

---

# Newton-LESS: Sparsification without Trade-offs for the Sketched Newton Update

---

**Michał Dereziński**

Department of Statistics  
University of California, Berkeley  
mderezin@berkeley.edu

**Jonathan Lacotte**

Department of Electrical Engineering  
Stanford University  
lacotte@stanford.edu

**Mert Pilanci**

Department of Electrical Engineering  
Stanford University  
pilanci@stanford.edu

**Michael W. Mahoney**

ICSI and Department of Statistics  
University of California, Berkeley  
mmahoney@stat.berkeley.edu

## Abstract

In second-order optimization, a potential bottleneck can be computing the Hessian matrix of the optimized function at every iteration. Randomized sketching has emerged as a powerful technique for constructing estimates of the Hessian which can be used to perform approximate Newton steps. This involves multiplication by a random sketching matrix, which introduces a trade-off between the computational cost of sketching and the convergence rate of the optimization algorithm. A theoretically desirable but practically much too expensive choice is to use a dense Gaussian sketching matrix, which produces unbiased estimates of the exact Newton step and which offers strong problem-independent convergence guarantees. We show that the Gaussian sketching matrix can be drastically sparsified, significantly reducing the computational cost of sketching, without substantially affecting its convergence properties. This approach, called Newton-LESS, is based on a recently introduced sketching technique: LEverage Score Sparsified (LESS) embeddings. We prove that Newton-LESS enjoys nearly the same problem-independent local convergence rate as Gaussian embeddings, not just up to constant factors but even down to lower order terms, for a large class of optimization tasks. In particular, this leads to a new state-of-the-art convergence result for an iterative least squares solver. Finally, we extend LESS embeddings to include uniformly sparsified random sign matrices which can be implemented efficiently and which perform well in numerical experiments.

## 1 Introduction

Consider the task of minimizing a twice-differentiable convex function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$ :

$$\text{find } \mathbf{x}^* = \underset{\mathbf{x} \in \mathbb{R}^d}{\operatorname{argmin}} f(\mathbf{x}).$$

One of the most classical iterative algorithms for solving this task is the Newton's method, which takes steps of the form  $\mathbf{x}_{t+1} = \mathbf{x}_t - \mu_t \nabla^2 f(\mathbf{x}_t)^{-1} \nabla f(\mathbf{x}_t)$ , and which leverages second-order information in the  $d \times d$  Hessian matrix  $\nabla^2 f(\mathbf{x}_t)$  to achieve rapid convergence, especially locally as it approaches the optimum  $\mathbf{x}^*$ . However, in many settings, the cost of forming the exact Hessian is prohibitively expensive, particularly when the function  $f$  is given as a sum of  $n \gg d$  components, i.e.,  $f(\mathbf{x}) = \sum_{i=1}^n f_i(\mathbf{x})$ . This commonly arises in machine learning when  $f$  represents the training loss over a

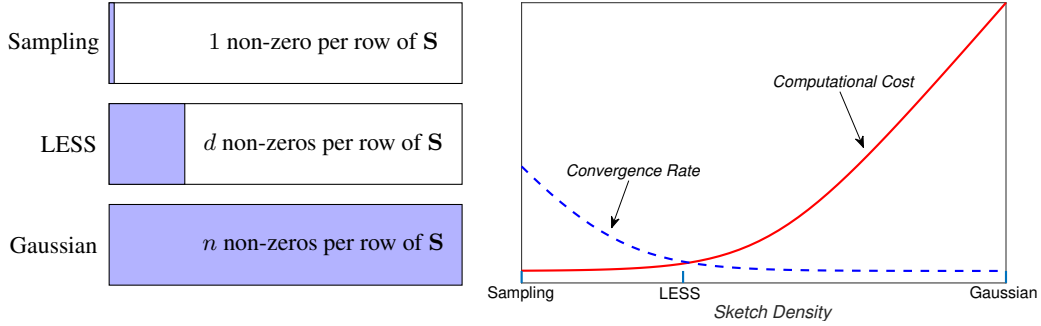


Figure 1: The effect of the density of  $m \times n$  sketching matrix  $\mathbf{S}$  applied to an  $n \times d$  matrix  $\mathbf{A}$  (with  $d, m \ll n$ ) on the convergence rate of Newton Sketch and the computational cost of constructing the Hessian estimate. LESS embeddings “interpolate” between Sub-Sampled Newton methods and Gaussian Newton Sketches, achieving a “sweet spot” in the computation-per-iteration versus number-of-iterations tradeoff.

dataset of  $n$  elements, as well as in solving semi-definite programs, portfolio optimization, and other tasks. In these contexts, we can represent the Hessian via a decomposition  $\nabla^2 f(\mathbf{x}) = \mathbf{A}_f(\mathbf{x})^\top \mathbf{A}_f(\mathbf{x})$ , where  $\mathbf{A}_f(\mathbf{x})$  is a tall  $n \times d$  matrix, which can be easily formed, and the main bottleneck is the matrix multiplication which takes  $O(nd^2)$  arithmetic operations. To avoid this bottleneck, many randomized second-order methods have been proposed which use a Hessian estimate in place of the exact Hessian (e.g., [BCNN11, EM15, ABH17, RKM19]). This naturally leads to a trade-off between the per-iteration cost of the method and the number of iterations needed to reach convergence. We develop Newton-LESS, a randomized second-order method which eliminates the computational bottleneck while minimizing the convergence trade-offs.

An important family of approximate second-order methods is known as the Newton Sketch [PW17]:

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t - \mu_t (\mathbf{A}_f(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_t^\top \mathbf{S}_t \mathbf{A}_f(\tilde{\mathbf{x}}_t))^{-1} \nabla f(\tilde{\mathbf{x}}_t), \quad (1)$$

where  $\mu_t$  is the step size, and  $\mathbf{S}_t$  is a random  $m \times n$  sketching matrix, with  $m \ll n$ , that is used to reduce  $\mathbf{A}_f(\tilde{\mathbf{x}}_t)$  to a small  $m \times d$  sketch  $\mathbf{S}_t \mathbf{A}_f(\tilde{\mathbf{x}}_t)$ . This brings the complexity of forming the Hessian down to  $O(md^2)$  time plus the cost of forming the sketch.

Naturally, sketching methods vary in their computational cost and they can affect the convergence rate, so the right choice of  $\mathbf{S}_t$  depends on the computation-convergence trade-off. On one end of this spectrum are the so-called Sub-Sampled Newton methods [RKM19, XRKM17, YXRKM18], where  $\mathbf{S}_t$  simply selects a random sample of  $m$  rows of  $\mathbf{A}_f(\mathbf{x})$  (e.g., a sample of data points in a training set) to form the sketch. Here the sketching cost is negligible, since  $\mathbf{S}_t$  is extremely sparse, but the convergence rate can be highly variable and problem-dependent. On the other end, we have what we will refer to as the Gaussian Newton Sketch, where  $\mathbf{S}_t$  is a dense matrix with i.i.d. scaled Gaussian entries (a.k.a. a Gaussian embedding). While the  $O(mnd)$  cost of performing this sketch limits its practical appeal, Gaussian Newton Sketch has a number of unique and desirable properties [LP19]: it enjoys strong problem-independent convergence rates; it produces unbiased estimates of the exact Newton update (useful in distributed settings); and it admits analytic expressions for the optimal step size.

A natural way to interpolate between these two extremes is to vary the sparsity  $s$  of the sketching matrix  $\mathbf{S}_t$ , from  $s = 1$  non-zero element per row (Sub-Sampling) to  $s = n$  non-zero elements (Gaussian embedding), with the sketching complexity  $O(mds)$ .<sup>1</sup> Motivated by this, we ask:

Can we sparsify the Gaussian embedding, making its sparsity closer to that of Sub-Sampling, without suffering *any* convergence trade-offs?

In this paper, we provide an affirmative answer to this question. We show that it is possible to drastically sparsify the Gaussian embedding so that two key statistics of the sketches, namely first

<sup>1</sup>For a more detailed discussion of other sketching techniques that may not fit this taxonomy, such as the Subsampled Randomized Hadamard Transform and the CountSketch, see Section 1.2.

and second inverse moments of the sketched Hessian, are nearly preserved in a very strong sense. Namely, the two inverse moments of the sparse sketches can be upper and lower bounded by the corresponding quantities for the dense Gaussian embeddings, where the upper/lower bounds are matching not just up to constant factors, but down to lower order terms (see Theorem 6). We use this to show that the Gaussian Newton Sketch can be sparsified to the point where the cost of sketching is proportional to the cost of other operations, while nearly preserving the convergence rate (again, down to lower order terms; see Theorem 1). This is illustrated conceptually in Figure 1, showing how the sparsity of the sketch affects the per-iteration convergence rate as well as the computational cost of the Newton Sketch. We observe that while the convergence rate improves as we increase the density, it eventually flattens out. On the other hand, the computational cost stays largely flat until some point when it starts increasing at a linear rate. As a result, there is a sparsity regime where we achieve the best of both worlds: the convergence rate and the computational cost are both nearly at their optimal values, thereby avoiding any trade-off.

To establish our results, we build upon a recently introduced sketching technique called LEverage Score Sparsified (LESS) embeddings [DLDM21]. LESS embeddings use leverage score techniques [DMIMW12a] to provide a carefully-constructed (random) sparsification pattern. This is used to produce a sub-Gaussian embedding with  $d$  non-zeros per row of  $\mathbf{S}_t$  (as opposed to  $n$  for a dense matrix), so that the cost of forming the sketch  $\mathbf{S}_t \mathbf{A}_f(\tilde{\mathbf{x}}_t)$  matches the cost of constructing the Hessian estimate, i.e.,  $O(md^2)$  (see Section 2). [DLDM21] analyzed the first inverse moment of the sketch to show that LESS embeddings retain certain unbiasedness properties of Gaussian embeddings. In our setting, this captures the bias of the Newton Sketch, but it does not capture the variance, which is needed to control the convergence rate.

**Contributions.** In this paper, we analyze both the bias and the variance of Newton Sketch with LESS embeddings (Newton-LESS; see Definition 2 and Lemma 7), resulting in a comprehensive convergence analysis. The following are our key contributions:

1. Characterization of the second inverse moment of the sketched Hessian for a class of sketches including sub-Gaussian matrices and LESS embeddings;
2. Precise problem-independent local convergence rates for Newton-LESS, matching the Gaussian Newton Sketch down to lower order terms;
3. Extension of Newton-LESS to *regularized* minimization tasks, with improved dimension-independent guarantees for the sketch sparsity and convergence rate;
4. *Notable corollary:* Best known global convergence rate for an iterative least squares solver, which translates to state-of-the-art numerical performance.

## 1.1 Main results

As our main contribution, we show that, under standard assumptions on the function  $f(\mathbf{x})$ , Newton-LESS achieves the same problem-independent local convergence rate as the Gaussian Newton Sketch, despite drastically smaller per-iteration cost.

**Theorem 1.** *Assume that  $f(\mathbf{x})$  is (a) self-concordant, or (b) has a Lipschitz continuous Hessian. Also, let  $\mathbf{H} = \nabla^2 f(\mathbf{x}^*)$  be positive definite. There is a neighborhood  $U$  containing  $\mathbf{x}^*$  such that if  $\tilde{\mathbf{x}}_0 \in U$ , then Newton-LESS with sketch size  $m \geq Cd \log(dT/\delta)$  and step size  $\mu_t = 1 - \frac{d}{m}$  satisfies:*

$$\left( \mathbb{E}_\delta \frac{\|\tilde{\mathbf{x}}_T - \mathbf{x}^*\|_{\mathbf{H}}^2}{\|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_{\mathbf{H}}^2} \right)^{1/T} \approx_\epsilon \frac{d}{m} \quad \text{for } \epsilon = O\left(\frac{1}{\sqrt{d}}\right),$$

where  $\mathbb{E}_\delta X$  is expectation conditioned on an event that holds with a  $1 - \delta$  probability,  $\|\mathbf{v}\|_{\mathbf{M}} = \sqrt{\mathbf{v}^\top \mathbf{M} \mathbf{v}}$ , and  $a \approx_\epsilon b$  means that  $|a - b| \leq \epsilon b$ .

**Remark 2.** *The same guarantee holds for the Gaussian Newton Sketch, but it is not known for any fast sketching method other than Newton-LESS (see Section 1.2). The alternative assumptions of self-concordance and Lipschitz continuous Hessian are standard in the local convergence analysis of the classical Newton's method, and they only affect the size of the neighborhood  $U$  (see Section 4). Global convergence of Newton-LESS follows from existing analysis of the Newton Sketch [PW17].*

The notion of expectation  $\mathbb{E}_\delta$  allows us to accurately capture the average behavior of a randomized algorithm over a moderate (i.e., polynomial in  $d$ ) number of trials even when the true expectation

is not well behaved. Here, this guards against the (very unlikely, but non-zero) possibility that the Hessian estimate produced by a sparse sketch will be ill-conditioned.

To illustrate this result in a special case (of obvious independent interest), we provide a simple corollary for the least squares regression task, i.e.,  $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|^2$ . Importantly, here the convergence rate of  $(\frac{d}{m})^T$  holds *globally*. Also, for this task we have  $\frac{1}{2}\|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{H}}^2 = f(\mathbf{x}) - f(\mathbf{x}^*)$ , so the convergence can be stated in terms of the excess function value. To our knowledge, this is the best known convergence guarantee for a fast iterative least squares solver.

**Corollary 3.** *Let  $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|^2$  for  $\mathbf{A} \in \mathbb{R}^{n \times d}$  and  $\mathbf{b} \in \mathbb{R}^n$ . Then, given any  $\tilde{\mathbf{x}}_0 \in \mathbb{R}^d$ , Newton-LESS with sketch size  $m \geq Cd \log(dT/\delta)$  and step size  $\mu_t = 1 - \frac{d}{m}$  satisfies:*

$$\left( \mathbb{E}_\delta \frac{f(\tilde{\mathbf{x}}_T) - f(\mathbf{x}^*)}{f(\tilde{\mathbf{x}}_0) - f(\mathbf{x}^*)} \right)^{1/T} \approx_\epsilon \frac{d}{m} \quad \text{for } \epsilon = O\left(\frac{1}{\sqrt{d}}\right).$$

Prior to this work, a convergence rate of  $(\frac{d}{m})^T$  was known only for dense Gaussian embeddings, and only for the least squares task [LP19]. On the other hand, our results apply as generally as the standard local convergence analysis of the Newton’s method, and they include a broad class of sketches. In Section 3, we provide general structural conditions on a randomized sketching matrix that are needed to enable our analysis. These conditions are satisfied by a wide range of sketching methods, including all sub-Gaussian embeddings (e.g., using random sign entries instead of Gaussians), the original LESS embeddings, and other choices of sparse random matrices (see Lemma 7). Moreover, we develop an improved local convergence analysis of the Newton Sketch, which allows us to recover the precise convergence rate and derive the optimal step size. In Appendix D, we also discuss a distributed variant of Newton-LESS, which takes advantage of the near-unbiasedness properties of LESS embeddings, extending the results of [DLDM21].

The performance of Newton-LESS can be further improved for regularized minimization tasks. Namely, suppose that function  $f$  can be decomposed as follows:  $f(\mathbf{x}) = f_0(\mathbf{x}) + g(\mathbf{x})$ , where  $g(\mathbf{x})$  has a Hessian that is easy to evaluate (e.g.,  $l_2$ -regularization,  $g(\mathbf{x}) = \frac{\lambda}{2}\|\mathbf{x}\|^2$ ). In this case, a modified variant of the Newton Sketch has been considered, where only the  $f_0$  component is sketched:

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t - \mu_t (\mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_t^\top \mathbf{S}_t \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) + \nabla^2 g(\tilde{\mathbf{x}}_t))^{-1} \nabla f(\tilde{\mathbf{x}}_t), \quad (2)$$

where, again, we let  $\mathbf{A}_{f_0}(\mathbf{x})$  be an  $n \times d$  matrix that encodes the second-order information in  $f_0$  at  $\mathbf{x}$ . For example, in the case of regularized least squares,  $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|^2 + \frac{\lambda}{2}\|\mathbf{x}\|^2$ , we have  $\mathbf{A}_{f_0}(\mathbf{x}) = \mathbf{A}$  and  $\nabla^2 g(\mathbf{x}) = \lambda \mathbf{I}$  for all  $\mathbf{x}$ . We show that the convergence rate of both Newton-LESS and the Gaussian Newton Sketch can be improved in the presence of regularization, by replacing the dimension  $d$  with an *effective* dimension  $d_{\text{eff}}$ . This can be significantly smaller than  $d$  when the Hessian of  $f_0$  at the optimum exhibits rapid spectral decay or is approximately low-rank:

$$d_{\text{eff}} = \text{tr}(\nabla^2 f_0(\mathbf{x}^*) \nabla^2 f(\mathbf{x}^*)^{-1}) \leq d.$$

**Theorem 4.** *Assume that  $f_0$  and  $f$  are (a) self-concordant, or (b) have a Lipschitz continuous Hessian. Also, let  $\nabla^2 f_0(\mathbf{x}^*)$  be positive definite and let  $\nabla^2 g(\mathbf{x}^*)$  be positive semidefinite, with  $\mathbf{H} = \nabla^2 f(\mathbf{x}^*)$ . There is a neighborhood  $U$  containing  $\mathbf{x}^*$  such that if  $\tilde{\mathbf{x}}_0 \in U$ , then Regularized Newton-LESS (2), with sketch size  $m \geq Cd_{\text{eff}} \log(d_{\text{eff}}T/\delta)$  and step size  $\mu_t = 1 - \frac{d_{\text{eff}}}{m}$ , satisfies:*

$$\left( \mathbb{E}_\delta \frac{\|\tilde{\mathbf{x}}_T - \mathbf{x}^*\|_{\mathbf{H}}^2}{\|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_{\mathbf{H}}^2} \right)^{1/T} \leq \frac{d_{\text{eff}}}{m} \cdot (1 + \epsilon) \quad \text{for } \epsilon = O\left(\frac{1}{\sqrt{d_{\text{eff}}}}\right).$$

**Remark 5.** *The same guarantee holds for the Gaussian Newton Sketch. Unlike in Theorem 1, here we can only obtain an upper-bound on the local convergence rate, because the exact rate may depend on the starting point  $\tilde{\mathbf{x}}_0$  (see Section 4). For regularized least squares,  $f(\mathbf{x}) = \frac{1}{2}\|\mathbf{Ax} - \mathbf{b}\|^2 + \frac{\lambda}{2}\|\mathbf{x}\|^2$ , the above convergence guarantee holds globally, i.e.,  $U = \mathbb{R}^d$ . Note that  $d_{\text{eff}}$  can be efficiently estimated using sketching-based trace estimators [ACW16, CEM<sup>+</sup>15].*

Finally, our numerical results show that Newton-LESS can be implemented very efficiently on modern hardware platforms, improving on the optimization cost over not only dense Gaussian embeddings, but also state-of-the-art sketching methods such as the Subsampled Randomized Hadamard Transform, as well as other first-order and second-order methods. Moreover, we demonstrate that our theoretical predictions for the optimal sparsity level and convergence rate are extremely accurate in practice.

