
KrADagrad: Kronecker Approximation-Domination Gradient Preconditioned Stochastic Optimization

Jonathan Mei^{1†}

Alexander Moreno^{1†}

Luke Walters^{1†}

¹Independent Researcher

Abstract

Second order stochastic optimizers allow parameter update step size and direction to adapt to loss curvature, but have traditionally required too much memory and compute for deep learning. Recently, Shampoo [Gupta et al., 2018] introduced a Kronecker factored preconditioner to reduce these requirements: it is used for large deep models [Anil et al., 2020] and in production [Anil et al., 2022]. However, it takes inverse matrix roots of ill-conditioned matrices. This requires 64-bit precision, imposing strong hardware constraints. In this paper, we propose a novel factorization, Kronecker Approximation-Domination (KrAD). Using KrAD, we update a matrix that directly approximates the inverse empirical Fisher matrix (like full matrix AdaGrad), avoiding inversion and hence 64-bit precision. We then propose KrADagrad*, with similar computational costs to Shampoo and the same regret. Synthetic ill-conditioned experiments show improved performance over Shampoo for 32-bit precision, while for several real datasets we have comparable or better generalization.

1 INTRODUCTION

Second order stochastic optimization methods adapt to loss curvature, allowing for smaller parameter update steps in regions where the gradient changes quickly, avoiding bouncing behavior, and larger ones in flat regions. Traditionally, they required storing and inverting the Hessian to update parameters: this requires quadratic memory and cubic computation in the number of parameters. Thus, methods using only a diagonal Hessian/Fisher approximation [Duchi et al., 2011, Kingma and Ba, 2015] have dominated the field. However, diagonal preconditioners only scale gradients in the

canonical basis, while full preconditioners can potentially perform scaling in a rotated basis aligning more closely with loss curvature.

Recently, Shampoo [Gupta et al., 2018, Anil et al., 2020] proposed approximating the (empirical) Fisher matrices using Kronecker factorized matrices. The matrix version of Shampoo factorizes the full preconditioner matrix into left and right Kronecker factors, which allows storing and inverting the smaller factors instead of the full matrix. For parameters $\mathbf{W} \in \mathbb{R}^{m \times n}$, this reduces computation costs from $O(m^3 n^3)$ to $O(m^3 + n^3)$ and storage costs from $O(m^2 n^2)$ to $O(m^2 + n^2)$. AdaGrad [Duchi et al., 2011] uses regret bound techniques based on Online Mirror Descent (OMD) Srebro et al. [2011] designed for vector updates. To use these techniques for matrix/tensor updates, [Gupta et al., 2018] exploit domination results that relate vector update preconditioners to their matrix/tensor counterparts. However, Shampoo still requires inverse matrix roots, which are numerically unstable or inaccurate for ill-conditioned matrices in 32-bit precision. For preconditioned gradient descent, it is important for Shampoo to maintain the accuracy of the smallest eigenvalues (largest when inverted). It thus needs 64-bit precision, which requires some combination of slow TPU-CPU data transfers, stale preconditioner matrices, or even new machine learning accelerator hardware [Anil et al., 2020] supporting fast 64-bit matrix multiplication or fast accurate eigendecomposition.

A primary motivator to use any optimizer is to reach the same quality solution in less time or reach a better solution that other optimizers fail to reach. Second order optimizers are currently not as popular as first order methods due to: a) inertia to adoption, with a lack of highly optimized implementations in all major ML frameworks; b) the added compute and memory requirements; c) numerical stability and consequently the additional considerations required to get them to work or to debug them (e.g. numerical linear algebra, computer number formats); and d) even though they sometimes reach a solution unreachable by 1st order methods (or the same solution with lower wall clock time

²Funded by and conducted at Luminous Computing.

(WCT) if properly optimized), they don't consistently for every task or architecture.

The key tradeoff is (b, c) vs (d). Shukla [2022] of Weights and Biases noted that some of their customers using Shampoo do find solutions that generalize better than those found with ADAM for their real world tasks (d), but that 2nd order optimizers are still more expensive (b) and have additional considerations (c).

In this paper, we address these limitations (c) by introducing a novel factorization, Kronecker Approximation-Domination (KrAD): it has a simple form that updates the preconditioning matrix without explicitly inverting it. Shampoo constructs Kronecker factors of intermediate statistics such that their Kronecker product dominates the gradient outer product matrix. Our key idea is to construct factors for those statistics such that the *inverse* of their Kronecker product dominates the gradient outer product matrix. This leads to preconditioners that require fractional powers rather than inverse fractional powers of factors. While this does not decrease the computational complexity compared to Shampoo, it avoids needing 64-bit precision.

