

Second-order Optimization for Non-convex Machine Learning: an Empirical Study

Peng Xu *

Fred Roosta †

Michael W. Mahoney ‡

Abstract

While first-order optimization methods such as SGD are popular in machine learning (ML), they come with well-known deficiencies, including relatively-slow convergence, sensitivity to the settings of hyper-parameters such as learning rate, stagnation at high training errors, and difficulty in escaping flat regions and saddle points. These issues are particularly acute in highly non-convex settings such as those arising in neural networks. Motivated by this, there has been recent interest in second-order methods that aim to alleviate these shortcomings by capturing curvature information. In this paper, we report detailed empirical evaluations of a class of Newton-type methods, namely sub-sampled variants of trust region (TR) and adaptive regularization with cubics (ARC) algorithms, for non-convex ML problems. In doing so, we demonstrate that these methods not only can be computationally competitive with hand-tuned SGD with momentum, obtaining comparable or better generalization performance, but also they are highly robust to hyper-parameter settings. Further, we show that the manner in which these Newton-type methods employ curvature information allows them to seamlessly escape flat regions and saddle points.

1 Introduction

The large-scale nature of many modern ML problems poses computational challenges which have rendered many classical optimization methods (developed for scientific computing and other related areas) inefficient or inapplicable. In this light, first-order optimization methods, such as SGD and its variants, have been the workhorse for ML applications due to their simplicity and versatility. On the other hand, since they consider only first-order gradient information, these methods come with well known deficiencies. These include

relatively-slow convergence and sensitivity to hyper-parameter settings such as learning rate. Furthermore, without dedicated “baby-sitting” of these methods, they can stagnate at high training loss and have difficulty in escaping saddle points and flat regions. These deficiencies are particularly problematic in highly non-convex ML problems such as those that arise in neural network applications. By incorporating curvature information, second-order optimization methods hold the promise to solve these well-known deficiencies. When implemented naïvely, however, second-order methods are clearly not computationally competitive with first order alternatives. This has unfortunately led to the conventional wisdom that second-order methods are inappropriate for large-scale ML applications.

Within the scientific computing community, more so than the ML community, it is well-known that not only stochastic Newton-type methods in general, and Gauss-Newton in particular, can be made scalable [19], but more importantly, and unlike first-order methods, they are also very *resilient* to a variety of adversarial effects [20, 25]. Perhaps the most well known example is their resilience to *ill-conditioning*. A subtle, yet potentially more severe, example is that the success of most first-order methods is tightly intertwined with *fine-tuning* (often many) *hyper-parameters*, most importantly the step size [1]. It is rare that these methods exhibit acceptable performance on first try, and it often takes many trials and errors before one can see reasonable results. In fact, the “true training time”, which almost always includes the time it takes to appropriately tune these parameters, can be frustratingly long. This sort of brute force hyper-parameter tuning is naturally computationally as well as financially expensive. In contrast, second-order optimization algorithms involve much less parameter tuning, and are less sensitive to specific choices of their hyper-parameters [1]. A third example that has received attention due to the popularity of non-convex deep learning problems [9, 14] has to do with avoiding (possibly degenerate) *saddle points* and finding a *local minimum*. While some first-order algorithms can guarantee convergence to an ap-

*Institute for Computational and Mathematical Engineering, Stanford University, Email: pengxu@stanford.edu

†School of Mathematics and Physics, University of Queensland, Brisbane, Australia, and International Computer Science Institute, Berkeley, USA, Email: fred.roosta@uq.edu.au

‡International Computer Science Institute and Department of Statistics, University of California at Berkeley, Email: mmahoney@stat.berkeley.edu

proximate *second-order critical point*¹, e.g., [12], the vast majority of them lack such performance guarantees. Instead, their convergence can, at best, only be ensured to *first-order critical points*, which include saddle points. It has been argued that converging to saddle points can be undesirable for obtaining low generalization errors [6, 8]. In addition, important cases have been demonstrated where SGD stagnates at such high training error points [11]. In contrast, employing the curvature information in the form of Hessian, in addition to the advantages mentioned above, can help with achieving convergence to second-order criticality.

Our Objective Here, we aim to provide an empirical evaluation of variants of Newton-type methods for non-convex ML problems, and study whether they can address, in a computationally competitive manner, the main aforementioned challenges associated with first order methods, i.e., relatively-slow convergence, sensitivity to hyper-parameter settings, stagnation, and entrapment near saddle points. To do so, we exploit recent theoretical developments in [24], and focus on variants of trust-region [7, 24] and (adaptive) cubic regularization [3, 24] algorithms, in which the Hessian is suitably approximated. More specifically, in the context of several non-convex machine learning applications, we study the empirical performance of *sub-sampled* versions of these algorithms and set out to paint a more complete picture of their *practical impact*. In the process, we highlight the above-mentioned shortcomings of first-order methods and, in their light, assess the *effectiveness of employing curvature information as remedy*.

Main Questions To accomplish our objective, we set out to answer the following questions, all of which are designed to shed light on various *practical* aspects of this paper’s underlying thesis.

Q.1 (*Computational Efficiency*) Can these sub-sampled Newton-type methods be *computationally efficient* enough to be competitive with hand-tuned SGD with momentum?

Q.2 (*Robustness to Hyper-parameters*) Does the performance of such Newton-type methods exhibit *robustness to hyper-parameter tuning*?