## 1.2 Related work

Leverage Score Sparsified (LESS) embeddings were proposed by [DLDM21] as a way of addressing the phenomenon of inversion bias, which arises in distributed second-order methods [WRKXM18, DM19, DBPM20, GGD<sup>+</sup>21]. Their results only establish the *near-unbiasedness* of Newton-LESS iterates (i.e., that  $\mathbb{E}[\tilde{\mathbf{x}}_{t+1}] \approx \mathbf{x}_{t+1}$ ), but they did not provide any improved guarantees on the convergence rate. Also, their notion of LESS embeddings is much narrower than ours, and so it does not capture Regularized Newton-LESS or uniformly sparsified sketches (LESS-uniform).

Convergence analysis of the Newton Sketch [PW17, LP19, LLDP20] and other randomized second-order methods [BCNN11, BCNW12, EM15, RKM19] has been extensively studied in the machine learning community, often using techniques from Randomized Numerical Linear Algebra (RandNLA) [DM16, DM21]. Some of the popular RandNLA methods include the Subsampled Randomized Hadamard Transform (SRHT, [AC09]) and several variants of sparse sketches, such as the CountSketch [CW17, MM13] and OSNAP [NN13, Coh16]. Also, row sampling based on Leverage Scores [DMM06, AM15] and Determinantal Point Processes [Der19, DCV19, DM21] has been used for sketching. Note that CountSketch and OSNAP sparse sketches differ from LESS embeddings in several ways, and in particular, they use a fixed number of non-zeros per column of the sketching matrix (as opposed to per row), so unlike LESS, their rows are not independent. While all of the mentioned methods, when used in conjunction with the Newton Sketch, exhibit similar per-iteration complexity as LESS embeddings (see Section 2), their existing convergence analysis is fundamentally limited: The best known rate is  $(C \log d \cdot \frac{d}{m})^T$ , which is worse than our result of  $(\frac{d}{m})^T$ , by a factor of  $C \log d$ , where  $C > 1$  is a non-negligible constant that arises in the measure concentration analysis.

In the specific context of least squares regression where  $f(\mathbf{x}) = \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$ , the Hessian  $\mathbf{A}^\top \mathbf{A}$  remains constant, and an alternative strategy is to keep the random sketch  $\mathbf{S}\mathbf{A}$  fixed at every iteration. Many efficient randomized iterative solvers are based on this precondition-and-solve approach [RT08, AMT10, MSM14]: form the sketch  $\mathbf{S}\mathbf{A}$ , compute an easy-to-invert square-root matrix  $\tilde{\mathbf{H}}^{\frac{1}{2}}$  of  $\mathbf{A}^\top \mathbf{S}^\top \mathbf{S}\mathbf{A}$  and apply an iterative least squares solver to the preconditioned objective  $\min_{\mathbf{z}} \frac{1}{2} \|\mathbf{A}\tilde{\mathbf{H}}^{-\frac{1}{2}}\mathbf{z} - \mathbf{b}\|^2$ , e.g., Chebyshev iterations or the preconditioned conjugate gradient. In contrast to the Newton Sketch, these methods do not naturally extend to more generic convex objectives for which the Hessian matrix changes at every iteration. Also, similarly as the Newton Sketch, their convergence guarantees are limited to  $(C \log d \cdot \frac{d}{m})^T$  when used in conjunction with fast sketching methods such as SRHT, OSNAP, or leverage score sampling.

## 2 Preliminaries

**Notation.** We let  $\|\mathbf{v}\|_{\mathbf{M}} = \sqrt{\mathbf{v}^\top \mathbf{M}\mathbf{v}}$ . We define  $a \approx_\epsilon b$  to mean  $|a - b| \leq \epsilon b$ , whereas  $a = b \pm \epsilon$  means that  $|a - b| \leq \epsilon$ , and  $C$  denotes a large absolute constant. We use  $\mathbb{E}_{\mathcal{E}}$  to denote expectation conditioned on  $\mathcal{E}$ , and for a  $\delta \in (0, 1)$ , we use  $\mathbb{E}_\delta$  as a short-hand for: ‘‘There is an event  $\mathcal{E}$  with probability at least  $1 - \delta$  s.t.  $\mathbb{E}_{\mathcal{E}}$  ...’’. Let pd and psd mean positive definite and positive semidefinite. Random variable  $X$  is sub-Gaussian if  $\Pr\{|X| \geq t\} \leq \exp(-ct^2)$  for all  $t \geq 0$  and some  $c = \Omega(1)$ .

We next introduce some concepts related to LEverage Score Sparsified (LESS) embeddings. We start with the notion of statistical leverage scores [DMIMW12a], which are importance weights assigned to the rows of a matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$ . The definition below for leverage scores is somewhat more general than standard definitions, because it allows for a *regularized* leverage score, where the regularization depends on a  $d \times d$  matrix  $\mathbf{C}$ . When  $\mathbf{C}$  is a scaled identity matrix, this matches the definition of ridge leverage scores [AM15].

**Definition 1** (Leverage scores). *For given matrices  $\mathbf{A} \in \mathbb{R}^{n \times d}$  and psd  $\mathbf{C} \in \mathbb{R}^{d \times d}$ , we define the  $i$ th leverage score  $l_i(\mathbf{A}, \mathbf{C})$  as the squared norm of the  $i$ th row of  $\mathbf{U} = \mathbf{A}\mathbf{H}^{-\frac{1}{2}}$ , where  $\mathbf{H} = \mathbf{A}^\top \mathbf{A} + \mathbf{C}$  is assumed to be invertible. The effective dimension of  $\mathbf{A}$  (given  $\mathbf{C}$ ) is defined as  $d_{\text{eff}} = \sum_i l_i(\mathbf{A}, \mathbf{C}) = \text{tr}(\mathbf{U}^\top \mathbf{U})$ , whereas the coherence of  $\mathbf{A}$  (given  $\mathbf{C}$ ) is  $\tau = \frac{n}{d_{\text{eff}}} \max_i l_i(\mathbf{A}, \mathbf{C}) \in [1, \frac{n}{d_{\text{eff}}}]$ .*

Next, we define a class of sparsified sub-Gaussian sketching matrices which will be used in our results. This captures LESS embeddings, as well as other sketching matrices that are supported by the analysis. To that end, we define what we call a *sparsifier*, which is an  $n$ -dimensional random vector  $\xi$  that specifies the sparsification pattern for one row of the  $m \times n$  sketching matrix  $\mathbf{S}$ .

**Definition 2** (LESS embeddings). Let  $t_1, \dots, t_s$  be sampled i.i.d. from a distribution  $p = (p_1, \dots, p_n)$ . Then, the random vector  $\boldsymbol{\xi}^\top = (\sqrt{\frac{b_1}{sp_1}}, \dots, \sqrt{\frac{b_n}{sp_n}})$ , where  $b_i = \sum_{j=1}^s 1_{[t_j=i]}$ , is a  $(p, s)$ -sparsifier. A  $(p, s)$ -sparsified sub-Gaussian sketch is a random matrix  $\mathbf{S}$  that consists of i.i.d. row vectors distributed as  $c \cdot (\mathbf{x} \circ \boldsymbol{\xi})^\top$ , where  $\circ$  denotes the entry-wise product,  $\mathbf{x}$  has i.i.d. mean zero, unit variance and sub-Gaussian entries, and  $c$  is some constant. For given matrices  $\mathbf{A} \in \mathbb{R}^{n \times d}$  and psd  $\mathbf{C} \in \mathbb{R}^{d \times d}$ , we focus on two variants of Leverage Score Sparsified embeddings:

1. LESS. We assume that  $p_i \approx_{1/2} l_i(\mathbf{A}, \mathbf{C})/d_{\text{eff}}$ , and let  $s \approx_{1/2} d_{\text{eff}}$ . For  $\mathbf{C} = \mathbf{0}$  and  $s = d$ , we recover the LESS embeddings proposed by [DLDM21].
2. LESS-uniform. We simply let  $p_i = 1/n$  (denoted as  $p = \text{unif}$ ). This avoids the preprocessing needed for approximating the  $l_i(\mathbf{A}, \mathbf{C})$ , but we may need larger  $s$  to recover the theory.

**Computational cost.** To implement LESS, we must first approximate (i.e., there is no need to compute exactly [DMIMW12a]) the leverage scores of  $\mathbf{A}$ . This can be done in time  $O(\text{nnz}(\mathbf{A}) \log n + d^3 \log d)$  by using standard RandNLA techniques [DMIMW12a, CW17], where  $\text{nnz}(\mathbf{A})$  is the number of non-zero entries in  $\mathbf{A}$  and it is bounded by  $nd$ . Since the prescribed sparsity for LESS satisfies  $s = O(d)$ , the sketching cost is at most  $O(md^2)$ . Thus, the total cost of constructing the sketched Hessian with LESS is  $O(\text{nnz}(\mathbf{A}) \log n + md^2)$ , which up to logarithmic factors matches other sparse sketching methods such as leverage score sampling (when implemented with approximate leverage scores [DMIMW12a]), CountSketch, and OSNAP. In comparison, using the SRHT leads to  $O(nd \log m + md^2)$  complexity, since this method does not take advantage of data sparsity. Note that, in practice, the computational trade-offs between sketching methods are quite different, and significantly hardware-dependent (see Section 5). In particular, the cost of approximating the leverage scores in LESS embeddings can be entirely avoided by using LESS-uniform. Here, the total cost of sketching is  $O(mds)$ , but the sparsity of the sketch that is needed for the theory depends on  $\mathbf{A}$  and  $\mathbf{C}$ . Yet, in Section 5, we show empirically that this approach works well even for  $s = d$ .

### 3 Equivalence between LESS and Gaussian Embeddings

In this section, we derive the basic quantities that determine the convergence properties of the Newton Sketch, namely, the first and second moments of the normalized sketched Hessian inverse. Our key technical contribution is a new analysis of the second moment for a wide class of sketching matrices that includes LESS embeddings and sub-Gaussian sketches.

Consider the Newton Sketch update as in (2), and let  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$ . Denoting  $\mathbf{H}_t = \nabla^2 f(\tilde{\mathbf{x}}_t)$ ,  $\mathbf{g}_t = \nabla f(\tilde{\mathbf{x}}_t)$ , and using  $\mathbf{p}_t = -\mu_t \mathbf{H}_t^{-1} \mathbf{g}_t$  to denote the exact Newton direction with step size  $\mu_t$ , a simple calculation shows that:

$$\|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 - \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 = 2\tilde{\Delta}_t^\top \mathbf{H}_t^{\frac{1}{2}} \tilde{\mathbf{Q}} \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}_t + \mathbf{p}_t^\top \mathbf{H}_t^{\frac{1}{2}} \tilde{\mathbf{Q}}^2 \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}_t, \quad (3)$$

where  $\tilde{\mathbf{Q}} = \mathbf{H}_t^{\frac{1}{2}} (\mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_t^\top \mathbf{S}_t \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) + \nabla^2 g(\tilde{\mathbf{x}}_t))^{-1} \mathbf{H}_t^{\frac{1}{2}}$ . From this, we have that the expected decrease in the optimization error is determined by the first two moments of the matrix  $\tilde{\mathbf{Q}}$ , i.e.,  $\mathbb{E}[\tilde{\mathbf{Q}}]$  and  $\mathbb{E}[\tilde{\mathbf{Q}}^2]$ . In the unregularized case, i.e.,  $g(\mathbf{x}) = 0$ , these moments can be derived exactly for the Gaussian embedding. For instance if we let  $\mathbf{S}_t$  be an  $m \times n$  matrix with i.i.d. standard normal entries scaled by  $\frac{1}{\sqrt{m-d-1}}$ , then we obtain that:

$$\mathbb{E}[\tilde{\mathbf{Q}}] = \mathbf{I}, \quad \mathbb{E}[\tilde{\mathbf{Q}}^2] = \frac{(m-1)(m-d-1)}{(m-d)(m-d-3)} \cdot \mathbf{I} \approx_\epsilon \frac{m}{m-d} \cdot \mathbf{I},$$

for  $\epsilon = O(1/d)$ . This choice of scaling for the Gaussian Newton Sketch ensures that each iterate  $\tilde{\mathbf{x}}_{t+1}$  is an unbiased estimate of the corresponding exact Newton update with the same step size, i.e., that  $\mathbb{E}[\tilde{\mathbf{x}}_{t+1}] = \tilde{\mathbf{x}}_t + \mathbf{p}_t$ . For most other sketching techniques, neither of the two moments is analytically tractable because of the bias coming from matrix inversion. Moreover, if we allow for regularization, e.g.,  $g(\mathbf{x}) = \frac{\lambda}{2} \|\mathbf{x}\|^2$ , then even the Gaussian embedding does not enjoy tractable formulas for the moments of  $\tilde{\mathbf{Q}}$ . However, using ideas from asymptotic random matrix theory, [DLDM21] showed that in the unregularized case, the exact Gaussian formula for the first moment holds approximately for sub-Gaussian sketches and LESS embeddings:  $\mathbb{E}_\delta[\tilde{\mathbf{Q}}] \approx_\epsilon \mathbf{I}$ . This implies near-unbiasedness of



the unregularized Newton-LESS iterates relative to the exact Newton step, but it is not sufficient to ensure any convergence guarantees.

In this work, we develop a general characterization of the first and second moments of  $\tilde{\mathbf{Q}}$  for a wide class of sketching matrices, both in the unregularized and in regularized settings. For the sake of generality, we will simplify the notation here, and analyze the first and second moment of  $\mathbf{Q} = \mathbf{H}^{\frac{1}{2}}(\mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A} + \mathbf{C})^{-1} \mathbf{H}^{\frac{1}{2}}$  for some matrices  $\mathbf{A} \in \mathbb{R}^{n \times d}$  and  $\mathbf{C} \in \mathbb{R}^{d \times d}$  such that  $\mathbf{A}^\top \mathbf{A} + \mathbf{C} = \mathbf{H}$ . In the context of Newton Sketch (2), these quantities correspond to  $\mathbf{A} = \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)$  and  $\mathbf{C} = \nabla^2 g(\tilde{\mathbf{x}}_t)$ . Also, as a shorthand, we will define the normalized version of matrix  $\mathbf{A}$  as  $\mathbf{U} = \mathbf{A} \mathbf{H}^{-\frac{1}{2}}$ . The following are the two structural conditions that need to be satisfied by a sketching matrix to enable our analysis.

The first condition is standard in the sketching literature. Essentially, it implies that the sketching matrix  $\mathbf{S}$  produces a useful approximation of the Hessian with high probability (although this guarantee is still far too coarse by itself to obtain our results).

**Condition 1** (Property of random matrix  $\mathbf{S}$ ). *Given  $\mathbf{U} \in \mathbb{R}^{n \times d}$ , the  $m \times n$  random matrix  $\mathbf{S}$  satisfies  $\|\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} - \mathbf{U}^\top \mathbf{U}\| \leq \eta$  with probability  $1 - \delta$ .*

This property is known as the *subspace embedding property*. Subspace embeddings were first used by [DMM06], where they were used in a data-aware context to obtain relative-error approximations for  $\ell_2$  regression and low-rank matrix approximation [DMM08]. Subsequently, data-oblivious subspace embeddings were used by [Sar06] and popularized by [Woo14]. Both data-aware and data-oblivious subspace embeddings can be used to derive bounds for the accuracy of various algorithms [DM16, DM18].

For our analysis, it is important to assume that  $\mathbf{S}$  has i.i.d. row vectors  $c\mathbf{s}_i^\top$ , where  $c$  is an appropriate scaling constant. The second condition is defined as a property of those row vectors, which makes them sufficiently similar to Gaussian vectors. This is a relaxation of the Restricted Bai-Silverstein condition, proposed by [DLDM21], which leads to significant improvements in the sparsity guarantee for LESS embeddings when the Newton Sketch is regularized.

**Condition 2** (Property of random vector  $\mathbf{s}$ ). *Given  $\mathbf{U} \in \mathbb{R}^{n \times d}$ , the  $n$ -dimensional random vector  $\mathbf{s}$  satisfies  $\text{Var}[\mathbf{s}^\top \mathbf{U} \mathbf{B} \mathbf{U}^\top \mathbf{s}] \leq \alpha \cdot \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top)$  for all p.s.d. matrices  $\mathbf{B}$  and some  $\alpha = O(1)$ .*

Given these two conditions, we are ready to derive precise non-asymptotic analytic expressions for the first two moments of the regularized sketched inverse matrix, which is the main technical contribution of this work (proof in Appendix A).