This paper has three primary contributions: 1) we introduce two new algorithms, both with $O(m^3 + n^3)$ and $O(m^2 + n^2)$ computational and memory complexity, respectively and only requiring positive matrix roots, in contrast to previous work requiring inverse matrix roots; 2) we show domination properties and use them to prove that our algorithm, which has similar computation cost to Shampoo, achieves optimal regret; 3) we show empirically that in 32-bit precision, we outperform Shampoo in synthetic experiments and perform similarly on some real experiments. We first describe some mathematical tools, set up the problem, and describe second order optimization and established results related to our method in Section 2. Then we present our method and its theoretical properties in Section 3. Next, we consider the practical implementation of our method in Section 4. Then, we show empirical results in Section 5. Finally, we discuss implications in Section 6.

2 BACKGROUND AND RELATED WORK

Here, we set up the notation and the problem and then describe relevant related works. We briefly describe optimization for vector-valued parameters then extend the discussion to matrix-valued parameters. We can further generalize to tensors, but continue with matrices in the main text for clarity and leave the tensor formulation for Appendix C.5.

2.1 NOTATION AND PRELIMINARIES

We use bold lower case letters to denote column vectors (e.g. $\mathbf{g} \in \mathbb{R}^n$), bold upper case letters to denote matrices (e.g. $\mathbf{G} \in \mathbb{R}^{m \times n}$), and calligraphic letters to denote

matrices composed of stacking vectors of interest (e.g. $\mathcal{G}_k = (\mathbf{g}_k \ \mathbf{g}_{k-1} \ \dots \ \mathbf{g}_1) \in \mathbb{R}^{n \times k}$). For square matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, let the trace be $\text{tr}(\mathbf{A}) = \sum_{i=1}^n a_{i,i}$, where $a_{i,j}$ denotes the element in row i and column j of \mathbf{A} . For matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{m \times n}$, let $\mathbf{A} \cdot \mathbf{B} = \text{tr}(\mathbf{A}^\top \mathbf{B})$ be the matrix (Frobenius) inner product, and the induced Frobenius norm $\|\mathbf{A}\|_F = (\mathbf{A} \cdot \mathbf{A})^{1/2}$. We use $\|\mathbf{A}\|_2$ to denote the spectral norm, the largest singular value of \mathbf{A} . We write $\mathbf{A} \succeq 0$ to mean \mathbf{A} is symmetric positive semi-definite (PSD), while $\mathbf{A} \succ 0$ means \mathbf{A} is symmetric positive definite (PD). For two PSD matrices $\mathbf{B} \succeq \mathbf{A}$ means that $\mathbf{B} - \mathbf{A} \succeq 0$ (similarly for \succ). For a PSD matrix, take $\mathbf{A} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top$ to be the eigenvalue decomposition, which results in orthonormal \mathbf{V} (i.e. $\mathbf{V}^{-1} = \mathbf{V}^\top$). Define $f(\mathbf{\Lambda})$ such that the diagonal elements are $(f(\mathbf{\Lambda}))_{i,i} = f(\lambda_i)$, the principal values of the function applied to the scalar eigenvalues. Then we take $f(\mathbf{A}) = \mathbf{V}f(\mathbf{\Lambda})\mathbf{V}^\top$. In this way, we define a unique value for functions applied to PSD matrices with eigenvalues within the domain of the function. In particular, we have a definition for real powers of PD matrices.

Let \otimes denote the Kronecker product, for matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{q \times r}$ defined as

$$\mathbf{B} \otimes \mathbf{A} = \begin{pmatrix} b_{1,1}\mathbf{A} & \dots & b_{1,r}\mathbf{A} \\ \vdots & \ddots & \vdots \\ b_{q,1}\mathbf{A} & \dots & b_{q,r}\mathbf{A} \end{pmatrix} \in \mathbb{R}^{mq \times nr}.$$

Let the vectorization operation for a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ be

$$\text{vec}(\mathbf{A}) = (\mathbf{a}_1^\top \ \dots \ \mathbf{a}_n^\top)^\top \in \mathbb{R}^{mn};$$

\mathbf{a}_i is the i -th column of \mathbf{A} , and the corresponding inverse vectorization for $\mathbf{a} \in \mathbb{R}^{mn}$ given a target matrix in $\mathbb{R}^{m \times n}$ is

$$\text{vec}_{m,n}^{-1}(\mathbf{a}) = (\mathbf{a}_{1:m} \ \dots \ \mathbf{a}_{m(n-1)+1:mn}) \in \mathbb{R}^{m \times n},$$

where $\mathbf{a}_{i:j} = (a_i \ \dots \ a_j)^\top$.

Several properties [Bellman, 1980, Van Loan, 2000, Boyd and Vandenberghe, 2004, Baumgartner, 2011] of trace and Kronecker products are important for our results. We list them in Appendix D due to space constraints.