Q.3 (*Escaping Saddle Point*) Does employing Hessian information help with *avoiding saddle points* and converging to *lower training errors* in highly non-convex problems?

Q.4 (*Generalization Performance*) Can second-order methods be beneficial for obtaining *low generalization*

error in machine learning problems?

Q.5 (*Benefits of Sub-sampling*) How does *sub-sampling* in general, and various sub-sampling *schemes* in particular, affect the performance of TR and ARC methods?

Q.6 (*Comparison Among Second-Order Methods*) How do sub-sampled TR and ARC *compare with other second-order methods* like (sub-sampled) Gauss-Newton (GN) and limited-memory BFGS (L-BFGS) methods?

Methodology To minimize the effect of many confounding factors involved in almost all empirical evaluations, rather than chasing to beat the state-of-the-art performances on benchmark tasks, we focus our attention on two simple, yet illustrative, classes of non-convex ML problems, i.e., *multi-layer perceptron (MLP) networks* and *non-linear least squares (NLS)*. More specifically, to address Q.1, Q.2, Q.3, and Q.4, which are concerned with comparisons among first and second-order methods, we consider training (deep) MLP networks with no additional bells and whistles; and to address Q.5 and Q.6, which involve comparisons among various second-order methods, we consider a simpler NLS problem. These questions, references to their theoretical studies, and sections of this paper involving their empirical treatment with relevant figures, are all gathered in Table 1.

In Section 3.2, we address Q.1–Q.4. In particular, in the context of image classification (Section 3.2.1) and deep auto-encoder (Section 3.2.2), we study the efficiency of sub-sampled TR method, which incorporates inexactness in both Hessian and the sub-problem solver, as compared with hand-tuned SGD with momentum. We answer item Q.1 by measuring the convergence speed over total computational cost. We then treat Q.2 by demonstrating the resiliency/sensitivity of various algorithms with respect to their main hyper-parameter. We do this through multiple simulations of all these examples with several choices of the main hyper-parameter for each algorithm. The main hyper-parameter is understood as the one for which, in practice, there is no “typical” value. For SGD with momentum, the *learning rate* is considered as the main parameter, since the momentum parameter is typically set to ≈ 0.9 . For trust region, the main hyper-parameter is the *initial trust region*, as there are typical values for other parameters of the algorithm. For Q.3, we consider various initialization schemes, including those that are close to high-level saddle points. We study the behavior of SGD-based methods near such regions and evaluate whether an appropriate use of curvature can indeed help with escaping such undesirable saddle points and making continued progress towards areas with lower training error. To address Q.4, we present test performances for

¹where gradient is very small and Hessian is almost positive semi-definite.

Table 1: Fundamental questions related to the underlying thesis of this paper.

| Question | Theoretical Results | Empirical Results | Figure |
|------------|------------------------|---------------------------|--------|
| Q.1 | [24, Theorems 4-6] | Sections 3.2.1 and 3.2.2 | 1,2,3 |
| Q.2 | [24, Theorems 1-6] | Sections 3.2.1 and 3.2.2 | 1,2,3 |
| Q.3 | – | Sections 3.2.1 and 3.2.2 | 1,2,3 |
| Q.4 | – | Sections 3.2.1, and 3.2.2 | 1,2,3 |
| Q.5 | [24, Lemmas 16 and 17] | Section 3.3 | 5 |
| Q.6 | – | Section 3.3 | 5 |

all our experiments and assess the generalization errors obtained by all the methods considered.

In Section 3.3, we turn our attention to **Q.5–Q.6**, where we consider the NLS problem arising from binary classification task with least-squares loss. On several real datasets, we then demonstrate the effect of various sub-sampling strategies on the performance of sub-sampled TR and ARC methods, as compared with other second-order algorithms.

2 Background

In this section, we give a brief review of the general formulation of the optimization problems considered in our study as well as the Newton-type algorithms in question.

2.1 Non-Convex Finite-Sum Minimization Following many machine learning applications, we consider the “finite-sum” optimization problem

$$(P1) \quad \min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) \triangleq \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{x}),$$

where each $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is a smooth but possibly non-convex function. Many machine learning and scientific computing applications involve optimization problems of the form (P1) where each f_i is a loss (or misfit) function corresponding to i^{th} observation (or measurement), e.g., see [22]. In particular, in machine learning applications, F in (P1) corresponds to the *empirical risk* [21] and the goal of solving (P1) is to obtain a solution with small generalization error, i.e., high predictive accuracy on “unseen” data.

2.2 Main Algorithms: Sub-Sampled TR and ARC Arguably, line-search is the most straightforward approach for *globalization* of many Newton-type algorithms. However, near saddle points where the gradient magnitude can be small, traditional line search methods can be very ineffective and in fact produce iterates that can get stuck at a saddle point [17]. Trust region [7] and cubic regularization methods [3, 4] are two elegant globalization alternatives that, specially recently, have attracted much attention.