**Theorem 6.** *Fix  $\mathbf{A}$  and assume that  $\mathbf{C}$  is psd. Define  $\mathbf{H} = \mathbf{A}^\top \mathbf{A} + \mathbf{C}$  and  $\mathbf{U} = \mathbf{A} \mathbf{H}^{-\frac{1}{2}}$ . Let  $\mathbf{S}$  consist of  $m$  i.i.d. rows distributed as  $\frac{1}{\sqrt{m-d_{\text{eff}}}} \mathbf{s}^\top$ , where  $\mathbb{E}[\mathbf{s} \mathbf{s}^\top] = \mathbf{I}_n$  and  $d_{\text{eff}} = \text{tr}(\mathbf{U}^\top \mathbf{U})$ . Also, let  $\tilde{d}_{\text{eff}} = \text{tr}((\mathbf{U}^\top \mathbf{U})^2)$ . Suppose that the matrix consisting of the first  $m/3$  rows of  $\mathbf{S}$  scaled by  $\sqrt{3}$  satisfies Condition 1 w.r.t.  $\mathbf{U}$ , for  $\eta \leq 1/2$  and probability  $1 - \delta/3$ , where  $\delta \leq 1/m^3$ . Suppose also that  $\mathbf{s}$  satisfies Condition 2 w.r.t.  $\mathbf{U}$ . If  $m \geq O(d_{\text{eff}})$ , then, conditioned on event  $\mathcal{E}$  that holds with probability  $1 - \delta$ , matrix  $\mathbf{Q} = \mathbf{H}^{\frac{1}{2}}(\mathbf{A}^\top \mathbf{S}^\top \mathbf{S} \mathbf{A} + \mathbf{C})^{-1} \mathbf{H}^{\frac{1}{2}}$  satisfies  $\|\mathbf{Q} - \mathbf{I}\| \leq O(\eta)$  and:*

$$\|\mathbb{E}_{\mathcal{E}}[\mathbf{Q}] - \mathbf{I}\| \leq O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right), \quad \|\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2] - \left(\mathbf{I} + \frac{d_{\text{eff}}}{m-d_{\text{eff}}} \mathbf{U}^\top \mathbf{U}\right)\| \leq O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right).$$

Theorem 6 shows that, for a wide class of sketching matrices, we can approximately write  $\mathbb{E}[\mathbf{Q}] \approx \mathbf{I}$  and  $\mathbb{E}[\mathbf{Q}^2] \approx \mathbf{I} + \frac{d_{\text{eff}}}{m-d_{\text{eff}}} \mathbf{U}^\top \mathbf{U}$ , with the error term scaling as  $O(\frac{\sqrt{d_{\text{eff}}}}{m})$ . In the case of the first moment, this is a relatively straightforward generalization of the unregularized formula for the Gaussian case. However, for the second moment this expression is considerably more complicated, including not one but two notions of effective dimension,  $\tilde{d}_{\text{eff}} \neq d_{\text{eff}}$ . To put this in context, in the unregularized case, i.e.,  $\mathbf{C} = \mathbf{0}$ , we have  $d_{\text{eff}} = \tilde{d}_{\text{eff}} = d$  and  $\mathbf{U}^\top \mathbf{U} = \mathbf{I}$ , so we get  $\mathbf{I} + \frac{d_{\text{eff}}}{m-d_{\text{eff}}} \mathbf{U}^\top \mathbf{U} = \frac{m}{m-d} \mathbf{I}$ , which matches the second moment for the Gaussian sketch (up to lower order terms).

In the case of the first moment, the proof of Theorem 6 follows along the same lines as in [DLDM21], using a decomposition of  $\mathbb{E}[\mathbf{Q}]$  that is based on the Sherman-Morrison rank-one update of the inverse. This approach was originally inspired by the analysis of Stieltjes transforms that are used to establish the limiting spectral distribution in asymptotic random matrix theory (e.g., [BS10, CD11]), and applied to sketching by [DLLM20, DLDM21]. Our key contribution lies in deriving the bound for the second moment, which requires a substantially more elaborate decomposition of  $\mathbb{E}[\mathbf{Q}^2]$ .

In the following lemma, we establish that the assumptions of Theorem 6 are satisfied not only by Gaussian, but also sub-Gaussian and LESS embedding matrices (proof in Appendix C).

**Lemma 7.** Fix  $\mathbf{A}$  and assume that  $\mathbf{C}$  is psd. Let  $\mathbf{S}$  be a sketching matrix with  $m$  i.i.d. rows distributed as  $\frac{1}{\sqrt{m-d_{\text{eff}}}}\mathbf{s}^\top$ . Then,  $\mathbf{S}$  satisfies Conditions 1 and 2 as long as one of the following holds:

1. Sub-Gaussian:  $\mathbf{s}$  is an i.i.d. sub-Gaussian random vector and  $m \geq C(d_{\text{eff}} + \log(1/\delta))/\eta^2$ ;
2. LESS:  $\mathbf{s}$  is a  $(p, d_{\text{eff}})$ -sparsified i.i.d. sub-Gaussian random vector with  $p_i \approx_{1/2} l_i(\mathbf{A}, \mathbf{C})/d_{\text{eff}}$ , and  $m \geq Cd_{\text{eff}} \log(d_{\text{eff}}/\delta)/\eta^2$ ;
3. LESS-uniform:  $\mathbf{s}$  is a  $(\text{unif}, \tau d_{\text{eff}})$ -sparsified i.i.d. sub-Gaussian random vector, where  $\tau$  is the coherence of  $\mathbf{A}$ , i.e.,  $\frac{n}{d_{\text{eff}}} \max_i l_i(\mathbf{A}, \mathbf{C})$ , and  $m \geq Cd_{\text{eff}} \log(d_{\text{eff}}/\delta)/\eta^2$ .

## 4 Convergence Analysis for Newton-LESS

In this section, we demonstrate how our main technical results can be used to provide improved convergence guarantees for Newton-LESS (and, more broadly, any sketching methods that satisfy the conditions of Theorem 6). Here, we will focus on the more general regularized setting (2), where we can only show an upper bound on the convergence rate (Theorem 4). The unregularized result (Theorem 1) with matching upper/lower bounds follows similarly.

To start, we introduce the standard assumptions on the function  $f$ , which are needed to ensure strong local convergence guarantees for the classical Newton's method [BV04].

**Assumption 8.** Function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  has a Lipschitz continuous Hessian with constant  $L$ , i.e.,  $\|\nabla^2 f(\mathbf{x}) - \nabla^2 f(\mathbf{x}')\| \leq L \|\mathbf{x} - \mathbf{x}'\|$  for all  $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ .

**Assumption 9.** Function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  is self-concordant, i.e., for all  $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ , the function  $\phi(t) = f(\mathbf{x} + t\mathbf{x}')$  satisfies:  $|\phi'''(t)| \leq 2(\phi''(t))^{3/2}$ .

Only one of those two assumptions needs to be satisfied for our analysis to go through, and the choice of the assumption only affects the size of the neighborhood around the optimum  $\mathbf{x}^*$  for which our local convergence guarantee is satisfied. To clarify this, below we give an expanded version of Theorem 4 (proof in Appendix B).

**Theorem 10** (Expanded Theorem 4). Let  $\mathbf{H}_0 = \nabla^2 f_0(\mathbf{x}^*)$  be pd and  $\mathbf{C} = \nabla^2 g(\mathbf{x}^*)$  be psd. Define  $d_{\text{eff}} = \text{tr}(\mathbf{H}_0 \mathbf{H}^{-1})$  and  $\tilde{d}_{\text{eff}} = \text{tr}((\mathbf{H}_0 \mathbf{H}^{-1})^2)$  for  $\mathbf{H} = \mathbf{H}_0 + \mathbf{C}$ . Assume one of the following:

1.  $f_0$  and  $f$  satisfy Assumption 8, and  $U = \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{H}} < \frac{\sqrt{d_{\text{eff}}}}{m} (\lambda_{\min})^{3/2}/L\}$ , where  $\lambda_{\min}$  is the smallest eigenvalue of  $\mathbf{H}_0$ ;
2.  $f_0$  and  $f$  satisfy Assumption 9, and  $U = \{\mathbf{x} : \|\mathbf{x} - \mathbf{x}^*\|_{\mathbf{H}} < \frac{\sqrt{d_{\text{eff}}}}{m}\}$ .

Then, Newton Sketch (2) starting from  $\tilde{\mathbf{x}}_0 \in U$ , using any  $\mathbf{S}_t$  from Lemma 7 (i.e., sub-Gaussian, LESS, or LESS-uniform) with  $\delta$  replaced by  $\delta/T$ , and step size  $\mu_t = 1 - \frac{d_{\text{eff}}}{m+d_{\text{eff}}-\tilde{d}_{\text{eff}}}$ , satisfies:

$$\left( \mathbb{E}_\delta \frac{\|\tilde{\mathbf{x}}_T - \mathbf{x}^*\|_{\mathbf{H}}^2}{\|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_{\mathbf{H}}^2} \right)^{1/T} \leq \frac{d_{\text{eff}} \cdot (1 + \epsilon)}{m + d_{\text{eff}} - \tilde{d}_{\text{eff}}} \leq \frac{d_{\text{eff}}}{m} \cdot (1 + \epsilon), \quad \text{for } \epsilon = O\left(\frac{1}{\sqrt{d_{\text{eff}}}}\right).$$

Note that, compared to Theorem 4, here we present a slightly sharper bound which uses both types of effective dimension,  $\tilde{d}_{\text{eff}} \leq d_{\text{eff}}$ , that are present in Theorem 6. The statement from Theorem 4 is recovered by replacing  $\tilde{d}_{\text{eff}}$  with  $d_{\text{eff}}$  in the step size and in the bound. The key step in the proof of the result is the following lemma, which uses Theorem 6 to characterize the Newton Sketch iterate  $\tilde{\mathbf{x}}_{t+1}$  in terms of the corresponding Newton iterate  $\mathbf{x}_{t+1}$ . Note that this result holds globally for arbitrary  $\tilde{\mathbf{x}}_t$  and without the smoothness assumptions on  $f$ . Recall that we let  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$  denote the error residual at step  $t$ .

**Lemma 11.** Fix  $\mathbf{H}_t = \nabla^2 f(\tilde{\mathbf{x}}_t)$  and let  $\tilde{\mathbf{x}}_{t+1}$  be the Newton Sketch iterate with  $\mathbf{S}_t$  as in Lemma 7. If the exact Newton step  $\mathbf{x}_{t+1} = \tilde{\mathbf{x}}_t - \mu_t \mathbf{H}_t^{-1} \mathbf{g}_t$  is a descent direction, i.e.,  $\|\Delta_{t+1}\|_{\mathbf{H}_t} \leq \|\tilde{\Delta}_t\|_{\mathbf{H}_t}$  where  $\Delta_{t+1} = \mathbf{x}_{t+1} - \mathbf{x}^*$ , then

$$\mathbb{E}_\delta \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 = \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 + \frac{d_{\text{eff}}(\tilde{\mathbf{x}}_t)}{m - \tilde{d}_{\text{eff}}(\tilde{\mathbf{x}}_t)} \|\mathbf{x}_{t+1} - \tilde{\mathbf{x}}_t\|_{\nabla^2 f_0(\tilde{\mathbf{x}}_t)}^2 \pm O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2.$$



Importantly, the second term on the right-hand side uses the norm  $\|\cdot\|_{\nabla^2 f_0(\tilde{\mathbf{x}}_t)}$ , which is different than the norm  $\|\cdot\|_{\mathbf{H}_t}$  used for the remaining terms. As a result, in the regularized setting, it is possible that the second term will be much smaller than the last term (the approximation error). This prevents us from obtaining a matching lower-bound for the convergence rate of Regularized Newton-LESS. On the other hand, when  $g(\mathbf{x}) = 0$ , then  $f_0 = f$  and we obtain matching upper/lower bounds.

The remainder of the proof of Theorem 10 is essentially a variant of the local convergence analysis of the Newton’s method. Here, note that typically, we would set the step size to  $\mu_t = 1$  and we would expect superlinear (specifically, quadratic) convergence rate. However, for Newton Sketch, convergence is a mixture of linear rate and superlinear rate, where the linear part is due to the approximation error in sketching the Hessian. Sufficiently close to the optimum, the linear rate will dominate, and so this is what we focus on in our local convergence analysis. The key novelty here is that, unlike prior work, we strive to describe the linear rate precisely, down to lower order terms. As a key step, we observe that the convergence of exact Newton with step size  $\mu_t < 1$ , letting  $\Delta_{t+1} = \mathbf{x}_{t+1} - \mathbf{x}^*$ , is given by:

$$\|\Delta_{t+1}\|_{\mathbf{H}_t}^2 = \underbrace{(1 - \mu_t)^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2}_{\text{linear rate}} + \underbrace{\mu_t (\Delta_{t+1} + (1 - \mu_t)\tilde{\Delta}_t)^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t)}_{\text{superlinear rate}}, \quad (4)$$

where recall that  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$  corresponds to the previous iterate, and  $\mathbf{g}_t = \nabla f(\tilde{\mathbf{x}}_t)$ . Here,  $(1 - \mu_t)^2$  represents the linear convergence rate. The superlinear term vanishes near the optimum  $\mathbf{x}^*$ , because of the presence of  $\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t$ , which (under the smoothness assumptions on  $f$ ) vanishes at the same rate as  $\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2$ . Interestingly, with this precise analysis, the superlinear term does not have a quadratic rate, but rather a 3/2 rate. Entering (4) into the guarantee from Lemma 11, we obtain that the local rate of Newton Sketch can be expressed as  $(1 - \mu_t)^2 + \frac{d_{\text{eff}}}{m - d_{\text{eff}}} \mu_t^2$ , and minimizing this expression over  $\mu_t$  we obtain the desired quantities from Theorem 10. Finally, we note that while the norms and effective dimensions in the above exact calculations are stated with respect to the Hessian  $\mathbf{H}_t$  at the current iterate  $\tilde{\mathbf{x}}_t$ , these can all be approximated by the corresponding quantities computed using the Hessian at the optimum  $\mathbf{x}^*$  (as in Theorem 10), relying on smoothness of  $f$ .




Figure 2: LESS-uniform embeddings: convergence rate of the Newton Sketch for least squares regression and wall-clock time of forming SA versus row density, with sketch size  $m = 8d$ . The results were averaged over 100 trials and error bars show twice the empirical standard deviation.

## 5 Numerical Experiments

We evaluated our theory on a range of different problems, and we have found that the more precise analysis that our theory provides describes well the convergence behavior for a range of optimization problems. In this section, we present numerical simulations illustrating this for regularized logistic regression and least squares regression, with different datasets ranging from medium to large scale: the CIFAR-10 dataset, the Musk dataset, and WESAD [SRD<sup>+</sup>18]. Data preprocessing and implementation details, as well as additional numerical results for least squares and regularized least squares, can be found in Appendix E.

We investigate first the effect of the row density of a LESS-uniform embedding on the Newton Sketch convergence rate and on the time for computing the sketch SA for least squares regression. In Figure 2, we report these two performance measures versus the row density. (This is the empirical analog for real data of Figure 1). Note that here we also consider a synthetic data matrix with high-coherence, which aims to be more challenging for a uniform sparsifier. Remarkably, our prescribed row density of  $d$  non-zeros per row offers an excellent empirical trade-off: except for the CIFAR-10

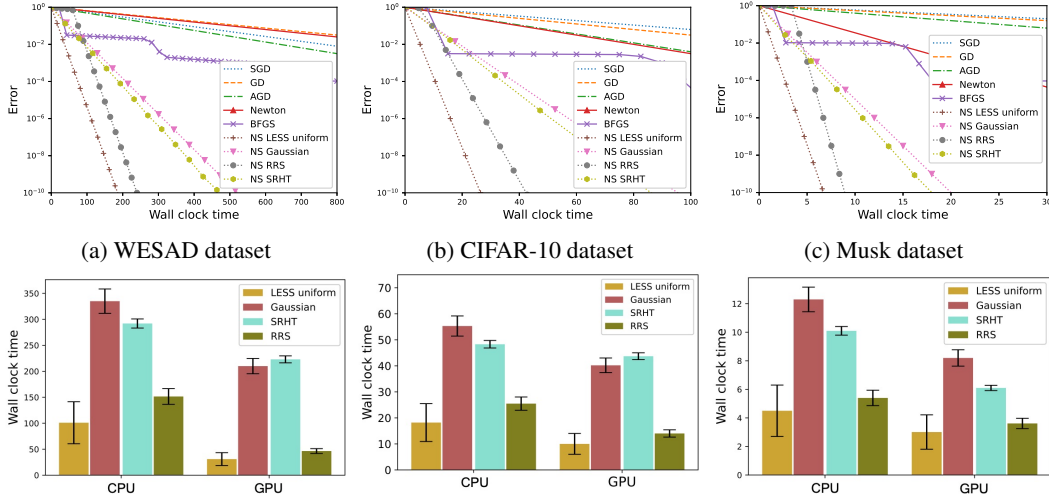


Figure 3: Top plots show the convergence of Newton Sketch (NS) and baselines for logistic regression. We use a sketch size  $m = d/2$  for NS. In the bottom plots, we report the CPU and GPU wall-clock times to reach a  $10^{-6}$  accurate solution for NS with different sketching methods.

dataset, for which random row sampling performs equally well to Gaussian embeddings, we observe that one can drastically decrease the row density without significantly impairing the convergence rate.

Next, in Figure 3, we investigate a minimization task for the regularized logistic regression loss. Namely, given a data matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$  with rows  $\mathbf{a}_i$ , a target vector  $\mathbf{b} \in \{\pm 1\}^n$  and a regularization parameter  $\lambda > 0$ , the goal is to solve:

$$\min_{\mathbf{x} \in \mathbb{R}^d} \frac{1}{n} \sum_{i=1}^n \log(1 + \exp(-b_i \mathbf{a}_i^\top \mathbf{x})) + \frac{\lambda}{2} \|\mathbf{x}\|_2^2.$$

For each dataset, we choose the value of  $\lambda$  among  $\{10^{-j} \mid j = 0, \dots, 8\}$  that minimizes the error on a hold out validation set. For CIFAR-10 and Musk, we pick  $\lambda = 10^{-4}$ . For WESAD, we pick  $\lambda = 10^{-5}$ . We plot the error versus wall-clock time for the Newton Sketch with LESS-uniform, Gaussian, Subsampled Randomized Hadamard Transform (SRHT) and Random Row Sampling (RRS) matrices, and we compare it with two standard second-order optimization baselines: Newton’s method and BFGS. We also included three first-order baselines: Gradient Descent (GD), Accelerated GD (AGD), and Stochastic GD (SGD), observing much worse performance. Based on the top plots in Figure 3, we conclude that Newton-LESS (i.e., NS LESS uniform) offers significant time speed-ups over all baselines.

Finally, we compare wall-clock time on different hardware (CPU versus GPU) for the Newton Sketch to reach a  $10^{-6}$ -accurate solution (see Appendix E for hardware details). From Figure 3 (bottom plots), we conclude the following: first, when switching from CPU to GPU for WESAD and CIFAR-10, Gaussian embeddings become more efficient than SRHT, despite a worse time complexity, by taking better advantage of the massively parallel architecture; second, Random Row Sampling performs better than either of them, despite having much weaker theoretical guarantees; and third, LESS-uniform is more efficient than all three other methods, on both CPU and GPU hardware platforms, observing a significant speed-up when switching to the parallel GPU architecture.