2.2 OPTIMIZATION IN MACHINE LEARNING

We are interested in iterative empirical risk minimization under loss f , with $\mathbf{x} \sim p_n(x)$ an empirical density and parameters $\mathbf{w} \in \mathbb{R}^N$

$$\mathbf{w}^* = \underset{\mathbf{w}}{\text{argmin}} \mathbb{E}_{p_n} [f(\mathbf{w}, \mathbf{x})].$$

We assume access to gradients $\nabla_{\mathbf{w}} f$. Let $\mathbf{g}_k = \sum_{i \in B_k} \nabla_{\mathbf{w}} f(\mathbf{w}_k, \mathbf{x}_i) \in \mathbb{R}^N$ be the estimated gradient of f w.r.t. \mathbf{w} evaluated at \mathbf{w}_k with data from batch B_k at iteration k . From here, we omit but imply ‘‘stochastic’’ gradients.

For step size $\eta_k \in \mathbb{R}^+$, gradient-based methods update

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k \mathbf{P}_k \mathbf{g}_k \quad (1)$$

where $\mathbf{P}_k \in \mathbb{R}^{N \times N}$ is a *preconditioner* matrix (some sources instead refer to \mathbf{P}_k^{-1} as the preconditioner). In some algorithms, additional intermediate *statistics* are stored and updated to aid preconditioner computation. Gradient descent uses $\mathbf{P}_k = \mathbf{I}_N$ (no preconditioning); Newton’s method takes $\mathbf{P}_k = \mathbf{H}_k^{-1}$ to be the (pseudo)inverse of the Hessian.

While vanilla gradient descent updates are trivial to compute, convergence can require many iterations. Newton updates are more expensive, but may require far fewer iterations. In practice, the chosen form of the preconditioner matrix appears to exist along a trade-off between computational tractability and improved convergence properties.

2.3 ADAPTIVE GRADIENT PRECONDITIONERS

We can collect gradients through iteration k ,

$$\mathcal{G}'_k = (\mathbf{g}_k \quad \mathbf{g}_{k-1} \quad \dots \quad \mathbf{g}_i \quad \dots \quad \mathbf{g}_1) \in \mathbb{R}^{N \times k}$$

and augment this with a scaled identity,

$$\mathcal{G}_k = (\mathcal{G}'_k \quad \epsilon \mathbf{I}_N) \in \mathbb{R}^{N \times (k+N)}.$$

One form of adaptive gradient update is

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k (\mathcal{G}_k \mathcal{G}_k^\top)^{-1/2} \mathbf{g}_k.$$

Expressing this in terms of the non-augmented \mathcal{G}'_k and taking $\delta = \epsilon^2$ gives the full version of AdaGrad [Duchi et al., 2011]: $\mathcal{G}_k \mathcal{G}_k^\top$ can be seen as the statistic that is stored and updated in each iteration, and $(\delta \mathbf{I}_N + \mathcal{G}'_k \mathcal{G}'_k{}^\top)^{-1/2}$ is the preconditioner computed from the statistic. Unfortunately, storing the full matrix $\mathcal{G}_k \mathcal{G}_k^\top$ is memory intensive at $O(N^2)$, and taking the inverse square root is computationally expensive at $O(N^3)$ to compute the SVD.

Diagonal AdaGrad reduces computational complexity

$$\mathbf{w}_{k+1} = \mathbf{w}_k - \eta_k \left(\delta \mathbf{I}_N + \text{diag} \left(\mathcal{G}'_k \mathcal{G}'_k{}^\top \right) \right)^{-1/2} \mathbf{g}_k.$$

This is $O(N)$ complexity in both memory and computation.

2.4 MATRIX VARIABLES AND SHAMPOO

Now, we consider $N = mn$ and optimize w.r.t. a matrix $\mathbf{W} \in \mathbb{R}^{m \times n}$. We consider a single matrix for clarity, but note that the derivations and analyses can be extended to tensors and applied individually to each tensor-valued parameter in a given model (e.g. to compute the total costs). Then we can use the same optimization framework, now taking $\mathbf{w} = \text{vec}(\mathbf{W}) \in \mathbb{R}^{mn}$. However, in forming a preconditioner, we utilize the fact that our parameter now has the additional structure of being a matrix. Consider

$$\mathbf{G}_k = \text{vec}_{m,n}^{-1}(\mathbf{g}_k).$$

One convenient factorized form of a preconditioner is

$$\mathbf{P}_k = \mathbf{R}_k \otimes \mathbf{L}_k,$$

where $\mathbf{L}_k \in \mathbb{R}^{m \times m}$ and $\mathbf{R}_k \in \mathbb{R}^{n \times n}$ are symmetric. This reduces storage and computation while not necessarily being low-rank. To see this, we return to Equation (1) and simplify

$$\begin{aligned} \mathbf{w}_{k+1} &= \mathbf{w}_k - \eta_k \mathbf{P}_k \mathbf{g}_k \\ &= \mathbf{w}_k - \eta_k \text{vec}(\mathbf{L}_k \mathbf{G}_k \mathbf{R}_k) \quad \because \text{(P18)} \\ \Rightarrow \mathbf{W}_{k+1} &= \mathbf{W}_k - \eta_k \mathbf{L}_k \mathbf{G}_k \mathbf{R}_k. \quad \because (\text{vec}_{m,n}^{-1}) \end{aligned}$$

