text{tr}(r^T dr) = \text{tr}(r^T B^T dW A) \quad (33)$$

By using the trace properties we find the derivative of W

$$G = dW = -B r A^T \quad (34)$$

Taking the derivative of G we have:

$$dG = -B B^T dW A A^T \quad (35)$$

We vectorize both sides and apply the Kronecker-Vector transform rule:

$$\text{vec}(dG) = (A A^T \otimes B B^T) \text{vec}(dW) \quad (36)$$

We can then by inspection see

$$H = (A A^T \otimes B B^T) \quad (37)$$

Since we have extracted the Hessian from an equation that was vectorized we write our update rule as:

$$\text{vec}(W) - H^{-1} \text{vec}(G) \quad (38)$$

To calculate this inverse we note that the inverse distributes over Kronecker product and the fact that vec distributes over Kronecker product we have

$$\text{vec}(W - (B B^T)^{-1} G (A A^T)^{-1}). \quad (39)$$

Which gives us the Newton update step for the unvectorized form of the problem.

A.6.2. PIECEWISE LINEAR NETWORK

For a piecewise linear neural network, the loss for each example is locally linear, and we can write our total loss as sum of per-example losses

$$f = \sum_i f_i \tag{40}$$

where each example i is associated with its own versions of A_i, B_i and r_i . And so we have the per-example derivative of the loss:

$$df_i = -tr(r_i^T B_i^T dW_i A_i) \tag{41}$$

Where the total Hessian is the sum of H_i which we can approximate using equation 1 in the KFAC paper

$$H = \sum_i (A_i A_i^T \otimes B_i B_i^T) \approx \sum_i (A_i A_i^T) \otimes \sum_i (B_i B_i^T) \tag{42}$$

Since each example corresponds to a vector we can stack these vectors into a matrix resulting in $H \approx AA^T \otimes BB^T$. Again, our Hessian inverse can be written as

$$H^{-1} \approx (AA^T)^{-1} \otimes (BB^T)^{-1} \tag{43}$$

A.6.3. INTUITION AND DISTINGUISHING FACTORS OF KFAC

The core of the KFAC method is to use an approximation of the Fisher Information matrix which assumes each of the layers are independent. The independence can be seen by the authors formulation in which all the weights of a each layer of a NN are broken down into groups by rows and columns in the Fisher matrix. This allows for the Fisher matrix to be factored into the multiplication of much smaller matrix multiplications via Kronecker factorization. Such a factorization allows for very efficient inversion of the Fisher matrix. Online estimation of the matrix is possible using arbitrarily large subsets of the training data (without increasing the cost of inversion). This assumption, although a strong one, allows for the math to be factored in a computationally efficient way. While the assumption that the layers of a NN are independent is a strong one, it allows for the math to be factored in a computationally efficient way.

Another distinguishing factor that sets the KFAC method apart from the rest of the optimization methods looked at thus far, is that the KFAC method can estimate the curvature matrix from a lot of data by using an online exponentially-decayed average, as opposed to relatively small-sized fixed mini-batches used by the other stochastic methods. This allows for curvature estimates to depend on much more data than can be reasonably processed in a single mini-batch.

Notably, for methods like LM which deals with an approximate Hessian directly or HF which deal with the exact Fisher indirectly via matrix-vector products, such a scheme would be impossible to implement efficiently, as the exact Fisher matrix (or Generalized-Gauss-Newton) seemingly cannot be summarized using a compact data structure whose size is independent of the amount of data used to estimate it. Indeed, it seems that the only representation of the exact Fisher which would be independent of the amount of data used to estimate it would be an explicit $n \times n$ matrix (which is far too big to be practical). Because of this, HF and related methods must base their curvature estimates only on subsets of data that can be reasonably processed all at once, which limits their effectiveness in the stochastic optimization regime [11].

A.7. Learning Rate Line Search

Algorithm 2 Learning Rate Line Search

Initialize $lrs = (1e-6, 9)$ with step size of 0.125, empty list $loss_list = []$

for $lr \in lrs$ **do**

Update network: $\theta^{(t+1)} = \theta^{(t)} + lr * v_{gn}$
 Evaluate network: $loss_{lr} = f(\theta^{(t+1)})$
 Add current loss to list: $loss_list.append(loss_{lr})$
 Reset network to θ^t

end

$lr_best = \text{argmin}(loss_list)$

Update network with lr of least loss: $\theta^{(t+1)} = \theta^{(t)} + lr_best * v_{gn}$

Appendix B. Experimental Setup

For the regression task, we fit the Sine function with a simple two-layer fully-connected MLP, ELU as the non-linear activation function, and mean square error as the loss function. For the classification task, we use the MNIST datasets, with 28×28 resolution, 60000 training samples and 10000 testing samples. We train a convolutional NN with two convolution layers, and Randomized Leaky ReLU as the non-linear activation function, two max-pooling layers and one fully-connected layers. We compare LM method with popular first-order methods (SGD, Adam) and second-order methods (L-BFGS, Hessian-Free, KFAC). We use the best hyper-parameters per optimizer. See appendix for more details.