## 6 Conclusions

We showed that, when constructing randomized Hessian estimates for second-order optimization, we can get the best of both worlds: the efficiency of Sub-Sampling methods and the precision of Gaussian embeddings, by using sparse sketching matrices known as LEverage Score Sparsified (LESS) embeddings. Our algorithm, called Newton-LESS, enjoys both strong theoretical convergence guarantees and excellent empirical performance on a number of hardware platforms. An important future direction is to explain the surprising effectiveness of the simpler LESS-uniform method, particularly on high-coherence matrices, which goes beyond the predictions of our current theory.

## Acknowledgements

We would like to acknowledge DARPA, IARPA, NSF, and ONR via its BRC on RandNLA for providing partial support of this work. MP and JL were partially supported by the National Science Foundation under grants IIS-1838179, ECCS-2037304 and the Army Research Office. Our conclusions do not necessarily reflect the position or the policy of our sponsors, and no official endorsement should be inferred.

## References

- [ABH17] Naman Agarwal, Brian Bullins, and Elad Hazan. Second-order stochastic optimization for machine learning in linear time. *Journal of Machine Learning Research*, 18(116):1–40, 2017.
- [AC09] Nir Ailon and Bernard Chazelle. The fast Johnson–Lindenstrauss transform and approximate nearest neighbors. *SIAM Journal on computing*, 39(1):302–322, 2009.
- [ACW16] Haim Avron, Kenneth L Clarkson, and David P Woodruff. Sharper bounds for regularized data fitting. *arXiv preprint arXiv:1611.03225*, 2016.
- [AM15] Ahmed El Alaoui and Michael W. Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, pages 775–783, December 2015.
- [AMT10] Haim Avron, Petar Maymounkov, and Sivan Toledo. Blendenpik: Supercharging lapack’s least-squares solver. *SIAM Journal on Scientific Computing*, 32(3):1217–1236, 2010.
- [BCNN11] Richard H Byrd, Gillian M Chin, Will Neveitt, and Jorge Nocedal. On the use of stochastic hessian information in optimization methods for machine learning. *SIAM Journal on Optimization*, 21(3):977–995, 2011.
- [BCNW12] Richard H Byrd, Gillian M Chin, Jorge Nocedal, and Yuchen Wu. Sample size selection in optimization methods for machine learning. *Mathematical programming*, 134(1):127–155, 2012.
- [BS10] Zhidong Bai and Jack W Silverstein. *Spectral analysis of large dimensional random matrices*, volume 20. Springer, 2010.
- [BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [CD11] Romain Couillet and Mrouane Debbah. *Random Matrix Methods for Wireless Communications*. Cambridge University Press, USA, 2011.
- [CEM<sup>+</sup>15] Michael B Cohen, Sam Elder, Cameron Musco, Christopher Musco, and Madalina Persu. Dimensionality reduction for k-means clustering and low rank approximation. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 163–172, 2015.
- [Coh16] Michael B Cohen. Nearly tight oblivious subspace embeddings by trace inequalities. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 278–287. SIAM, 2016.
- [CW17] Kenneth L. Clarkson and David P. Woodruff. Low-rank approximation and regression in input sparsity time. *J. ACM*, 63(6):54:1–54:45, January 2017.
- [DBPM20] Michał Dereziński, Burak Bartan, Mert Pilanci, and Michael W Mahoney. Debiasing distributed second order optimization with surrogate sketching and scaled regularization. In *Advances in Neural Information Processing Systems*, volume 33, pages 6889–6899, 2020.

- [DCV19] Michał Dereziński, Daniele Calandriello, and Michal Valko. Exact sampling of determinantal point processes with sublinear time preprocessing. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 11542–11554. 2019.
- [Der19] Michał Dereziński. Fast determinantal point processes via distortion-free intermediate sampling. In *Proceedings of the 32nd Conference on Learning Theory*, 2019.
- [DLDM21] Michał Dereziński, Zhenyu Liao, Edgar Dobriban, and Michael W Mahoney. Sparse sketches with small inversion bias. In *Proceedings of the 34th Conference on Learning Theory*, 2021.
- [DLLM20] Michał Dereziński, Feynman Liang, Zhenyu Liao, and Michael W Mahoney. Precise expressions for random projections: Low-rank approximation and randomized Newton. In *Advances in Neural Information Processing Systems*, volume 33, pages 18272–18283, 2020.
- [DM16] Petros Drineas and Michael W. Mahoney. RandNLA: Randomized numerical linear algebra. *Communications of the ACM*, 59:80–90, 2016.
- [DM18] P. Drineas and M. W. Mahoney. Lectures on randomized numerical linear algebra. In M. W. Mahoney, J. C. Duchi, and A. C. Gilbert, editors, *The Mathematics of Data*, IAS/Park City Mathematics Series, pages 1–48. AMS/IAS/SIAM, 2018.
- [DM19] Michał Dereziński and Michael W Mahoney. Distributed estimation of the inverse hessian by determinantal averaging. In *Advances in Neural Information Processing Systems 32*, pages 11401–11411. Curran Associates, Inc., 2019.
- [DM21] Michał Dereziński and Michael W Mahoney. Determinantal point processes in randomized numerical linear algebra. *Notices of the American Mathematical Society*, 68(1):34–45, 2021.
- [DMIMW12a] Petros Drineas, Malik Magdon-Ismail, Michael W. Mahoney, and David P. Woodruff. Fast approximation of matrix coherence and statistical leverage. *J. Mach. Learn. Res.*, 13(1):3475–3506, December 2012.
- [DMIMW12b] Petros Drineas, Malik Magdon-Ismail, Michael W Mahoney, and David P Woodruff. Fast approximation of matrix coherence and statistical leverage. *The Journal of Machine Learning Research*, 13(1):3475–3506, 2012.
- [DMM06] Petros Drineas, Michael W Mahoney, and S Muthukrishnan. Sampling algorithms for  $\ell_2$  regression and applications. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pages 1127–1136, 2006.
- [DMM08] Petros Drineas, Michael W. Mahoney, and S. Muthukrishnan. Relative-error CUR matrix decompositions. *SIAM J. Matrix Anal. Appl.*, 30(2):844–881, September 2008.
- [EM15] Murat A Erdogdu and Andrea Montanari. Convergence rates of sub-sampled Newton methods. *Advances in Neural Information Processing Systems*, 28:3052–3060, 2015.
- [GGD<sup>+</sup>21] Vipul Gupta, Avishek Ghosh, Michał Dereziński, Rajiv Khanna, Kannan Ramchandran, and Michael Mahoney. LocalNewton: Reducing communication bottleneck for distributed learning. *arXiv preprint arXiv:2105.07320*, 2021.
- [KL17] Vladimir Koltchinskii and Karim Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133, 2017.
- [LLDP20] Jonathan Lacotte, Sifan Liu, Edgar Dobriban, and Mert Pilanci. Limiting spectrum of randomized Hadamard transform and optimal iterative sketching methods. In *Conference on Neural Information Processing Systems*, 2020.

- [LP19] Jonathan Lacotte and Mert Pilanci. Faster least squares optimization. *arXiv preprint arXiv:1911.02675*, 2019.
- [MM13] Xiangrui Meng and Michael W. Mahoney. Low-distortion subspace embeddings in input-sparsity time and applications to robust linear regression. In *Proceedings of the Forty-fifth Annual ACM Symposium on Theory of Computing, STOC '13*, pages 91–100, New York, NY, USA, 2013. ACM.
- [MSM14] Xiangrui Meng, Michael A Saunders, and Michael W Mahoney. LSRN: A parallel iterative solver for strongly over- or underdetermined systems. *SIAM Journal on Scientific Computing*, 36(2):C95–C118, 2014.
- [MZ20] Shahar Mendelson and Nikita Zhivotovskiy. Robust covariance estimation under  $L_4 - L_2$  norm equivalence. *Annals of Statistics*, 48(3):1648–1664, 2020.
- [NN13] Jelani Nelson and Huy L Nguyễn. Osnap: Faster numerical linear algebra algorithms via sparser subspace embeddings. In *2013 IEEE 54th annual symposium on foundations of computer science*, pages 117–126. IEEE, 2013.
- [PW17] Mert Pilanci and Martin J Wainwright. Newton sketch: A near linear-time optimization algorithm with linear-quadratic convergence. *SIAM Journal on Optimization*, 27(1):205–245, 2017.
- [RKM19] Farbod Roosta-Khorasani and Michael W. Mahoney. Sub-Sampled Newton methods. *Math. Program.*, 174(1–2):293–326, March 2019.
- [RT08] Vladimir Rokhlin and Mark Tygert. A fast randomized algorithm for overdetermined linear least-squares regression. *Proceedings of the National Academy of Sciences*, 105(36):13212–13217, 2008.
- [RV13] Mark Rudelson and Roman Vershynin. Hanson-Wright inequality and sub-gaussian concentration. *Electronic Communications in Probability*, 18, 2013.
- [Sar06] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science, FOCS '06*, pages 143–152, Washington, DC, USA, 2006. IEEE Computer Society.
- [SRD<sup>+</sup>18] Philip Schmidt, Attila Reiss, Robert Duerichen, Claus Marberger, and Kristof Van Laerhoven. Introducing wesad, a multimodal dataset for wearable stress and affect detection. In *Proceedings of the 20th ACM international conference on multimodal interaction*, pages 400–408, 2018.
- [Tro11] Joel A Tropp. Improved analysis of the subsampled randomized Hadamard transform. *Advances in Adaptive Data Analysis*, 3(01n02):115–126, 2011.
- [Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, August 2012.
- [Tro15] Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.
- [Woo14] David P Woodruff. Sketching as a tool for numerical linear algebra. *Foundations and Trends® in Theoretical Computer Science*, 10(1–2):1–157, 2014.
- [WRKXM18] Shusen Wang, Farbod Roosta-Khorasani, Peng Xu, and Michael W Mahoney. GI-ANT: Globally improved approximate Newton method for distributed optimization. In *Advances in Neural Information Processing Systems 31*, pages 2332–2342. 2018.
- [XRKM17] Peng Xu, Farbod Roosta-Khorasani, and Michael Mahoney. Newton-type methods for non-convex optimization under inexact hessian information. *Mathematical Programming*, August 2017.
- [YXRKM18] Z. Yao, P. Xu, F. Roosta-Khorasani, and M. W. Mahoney. Inexact non-convex Newton-type methods. Technical report, 2018. Preprint: arXiv:1802.06925.



## A Characterization of Inverse Moments (Proof of Theorem 6)

In this section, we prove Theorem 6. The proof consists of two parts corresponding to the first and second moment of  $\mathbf{Q}$ . The analysis of the first moment bound is nearly the same as in [DLDM21], so we only outline it here, highlighting the differences coming from the regularization matrix  $\mathbf{C}$ . The analysis of the second moment is our main contribution in this proof, and we discuss it in detail. First, however, we define the high probability event  $\mathcal{E}$  which is common to both parts.

To simplify the proof, we will let  $m$  be divisible by 3. Note that we have  $\mathbf{Q} = (\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} + \mathbf{D})^{-1}$  for  $\mathbf{D} = \mathbf{H}^{-\frac{1}{2}} \mathbf{C} \mathbf{H}^{-\frac{1}{2}}$ . Moreover, let  $\mathbf{S}_{-i}$  denote  $\mathbf{S}$  without the  $i$ th row, and let  $\mathbf{Q}_{-i} = (\mathbf{U}^\top \mathbf{S}_{-i}^\top \mathbf{S}_{-i} \mathbf{U} + \mathbf{D})^{-1}$ . Also, define  $\tilde{\mathbf{S}}_1, \tilde{\mathbf{S}}_2, \tilde{\mathbf{S}}_3$  as the matrices consisting of the first, second and third group of  $m/3$  rows in  $\mathbf{S}$ , all scaled by  $\sqrt{3}$ , so that  $\mathbf{S}^\top \mathbf{S} = \frac{1}{3} \sum_{j=1}^3 \tilde{\mathbf{S}}_j^\top \tilde{\mathbf{S}}_j$ . Next, using  $\Sigma = \mathbf{U}^\top \mathbf{U}$ , similarly as in [DLDM21] we let  $t = m/3$  and define three independent events:

$$\mathcal{E}_j : \quad \|\Sigma - \mathbf{U}^\top \tilde{\mathbf{S}}_j^\top \tilde{\mathbf{S}}_j \mathbf{U}\| \leq \eta, \quad \text{for } j = 1, 2, 3, \quad (5)$$

with  $\mathcal{E} = \bigwedge_{j=1}^3 \mathcal{E}_j$  defined as the intersection of the events. Conditioned on  $\mathcal{E}$ , we have:

$$\|\mathbf{I} - (\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} + \mathbf{D})\| = \left\| \frac{1}{3} \sum_{j=1}^3 (\Sigma - \mathbf{U}^\top \tilde{\mathbf{S}}_j^\top \tilde{\mathbf{S}}_j \mathbf{U}) \right\| \leq \frac{1}{3} \sum_{j=1}^3 \|\Sigma - \mathbf{U}^\top \tilde{\mathbf{S}}_j^\top \tilde{\mathbf{S}}_j \mathbf{U}\| \leq \eta,$$

which implies that  $\|\mathbf{Q} - \mathbf{I}\| \leq \frac{\eta}{1-\eta} \leq 2\eta$ . Furthermore, an important property of the definition of  $\mathcal{E}$  is that for each  $i \in \{1, \dots, m\}$  there is a  $j \in \{1, 2, 3\}$  such that  $\mathcal{E}_j$  is independent of  $\mathbf{x}_i$ , and after conditioning only on  $\mathcal{E}_j$  we get  $\|\mathbf{Q}_{-i}\| \leq 6$ . From Condition 1 and the union bound we conclude that  $\Pr(\mathcal{E}) \geq 1 - \delta$ .

The analysis of both the first and second moment uses the Sherman-Morrison formula, to separate one of the rows from the rest of the sketch. We state this formula in the following lemma.

**Lemma 12** (Sherman-Morrison). *For  $\mathbf{A} \in \mathbb{R}^{n \times n}$  invertible and  $\mathbf{u}, \mathbf{v} \in \mathbb{R}^n$ ,  $\mathbf{A} + \mathbf{u}\mathbf{v}^\top$  is invertible if and only if  $1 + \mathbf{v}^\top \mathbf{A}^{-1} \mathbf{u} \neq 0$  and*

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^\top)^{-1} = \mathbf{A}^{-1} - \frac{\mathbf{A}^{-1} \mathbf{u} \mathbf{v}^\top \mathbf{A}^{-1}}{1 + \mathbf{v}^\top \mathbf{A}^{-1} \mathbf{u}}.$$

From the above formula, it follows that:

$$(\mathbf{A} + \mathbf{u}\mathbf{v}^\top)^{-1} \mathbf{u} = \frac{\mathbf{A}^{-1} \mathbf{u}}{1 + \mathbf{v}^\top \mathbf{A}^{-1} \mathbf{u}}.$$

### A.1 Proof of first moment bound

In this part of the proof we recall the decomposition of  $\mathbb{E}_{\mathcal{E}}[\mathbf{Q}]$  used by [DLDM21]. Most of their analysis is unaffected by the presence of the regularization matrix  $\mathbf{D}$ , so we will focus on the steps that will also be needed for our analysis of the second moment. Let the  $i$ th row of  $\mathbf{S}$  be  $\frac{1}{\sqrt{m-d_{\text{eff}}}} \mathbf{s}_i^\top$ , and define  $\mathbf{x}_i = m \mathbf{U}^\top \mathbf{s}_i$ , so that  $\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} = \frac{\gamma}{m} \sum_i \mathbf{x}_i \mathbf{x}_i^\top$ , where  $\gamma = \frac{m}{m-d_{\text{eff}}}$ . Using  $\gamma_i = 1 + \frac{\gamma}{m} \mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{x}_i$ , we have:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\mathbf{Q}] - \mathbf{I} &= \mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\Sigma + \mathbf{D}) - \mathbf{Q}(\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} + \mathbf{D})] \\ &= \mathbb{E}_{\mathcal{E}}[\mathbf{Q}\Sigma] - \mathbb{E}_{\mathcal{E}}[\mathbf{Q}\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U}] \\ &\stackrel{(*)}{=} \mathbb{E}_{\mathcal{E}}[\mathbf{Q}\Sigma] - \mathbb{E}_{\mathcal{E}}\left[\frac{\gamma}{\gamma_i} \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top\right] \\ &= \mathbb{E}_{\mathcal{E}}[\mathbf{Q} - \mathbf{Q}_{-i}]\Sigma + \mathbb{E}_{\mathcal{E}}[\mathbf{Q}_{-i}(\Sigma - \mathbf{x}_i \mathbf{x}_i^\top)] + \mathbb{E}_{\mathcal{E}}\left[\left(1 - \frac{\gamma}{\gamma_i}\right) \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top\right], \end{aligned}$$

where  $(*)$  follows from the Sherman-Morrison formula. From this point, the analysis of [DLDM21] proceeds to bound the spectral norm of the first two terms by  $O(1/m)$ , and the spectral norm of the last term by  $O(\sqrt{\text{tr}(\mathbf{U}^\top \mathbf{U})/m})$ . In their setup,  $\mathbf{C} = \mathbf{0}$ , which means that  $\text{tr}(\mathbf{U}^\top \mathbf{U}) = d$ , whereas in our more general statement, we let  $d_{\text{eff}} = \text{tr}(\mathbf{U}^\top \mathbf{U})$ . This does not affect the proofs. For the sake of our analysis of the second moment, we separate out the following guarantees obtained by [DLDM21], given here in a slightly more general form than originally.