This requires $O(m^2 + n^2) = O(N(\frac{m}{n} + \frac{n}{m}))$ storage and $O(N(m+n))$ compute. Unless otherwise specified, we assume w.l.o.g. that $m \leq n$. Shampoo [Gupta et al., 2018] tracks statistics

$$\mathbf{B}_k = \epsilon \mathbf{I}_m + \sum_{i=1}^k \mathbf{G}_i \mathbf{G}_i^\top, \quad \mathbf{C}_k = \epsilon \mathbf{I}_n + \sum_{i=1}^k \mathbf{G}_i^\top \mathbf{G}_i, \quad (2)$$

and forms the preconditioner from the Kronecker factors

$$(\mathbf{L}_k, \mathbf{R}_k) = (\mathbf{B}_k^{-1/4}, \mathbf{C}_k^{-1/4}). \quad (3)$$

Gupta et al. [2018] relies on 3 key conditions to prove Shampoo achieves optimal regret. First,

$$\mathbf{R}_k^{-2} \otimes \mathbf{L}_k^{-2} \succeq \epsilon \mathbf{I}_{mn} + \sum_{i=1}^k \mathbf{g}_i \mathbf{g}_i^\top. \quad (4)$$

It secondly requires that

$$\mathbf{R}_k^{-2} \otimes \mathbf{L}_k^{-2} \succeq \mathbf{R}_{k-1}^{-2} \otimes \mathbf{L}_{k-1}^{-2}. \quad (5)$$

Finally it requires that under mild conditions,

$$\text{tr}(\mathbf{L}_k), \text{tr}(\mathbf{R}_k) = O(k^{1/4}). \quad (6)$$

An additional $O(n^3 + m^3)$ cost comes from taking the inverse fractional powers via a high-precision (64-bit) Newton iteration (which involves repeated matrix multiplications; see Equation (25) in Appendix C.3 for further details) or SVD, which dominates the previous $O(N(m+n))$.

The motivation for these properties stems from OMD analysis. The regret vector parameter \mathbf{w}_k updates and general preconditioners $\mathbf{P}_k \succ 0$ is initially bounded by sums of quadratic forms $\mathbf{w}_k^\top \mathbf{P}_k \mathbf{w}_k$. If the domination property holds, these can be bounded in terms of $\text{tr}(\mathbf{P}_k)$. Using Property (P13), this can be further expressed in terms of $\text{tr}(\mathbf{L}_k)$ and $\text{tr}(\mathbf{R}_k)$. The trace growth rates give the final bound.

2.5 RELATED WORK

Recently, there has been a surge in interest in tractable preconditioned gradient methods. We briefly contrast some of the most similar or otherwise notable methods to ours.

Kronecker Factored: KFAC [Martens and Grosse, 2015] and TNT [Ren and Goldfarb, 2021] use Kronecker factors to approximate Fisher matrices while reducing storage and computation costs. KFAC requires knowledge of network architecture and thus modifications or even re-implementations corresponding to each parametric layer

type within the network. KFAC and TNT require an additional backward pass and matrix inversion. KBFGS [Goldfarb et al., 2020] does not require matrix inversion but requires an additional forward and backward pass. We note that the empirical performance achieved by KBFGS is partially due to initialization using curvature estimated from the entire training set [Goldfarb et al., 2020], which is not available in a truly online setting. Shampoo is the most closely related work, relying on the empirical Fisher matrix rather than estimating the Fisher matrix, thus not requiring additional sampling or forward/backward passes. In addition, only tensor shape knowledge is required. While estimating the Fisher matrix may have intuitively desirable properties over the empirical Fisher, the empirical Fisher is more practical to compute [Martens, 2020], since in distributed data or model parallel settings, additional forward/backward passes become prohibitive (both in terms of computation and engineering cost). To our knowledge, Shampoo is the only second order optimizer that has been successfully implemented in a large-scale, production, deep learning setting [Anil et al., 2022], which makes it of primary interest.

Limited Memory: GGT [Agarwal et al., 2019] uses a limited history of h past gradients to form a low-rank approximation to the full AdaGrad matrix, reducing storage costs to $O(Nh)$ and compute costs to $O(Nh^2)$; however, this requires many copies (200, for the problems they consider) of the full gradient to be stored as statistics (h still scales as a function of N), which can become prohibitive without modifications to reduce N .