In large-scale non-convex settings, where the application of exact Hessian can be computationally infeasible, recently, [24] theoretically studied the variants of TR and ARC algorithms in which the Hessian is suitably approximated. For ease of reference, Appendix A gathers the descriptions of TR and ARC algorithms, henceforth referred to as Algorithms 1 and 2, respectively. Iterations of these algorithms involve the following sub-problems:

$$(2.1a)$$

TR Sub-Problem:

$$\mathbf{s}_t \approx \underset{\|\mathbf{s}\| \leq \Delta_t}{\operatorname{argmin}} m_t(\mathbf{s}) \triangleq \langle \nabla F(\mathbf{x}_t), \mathbf{s} \rangle + \frac{1}{2} \langle \mathbf{s}, \mathbf{H}_t \mathbf{s} \rangle,$$

$$(2.1b)$$

ARC Sub-Problem:

$$\mathbf{s}_t \approx \underset{\mathbf{s} \in \mathbb{R}^d}{\operatorname{argmin}} m_t(\mathbf{s}) \triangleq \langle \nabla F(\mathbf{x}_t), \mathbf{s} \rangle + \frac{1}{2} \langle \mathbf{s}, \mathbf{H}_t \mathbf{s} \rangle + \frac{\sigma_t}{3} \|\mathbf{s}\|^3,$$

where \mathbf{H} is an approximation of the the exact Hessian. **Hessian Sub-Sampling** We consider (P1) in large-scale regime where $n, d \gg 1$. In such settings, the mere evaluations of the Hessian and the gradient increase linearly in n . As studied in [24], given a sampling distribution $\{p_i\}_{i=1}^n$ over the set of indices $\{1, 2, \dots, n\}$, the sub-sampled Hessian has the form

$$(2.2) \quad \mathbf{H}(\mathbf{x}) \triangleq \frac{1}{n|\mathcal{S}|} \sum_{j \in \mathcal{S}} \frac{1}{p_j} \nabla^2 f_j(\mathbf{x}),$$

where $\mathcal{S} \subset \{1, 2, \dots, n\}$ is a random sample collection. When $|\mathcal{S}| \ll n$, sub-sampling can offer significant computational savings; see [2, 20, 25] for examples of studies in convex settings. Recently [24] theoretically showed that randomized sub-sampling can also be seamlessly extended to non-convex settings.

[24] further showed that, in certain settings, one can construct more “informative” distributions over the indices in $\{1, 2, \dots, n\}$, as opposed to oblivious uniform sampling. Indeed, it is typically advantageous to bias the probability distribution towards picking indices corresponding to those f_i ’s which are more *relevant*, in certain sense, in forming the Hessian. One

such setting where this is possible is the finite-sum optimization of the form,

$$(P2) \quad \min_{\mathbf{x} \in \mathbb{R}^d} F(\mathbf{x}) \triangleq \frac{1}{n} \sum_{i=1}^n f_i(\mathbf{a}_i^T \mathbf{x}),$$

for some given data vectors $\{\mathbf{a}_i\}_{i=1}^n \subset \mathbb{R}^d$. Problems of the form (P2), which is a special case of (P1), arise often in many machine learning problems [21], e.g., logistic regression with least squares loss as in Example 3.3. For problems of the form (P2), one can construct a more informative sampling scheme by considering the sampling distribution as $p_i = |f_i''(\mathbf{a}_i^T \mathbf{x})| \|\mathbf{a}_i\|^2 / (\sum_{j=1}^n |f_j''(\mathbf{a}_j^T \mathbf{x})| \|\mathbf{a}_j\|^2)$. In [24], it was shown that, in order to obtain similar approximation guarantee, such nonuniform sub-sampling scheme yields a sample size which can be significantly smaller than that required by oblivious uniform sampling.

Inexact Sub-problem Solver In Algorithms 1 and 2, it is imperative that the sub-problems (2.1a) and (2.1b), respectively, are solved only approximately. Indeed, in large-scale problems, where the exact solution of sub-problems is a computational bottleneck, this relaxation is crucial; see [24] for precise definitions as well as ways to obtain the approximate solution of the sub-problems (2.1a) and (2.1b).

3 Numerical Experiments

We are now ready to empirically evaluate the performance of the Newton-type methods considered in this paper (and studied theoretically in [24]) in several settings. In particular, we study the answers to Q.1–Q.6 posed at the outset in Section 1 in the context of two simple, yet illustrative, classes of non-convex optimization problems, i.e. (*deep*) *MLP networks* (Section 3.2) and *NLS* (Section 3.3). See the supplementary materials for more experiments.

3.1 General Experimental Settings

Complexity Measure In all of our experiments, we plot various quantities vs. total *number of propagations* [9], which is equivalent to measuring the number of oracle calls of function, gradient and Hessian-vector product. This is so since comparing algorithms in terms of “wall-clock” time can be highly affected by their particular implementation details as well as system specifications. In contrast, counting the number of propagations, as an implementation and system independent unit of complexity, is most appropriate and fair. Specifically, in neural nets, for a given data at the input layer, evaluation of network’s output layer involves one forward propagation. Performing one additional backward propagation gives the corresponding gradient. Each of the Hessian-vector products, required to solve the respective sub-problems of the second-order methods, is

equivalent to two gradient evaluations, i.e., compared to the gradient, it involves one additional forward and backward propagations [18]. Combining the batch size in each iteration, we summarize the number of propagations per iteration for each algorithm in Table 2.

Table 2: Total number of propagations per iteration for various algorithms. “ r ” denotes the number of Hessian-vector products for solving the respective subproblems, and hence can be different for each algorithm. In the “Full” cases (e.g., Full TR), $|\mathcal{S}| = n$.

| SUB-SAMPLED TR,ARC,GN | L-BFGS | SGD |
|-------------------------|--------|------------------|
| $2(n + \mathcal{S} r)$ | $2n$ | $2 \mathcal{S} $ |

In all of the following experiments, to approximately solve the sub-problems (2.1a) and (2.1b), respectively, we use CG-Steihaug method [17] and the generalized Lanczos method [3] with a maximum of 250 Lanczos iterations.