We set the hyper-parameters to be the best for each optimizer. For SGD we set the learning rate (lr) to be 0.05, with momentum=0.9 and weight decay= 5×10^{-4} , for Adam, we set the lr to be 0.01, for L-BFGS, we set the lr to 0.005, for Hessian Free, we set the lr to be 0.1, for KFAC, we set the lr to 0.05. For LM, we set the initial lr to be 1 and initial λ to 1 as well. For the full sized MNIST problem we use KFAC with a lr of 1, Adam with a lr of 0.05, SGD with a lr of 0.1. Keeping the rest of the learning rates the same.

Appendix C. Further Experiments

We see that LM can learn the regression problem with a single iteration compared to the results of Adam and KFAC.




Figure 3: Comparison of optimizers on simple regression problem a) LM b) KFAC c) Adam

To then see how the optimizers would do with a limited dataset, we switched the train and the test set and reran the algorithms. We again found that the LM method outperformed the rest, both in the few number of iterations it took to reach a high accuracy and overall accuracy at the end of an epoch.

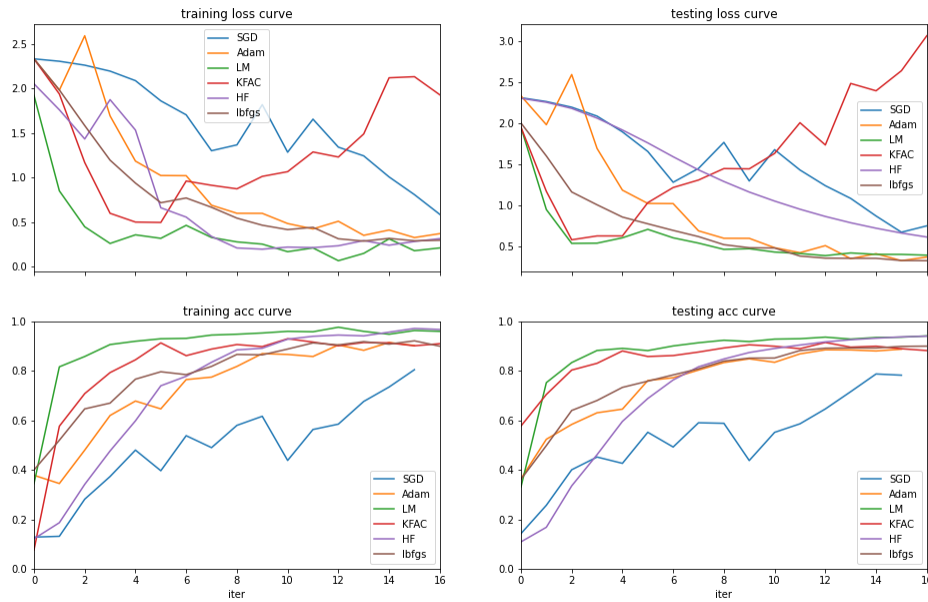


Figure 4: CNN on Full MNIST 1 Epoch Train and Test set switched

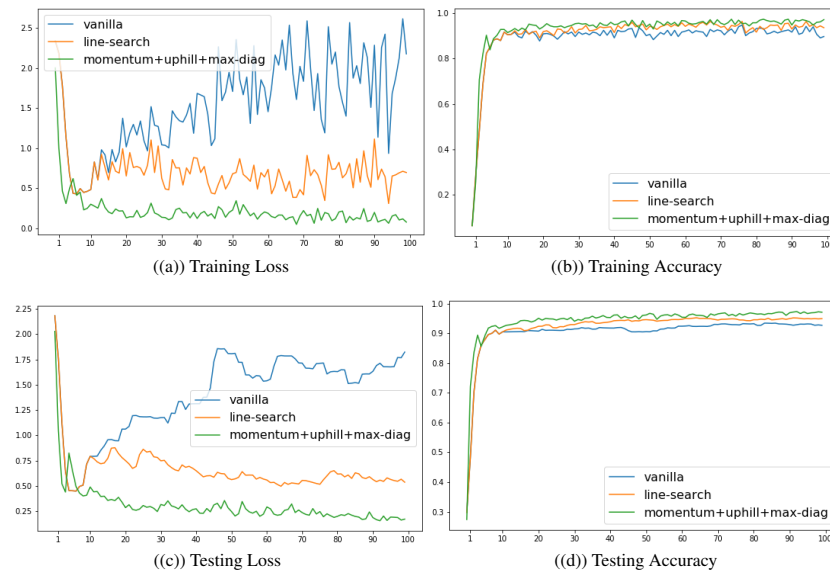


Figure 5: Ablation Study: CNN on Full MNIST (100 Epochs)

Appendix D. Future Work

Very recently [14] has shown that one can efficiently learn from large datasets by selecting subsets based on curvature information to train on. In the experiments they see that using curvature informed optimizers can significantly boost performance. We believe by applying our LM method or PSGD to optimize networks with these AdaCore subsets, we can achieve better generalization compared to first order methods.

Furthermore very recently [13] sees the use of second order methods to greatly benefit optimization of causal models. We believe there is potential to combine ideas from this paper with more efficient methods like PSGD and apply them to causal models and many other machine learning frameworks.