Sketching: AdaHessian [Yao et al., 2021], SENG [Yang et al., 2022], and SketchySGD [Frangella et al., 2022] estimate the Hessian (either the diagonal or a low-rank approximation) via automatic differentiation to compute Hessian-vector products (HVP), which require additional backpropagation steps and two batches of data per step, one for gradient computation and one for Hessian sketching. This is less expensive than a full forward/backward pass, but can still be expensive in distributed settings. In addition, the low-rank factorization still requires many times the storage of the full gradient of the model (100 – 200× in the problems they consider). While KrADagrad does not use sketching or HVP, these methods could potentially be combined with KrAD factorization as future work.

3 KRADAGRAD

Here, we derive a pair of new optimization algorithms, KrADagrad and KrADagrad*, presenting along the way intermediate results that allow us to attain domination properties analogous to those in Equations (4)-(5) and trace growth rates in Equation (6) required for good regret. Simultaneously, this maintains low computational complexity achieved by Kronecker factorized methods. To derive the algorithms, we present

1. KrAD, a method for producing a Kronecker factorization approximating a matrix that yields the property in Equation (4)
2. Derivation of the basic form of KrADagrad, applying to AdaGrad style updates KrAD and the Woodbury matrix identity Woodbury [1950] as the key tricks
3. Statements confirming the property in Equation (5)
4. Extension to KrADagrad*.

KrADagrad alternates updates of Kronecker factors of the statistics and has within ε tolerance of optimal regret. With additional insights from KrADagrad regret analysis, we formulate KrADagrad*, which can be seen as an “average” of two KrADagrad estimators. KrADagrad* updates both Kronecker factors of the statistics simultaneously and achieves optimal regret. We present theoretical results along the way in this section as they are needed but defer the proofs to Appendix B. We derive the algorithm for matrix-valued parameters for clarity and again leave the extension to tensor-valued parameters for Appendix C.5.

3.1 KRONECKER APPROXIMATION-DOMINATION

First, we state a Lemma as the goal of KrAD, which is needed to achieve the condition in Equation (4) that allows us to prove optimal regret.

Lemma 3.1. *Let PD matrix $\mathbf{C} \in \mathbb{R}^{n \times n}$, $\mathbf{U} \in \mathbb{R}^{mn \times r}$, \mathbf{u}_i be the i -th column of \mathbf{U} , $\mathbf{U}_i = \text{vec}_{m,n}^{-1}(\mathbf{u}_i)$. Then*

$$\begin{aligned} \mathbf{B} &= \sum_{i=1}^r \mathbf{U}_i \mathbf{C}^{-1} \mathbf{U}_i^\top \succeq 0 \\ \Rightarrow \mathbf{C} \otimes \mathbf{B} &\succeq \frac{1}{n} \mathbf{U} \mathbf{U}^\top. \end{aligned}$$

These matrices \mathbf{U} and \mathbf{U} are fairly general. In our setting, we will use KrAD on gradient matrices. In context, this result states that given a PD right matrix, we can express a left matrix in a quadratic form, and the Kronecker product of the right and left matrices will dominate the scaled gradient outer product matrix. We present the proof in Appendix B.1.

3.2 KRADAGRAD UPDATES: DERIVATION

We start with deriving the general statistic update. We outline it here and fill in details in Appendix D. Suppose we have the previous Kronecker factorization of a statistic \mathbf{Q}_{k-1} that dominates the gradient outer product matrix by a factor t (which we will clarify later) and its inverse \mathbf{P}_{k-1}

$$\mathbf{Q}_{k-1} = \mathbf{C}_{k-1} \otimes \mathbf{B}_{k-1} \succeq \frac{1}{nt} \mathcal{G}_{k-1} \mathcal{G}_{k-1}^\top \quad (7)$$

$$\mathbf{P}_{k-1} = \mathbf{R}_{k-1} \otimes \mathbf{L}_{k-1} = \mathbf{C}_{k-1}^{-1} \otimes \mathbf{B}_{k-1}^{-1} = \mathbf{Q}_{k-1}^{-1}, \quad (8)$$

where \mathbf{B}_{k-1} , \mathbf{C}_{k-1} , \mathbf{L}_{k-1} , and \mathbf{R}_{k-1} are PD. Our initial intermediate update, which we will apply our KrAD factorization to, is

$$\tilde{\mathbf{Q}}_k = \frac{1}{nt_k} \mathbf{g}_k \mathbf{g}_k^\top + \mathbf{Q}_{k-1},$$

for some $t_k \leq t$. We will then compute an intermediate version of the update on the inverse of the left statistic $\tilde{\mathbf{B}}_k = \tilde{\mathbf{L}}_k^{-1}$. Letting $\tilde{\mathbf{B}}_k = \mathbf{B}_{k-1} + \Delta \mathbf{B}_k$, we can apply KrAD to $\tilde{\mathbf{Q}}_k$ with a fixed $\mathbf{C}_k = \mathbf{C}_{k-1}$ (i.e., we use the old right statistic to update the current left one),

$$\Delta \mathbf{B}_k = \frac{1}{t_k} \mathbf{G}_k \mathbf{C}_k^{-1} \mathbf{G}_k^\top \quad (9)$$

$$\Rightarrow \mathbf{C}_k \otimes \Delta \mathbf{B}_k \succeq \frac{1}{nt_k} \mathbf{g}_k \mathbf{g}_k^\top \quad \because \text{Lemma 3.1} \quad (10)$$

$$\Rightarrow \mathbf{C}_k \otimes (\mathbf{B}_{k-1} + \Delta \mathbf{B}_k) \succeq \frac{1}{nt} \mathbf{g}_k \mathbf{g}_k^\top \quad \because \text{(P15)}. \quad (11)$$