**Lemma 13** ([DLDM21]). *The following bounds hold for  $k \in \{1, 2\}$ :*

$$\begin{aligned}\|\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^k - \mathbf{Q}_{-i}^k] \boldsymbol{\Sigma}\| &= O(1/m), \\ \|\mathbb{E}_{\mathcal{E}}[\mathbf{Q}_{-i}^k (\mathbf{x}_i \mathbf{x}_i^\top - \boldsymbol{\Sigma})]\| &= O(1/m), \\ \|\mathbb{E}_{\mathcal{E}}[(\frac{\gamma}{\gamma_i} - 1) \mathbf{Q}_{-i}^k \mathbf{x}_i \mathbf{x}_i^\top]\| &= O(\sqrt{d_{\text{eff}}}/m).\end{aligned}$$

## A.2 Proof of second moment bound

We next present the analysis of the second moment,  $\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2]$ , which requires a considerably more elaborate decomposition. Using  $\rho = \frac{d_{\text{eff}}}{m - d_{\text{eff}}}$  and  $\boldsymbol{\Sigma} = \mathbf{U}^\top \mathbf{U}$ , we have:

$$\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2] - (\mathbf{I} + \rho \boldsymbol{\Sigma}) = \underbrace{(\mathbb{E}_{\mathcal{E}}[\mathbf{Q}] - \mathbf{I})}_{\mathbf{T}_1} + (\mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\mathbf{Q} - \mathbf{I})] - \rho \boldsymbol{\Sigma}).$$

Recalling that  $\boldsymbol{\Sigma} + \mathbf{D} = \mathbf{I}$ , we can rewrite the last term as:

$$\begin{aligned}\mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\mathbf{Q} - \mathbf{I})] &= \mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\mathbf{Q}(\boldsymbol{\Sigma} + \mathbf{D}) - \mathbf{Q}(\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} + \mathbf{D}))] \\ &= \mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\mathbf{Q}\boldsymbol{\Sigma} - \mathbf{Q}\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U})] \\ &\stackrel{(a)}{=} \mathbb{E}_{\mathcal{E}}[\mathbf{Q}(\mathbf{Q}\boldsymbol{\Sigma} - \frac{\gamma}{\gamma_i} \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top)] \\ &\stackrel{(b)}{=} \mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2] \boldsymbol{\Sigma} - \mathbb{E}_{\mathcal{E}}[\frac{\gamma}{\gamma_i} \mathbf{Q}_{-i}^2 \mathbf{x}_i \mathbf{x}_i^\top] + \mathbb{E}_{\mathcal{E}}\left[\frac{\mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i}{m} \frac{\gamma^2}{\gamma_i^2} \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top\right] \\ &= \underbrace{\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2 - \mathbf{Q}_{-i}^2] \boldsymbol{\Sigma}}_{\mathbf{T}_2} + \underbrace{\mathbb{E}_{\mathcal{E}}[\mathbf{Q}_{-i}^2 (\boldsymbol{\Sigma} - \mathbf{x}_i \mathbf{x}_i^\top)]}_{\mathbf{T}_3} + \underbrace{\mathbb{E}_{\mathcal{E}}[(1 - \frac{\gamma}{\gamma_i}) \mathbf{Q}_{-i}^2 \mathbf{x}_i \mathbf{x}_i^\top]}_{\mathbf{T}_4} + \mathbb{E}_{\mathcal{E}}\left[\frac{\mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i}{m} \frac{\gamma^2}{\gamma_i^2} \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top\right],\end{aligned}$$

for a fixed  $i$ , where we denote  $\gamma_i = 1 + \frac{\gamma}{m} \mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{x}_i$ . Note that we used the Sherman-Morrison formula twice, in steps (a) and (b). We can put everything together as follows:

$$\begin{aligned}\mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2] - (\mathbf{I} + \rho \boldsymbol{\Sigma}) &= \mathbf{T}_1 + \mathbf{T}_2 + \mathbf{T}_3 + \mathbf{T}_4 \\ &\quad + \underbrace{\rho(\mathbb{E}_{\mathcal{E}}[\mathbf{Q}_{-i}] - \mathbf{I}) \boldsymbol{\Sigma}}_{\mathbf{T}_5} + \underbrace{\rho \mathbb{E}_{\mathcal{E}}[\mathbf{Q}_{-i} (\mathbf{x}_i \mathbf{x}_i^\top - \boldsymbol{\Sigma})]}_{\mathbf{T}_6} + \underbrace{\mathbb{E}_{\mathcal{E}}\left[\left(\frac{\mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i}{m} \frac{\gamma^2}{\gamma_i^2} - \rho\right) \mathbf{Q}_{-i} \mathbf{x}_i \mathbf{x}_i^\top\right]}_{\mathbf{T}_7}.\end{aligned}$$

From the bound on the first moment of  $\mathbf{Q}$ , we conclude that  $\|\mathbf{T}_1\| = O(\sqrt{d_{\text{eff}}}/m)$  and that  $\|\mathbf{T}_5\| = O(\sqrt{d_{\text{eff}}}/m)$ . Without loss of generality, assume that events  $\mathcal{E}_1$  and  $\mathcal{E}_2$  are both independent of  $\mathbf{x}_i$ , and let  $\mathcal{E}' = \mathcal{E}_1 \wedge \mathcal{E}_2$  as well as  $\delta_3 = \Pr(\neg \mathcal{E}_3)$ . Next, we will use the fact that for a p.s.d. random matrix  $\mathbf{M}$  in the probability space of  $\mathbf{S}$ , we have  $\mathbb{E}_{\mathcal{E}}[\mathbf{M}] \preceq \frac{1}{1-\delta} \mathbb{E}_{\mathcal{E}'}[\mathbf{M}] \preceq 2 \cdot \mathbb{E}_{\mathcal{E}'}[\mathbf{M}]$ .

Using  $k = 2$  in Lemma 13, we can bound  $\|\mathbf{T}_2\|$ ,  $\|\mathbf{T}_3\|$  and  $\|\mathbf{T}_4\|$  by  $O(\sqrt{d_{\text{eff}}}/m)$ , and setting  $k = 1$ , we can do the same for  $\|\mathbf{T}_6\|$ .

Thus, it remains to bound  $\|\mathbf{T}_7\|$ . Let  $\tilde{\gamma} = \frac{m}{m - d_{\text{eff}}}$ . We first use the Cauchy-Schwartz inequality twice, obtaining that:

$$\|\mathbf{T}_7\| \leq \frac{1}{m} \sqrt{\mathbb{E}_{\mathcal{E}}\left[\left(\tilde{\gamma} d_{\text{eff}} - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i \cdot \frac{\gamma^2}{\gamma_i^2}\right)^2\right]} \cdot \sup_{\|\mathbf{u}\|=1} \sqrt{\mathbb{E}_{\mathcal{E}}[(\mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{x}_i)^4]} \cdot \sup_{\|\mathbf{u}\|=1} \sqrt[4]{\mathbb{E}_{\mathcal{E}}[(\mathbf{x}_i^\top \mathbf{u})^4]}.\tag{6}$$

The latter two terms can each be bounded easily by  $O(\sqrt[4]{\alpha + 1})$  using Condition 2. For instance, considering the middle term, we have:

$$\begin{aligned}\mathbb{E}_{\mathcal{E}}[(\mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{x}_i)^4] &\leq 2 \mathbb{E}_{\mathcal{E}'}\left[\mathbb{E}[(\mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{u} \mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{x}_i)^2 \mid \mathbf{Q}_{-i}]\right] \\ &= 2 \mathbb{E}_{\mathcal{E}'}\left[\text{Var}[\mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{u} \mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{x}_i \mid \mathbf{Q}_{-i}] + (\mathbb{E}[\mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{u} \mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{x}_i \mid \mathbf{Q}_{-i}])^2\right] \\ &\leq 2 \mathbb{E}_{\mathcal{E}'}\left[\alpha \text{tr}(\mathbf{U}(\mathbf{Q}_{-i} \mathbf{u} \mathbf{u}^\top \mathbf{Q}_{-i})^2 \mathbf{U}^\top) + 2(\text{tr}(\mathbf{U} \mathbf{Q}_{-i} \mathbf{u} \mathbf{u}^\top \mathbf{Q}_{-i} \mathbf{U}^\top))^2\right] \\ &\leq 2 \mathbb{E}_{\mathcal{E}'}\left[O(\alpha) \mathbf{u}_i^\top \mathbf{Q}_{-i} \mathbf{U}^\top \mathbf{U} \mathbf{Q}_{-i} \mathbf{u}_i + (\mathbf{u}_i^\top \mathbf{Q}_{-i} \mathbf{U}^\top \mathbf{U} \mathbf{Q}_{-i} \mathbf{u}_i)^2\right] \\ &\leq 2 \mathbb{E}_{\mathcal{E}'}\left[O(\alpha) \|\mathbf{Q}_{-i} \boldsymbol{\Sigma} \mathbf{Q}_{-i}\| + \|\mathbf{Q}_{-i} \boldsymbol{\Sigma} \mathbf{Q}_{-i}\|^2\right] = O(\alpha + 1),\end{aligned}$$

where we also used that matrices  $\mathbf{Q}_{-i}$ ,  $\Sigma$  and  $\mathbf{u}\mathbf{u}^\top$  have spectral norms bounded by  $O(1)$ . Similarly, we obtain that  $\mathbb{E}_{\mathcal{E}}[(\mathbf{x}_i^\top \mathbf{u})^4] = O(\alpha + 1)$ . For the first term in (6), we have:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}} \left[ \left( \tilde{\gamma} d_{\text{eff}} - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i \cdot \frac{\gamma^2}{\gamma_i^2} \right)^2 \right] &\leq 2 \cdot \mathbb{E}_{\mathcal{E}'} \left[ \left( \tilde{\gamma} d_{\text{eff}} - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i \cdot \frac{\gamma^2}{\gamma_i^2} \right)^2 \right] \\ &\leq 4 \cdot \mathbb{E}_{\mathcal{E}'} [(\tilde{\gamma} d_{\text{eff}} - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i)^2] + 4 \cdot \mathbb{E}_{\mathcal{E}'} \left[ (\mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i)^2 \left( \frac{\gamma^2}{\gamma_i^2} - 1 \right)^2 \right]. \end{aligned}$$

We can further break down the first term as follows:

$$\mathbb{E}_{\mathcal{E}'} [(\tilde{\gamma} d_{\text{eff}} - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i)^2] = (\tilde{\gamma} d_{\text{eff}} - \mathbb{E}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2 \Sigma)])^2 + \text{Var}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2 \Sigma)] + \mathbb{E}_{\mathcal{E}'} [(\text{tr}(\mathbf{Q}_{-i}^2 \Sigma) - \mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i)^2] \quad (7)$$

The latter term can be bounded immediately using Condition 2. The middle term is handled by a separate lemma, which is an immediate extension of Lemma 25 in [DLDM21].

**Lemma 14** ([DLDM21]). *Let  $\text{Var}_{\mathcal{E}'}[\cdot]$  be the conditional variance with respect to event  $\mathcal{E}' = \mathcal{E}_1 \wedge \mathcal{E}_2$ . Then, for  $k \in \{1, 2\}$ ,*

$$\text{Var}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^k \Sigma)] = O(d_{\text{eff}}).$$

Next, note that  $|\frac{\gamma^2}{\gamma_i^2} - 1| = |\gamma - \gamma_i| \cdot \frac{\gamma + \gamma_i}{\gamma_i^2} \leq |\gamma - \gamma_i| \cdot \frac{\gamma + 1}{\gamma_i}$ , since  $\gamma_i > 1$ , so we get:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}'} \left[ (\mathbf{x}_i^\top \mathbf{Q}_{-i}^2 \mathbf{x}_i)^2 \left( \frac{\gamma^2}{\gamma_i^2} - 1 \right)^2 \right] &\leq 6^2 (\gamma + 1)^2 \cdot \mathbb{E}_{\mathcal{E}'} \left[ (\mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{x}_i)^2 \frac{(\gamma - \gamma_i)^2}{\gamma_i^2} \right] \\ &\leq 6^2 (\gamma + 1)^2 \cdot \mathbb{E}_{\mathcal{E}'} \left[ \frac{(\mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{x}_i)^2}{(1 + \frac{\gamma}{m} \mathbf{x}_i^\top \mathbf{Q}_{-i} \mathbf{x}_i)^2} (\gamma - \gamma_i)^2 \right] \\ &\leq O(m^2) \cdot \mathbb{E}_{\mathcal{E}'} [(\gamma - \gamma_i)^2] \\ &\leq O(m^2) \cdot O(\alpha d_{\text{eff}}/m^2) = O(\alpha d_{\text{eff}}). \end{aligned}$$

Finally, we analyze the first term in (7) as follows:

$$\begin{aligned} |\tilde{\gamma} d_{\text{eff}} - \mathbb{E}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2 \Sigma)]| &= |\text{tr}((\mathbb{E}_{\mathcal{E}} - \mathbb{E}_{\mathcal{E}'})[\mathbf{Q}_{-i}^2 \Sigma]) - \text{tr}(\mathbf{T}_2) + \text{tr}(\tilde{\gamma} \Sigma - \mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2 \Sigma])| \\ &= |\text{tr}((\mathbb{E}_{\mathcal{E}} - \mathbb{E}_{\mathcal{E}'})[\mathbf{Q}_{-i}^2 \Sigma]) - \text{tr}(\mathbf{T}_2) + \text{tr}((\mathbf{I} + \rho \Sigma - \mathbb{E}_{\mathcal{E}}[\mathbf{Q}^2]) \Sigma)| \\ &\leq O(d_{\text{eff}}/m^3) + d_{\text{eff}} \cdot O(\alpha \sqrt{d_{\text{eff}}}/m) + |\text{tr}(\mathbf{T}_7 \Sigma)|, \end{aligned}$$

where to bound the first term we used the fact that  $\Pr(-\mathcal{E}_3 | \mathcal{E}') \leq 1/m^3$  and for the last term, recall that  $d_{\text{eff}} = \text{tr}(\Sigma)$  and  $\tilde{d}_{\text{eff}} = \text{tr}(\Sigma^2)$ , which leads to the following identity:

$$\text{tr}(\tilde{\gamma} \Sigma - (\mathbf{I} + \rho \Sigma) \Sigma) = \text{tr} \left( \frac{\tilde{d}_{\text{eff}}}{m - \tilde{d}_{\text{eff}}} \Sigma - \frac{d_{\text{eff}}}{m - d_{\text{eff}}} \Sigma^2 \right) = 0.$$

Further, note that from the analysis of  $\mathbf{T}_7$  we have:

$$\begin{aligned} |\text{tr}(\mathbf{T}_7)| &\leq \frac{d_{\text{eff}}}{m} \sqrt{4(\tilde{\gamma} d_{\text{eff}} - \mathbb{E}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2)])^2 + O(\alpha d_{\text{eff}}) \cdot O(\sqrt{\alpha})} \\ &\leq O(\alpha d_{\text{eff}}/m) \cdot (|\tilde{\gamma} d_{\text{eff}} - \mathbb{E}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2)] + \sqrt{d_{\text{eff}}}). \end{aligned}$$

Putting this together with the previous inequality, we conclude that for sufficiently large  $m$ :

$$|\tilde{\gamma} d_{\text{eff}} - \mathbb{E}_{\mathcal{E}'}[\text{tr}(\mathbf{Q}_{-i}^2)]| \leq \frac{O(\alpha \sqrt{d_{\text{eff}}})}{1 - O(\alpha d_{\text{eff}}/m)} = O(\alpha \sqrt{d_{\text{eff}}}).$$

Plugging this back into the analysis of  $\|\mathbf{T}_7\|$ , we can bound it by  $O(\alpha \sqrt{d}/m)$ , which concludes the proof.

## B Local Convergence Rate of Newton-LESS

In this section, we present the convergence analysis of Newton Sketch for sketching matrices satisfying the structural conditions of Theorem 6. We start by proving Lemma 11, then we show how it can be used to establish the guarantee from Theorem 10. Finally, we discuss how the analysis needs to be adjusted to obtain the two-sided bound from Theorem 1.

## B.1 Proof of Lemma 11

Let  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$ ,  $\Delta_{t+1} = \mathbf{x}_{t+1} - \mathbf{x}^*$ , and  $\mathbf{p}_t = \mathbf{x}_{t+1} - \tilde{\mathbf{x}}_t$ . Also, define  $\rho = \frac{d_{\text{eff}}}{m-d_{\text{eff}}}$  as well as the matrices  $\tilde{\mathbf{Q}} = \mathbf{H}_t^{\frac{1}{2}} (\mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_t^\top \mathbf{S}_t \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) + \nabla^2 g(\tilde{\mathbf{x}}_t))^{-1} \mathbf{H}_t^{\frac{1}{2}}$  and  $\mathbf{U}_t = \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) \mathbf{H}_t^{-\frac{1}{2}}$ . We have:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}} \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 - \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 &= 2\Delta_{t+1}^\top \mathbf{H}_t \mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{x}}_{t+1} - \mathbf{x}_{t+1}] + \mathbb{E}_{\mathcal{E}} \|\tilde{\mathbf{x}}_{t+1} - \mathbf{x}_{t+1}\|_{\mathbf{H}_t}^2 \\ &= 2\Delta_{t+1}^\top \mathbf{H}_t^{\frac{1}{2}} (\mathbf{I} - \mathbb{E}_{\mathcal{E}} \tilde{\mathbf{Q}}) \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}_t + \mathbf{p}_t^\top \mathbf{H}_t^{\frac{1}{2}} \mathbb{E}_{\mathcal{E}} (\mathbf{I} - \tilde{\mathbf{Q}})^2 \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}_t \\ &\leq \rho \cdot \mathbf{p}_t^\top \mathbf{H}_t^{\frac{1}{2}} \mathbf{U}_t^\top \mathbf{U}_t \mathbf{H}_t^{\frac{1}{2}} \mathbf{p}_t + 2\|\mathbf{I} - \mathbb{E}_{\mathcal{E}} \tilde{\mathbf{Q}}\| \cdot (\|\Delta_{t+1}\|_{\mathbf{H}_t} \|\mathbf{p}_t\|_{\mathbf{H}_t} + \|\mathbf{p}_t\|_{\mathbf{H}_t}^2) \\ &\quad + \|\mathbf{I} + \rho \mathbf{U}_t^\top \mathbf{U}_t - \mathbb{E}_{\mathcal{E}} \tilde{\mathbf{Q}}^2\| \cdot \|\mathbf{p}_t\|_{\mathbf{H}_t}^2 \\ &\leq \rho \|\mathbf{p}_t\|_{\nabla^2 f_0(\tilde{\mathbf{x}}_t)}^2 + O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2, \end{aligned}$$

where the last step follows by applying Theorem 6 and observing that  $\|\mathbf{p}_t\|_{\mathbf{H}_t} \leq \|\tilde{\Delta}_t\|_{\mathbf{H}_t} + \|\Delta_{t+1}\|_{\mathbf{H}_t} \leq 2\|\tilde{\Delta}_t\|_{\mathbf{H}_t}$  and  $\|\Delta_{t+1}\|_{\mathbf{H}_t} \|\mathbf{p}_t\|_{\mathbf{H}_t} \leq \frac{1}{2}(\|\Delta_{t+1}\|_{\mathbf{H}_t}^2 + \|\mathbf{p}_t\|_{\mathbf{H}_t}^2) \leq 3\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2$ . The matching lower-bound follows identically.