3.2 Hessian-Free Optimization for MLPs In this section, we empirically study the main questions Q.1–Q.4 of Section 1. For this, we consider two simple but non-trivial MLP networks under various settings, namely (a) 1-hidden layer neural networks for image classification, and (b) deep auto-encoder. The empirical performance of the following algorithms, in light of Questions Q.1–Q.4, are compared²:

- (1) *TR Uniform*: Alg 1 with uniform sub-sampling,
- (2) *GN Uniform*: sub-sampled Gauss-Newton method with modifications introduced in [16], and
- (3) *SGD with Momentum*: mini-batch SGD with momentum term [23] and fixed step-size.

Parameter settings For Algorithm 1, we set $\eta_1 = 10^{-4}$, $\eta_2 = 0.8$, $\gamma_1 = 1.2$, and $\gamma_2 = 2$ (these are some values typical used in literature). The sample size used for Hessian sub-sampling is set to 5% of the total training set, i.e., $0.05n$. For momentum SGD, the mini-batch size is chosen as $0.05n$ and the momentum parameter is set as 0.9 (which is typically set in literature). For GN method, we use the same parameter settings as in [16].

3.2.1 One-Hidden Layer Neural Network Here we consider a one-hidden layer neural network for the task of image classification using *cifar10* [13] data set. The hidden layer size is 512, amounting to $d = 1,578,506$. Initializations are done by setting \mathbf{x}_0 to a normalized vector drawn from standard normal distribution and all-zeros vector. Figure 1 shows training loss, training error and test error for all methods.

²We excluded ARC variants for reasons alluded to in Section 4.

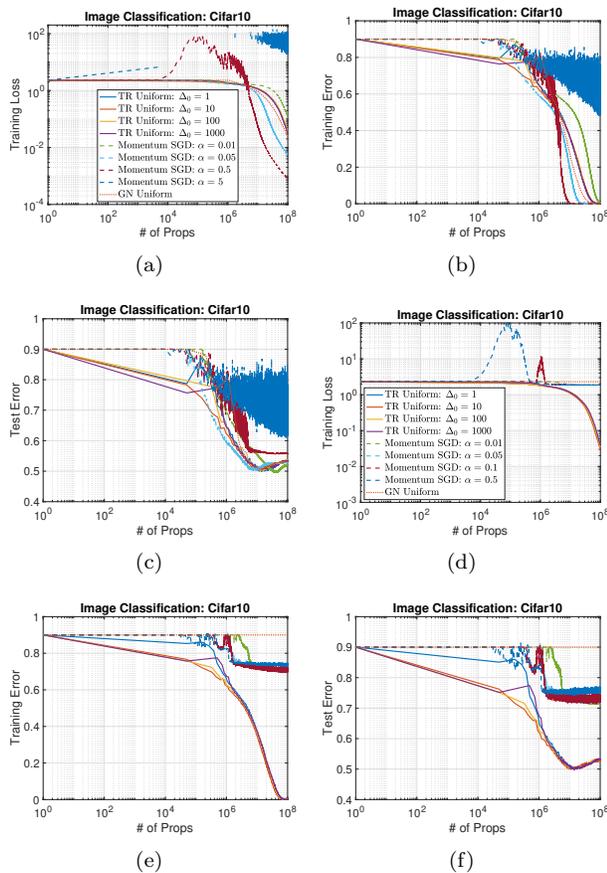


Figure 1: Cifar10 dataset on MLP using normalized random (a,b,c) and zero (d,e,f) initialization. Δ_0 is the initial trust-region radius of Algorithm 1 and α is the step size for SGD with momentum.

In light of Q.1–Q.4, we observe the following:

Re: Q.1 (Computational Efficiency) In Figures 1(a)(b)(c), all algorithms start from a normalized random vector. From Figure 1(a), we observe that in terms of training loss, sub-sampled TR algorithm, though very competitive, can be slightly slower than SGD as long as SGD’s step size is appropriately fine-tuned. However, this does not necessarily translate to better test error; see Figure 1(c). In particular, while all TR runs achieve similar test errors, SGD’s generalization performance does not mirror its training behavior and appears highly “chaotic”, i.e., the step-size that achieved the fastest training (red dashed line), generalizes very poorly (compare red and blue dashed lines).

Re: Q.2 (Robustness to Hyper-parameters) From Figures 1(a)(b)(c), we notice the dependence of SGD’s performance on the choice of its learning-rate. Well-tuned step size can give both fast convergence of training pro-

cess and good generalization performance. Too small a step size, however, can lead to slow convergence, while too large a step size, can cause SGD divergence or poor generalization performance. In contrast, Algorithm 1 with drastically different initial trust-region radii exhibit comparable performances; similar phenomenon are seen in Figures 1(d)(e)(f).

Re: Q.3 (Escaping Saddle Point) In Figure 1(d)(e)(f), all algorithms start from the origin. We can clearly see that SGD and GN easily get trapped at/near saddle points and/or flat regions (it is easy to check the gradients there are extremely small) and can barely make any progress. In contrast, sub-sampled TR, which effectively utilizes the Hessian, seamlessly escapes these regions and makes continued progress.