Then, letting $\hat{\mathbf{C}}_k = \mathbf{C}_k + \frac{1}{t_k} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k$ and applying the Woodbury matrix identity,

$$\tilde{\mathbf{L}}_k = \tilde{\mathbf{B}}_k^{-1} = (\mathbf{B}_{k-1} + \Delta \mathbf{B}_k)^{-1} \quad (12)$$

$$\succeq \underbrace{\mathbf{L}_{k-1} - \frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}}_{\mathbf{L}_k} \quad (13)$$

(we provide more detail in Appendix D). Note that our update for \mathbf{L}_k neither depends on \mathbf{B} , \mathbf{C} nor requires any other expensive matrix inverses to compute. This suggests that we do not need to actually store \mathbf{B} or \mathbf{C} to obtain a computationally tractable implementation.

Here, we state intermediate results that suggest that our proposed updates are reasonable (proof in Appendix B.2).

Proposition 3.2. *Taking $t_k = 1 + \|\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_2$ (or the looser but more computationally friendly $t_k = 1 + \|\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_F$), the PSD matrix*

$$\mathbf{M}_k \triangleq \frac{1}{t_k} \mathbf{L}_{k-1}^{1/2} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{L}_{k-1}^{1/2} \prec \mathbf{I}.$$

Corollary 3.3. *If $\mathbf{L}_{k-1} \succ 0$, the updated $0 \prec \mathbf{L}_k \preceq \mathbf{L}_{k-1}$.*

In Corollary 3.3, the second inequality $\mathbf{L}_k \preceq \mathbf{L}_{k-1}$ has the effect of not increasing the step size, while the first inequality $\mathbf{L}_k \succ 0$ guarantees that we do not reverse the direction of the gradient. These are both valued theoretical properties of useful preconditioners for avoiding divergent behavior. In practice, they may be lightly violated to great effect; for example, in Adam, it is technically possible to have increasing step size as shown by [Reddi et al., 2018]. The result from Corollary 3.3 also allows us to leverage existing techniques to bound the regret of our algorithm.

3.2.1 Update schemes

Thus far, we have glossed over the fact that we have actually only updated \mathbf{L}_k . We may update \mathbf{R}_k in the same way; since we have already achieved domination by just updating

\mathbf{L}_k , we are interested in the process for jointly updating $(\mathbf{L}_k, \mathbf{R}_k)$. Further, we have only hinted at a method for updating statistics; we must still compute the preconditioner.

In the next few subsections, we discuss these issues more concretely, proposing KrADagrad*, an algorithm that combines two sets of KrAD preconditioners to obtain optimal regret but requires updating two matrix statistics and involves higher order matrix roots. This algorithm is reminiscent of Shampoo, but avoids the numerical difficulty of inverting ill-conditioned matrices. We also propose, KrADagrad, a separate scheme that has suboptimal regret but is more intuitive, showing why we arrive at this form of update. Due to space constraints, we leave this to Appendix C.1.

3.3 KRADAGRAD*: COMBINING PRECONDITIONERS

Suppose we have two distinct sets of KrADagrad preconditioners. Here we overload notation a bit, holding the iteration k fixed and dropping it from the subscript, instead using the subscript to denote the index for the set of preconditioners to which each matrix belongs,

$$\mathbf{R}_1 \otimes \mathbf{L}_1 \succeq \mathcal{G} \mathcal{G}^\top, \quad \mathbf{R}_2 \otimes \mathbf{L}_2 \succeq \mathcal{G} \mathcal{G}^\top,$$

recalling that $\mathcal{G} = (\mathcal{G}' \ \epsilon \mathbf{I}_N)$, the augmented matrix collecting the history of observed gradients. Using the matrix geometric mean for two matrices [Ando et al., 2004],

$$\begin{aligned} M_g(\mathbf{A}, \mathbf{B}) &\triangleq \mathbf{A}^{1/2} (\mathbf{A}^{-1/2} \mathbf{B} \mathbf{A}^{-1/2})^{1/2} \mathbf{A}^{1/2} \\ &= \mathbf{A} (\mathbf{A}^{-1} \mathbf{B})^{1/2} \end{aligned}$$

and due to the geometry of the manifold of PD matrices [Bhatia, 2009],

$$\mathbf{R}_c \otimes \mathbf{L}_c \succeq \mathcal{G} \mathcal{G}^\top.$$

where $\mathbf{L}_c = M_g(\mathbf{L}_1, \mathbf{L}_2)$ and $\mathbf{R}_c = M_g(\mathbf{R}_1, \mathbf{R}_2)$ form another pair of KrAD estimates¹.