## B.2 Proof of Theorem 10

We start by analyzing the exact Newton step  $\mathbf{x}_{t+1} = \tilde{\mathbf{x}}_t - \mu_t \mathbf{H}_t^{-1} \mathbf{g}_t$  with step size  $\mu_t$ , gradient  $\mathbf{g}_t = \nabla f(\tilde{\mathbf{x}}_t)$ , and Hessian  $\mathbf{H}_t = \nabla^2 f(\tilde{\mathbf{x}}_t)$ . Letting  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$  and  $\Delta_{t+1} = \mathbf{x}_{t+1} - \mathbf{x}^*$ , we have:

$$\begin{aligned} \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 &= (1 - \mu_t) \Delta_{t+1}^\top \mathbf{g}_t + \Delta_{t+1}^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t) \\ &= (1 - \mu_t) \Delta_{t+1}^\top \mathbf{H}_t \tilde{\Delta}_t - (1 - \mu_t) \Delta_{t+1}^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t) + \Delta_{t+1}^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t) \\ &= (1 - \mu_t) (\tilde{\Delta}_t^\top \mathbf{H}_t \tilde{\Delta}_t - \mu_t \mathbf{g}_t^\top \tilde{\Delta}_t) + \mu_t \Delta_{t+1}^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t) \\ &= (1 - \mu_t)^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + \mu_t (\Delta_{t+1} + (1 - \mu_t) \tilde{\Delta}_t)^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t). \end{aligned}$$

Before we proceed, we make the following assumptions, which will be addressed later.

$$\text{Assume: } \quad \|\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t\|_{\mathbf{H}_t^{-1}} \leq \epsilon \beta \|\tilde{\Delta}_t\|_{\mathbf{H}_t}, \quad \mathbf{H}_t \approx_\epsilon \mathbf{H}, \quad (8)$$

where  $\epsilon = O\left(\frac{1}{\sqrt{d_{\text{eff}}}}\right)$  and  $\beta = \frac{\rho}{1+\rho}$  will become the convergence rate of Newton-LESS, and recall that  $\rho = \frac{d_{\text{eff}}}{m-d_{\text{eff}}}$ . Now, using the Cauchy-Schwartz inequality we obtain that:

$$\begin{aligned} \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 &\leq (1 - \mu_t)^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + \mu_t \|\Delta_{t+1}\|_{\mathbf{H}_t} + (1 - \mu_t) \|\tilde{\Delta}_t\|_{\mathbf{H}_t} \|\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t\|_{\mathbf{H}_t^{-1}} \\ &\leq (1 - \mu_t)^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + \epsilon \beta \mu_t \|\Delta_{t+1}\|_{\mathbf{H}_t} \|\tilde{\Delta}_t\|_{\mathbf{H}_t} + \epsilon \beta \mu_t (1 - \mu_t) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2. \end{aligned}$$

Solving for  $\|\Delta_{t+1}\|_{\mathbf{H}_t}$  (we use that if  $x^2 \leq ax + b$  then  $x^2 \leq a^2 + 2b$ ), we obtain the following:

$$\begin{aligned} \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 &\leq 2(1 - \mu_t)^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + 2\epsilon \beta \mu_t (1 - \mu_t) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + \epsilon^2 \beta^2 \mu_t^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \\ &\leq 2((1 - \mu_t)^2 + \epsilon \beta \mu_t) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2. \end{aligned}$$

Setting  $\mu_t = \frac{1}{1+\rho}$ , we conclude that  $\|\Delta_{t+1}\|_{\mathbf{H}_t}^2 \leq \frac{2\rho}{1+\rho} \frac{\rho+\epsilon}{1+\rho} \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \leq \beta \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2$  when  $m \geq 4d_{\text{eff}} + 2$  and  $\epsilon < 1/4$ . Next, we return to the Newton Sketch. Recall that using Lemma 11 with the event  $\mathcal{E}$  having failure probability  $\delta/T$ , we have:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}} \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 - \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 &= \rho \|\mathbf{p}_t\|_{\nabla^2 f_0(\tilde{\mathbf{x}}_t)}^2 \pm O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \\ &\leq \rho \|\mathbf{p}_t\|_{\mathbf{H}_t}^2 + O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2, \end{aligned} \quad (9)$$

where we also used the fact that  $\nabla^2 f_0(\tilde{\mathbf{x}}_t) \preceq \mathbf{H}_t$ . The leading term in the above decomposition can be written as follows:

$$\begin{aligned} \rho \|\mathbf{p}_t\|_{\mathbf{H}_t}^2 &= \rho \mu_t^2 (\mathbf{g}_t \tilde{\Delta}_t - \mathbf{g}_t^\top \mathbf{H}_t^{-1} (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t)) \\ &= \rho \mu_t^2 (\tilde{\Delta}_t^\top \mathbf{H}_t \tilde{\Delta}_t - \tilde{\Delta}_t^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t) - \mathbf{g}_t^\top \mathbf{H}_t^{-1} (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t)) \\ &= \rho \mu_t^2 \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 - \rho \mu_t^2 (\tilde{\Delta}_t + \mathbf{H}_t^{-1} \mathbf{g}_t)^\top (\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t). \end{aligned}$$

Putting everything together, and then setting  $\mu_t = \frac{1}{1+\rho}$ , we obtain:

$$\begin{aligned}\mathbb{E}_{\mathcal{E}} \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 &\leq \left( (1-\mu_t)^2 + \rho\mu_t^2 + O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \\ &\quad + \left( (1+\rho)\mu_t(\Delta_{t+1} - \tilde{\Delta}_t) + (1-\rho\mu_t^2)\tilde{\Delta}_t \right)^\top (\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t) \\ &= \left( \beta + O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 + (\Delta_{t+1} - \beta\mu_t\tilde{\Delta}_t)^\top (\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t).\end{aligned}$$

We can bound the second term by using Cauchy-Schwartz, the first assumption in (8) and  $\mu_t \leq 1$ :

$$\left| (\Delta_{t+1} - \beta\mu_t\tilde{\Delta}_t)^\top (\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t) \right| \leq \epsilon\beta\|\Delta_{t+1}\|_{\mathbf{H}_t}\|\tilde{\Delta}_t\|_{\mathbf{H}_t} + \epsilon\beta^2\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \leq 2\epsilon\beta\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2.$$

Combining this with the assumption  $\mathbf{H}_t \approx_\epsilon \mathbf{H}$ , which implies that  $\|\mathbf{v}\|_{\mathbf{H}_t}^2 \approx_\epsilon \|\mathbf{v}\|_{\mathbf{H}}^2$ , we obtain:

$$\mathbb{E}_{\delta/T} \frac{\|\tilde{\Delta}_{t+1}\|_{\mathbf{H}}^2}{\|\tilde{\Delta}_t\|_{\mathbf{H}}^2} \leq \beta \cdot \left( 1 + O\left(\frac{1}{\sqrt{d_{\text{eff}}}}\right) \right). \quad (10)$$

Note that since  $\rho = \frac{d_{\text{eff}}}{m-d_{\text{eff}}}$ , we have  $\beta = \frac{d_{\text{eff}}}{m+d_{\text{eff}}-d_{\text{eff}}}$  and  $\mu_t = \frac{d_{\text{eff}}}{m+d_{\text{eff}}-d_{\text{eff}}}$ . Alternatively, if throughout the analysis we use  $\rho = \frac{d_{\text{eff}}}{m-d_{\text{eff}}} \leq \frac{d_{\text{eff}}}{m-d_{\text{eff}}}$ , then we obtain the simpler (and slightly weaker) convergence rate  $\beta = \frac{d_{\text{eff}}}{m}$  with step size  $\mu_t = 1 - \frac{d_{\text{eff}}}{m}$ , as in Theorem 4.

It remains to address the assumptions from (8), and then carefully chain the expectations together. Next, we define the neighborhood  $U$  in which we can establish our convergence guarantee, and show that when the iterate lies in the neighborhood, then (8) is satisfied. This part of the proof will depend on what type of function  $f(\mathbf{x})$  we are minimizing.

**Lipschitz Hessian.** Suppose that function  $f(\mathbf{x})$  has a Lipschitz continuous Hessian with constant  $L$  (Assumption 8). We define the neighborhood  $U$  through the following condition:

$$\|\tilde{\Delta}_t\|_{\mathbf{H}} < \frac{\sqrt{d_{\text{eff}}}(\lambda_{\min})^{3/2}}{mL},$$

where  $\lambda_{\min}$  denotes the smallest eigenvalue of  $\mathbf{H}$ . Suppose that the condition holds for some  $t$ . Then, we have:

$$\|\mathbf{H}^{-\frac{1}{2}}(\mathbf{H}_t - \mathbf{H})\mathbf{H}^{-\frac{1}{2}}\| \leq \frac{1}{\lambda_{\min}}\|\mathbf{H}_t - \mathbf{H}\| \leq \frac{L}{\lambda_{\min}}\|\tilde{\Delta}_t\| \leq \frac{L}{(\lambda_{\min})^{3/2}}\|\tilde{\Delta}_t\|_{\mathbf{H}} \leq \frac{\sqrt{d_{\text{eff}}}}{m} \leq \epsilon,$$

for  $\epsilon = O\left(\frac{1}{\sqrt{d_{\text{eff}}}}\right)$ , showing that  $\mathbf{H}_t \approx_\epsilon \mathbf{H}$ . In particular, this implies that  $\|\mathbf{H}_t^{-1}\| \geq \frac{1}{\lambda_{\min}(1-\epsilon)}$ . To get the second assumption in (8), we first follow standard analysis of the Newton's method [BV04]:

$$\begin{aligned}\|\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t\| &= \left\| \mathbf{H}_t\tilde{\Delta}_t - \left( \int_0^1 \nabla^2 f(\mathbf{x}^* + \tau\tilde{\Delta}_t) d\tau \right) \tilde{\Delta}_t \right\| \\ &\leq \|\tilde{\Delta}_t\| \cdot \int_0^1 \|\nabla^2 f(\tilde{\mathbf{x}}_t) - \nabla^2 f(\mathbf{x}^* + \tau\tilde{\Delta}_t)\| d\tau \\ &\leq \|\tilde{\Delta}_t\| \cdot \int_0^1 (1-\tau)L\|\tilde{\Delta}_t\| d\tau \leq \frac{L}{2}\|\tilde{\Delta}_t\|^2.\end{aligned}$$

Then, we simply use the fact that  $\|\mathbf{H}_t^{-1}\| \geq \frac{1}{\lambda_{\min}(1-\epsilon)}$  to conclude:

$$\begin{aligned}\|\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t\|_{\mathbf{H}_t^{-1}} &\leq \frac{1}{\sqrt{\lambda_{\min}(1-\epsilon)}}\|\mathbf{H}_t\tilde{\Delta}_t - \mathbf{g}_t\| \\ &\leq \frac{1}{\sqrt{\lambda_{\min}(1-\epsilon)}} \frac{L}{2}\|\tilde{\Delta}_t\|^2 \\ &\leq \frac{L}{(\lambda_{\min})^{3/2}}\|\tilde{\Delta}_t\|_{\mathbf{H}}^2 \leq \epsilon\beta\|\tilde{\Delta}_t\|_{\mathbf{H}_t},\end{aligned}$$

since  $\beta = O\left(\frac{d_{\text{eff}}}{m}\right)$ , thus establishing the assumptions from (8).



**Self-concordant function.** Suppose that function  $f(\mathbf{x})$  is self-concordant (Assumption 9). We will define the neighborhood  $U$  through the following condition:

$$\|\tilde{\Delta}_t\|_{\mathbf{H}} < \frac{\sqrt{d_{\text{eff}}}}{m}.$$

Now, using a standard property of self-concordant functions [BV04, Chapter 9], we have:

$$(1 - \|\tilde{\Delta}_t\|_{\mathbf{H}})^2 \mathbf{H} \preceq \mathbf{H}_t \preceq \frac{1}{(1 - \|\tilde{\Delta}_t\|_{\mathbf{H}})^2} \mathbf{H},$$

and note that  $\frac{1}{(1 - \|\tilde{\Delta}_t\|_{\mathbf{H}})^2} \leq 1 + \epsilon$  for  $\epsilon = O(\frac{1}{\sqrt{d_{\text{eff}}}})$ , so it follows that  $\mathbf{H}_t \approx_{\epsilon} \mathbf{H}$ . Furthermore, for self-concordant functions, it follows that:

$$\begin{aligned} \|\mathbf{H}_t \tilde{\Delta}_t - \mathbf{g}_t\|_{\mathbf{H}_t^{-1}} &= \left\| \mathbf{H}_t^{\frac{1}{2}} \tilde{\Delta}_t - \mathbf{H}_t^{-\frac{1}{2}} \left( \int_0^1 \nabla^2 f(\mathbf{x}^* + \tau \tilde{\Delta}_t) d\tau \right) \tilde{\Delta}_t \right\| \\ &= \left\| \left( \int_0^1 (\mathbf{I} - \mathbf{H}_t^{-\frac{1}{2}} \nabla^2 f(\mathbf{x}^* + \tau \tilde{\Delta}_t) \mathbf{H}_t^{-\frac{1}{2}}) d\tau \right) \mathbf{H}_t^{\frac{1}{2}} \tilde{\Delta}_t \right\| \\ &\leq \|\tilde{\Delta}_t\|_{\mathbf{H}_t} \cdot \int_0^1 \|\mathbf{I} - \mathbf{H}_t^{-\frac{1}{2}} \nabla^2 f(\mathbf{x}^* + \tau \tilde{\Delta}_t) \mathbf{H}_t^{-\frac{1}{2}}\| d\tau \\ &\leq \|\tilde{\Delta}_t\|_{\mathbf{H}_t} \cdot \int_0^1 \frac{1}{(1 - \tau \|\tilde{\Delta}_t\|_{\mathbf{H}_t})^2} d\tau = \frac{\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2}{1 - \|\tilde{\Delta}_t\|_{\mathbf{H}_t}}. \end{aligned}$$

Using the neighborhood condition, we conclude that  $\frac{\|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2}{1 - \|\tilde{\Delta}_t\|_{\mathbf{H}_t}} \leq O(\frac{\sqrt{d_{\text{eff}}}}{m}) \|\tilde{\Delta}_t\|_{\mathbf{H}_t} \leq \epsilon \beta \|\tilde{\Delta}_t\|_{\mathbf{H}_t}$ .

**Chaining the expectations.** Let  $\mathcal{E}_t$  denote the high-probability event corresponding to the conditional expectation in (10) for the iteration  $t$ . It remains to show that after conditioning on event  $\mathcal{E} = \bigwedge_{t=0}^{T-1} \mathcal{E}_t$ , we maintain that  $\tilde{\mathbf{x}}_t \in U$  for all  $t$ . Assume that this holds for  $t = 0$ . Then, it suffices to show that  $\|\tilde{\Delta}_{t+1}\|_{\mathbf{H}} \leq \|\tilde{\Delta}_t\|_{\mathbf{H}}$  for every  $t$  almost surely (conditioned on  $\mathcal{E}$ ). Recall that Theorem 6 implies that conditioned on  $\mathcal{E}_t$  we have

$$\|\mathbf{I} - \tilde{\mathbf{Q}}\| \leq \eta,$$

where Lemma 7 ensures that  $\eta$  is small. We use this to show the following coarse convergence guarantee that holds almost surely conditioned on  $\mathcal{E}_t$ , but is substantially weaker than  $\beta$ . First, note that using the derivation as in the proof of Lemma 11 and the analysis of the exact Newton step,

$$\begin{aligned} \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 &\leq \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 + O(\|\mathbf{I} - \tilde{\mathbf{Q}}\|) \cdot \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2 \\ &\leq (\beta + O(\eta)) \cdot \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2. \end{aligned}$$

Using a sufficiently large constant  $C$  in Lemma 7 so that  $\beta + O(\eta)$  is small enough, and given the assumption  $\mathbf{H}_t \approx_{\epsilon} \mathbf{H}$ , we obtain  $\|\tilde{\Delta}_{t+1}\|_{\mathbf{H}}^2 \leq \|\tilde{\Delta}_t\|_{\mathbf{H}}^2$ . Thus, we conclude that all of the iterates will lie in the neighborhood  $U$ , and so (10) will hold for all  $t = 0, 1, \dots, T-1$ . Finally, note that by the union bound, event  $\mathcal{E}$  holds with probability  $1 - \delta$ , which completes the proof.

**Lower-bound from Theorem 1** The matching lower-bound from Theorem 1 holds only in the unregularized setting. In this case, we have  $\nabla^2 f_0(\tilde{\mathbf{x}}_t) = \mathbf{H}_t$ , so instead of an inequality in (9), we can obtain a two-sided approximation. The rest of the proof proceeds identically.

## C Sketches Satisfying Structural Conditions (Proof of Lemma 7)

In this section, we prove Lemma 7, showing that sub-Gaussian, LESS, and LESS-uniform embeddings all satisfy the assumptions of Theorem 6, which are derived from Conditions 1 and 2. This analysis follows along similar lines as in [DLDM21], except for extending LESS embeddings to LESS-uniform, and allowing for the presence of regularization.

## C.1 Sub-Gaussian embeddings

For sub-Gaussian embeddings, both conditions follow from existing results. To establish Condition 1, we rely on a covariance estimation result of [KL17], stated as in [MZ20]. Here, we say that a random vector  $\mathbf{x}$  is sub-Gaussian if  $\mathbf{v}^\top \mathbf{x}$  is a sub-Gaussian variable for all unit vectors  $\mathbf{v}$ .