Re: Q.4 (Generalization Performances) In all initialization schemes, as shown in Figures 1, sub-sampled TR obtains competitive, if not better, generalization performance.

3.2.2 Deep Auto-Encoder Here, we consider the deep auto-encoder problem [9] and use the same model architectures as well as loss functions as in [16]. The dataset and network architectures are given in Table 3. The experiments in Figures 2 and 3 are each done with initialization to a vector drawn from standard normal distribution as well as the all-zeros vector.

Table 3: Datasets for deep auto-encoder experiment.

| Dataset | Size | Encoder Network (# parameters) |
|---------|-------|------------------------------------|
| curves | 20000 | 784-400-200-100-50-25-6 (842, 340) |
| mnist | 60000 | 784-1000-500-250-30 (2, 837, 314) |

Although deep auto-encoder networks here are much more complex than 1-hidden layer network of Section 3.2.1, but we can still make very similar observations as they relate to Q.1–Q.4. In particular, on both datasets, Algorithm 1 converges comparably as fast, or faster than, other methods in terms of number of propagations. More importantly, it exhibits great robustness to its hyper-parameter, the initial trust-region radius, in contrast to SGD with momentum, which is heavily dependent on the choice of step size. Since these are rather complex networks, the optimization landscape is riddled with saddle and/or very flat regions. In this light, SGD and GN algorithms can both, rather easily, get trapped at/near these regions, which is quite contrary to the behavior of Algorithm 1 (Figures 2(b) and 3(b)). In terms of test error, Algorithm 1 is competitive to other methods on some experiments (Figures 2(a) and 3(a)) and clearly outperforms on others (Figures 2(b) and 3(b)).

We further examine two additional initialization strategies: normalized random initial point (Figure

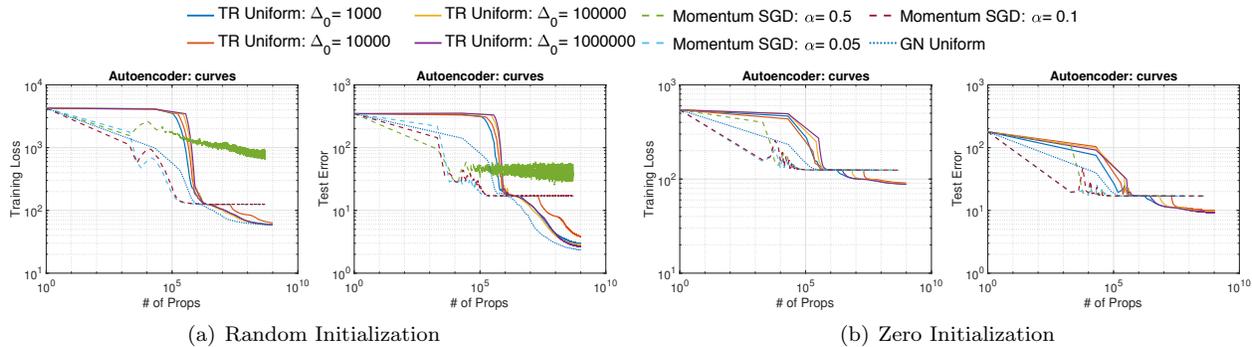


Figure 2: Deep autoencoder on *curves* dataset using random (a) and zero (b) initial point. Δ_0 is the initial trust-region radius of Algorithm 1 and α is the step size for SGD with momentum.

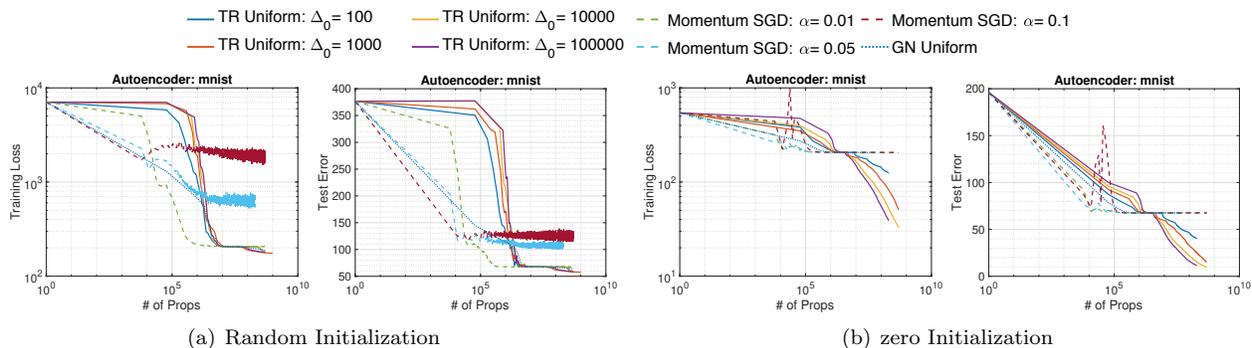


Figure 3: Deep autoencoder on *mnist* dataset using random (a) and zero (b) initial point. Δ_0 is the initial trust-region radius of Algorithm 1 and α is the step size for SGD with momentum.

4(a)) as well as a random vector scaled by 0.25 (Figure 4(b)). Similar observations as earlier can also be made here. In particular, unlike Section 3.2.1, for this problem normalized random initial point (Figure 4(a)) seems to paint a different picture, i.e., SGD with momentum as well as GN all get trapped at high training levels while Algorithm 1 makes continued progress. It is worth mentioning that among all the initialization schemes we considered here, only under the particular scaled random initialization SGD can obtain desirable performance. This further demonstrates the versatility of Algorithm 1.