In this case, we may need to take additional square roots of the quantities $(\mathbf{L}_1^{-1} \mathbf{L}_2, \mathbf{R}_1^{-1} \mathbf{R}_2)$ and of $(\mathbf{L}_c, \mathbf{R}_c)$ themselves. We will first show how these combined estimators are relevant to Shampoo, and use this insight to arrive at a second version of preconditioning, which we call KrADagrad*.

3.3.1 Shampoo combines inverse KrAD estimates

If we keep a pair of KrAD estimates for \mathbf{Q} (an integral power of the preconditioner inverse) instead of directly for \mathbf{P} , one in which we only update \mathbf{B}_1 and keep $\mathbf{C}_1 = \mathbf{I}$ fixed, while in the other we only update \mathbf{C}_2 and keep $\mathbf{B}_2 = \mathbf{I}$ fixed, we end up with exactly the Shampoo statistics updates

$$\Delta \mathbf{B}_1 = \mathbf{G} \mathbf{C}_1^{-1} \mathbf{G}^\top = \mathbf{G} \mathbf{G}^\top$$

$$\Delta \mathbf{C}_2 = \mathbf{G}^\top \mathbf{B}_2^{-1} \mathbf{G} = \mathbf{G}^\top \mathbf{G}.$$

¹the subscript c stands for ‘‘combined’’

Algorithm 1 KrADagrad*: Precondition *without* inversion

Input: Parameters $\mathbf{W}_0 \in \mathbb{R}^{m \times n}$, iterations K , step size η , exponent $\alpha = 1/2$.
Initialize: $(\mathbf{L}_0, \mathbf{R}_0) = (\mathbf{I}_m, \mathbf{I}_n)$
for $k = 1, \dots, K$ **do**
 Obtain gradient \mathbf{G}_k
 Compute $\Delta \mathbf{L}_k = \frac{1}{t_{k,L}} \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}$
 Compute $\Delta \mathbf{R}_k = \frac{1}{t_{k,R}} \mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{G}_k \mathbf{R}_{k-1}$
 Update $(\mathbf{L}_k, \mathbf{R}_k) \leftarrow (\mathbf{R}_{k-1} - \Delta \mathbf{R}_k, \mathbf{L}_{k-1} - \Delta \mathbf{L}_k)$
 Compute $(\mathbf{L}_k^{\alpha/2}, \mathbf{R}_k^{\alpha/2})$ from $(\mathbf{L}_k, \mathbf{L}_{k-1}^{\alpha/2}, \mathbf{R}_k, \mathbf{R}_{k-1}^{\alpha/2})$
 Apply preconditioned gradient step
 $\mathbf{W}_k = \mathbf{W}_{k-1} - \eta \mathbf{L}_k^{\alpha/2} \mathbf{G}_k \mathbf{R}_k^{\alpha/2}$
end for

Then, the full Shampoo statistics matrices are related to a combination of these two statistics,

$$\begin{aligned} \mathbf{B}_c &= M_g(\mathbf{B}_1, \mathbf{I}) = \mathbf{B}_1^{1/2} \\ \mathbf{C}_c &= M_g(\mathbf{I}, \mathbf{C}_2) = \mathbf{C}_2^{1/2}, \end{aligned}$$

These still need to be inverse square rooted to be applied as the preconditioner, since here we have constructed

$$\mathbf{C}_c \otimes \mathbf{B}_c \succeq \mathcal{G} \mathcal{G}^\top.$$

The inverse square root ultimately gets grouped together with the square roots in the expressions above to yield the inverse $1/4$ -th power in the preconditioning,

$$\begin{aligned} \mathbf{W}_k &= \mathbf{W}_{k-1} - \eta \mathbf{B}_c^{-1/2} \mathbf{G}_k \mathbf{C}_c^{-1/2} \\ &= \mathbf{W}_{k-1} - \eta \mathbf{B}_1^{-1/4} \mathbf{G}_k \mathbf{C}_2^{-1/4} \end{aligned}$$

(compare with Equations (2)-(3)). Each individual estimator keeps one Kronecker factor as identity.