**Lemma 15** ([MZ20, Theorem 1.4]). *For  $i \in \{1, \dots, m\}$ , let  $\mathbf{x}_i \in \mathbb{R}^n$  be independent sub-Gaussian random vectors such that  $\mathbb{E}[\mathbf{x}_i] = 0$  and  $\mathbb{E}[\mathbf{x}_i \mathbf{x}_i^\top] = \Sigma$ . Then, it holds with probability at least  $1 - 2 \exp(-t^2)$  that*

$$\left\| \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \mathbf{x}_i^\top - \Sigma \right\| \leq C \|\Sigma\| \left( \sqrt{\frac{\text{tr} \Sigma / \|\Sigma\|}{m}} + \frac{\text{tr} \Sigma / \|\Sigma\|}{m} + \frac{t}{\sqrt{m}} + \frac{t^2}{m} \right).$$

Now, to establish Condition 1 for the sub-Gaussian sketching matrix  $\mathbf{S}$ , i.e., where the  $m$  rows are distributed as  $\frac{1}{\sqrt{m-d_{\text{eff}}}} \mathbf{s}_i^\top$  for  $\mathbf{s}_i$  having i.i.d. zero mean, unit variance and sub-Gaussian entries, we let  $\mathbf{x}_i = \mathbf{U}^\top \mathbf{s}_i$ . Recall that  $\mathbf{U} = \mathbf{A} \mathbf{H}^{-\frac{1}{2}}$  for  $\mathbf{H} = \mathbf{A}^\top \mathbf{A} + \mathbf{C}$ , so  $\|\mathbf{U}\| \leq 1$ , so it follows that  $\mathbf{x}_i$  is a sub-Gaussian random vector. Therefore, letting  $\gamma = \frac{m}{m-d_{\text{eff}}}$  and  $\Sigma = \mathbf{U}^\top \mathbf{U}$ , we have  $\mathbb{E}[\mathbf{x}_i \mathbf{x}_i^\top] = \Sigma$  and with probability  $1 - \delta$ :

$$\begin{aligned} \|\mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} - \Sigma\| &\leq \gamma \cdot \left\| \frac{1}{m} \sum_{i=1}^m \mathbf{x}_i \mathbf{x}_i^\top - \Sigma \right\| + (\gamma - 1) \cdot \|\Sigma\| \\ &\leq C \left( \sqrt{\frac{d_{\text{eff}}}{m}} + \sqrt{\frac{\log(1/\delta)}{m}} \right) + \frac{d_{\text{eff}}}{m - d_{\text{eff}}}, \end{aligned}$$

thus setting  $m \geq O(1) \cdot (d_{\text{eff}} + \log(1/\delta))/\eta^2$ , we can bound the above by  $\eta$ , obtaining Condition 1.

To show Condition 2 for sub-Gaussian embeddings, we can again rely on a more general moment bound for quadratic forms, which is a special case of Lemma B.26 in [BS10].

**Lemma 16** ([BS10]). *Let  $\mathbf{M}$  be a  $n \times n$  matrix, and let  $\mathbf{x}$  be an  $n$ -dimensional random vector with independent, mean zero, unit variance entries such that  $\mathbb{E}[x_i^4] = O(1)$ . Then,*

$$\text{Var}[\mathbf{x}^\top \mathbf{M} \mathbf{x}] \leq O(1) \cdot \text{tr}(\mathbf{M} \mathbf{M}^\top).$$

To obtain Condition 2, we simply set  $\mathbf{x} = \mathbf{s}$  and  $\mathbf{M} = \mathbf{U} \mathbf{B} \mathbf{U}^\top$ . Note that since  $\|\mathbf{U}\| \leq 1$ , we have  $\text{tr}(\mathbf{M} \mathbf{M}^\top) = \text{tr}(\mathbf{U} \mathbf{B} \mathbf{U}^\top \mathbf{U} \mathbf{B} \mathbf{U}^\top) \leq \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top)$ .

## C.2 LESS embeddings: Condition 1

Now, we demonstrate that Condition 1 also holds for LESS embeddings. We will use the following matrix concentration inequality which is a straightforward combination of two standard results.

**Lemma 17** ([Tro12, Theorem 6.2] and [Tro15, Theorem 7.7.1]). *For  $i = 1, 2, \dots$ , consider a finite sequence  $\mathbf{X}_i$  of  $d \times d$  independent symmetric random matrices such that  $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$ , and one of the following holds for all  $i$ :*

1.  $\mathbb{E}[\mathbf{X}_i^p] \preceq \frac{p!}{2} \cdot R^{p-2} \mathbf{A}_i^2$  for  $p = 2, 3, \dots$ ;
2.  $\|\mathbf{X}_i\| \leq R$  and  $\mathbb{E}[\mathbf{X}_i^2] \preceq \mathbf{A}_i^2$ .

Then, defining the variance matrix  $\mathbf{V} = \sum_i \mathbf{A}_i^2$ , parameter  $\sigma^2 = \|\mathbf{V}\|$  and  $d_{\text{eff}} = \text{tr}(\mathbf{V})/\|\mathbf{V}\|$ , for any  $t \geq \sigma + R$  we have:

$$\Pr \left\{ \lambda_{\max} \left( \sum_i \mathbf{X}_i \right) \geq t \right\} \leq 4d_{\text{eff}} \cdot \exp \left( \frac{-t^2/2}{\sigma^2 + Rt} \right).$$

Before we can use matrix concentration, we must first establish high-probability concentration of the quadratic form  $\mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s}$ , for a leverage score sparsified sub-Gaussian random vector  $\mathbf{s}$ . This is an analog of the Hanson-Wright inequality, which holds for non-sparsified sub-Gaussian random vectors, as given below.

**Lemma 18** (Hanson-Wright inequality, [RV13, Theorem 1.1]). *Let  $\mathbf{x}$  have independent sub-Gaussian entries with mean zero and unit variance. Then, there is  $c = \Omega(1)$  such that for any  $n \times n$  matrix  $\mathbf{B}$  and  $t \geq 0$ ,*

$$\Pr \left\{ |\mathbf{x}^\top \mathbf{B} \mathbf{x} - \text{tr}(\mathbf{B})| \geq t \right\} \leq 2 \exp \left( -c \min \left\{ \frac{t^2}{\|\mathbf{B}\|_F^2}, \frac{t}{\|\mathbf{B}\|} \right\} \right).$$

Our version of this result for sparsified sub-Gaussian vectors is an extension of Lemma 31 of [DLDM21], introducing the effective dimension  $d_{\text{eff}}$  as opposed to the regular dimension  $d$ , and allowing a broader class of sparsifiers, so that we can cover the results for LESS-uniform embeddings.

**Lemma 19.** *Let  $\mathbf{U} = \mathbf{A}\mathbf{H}^{-\frac{1}{2}}$  for  $\mathbf{H} = \mathbf{A}^\top \mathbf{A} + \mathbf{C}$ , and let  $d_{\text{eff}} = \text{tr}(\mathbf{U}^\top \mathbf{U})$ . Let  $\boldsymbol{\xi}$  be a  $(p, s)$ -sparsifier and  $\mathbf{x}$  have independent sub-Gaussian entries with mean zero and unit variance. If  $p_i = \Omega(l_i(\mathbf{A}, \mathbf{C})/s)$  for all  $i$ , then for any  $t \geq Cd_{\text{eff}}$ , vector  $\mathbf{s} = \mathbf{x} \circ \boldsymbol{\xi}$  satisfies:*

$$\Pr \left\{ \mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s} \geq t \right\} \leq \exp \left( -c(\sqrt{t} + t/d_{\text{eff}}) \right).$$

*Proof.* The analysis follows along the same lines as the proof of Lemma 31 in [DLDM21]. First, we define the shorthand  $\bar{\mathbf{U}} = \text{diag}(\boldsymbol{\xi})\mathbf{U}$ , and use Lemma 17 to bound the spectral norm  $\|\bar{\mathbf{U}}\|$ . Observe that from the definition of the sparsifier  $\boldsymbol{\xi}$  we have the following decomposition:  $\bar{\mathbf{U}}^\top \bar{\mathbf{U}} = \sum_{i=1}^s \frac{1}{sp_{t_i}} \mathbf{u}_{t_i} \mathbf{u}_{t_i}^\top$ , where  $\mathbf{u}_{t_i}^\top$  denotes the  $i$ th row of  $\mathbf{U}$  and  $t_1, \dots, t_s$  are the independently sampled indices from  $p$ . Note that since  $l_i(\mathbf{A}, \mathbf{C}) = \|\mathbf{u}_{t_i}\|^2$ , we have that  $p_i = \Omega(\|\mathbf{u}_{t_i}\|^2/s)$ , so  $\mathbf{X}_i = \frac{1}{sp_{t_i}} \mathbf{u}_{t_i} \mathbf{u}_{t_i}^\top - \frac{1}{s} \mathbf{U}^\top \mathbf{U}$  satisfies  $\mathbb{E}[\mathbf{X}_i] = \mathbf{0}$ ,  $\|\mathbf{X}_i\| = O(1)$ , and  $\mathbb{E}[\mathbf{X}_i^2] = O(1/s) \cdot \mathbf{I}$ . So, using Lemma 17 with  $\sigma^2 = R = O(1)$ , for any  $t \geq Cd_{\text{eff}}$  we have  $\Pr\{\|\bar{\mathbf{U}}\|^2 \geq \sqrt{t}\} \leq \exp(-c\sqrt{t})$ , with  $c = \Omega(1)$ . Using the fact that  $\|\bar{\mathbf{U}}\|^2 \leq \text{tr}(\bar{\mathbf{U}}^\top \bar{\mathbf{U}}) \leq Cd_{\text{eff}}$  almost surely, it follows that the event  $\mathcal{E} : \|\bar{\mathbf{U}}\|^2 \leq \min\{\sqrt{t}, Cd_{\text{eff}}\}$  has probability  $1 - \exp(-c(\sqrt{t} + t/d_{\text{eff}}))$ . Now, it suffices to condition on  $\boldsymbol{\xi}$  and apply the Hanson-Wright inequality (Lemma 18), concluding that:

$$\begin{aligned} \Pr \left\{ \mathbf{x}^\top \bar{\mathbf{U}} \bar{\mathbf{U}}^\top \mathbf{x} \geq Cd_{\text{eff}} + t \mid \boldsymbol{\xi}, \mathcal{E} \right\} &\leq 2 \exp \left( -c \min \left\{ \frac{t^2}{\|\bar{\mathbf{U}}\|_F^2}, \frac{t}{\|\bar{\mathbf{U}}\|} \right\} \right) \\ &\leq 2 \exp(-\Omega(\sqrt{t} + t/d_{\text{eff}})), \end{aligned}$$

which completes the proof. ■

By appropriately integrating out the concentration inequality from Lemma 19, as in Lemma 30 of [DLDM21] but replacing  $d$  with  $d_{\text{eff}}$ , we can show the following matrix moment bound.

**Lemma 20.** *Under the assumptions of Lemma 19, for all  $p = 2, 3, \dots$  we have:*

$$\left\| \mathbb{E} \left[ \left( \mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U} - \mathbf{U}^\top \mathbf{U} \right)^p \right] \right\| \leq \frac{p!}{2} \cdot (Cd_{\text{eff}})^{p-1}.$$

*Proof.* First, we bound the expression in terms of the quadratic form  $\mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s}$ , so that we can use the concentration inequality from Lemma 19. To that end, we have:

$$\begin{aligned} \mathbb{E} \left[ \left( \mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U} - \mathbf{U}^\top \mathbf{U} \right)^p \right] &\preceq \mathbb{E} \left[ \left\| \mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U} - \mathbf{U}^\top \mathbf{U} \right\|^{p-2} \left( \mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U} - \mathbf{U}^\top \mathbf{U} \right)^2 \right] \\ &\stackrel{(*)}{\preceq} \mathbb{E} \left[ \left( \mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s} + d_{\text{eff}} \right)^{p-2} \left( 2(\mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U})^2 + 2(\mathbf{U}^\top \mathbf{U})^2 \right) \right] \\ &\preceq 2 \mathbb{E} \left[ \left( \mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s} + d_{\text{eff}} \right)^{p-1} \mathbf{U}^\top \mathbf{s} \mathbf{s}^\top \mathbf{U} \right] + 2 \mathbb{E} \left[ \left( \mathbf{s}^\top \mathbf{U} \mathbf{U}^\top \mathbf{s} + d_{\text{eff}} \right)^{p-2} \right] \cdot \mathbf{I}, \end{aligned}$$

where in  $(*)$  we used the fact that function  $f(x) = x^2$  is operator convex. Now, integrating out the concentration inequality from Lemma 19 for each of the two terms (following the steps of [DLDM21, Appendix D.2]), we obtain the desired bound. ■

We can now apply Lemma 17 with  $\mathbf{X}_i = \frac{1}{m} \mathbf{U}^\top \mathbf{s}_i \mathbf{s}_i^\top \mathbf{U} - \frac{1}{m} \mathbf{U}^\top \mathbf{U}$ , and  $\sigma^2 = R = O(d_{\text{eff}}/m)$ , obtaining that:

$$\Pr \{ \|\gamma^{-1} \mathbf{U}^\top \mathbf{S}^\top \mathbf{S} \mathbf{U} - \mathbf{U}^\top \mathbf{U}\| \geq \eta \} \leq d_{\text{eff}} \cdot \exp(-\Omega(\eta^2 m / d_{\text{eff}})).$$

Setting  $m \geq C d_{\text{eff}} \log(d_{\text{eff}}/\delta) / \eta^2$ , we obtain the desired bound. Note that we must account again for the scaling  $\gamma = \frac{m}{m-d_{\text{eff}}}$ , which gets absorbed into the error  $\eta$ .

Finally, observe that the conditions imposed on the sparsifier  $\xi$  in Lemma 19 encompass both LESS and LESS-uniform embeddings. In the case of LESS, we can simply let  $s \approx_{1/2} d_{\text{eff}}$ , and then the condition on sparsifying distribution is  $p_i = \Omega(l_i(\mathbf{A}, \mathbf{C}) / d_{\text{eff}})$ . On the other hand, for LESS-uniform, as long as  $s = \Omega(\tau d_{\text{eff}})$  where  $\tau = \frac{n}{d_{\text{eff}}} \max_i l_i(\mathbf{A}, \mathbf{C})$  is the coherence of  $\mathbf{A}$ , it follows that  $\frac{1}{n} = \Omega(l_i(\mathbf{A}, \mathbf{C}) / s)$  for all  $i$ 's, so a uniformly sparsifying distribution suffices.

### C.3 LESS embeddings: Condition 2

Here, we prove a result that is similar to the so-called Restricted Bai-Silverstein inequality from [DLDM21, Lemma 28]. Our assumptions on the sparsifier are somewhat weaker, to account for LESS-uniform embeddings and for the presence of regularization.

**Lemma 21.** *Let  $\mathbf{U} = \mathbf{A} \mathbf{H}^{-\frac{1}{2}}$  for  $\mathbf{H} = \mathbf{A}^\top \mathbf{A} + \mathbf{C}$ . Let  $\xi$  be a  $(p, s)$ -sparsifier and  $\mathbf{x}$  have independent sub-Gaussian entries with mean zero and unit variance. If  $p_i = \Omega(l_i(\mathbf{A}, \mathbf{C}) / s)$  for all  $i$ , then for all  $d \times d$  psd matrices  $\mathbf{B}$ , vector  $\mathbf{s} = \mathbf{x} \circ \xi$  satisfies:*

$$\text{Var}[\mathbf{s}^\top \mathbf{U} \mathbf{B} \mathbf{U}^\top \mathbf{s}] \leq O(1) \cdot \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top).$$

*Proof.* Let  $\bar{\mathbf{U}} = \text{diag}(\xi) \mathbf{U}$ . We start with a decomposition of the variance:

$$\begin{aligned} \text{Var}[\mathbf{s}^\top \mathbf{U} \mathbf{B} \mathbf{U}^\top \mathbf{s}] &= \mathbb{E}[(\mathbf{x}^\top \bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top \mathbf{x} - \text{tr}(\bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top) + \text{tr}(\bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top) - \text{tr}(\mathbf{B}))^2] \\ &= \mathbb{E}[\text{Var}[\mathbf{x}^\top \bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top \mathbf{x} \mid \bar{\mathbf{U}}] + \text{Var}[\text{tr}(\bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top)]]. \end{aligned}$$

Recall that  $\bar{\mathbf{U}}^\top \bar{\mathbf{U}} = \sum_{i=1}^s \frac{1}{sp_{t_i}} \mathbf{u}_{t_i} \mathbf{u}_{t_i}^\top$ , where  $\mathbf{u}_i^\top$  is the  $i$ th row of  $\mathbf{U}$  and  $p_i = \Omega(\|\mathbf{u}_i\|^2 / s)$ . Then

$$\text{Var}[\text{tr}(\bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top)] = s \text{Var}\left[\frac{\mathbf{u}_{t_1}^\top \mathbf{B} \mathbf{u}_{t_1}}{sp_{t_1}}\right] \leq \mathbb{E}\left[\frac{\|\mathbf{u}_{t_1}\|^2}{sp_{t_1}} \frac{\mathbf{u}_{t_1}^\top \mathbf{B}^2 \mathbf{u}_{t_1}}{p_{t_1}}\right] = O(1) \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top),$$

where we use that  $\mathbf{U}^\top \mathbf{U} \preceq \mathbf{I}$ . Next, we use the classical Bai-Silverstein inequality (Lemma 16):

$$\begin{aligned} \mathbb{E}[\text{Var}[\mathbf{x}^\top \bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top \mathbf{x} \mid \bar{\mathbf{U}}]] &\leq O(1) \cdot \mathbb{E}[\text{tr}((\bar{\mathbf{U}} \mathbf{B} \bar{\mathbf{U}}^\top)^2)] = O(1) \cdot \mathbb{E}\left[\text{tr}\left(\left(\sum_{i=1}^s \frac{1}{sp_{t_i}} \mathbf{B} \mathbf{u}_{t_i} \mathbf{u}_{t_i}^\top\right)^2\right)\right] \\ &= O(1) \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top) + O(1) \text{tr}((\mathbf{U} \mathbf{B} \mathbf{U}^\top)^2), \end{aligned}$$

where the last step follows by breaking down the expanded square into the diagonal part and the cross-terms. Since  $\text{tr}((\mathbf{U} \mathbf{B} \mathbf{U}^\top)^2) \leq \text{tr}(\mathbf{U} \mathbf{B}^2 \mathbf{U}^\top)$ , this completes the proof.  $\blacksquare$

## D Distributed Averaging for Newton-LESS

An important property of the Gaussian Newton Sketch is that it produces unbiased estimates of the exact Newton step. This is useful in distributed settings, where we can construct multiple independent estimates in parallel, and then produce an improved estimate by averaging them together. Newton-LESS retains this unbiasedness property, up to a small error, which also makes it amenable to distributed averaging. This near-unbiasedness of LESS embeddings follows from the characterization of the first inverse moment of the sketched Hessian (see [DLDM21] and the first part of Theorem 6).