3.3 Non-Linear Least Squares We now turn to study Questions Q.5 and Q.6. For this, we consider the class of NLS problems, and focus on the task of binary classification with square loss as a concrete instance. Since logistic loss, which is the “standard” loss used in this task, leads to a convex problem, we use square loss to obtain a non-convex objective. More specifically, suppose we are given training data $\{\mathbf{x}_i, y_i\}_{i=1}^n$, where $\mathbf{x}_i \in \mathbb{R}^d$ and $y_i \in \{0, 1\}$ are, respectively, i^{th} feature vector and the corresponding

label. Consider minimizing the empirical risk problem $\min_{\mathbf{w} \in \mathbb{R}^d} \sum_{i=1}^n (y_i - \phi(\langle \mathbf{x}_i, \mathbf{w} \rangle))^2 / n$, where $\phi(z)$ is the sigmoid function, i.e., $\phi(z) = 1/(1 + e^{-z})$. Since this is an example of (P2), we can apply both uniform and non-uniform sampling schemes.

Table 4 summarizes the real data sets used for the experiments of this section. All datasets are from LIBSVM library [5]. The following algorithms are compared (exact Hessian refers to $\mathbf{H}_t = \nabla^2 F(\mathbf{x}_t)$):

- i. *TR Full/Uniform/Non-Uniform*: Algorithm 1 with full or uniform/non-uniform estimation of Hessian, respectively,
- ii. *ARC Full/Uniform/Non-Uniform*: Algorithm 2 with full or uniform/non-uniform estimation of Hessian, respectively,
- iii. *GN Full/Uniform/Non-Uniform*: GN with full or uniform/non-uniform estimation of Hessian, respectively,
- iv. *LBFGS-100*: the standard L-BFGS method [15] with history size 100 and using line-search.

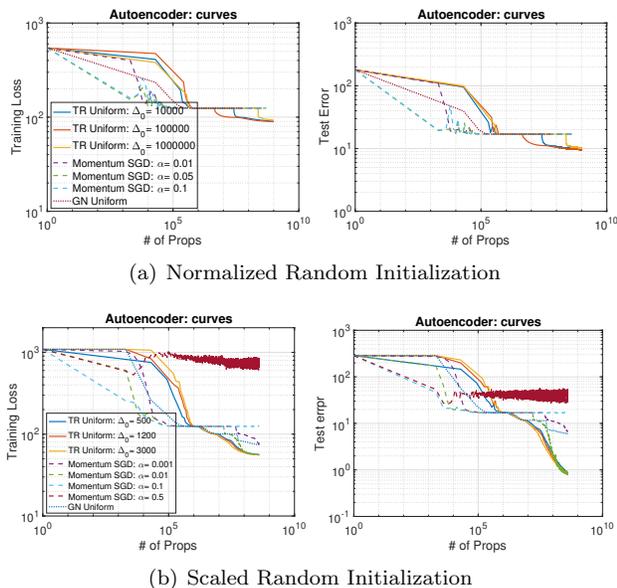


Figure 4: Deep autoencoder on `curves` dataset using (a) normalized and (b) scaled random initial points. Δ_0 is the initial trust-region radius of Algorithm 1 and α is the step size for SGD.

Table 4: Datasets used in binary linear classification.

| DATA | n | d | Test Size |
|----------------------|---------|-----|-----------|
| <code>covtype</code> | 464,810 | 54 | 116,202 |
| <code>mnist</code> | 60,000 | 784 | 10,000 |

Parameter settings For both Algorithms 1 and 2, we set η_1, η_2, γ_1 and γ_2 , the same as Section 3.2. The sampling ratios, i.e., $|\mathcal{S}|/n$, for uniform and non-uniform sampling are set to 1% and 0.1%, respectively. For all datasets, we set $\Delta_0 = 10$ for Algorithm 1 and $\sigma_0 = 10^{-4}$ for Algorithm 2. For GN, we do not regularize Hessian as in [16]. Figure 5 gathers all comparison results of this section.

Re: Q.5 (Benefits of Sub-Sampling) Overall, both Algorithms 1 and 2 compare well with the classical TR and ARC methods. In fact, sub-sampling can, at times, help increase the efficiency, e.g., TR variants for `covtype2`. However, with too small a sample, the performance can hurt; ARC Full vs. ARC Non-Uniform and ARC Uniform for `covtype2` and `mnist2`, respectively. The benefits of non-uniform sampling over uniform alternative are far more pronounced in the performance of Algorithm 2 than Algorithm 1. This can be attributed mainly to their respective sub-problem solvers in terms of total number of performed Hessian-vector products. In particular, CG-Steihaug used for the sub-problem (2.1a) of Algorithm 1 typically terminates in a handful

of iterations whereas the generalized Lanczos method for solving the sub-problem (2.1b) of Algorithm 2 usually exhausts the allotted 250 iterations.

Re: Q.6 (Comparison Among Second-Order Methods) One can observe the consistent poor performances of L-BFGS-100 (green dotted lines) methods on all datasets, in particular with all 1’s initialization vector. This is rather expected as contrary to popular belief, BFGS is not quite a “full-fledged” second-order method. Indeed, BFGS merely employs first-order information, i.e. gradients, to approximate the curvature, and starting from all 1’s vector, L-BFGD cannot capture enough curvature information to navigate its way out of this region effectively. Gauss-Newton (dash lines), which has been specifically designed to effectively solve NLS problems, performs very well with random initialization. Starting from all 1’s vector, however, where the gradient is very small, GN performs poorly. This is also expected because GN, similar in spirit to BFGS, does not fully utilize the Hessian information. In particular, in exchange for obtaining a positive definite approximation matrix, GN completely ignores the information from *negative curvature*, which is critical for allowing to escape from regions with small gradient. We can also observe that ARC is consistently no better than TR. This is an empirical evidence that the optimal worst-case complexity of ARC [4, 24], though theoretically highly interesting, might be hard to observe in many practical settings.

4 Conclusion

In this paper, we aimed at painting a more complete picture of the practical advantages of Newton-type algorithms in general, and sub-sampled variants of trust-region and adaptive cubic regularization methods in particular, as compared with first-order alternatives. In the context of (deep) multi-layer perceptron networks as well as non-linear least squares, two simple, yet illustrative, non-convex machine learning applications, by making the following observations, we empirically attempted to make a case for the application of such second-order methods for machine learning.

A.1 (Computational Efficiency) The randomized sub-sampling approaches described here, and studied in detail in [24], can effectively make Newton-type methods *computationally efficient* enough to be competitive with popular first-order methods, widely used in machine learning, e.g., SGD with momentum. This is indeed due to the amortized combination of (1) low per-iteration cost offered by randomized sampling, and (2) small number of overall iterations due to the application of curvature information.

A.2 (Robustness to Hyper-parameters) In contrast to

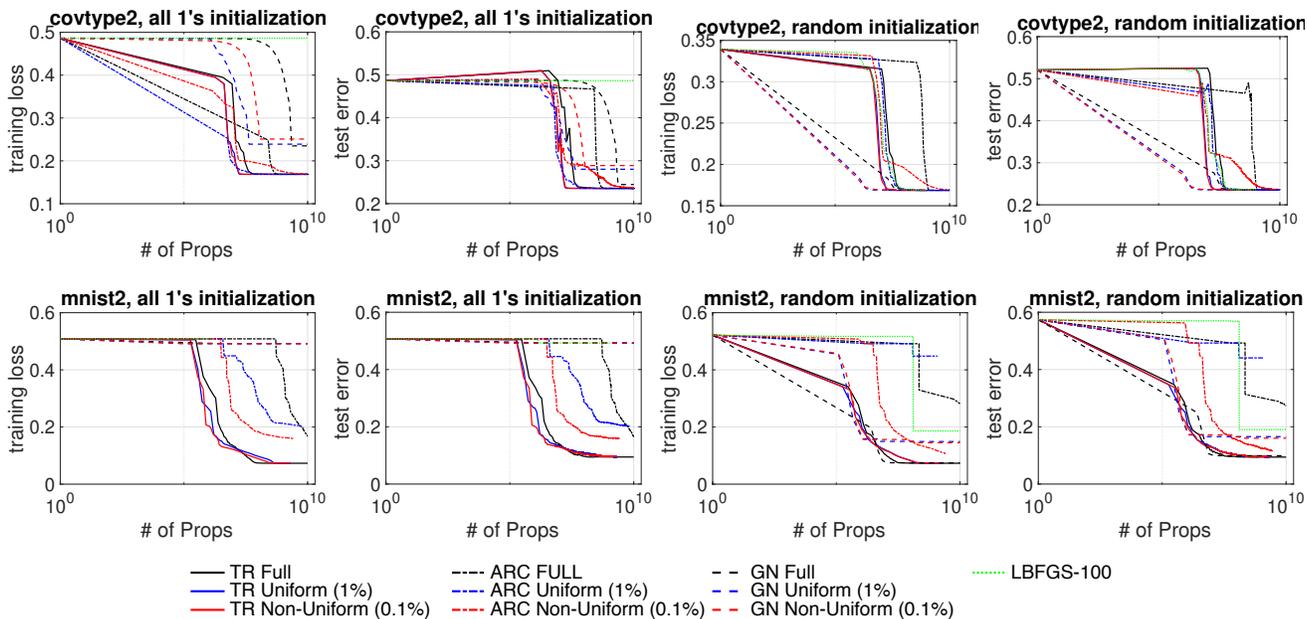


Figure 5: Binary Classification on different datasets with all 1's initialization (left two columns) and random initialization (right two columns). Each row corresponds to a dataset. The x-axis is the number of propagations in log-scale.

first-order algorithms whose performance is greatly affected by the choice of hyper-parameters, most notably *step-size*, the performance of the proposed Newton-type methods exhibit great *robustness* to such *parameter tuning*.

A.3 (Escaping Saddle Point) A greatly beneficial advantage of employing Hessian information is that it allows for such Newton-type algorithms, unlike many first-order alternatives, to seamlessly escape regions near *saddle points*.

A.4 (Generalization Performance) Second-order methods prove beneficial for the downstream machine learning objective of obtaining *good generalization error*. In particular, one can obtain very good levels of prediction accuracy only after a few iterations of such methods. This is highly beneficial in, say, distributed settings where the communication across the network is the main computational bottleneck.

A.5 (Benefits of Sub-sampling) On several real datasets, we validated the effectiveness of sub-sampled Newton-type methods, as compared with classical versions, in speeding up computations. Also, the advantages of non-uniform sampling over the oblivious uniform alternative was verified.