In this section, we show that the near-unbiasedness of LESS embeddings can be combined with our new convergence analysis to provide improved convergence rates for Distributed Newton-LESS:

$$\tilde{\mathbf{x}}_{t+1} = \tilde{\mathbf{x}}_t - \frac{\mu t}{q} \sum_{i=1}^q (\mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_{t,i}^\top \mathbf{S}_{t,i} \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) + \nabla^2 g(\tilde{\mathbf{x}}_t))^{-1} \nabla f(\tilde{\mathbf{x}}_t), \quad (11)$$

where  $\mathbf{S}_{t,i}$  are independently drawn LESS embedding matrices. To adapt our analysis for this algorithm, we extend the characterization from Lemma 11.

**Lemma 22.** Fix  $\mathbf{H}_t = \nabla^2 f(\tilde{\mathbf{x}}_t)$  and let  $\tilde{\mathbf{x}}_{t+1}$  be as in (11) with  $\mathbf{S}_{t,i}$  as in Lemma 7 (i.e., sub-Gaussian, LESS or LESS-uniform). Also, suppose that the exact Newton step  $\mathbf{x}_{t+1} = \tilde{\mathbf{x}}_t - \mu_t \mathbf{H}_t^{-1} \mathbf{g}_t$  is a descent direction, i.e.,  $\|\Delta_{t+1}\|_{\mathbf{H}_t} \leq \|\tilde{\Delta}_t\|_{\mathbf{H}_t}$  where  $\Delta_{t+1} = \mathbf{x}_{t+1} - \mathbf{x}^*$  and  $\tilde{\Delta}_t = \tilde{\mathbf{x}}_t - \mathbf{x}^*$ . Then, letting  $\rho = \frac{d_{\text{eff}}(\tilde{\mathbf{x}}_t)}{m - d_{\text{eff}}(\tilde{\mathbf{x}}_t)}$ , we have:

$$\mathbb{E}_{q\delta} \|\tilde{\Delta}_{t+1}\|_{\mathbf{H}_t}^2 = \|\Delta_{t+1}\|_{\mathbf{H}_t}^2 + \frac{\rho}{q} \|\mathbf{x}_{t+1} - \tilde{\mathbf{x}}_t\|_{\nabla^2 f_0(\tilde{\mathbf{x}}_t)}^2 \pm O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right) \|\tilde{\Delta}_t\|_{\mathbf{H}_t}^2.$$

*Proof.* The proof is analogous to the proof of Lemma 11, except we must replace  $\tilde{\mathbf{Q}}$  with

$$\bar{\mathbf{Q}} = \frac{1}{q} \sum_{i=1}^q \tilde{\mathbf{Q}}_i, \quad \text{for} \quad \tilde{\mathbf{Q}}_i = \mathbf{H}_t^{\frac{1}{2}} (\mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t)^\top \mathbf{S}_{t,i}^\top \mathbf{S}_{t,i} \mathbf{A}_{f_0}(\tilde{\mathbf{x}}_t) + \nabla^2 g(\tilde{\mathbf{x}}_t))^{-1} \mathbf{H}_t^{\frac{1}{2}}.$$

Each  $\tilde{\mathbf{Q}}_i$  satisfies the first and second moment characterizations from Theorem 6. Let  $\mathcal{E} = \bigwedge_{i=1}^q \mathcal{E}_i$  denote the intersection of the corresponding  $1 - \delta$  probability events. Then,  $\|\mathbb{E}_{\mathcal{E}}[\bar{\mathbf{Q}}] - \mathbf{I}\| \leq O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right)$  and also:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\bar{\mathbf{Q}}^2] - \mathbf{I} &= \frac{1}{q^2} \sum_{i=1}^q \mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_i^2] - \sum_{i \neq j} \mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_i] \mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_j] - \mathbf{I} \\ &= \frac{1}{q} (\mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_1^2] - \mathbf{I}) + \frac{q(q-1)}{q^2} (\mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_1]^2 - \mathbf{I}), \end{aligned}$$

so using that  $\|\mathbb{E}_{\mathcal{E}}[\tilde{\mathbf{Q}}_1^2] - (\mathbf{I} + \rho \mathbf{U}^\top \mathbf{U})\| \leq O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right)$ , where  $\mathbf{U} = \mathbf{A}_{f_0}(\tilde{\mathbf{x}}) \mathbf{H}_t^{-\frac{1}{2}}$ , we get:

$$\|\mathbb{E}_{\mathcal{E}}[\bar{\mathbf{Q}}^2] - (\mathbf{I} + \frac{\rho}{q} \mathbf{U}^\top \mathbf{U})\| \leq O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right).$$

The rest of the proof follows identically as in Lemma 11, using  $\bar{\mathbf{Q}}$  in place of  $\tilde{\mathbf{Q}}$ . Note that, using the union bound, we can show that the probability of  $\mathcal{E}$  is at least  $1 - q\delta$ .  $\blacksquare$

From this lemma, repeating the local convergence analysis of Theorem 10, we obtain that in the neighborhood of  $\mathbf{x}^*$ , setting  $\mu_t = \frac{q(m - d_{\text{eff}})}{d_{\text{eff}} + q(m - d_{\text{eff}})}$ , Distributed Newton-LESS achieves:

$$\left( \mathbb{E}_{qT\delta} \frac{\|\tilde{\mathbf{x}}_T - \mathbf{x}^*\|_{\mathbf{H}}^2}{\|\tilde{\mathbf{x}}_0 - \mathbf{x}^*\|_{\mathbf{H}}^2} \right)^{1/T} \leq \frac{d_{\text{eff}}}{d_{\text{eff}} + q(m - d_{\text{eff}})} + O\left(\frac{\sqrt{d_{\text{eff}}}}{m}\right),$$

and for the unregularized case, where  $d_{\text{eff}} = \tilde{d}_{\text{eff}} = d$ , we can obtain a matching lower bound on the convergence rate. This shows that the convergence rate of Newton-LESS can be substantially improved via distributed averaging.

## E Additional Numerical Experiments and Implementation Details

Experiments are implemented in Python using the Pytorch module on Amazon Sagemaker instances with CPUs with 256 gigabytes of memory and GPUs NVIDIA Tesla V100. The code is publicly available at <https://github.com/lessketching/newtonsketch>.

### E.1 Sketching matrices

Given a data matrix  $\mathbf{A} \in \mathbb{R}^{n \times d}$ , we follow the procedure described in [DMIMW12b] (see Algorithm 1 therein) for fast approximation of the leverage scores. We use these approximate leverage scores to compute the RSS-lev-score embedding.

For LESS embeddings, we report the performance of the computationally most efficient method between using the approximate leverage scores, or, pre-processing the data matrix  $\mathbf{A}$  by a Hadamard matrix  $\mathbf{H}$  and then using a uniformly sparsified sketching matrix. Preprocessing with a Hadamard matrix uniformizes the leverage scores, so this second option is a valid implementation of a LESS embedding (see [DLDM21] for a detailed discussion). We found this second option to be the fastest method in practice.



For a LESS-uniform embedding  $\mathbf{S} \in \mathbb{R}^{m \times n}$ , we fix a number of non-zero entries per row to be  $s$ . For each row  $\mathbf{s}_i^\top$ , we sample  $s$  indices  $\{i_1, \dots, i_s\}$  in  $\{1, \dots, n\}$  uniformly at random with replacement. Each entry  $S_{i_i}$  is then chosen uniformly at random in  $\{\pm(n/ms)^{1/2}\}$ . We choose to sample with replacement for maximal computational efficiency. In our experiments, the number of non-zero entries  $s$  is small in comparison to the sample size  $n$ , so the probability of sampling twice the same index remains very small.

## E.2 Datasets

The high-coherence synthetic data matrix  $\mathbf{A}$  is generated as follows. We construct a covariance matrix  $\Sigma \in \mathbb{R}^{d \times d}$  with entries  $\Sigma_{ij} = 2 \cdot 0.5^{|i-j|}$ . The rows  $\mathbf{a}_i$  of  $\mathbf{A}$  are then sampled independently as  $\mathbf{a}_i \sim \mathbf{g}_i / \sqrt{z_i}$  where  $\mathbf{g}_i \sim \mathcal{N}(0, \Sigma)$  and  $z_i$  follows a Gamma distribution with shape  $1/2$  and scale  $2$ . We use  $n = 16384$  and  $d = 256$ .

We downloaded the Musk and CIFAR-10 datasets from <https://www.openml.org/>. The Musk data matrix has size  $n = 4096$  and  $d = 256$ . The sample size of the CIFAR-10 dataset is  $n = 50000$ . We transform each image using a random features map that approximates the Gaussian kernel  $\exp(-\gamma x^2)$  with bandwidth  $\gamma = 0.02$ , and we use  $d = 2000$  random cosine components. We partition the ten classes of CIFAR-10 into two groups with labels 0 and 1.

For the WESAD dataset [SRD<sup>+</sup>18], we used the data obtained from the E4 Empatica device and we filtered the data over windows of one second.<sup>2</sup> This results in a sample size  $n = 262144$ . Then we applied a random features map that approximates the Gaussian kernel  $\exp(-\gamma x^2)$  with  $\gamma = 0.01$  and we use  $d = 2000$  components.

## E.3 Least squares regression

We first consider least squares regression. On Figure 4, we report the relative error versus number of iterations, as well as the relative error versus wall-clock time for the Newton Sketch. We compare to Gaussian embeddings, the SRHT, uniformly random row sampling matrices (RRS) and random row sampling based on approximate leverage scores (RRS-lev-scores). As predicted by our theory, LESS embeddings have convergence rate scaling as  $d/m$ . This is similar to the convergence rate of the Newton sketch with Gaussian embeddings [LP19]. We also observe similar convergence for the SRHT, which is not explained by existing worst-case theory [Tro11], but it matches the predictions based on high-dimensional asymptotic analysis of the SRHT [LLDP20]. Except for CIFAR-10, RSS and RSS-lev-scores have weaker convergence rates. This suggests that the CIFAR-10 data matrix has low coherence. Except for the high-coherence synthetic data matrix for which the convergence rate is slightly worse than  $d/m$ , using LESS with a uniformly random sparsifier does not affect the convergence rate. Here, we implement LESS-uniform with  $d$  non-zero entries per row subsampled uniformly at random. Importantly, LESS-uniform offers significant speed-ups over other sketching matrices.

Note that some curves stop earlier than others (e.g., RRS) on the wall-clock time versus error plots, because we run the Newton sketch for each embedding for a fixed number of iterations.

## E.4 Regularized least squares and effective dimension

In Figure 5, we report the error versus number of iterations of the Newton Sketch for regularized least squares regression. These results illustrate in particular our theoretical predictions: the convergence rate of Newton-LESS is upper bounded by  $d_{\text{eff}}/m$ . In fact, Newton-LESS has the same convergence rate as the Newton Sketch with dense Gaussian embeddings.

## E.5 Comparison with CountSketch

In Figure 6, we perform an empirical comparison of LESS-uniform with the Sparse Johnson Lindenstrauss Transform (SJLT) with one non-zero entry per column, also known as the CountSketch [CW17]. Here, we consider the wall clock time convergence for the Newton Sketch solving a least squares task on the high-coherence synthetic matrix. We present the results alongside Gaussian and Random Row Sampling (RRS) sketches.

<sup>2</sup>We refer to the public repository for implementation details about subsampling the signal, [https://github.com/WJMatthew/WESAD/blob/master/data\\_wrangling.py](https://github.com/WJMatthew/WESAD/blob/master/data_wrangling.py).

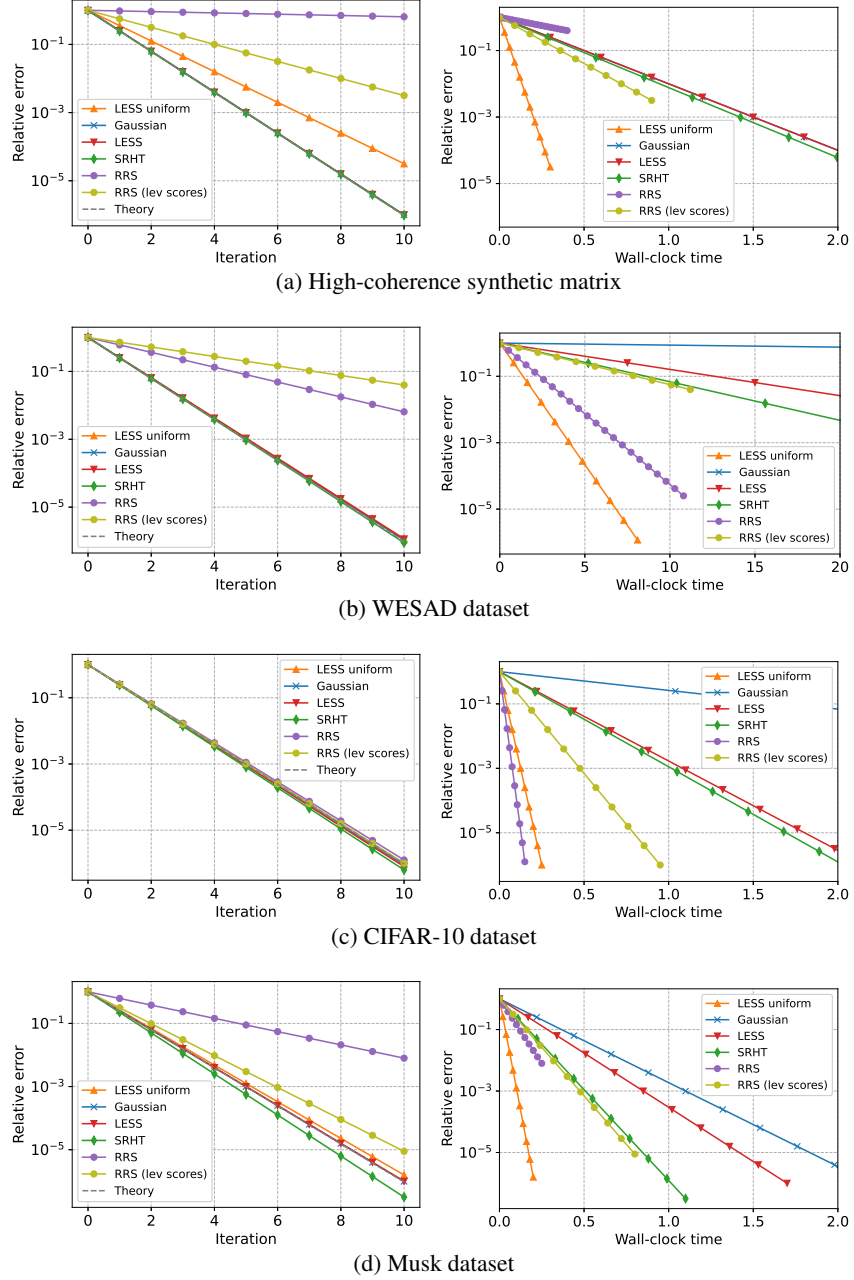


Figure 4: Newton sketch for least squares regression. We use the sketch size  $m = 4d$  for all experiments. Results are averaged over 10 trials.

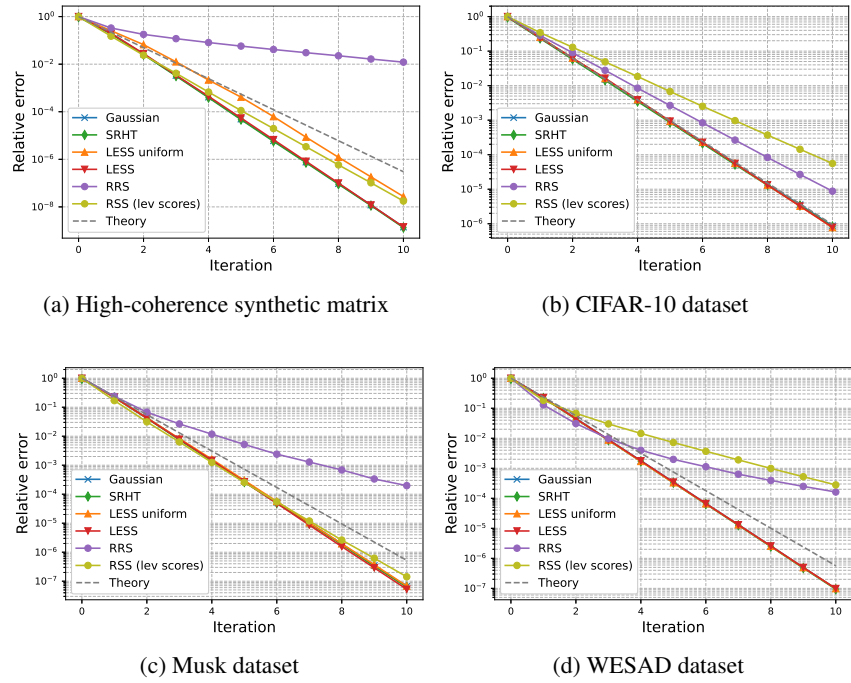


Figure 5: Newton Sketch for regularized least squares regression. We use the sketch size  $m = 4d_{\text{eff}}$  for all experiments. Results are averaged over 10 trials.

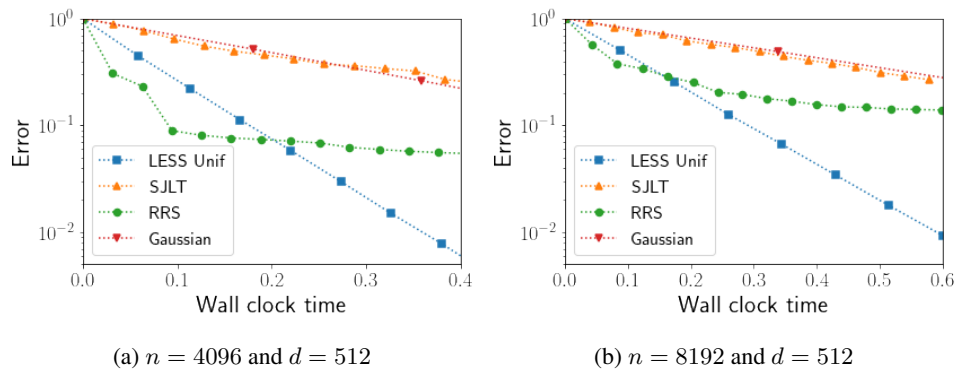


Figure 6: Relative error versus wall clock time: Newton Sketch for least squares regression on the high-coherence synthetic data matrix. We use sketch size  $m = 1024$ . For LESS-uniform, we use a number of non-zero entries  $s = 0.3d$  per row. For SJLT, we use the standard CountSketch implementation [CW17]. Results are averaged over 10 trials.