A.6 (Comparison Among Second-Order Methods) There are clear advantages in using (sub-sampled) TR and ARC methods over other second-order alternatives,

e.g., L-BFGS and GN, in terms of effective exploitation of curvature.

For the examples of Section 3.2, despite the best of our efforts, we were unable to obtain the expected performance of Algorithm 2 using a variety of implementations, e.g., our own hand-written code as well as some existing packages such as GALAHAD [10]. We believe that this is tightly connected to the choice of the sub-problem solver in all these implementations. Since we were unable to pinpoint the source of the problem, we did not include Algorithm 2 in examples of Section 3.2.

Finally, we acknowledge that, although we presented various experiments, the empirical study of methods such as the ones considered here, takes more than a single “proof-of-concept” paper, and the results presented here should be viewed as merely a glimpse into their various properties.

Acknowledgment FR and MM acknowledge the support by DARPA D3M (FA8750-17-2-0122) and Cray/77000. FR was partially supported by the Australian Research Council grants CE140100049 and DE180100923.

References

- [1] Albert S Berahas, Raghu Bollapragada, and Jorge Nocedal. “An Investigation of Newton-Sketch and Subsampled Newton Methods”. In: *arXiv preprint arXiv:1705.06211* (2017).

- [2] Raghu Bollapragada, Richard Byrd, and Jorge Nocedal. “Exact and Inexact Subsampled Newton Methods for Optimization”. In: *arXiv preprint arXiv:1609.08502* (2016).
- [3] Coralía Cartis, Nicholas IM Gould, and Philippe L Toint. “Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results”. In: *Mathematical Programming* 127.2 (2011), pp. 245–295.
- [4] Coralía Cartis, Nicholas IM Gould, and Philippe L Toint. “Adaptive cubic regularisation methods for unconstrained optimization. Part II: worst-case function-and derivative-evaluation complexity”. In: *Mathematical programming* 130.2 (2011), pp. 295–319.
- [5] Chih-Chung Chang and Chih-Jen Lin. “LIBSVM: A library for support vector machines”. In: *ACM Transactions on Intelligent Systems and Technology* 2 (3 2011). Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>, 27:1–27:27.
- [6] Anna Choromanska, Mikael Henaff, Michael Mathieu, Gérard Ben Arous, and Yann LeCun. “The loss surfaces of multilayer networks”. In: *Artificial Intelligence and Statistics*. 2015, pp. 192–204.
- [7] Andrew R Conn, Nicholas IM Gould, and Philippe L Toint. *Trust region methods*. SIAM, 2000.
- [8] Yann N Dauphin, Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Surya Ganguli, and Yoshua Bengio. “Identifying and attacking the saddle point problem in high-dimensional non-convex optimization”. In: *Advances in neural information processing systems*. 2014, pp. 2933–2941.
- [9] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.
- [10] Nicholas IM Gould, Dominique Orban, and Philippe L Toint. “GALAHAD, a library of thread-safe Fortran 90 packages for large-scale nonlinear optimization”. In: *ACM Transactions on Mathematical Software (TOMS)* 29.4 (2003), pp. 353–372.
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 770–778.
- [12] Chi Jin, Rong Ge, Praneeth Netrapalli, Sham M Kakade, and Michael I Jordan. “How to Escape Saddle Points Efficiently”. In: *arXiv preprint arXiv:1703.00887* (2017).
- [13] Alex Krizhevsky and Geoffrey Hinton. “Learning multiple layers of features from tiny images”. In: (2009). Data available at <https://www.cs.toronto.edu/~kriz/cifar.html>.
- [14] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *Nature* 521.7553 (2015), pp. 436–444.
- [15] Dong C. Liu and Jorge Nocedal. “On the limited memory BFGS method for large scale optimization”. In: 45 (1989), pp. 503–528.
- [16] James Martens. “Deep learning via Hessian-free optimization”. In: *International Conference on Machine Learning (ICML)*. 2010.
- [17] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [18] Barak A Pearlmutter. “Fast exact multiplication by the Hessian”. In: *Neural computation* 6.1 (1994), pp. 147–160.
- [19] Farbod Roosta-Khorasani, Kees van den Doel, and Uri Ascher. “Stochastic algorithms for inverse problems involving PDEs and many measurements”. In: *SIAM J. Scientific Computing* 36.5 (2014), S3–S22.
- [20] Farbod Roosta-Khorasani and Michael W Mahoney. “Sub-sampled Newton methods”. In: *Mathematical Programming* 174.1-2 (2019), pp. 293–326.
- [21] Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- [22] Suvrit Sra, Sebastian Nowozin, and Stephen J Wright. *Optimization for machine learning*. MIT Press, 2012.
- [23] Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. “On the importance of initialization and momentum in deep learning”. In: *International conference on machine learning*. 2013, pp. 1139–1147.
- [24] Peng Xu, Farbod Roosta-Khorasani, and Michael W Mahoney. “Newton-type methods for non-convex optimization under inexact Hessian information”. In: *Mathematical Programming* (2019). doi:10.1007/s10107-019-01405-z.
- [25] Peng Xu, Jiyan Yang, Farbod Roosta-Khorasani, Christopher Ré, and Michael W Mahoney. “Sub-sampled newton methods with non-uniform sampling”. In: *Advances in Neural Information Processing Systems*. 2016, pp. 3000–3008.