3.3.2 KrADagrad*

Inspired by Shampoo’s optimality, we can maintain a pair of KrADagrad preconditioners $(\mathbf{L}_{k,1}, \mathbf{R}_{k,1}) = (\mathbf{L}_k, \mathbf{I}_m)$ and $(\mathbf{L}_{k,2}, \mathbf{R}_{k,2}) = (\mathbf{I}_n, \mathbf{R}_k)$, where the second index in the LHS subscripts denotes the estimator. Since we will hold the identity matrices constant, we do not need to store or perform multiplication with them explicitly. In addition, this means we can unambiguously drop the second index in the subscripts. On iteration k , we update both \mathbf{L}_k and \mathbf{R}_k , just as in shampoo. With

$$\begin{aligned} t_{k,L} &\leftarrow 1 + \|\mathbf{G}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}\|_F \\ t_{k,R} &\leftarrow 1 + \|\mathbf{G}_k^\top \mathbf{G}_k \mathbf{R}_{k-1}\|_F, \end{aligned}$$

we summarize in Algorithm 1. In order to bound the regret, we need a few intermediate results.

Lemma 3.4. *Assume a 1-Lipschitz loss, implying $\|\mathbf{G}_k\|_2 \leq 1 \forall k$. Suppose by iteration k , \mathbf{L}_k is updated in $k - k'$ of those steps for $0 \leq k' \leq k$ (and thus \mathbf{R}_k is updated in the*

remaining k' steps). Letting $\mathbf{B}_k = \mathbf{L}_k^{-1}$, $\mathbf{C}_k = \mathbf{R}_k^{-1}$, and $s = 1 + \|\mathbf{L}_0\|_2 \|\mathbf{R}_0\|_2$,

$$\text{tr}(\mathbf{B}_k) \leq \text{tr}(\mathbf{B}_0) + k' m s \|\mathbf{R}_0\|_2 \quad (14)$$

$$\text{tr}(\mathbf{C}_k) \leq \text{tr}(\mathbf{C}_0) + (k - k') n s \|\mathbf{L}_0\|_2. \quad (15)$$

Next, we restate Theorem 7 from Gupta et al. [2018] in our notational setting. We additionally make one minor but straightforward substitution of the smaller matrix dimension instead of the rank, as the dimensions upper bound the rank.

Lemma 3.5 (From Gupta et al. [2018]). *Let $\mathbf{w}_* \in \mathbb{R}^{mn}$ with $m \geq n$, $t > 0$, and $D \triangleq \max_{1 \leq k \leq K} \|\mathbf{w}_k - \mathbf{w}_*\|_2$. The regret from using a Kronecker factorized preconditioner*

$$\mathbf{P}_k = \mathbf{C}_k^{-1/2} \otimes \mathbf{B}_k^{-1/2}$$

that dominates the empirical Fisher matrix

$$\mathbf{C}_k \otimes \mathbf{B}_k \succeq \frac{1}{nt} \mathcal{G}_k \mathcal{G}_k^\top$$

is bounded

$$\sum_{k=1}^K (f_k(\mathbf{w}_k) - f_k(\mathbf{w}_*)) \leq D \sqrt{2tn} \text{tr}(\mathbf{B}_K^{1/2}) \text{tr}(\mathbf{C}_K^{1/2})$$

Now, we have our regret result: a proof is in Appendix B.4.

Theorem 3.6. *Assuming a 1-Lipschitz loss, the regret from using KrADagrad* scales as $O(\sqrt{K})$.*

4 IMPLEMENTATION

Now we discuss the algorithmic considerations for actually implementing KrADagrad*. The main difficulty lies in efficiently computing the matrix roots. Differentiable matrix square roots for machine learning have been the subject of a substantial amount of research [Song et al., 2021]. For preconditioning, we do not require differentiability for our roots, so they can be computed using numerical linear algebra methods (e.g. SVD) without concern for the backward pass. While such algorithms are gaining hardware support on current GeMM accelerators and software frameworks, they are still not universal (Pytorch supports SVD on CUDA in double precision via cuSOLVER and MAGMA [PyTorch Team, 2021, NVidia Team, 2023, Dongarra et al., 2014], but TPUs to our knowledge do not [Jouppi et al., 2021]).

However, as [Anil et al., 2020] discovered, accurate inverse powers require high precision to retain the important contributions of the eigenvectors corresponding to the smallest eigenvalues, and current general matrix multiplication (GeMM) accelerators do not prioritize 64-bit computation (i.e. GPUs see significant speed reductions, while TPUs do not support it at all). Plus, they assume preconditioners are a slowly-varying sequence of matrices. Thus, they propose computing matrix roots iteratively on CPU due to these algorithmic and hardware architectural constraints.

²<http://github.com/jonathanmei/kradagrad>

In contrast, KrADagrad* deals with positive roots, and both these iterative matrix methods and the numerical linear algebra techniques are otherwise actually amenable to computation using GeMM accelerators. Ultimately, as a straightforward solution for the roots that appear in KrADagrad*, the SVD on GPU suffices. We note that avoiding inversion is thus not about reducing the computational complexity, rather we aim to avoid needing 64-bit computation. Further details of matrix roots are discussed in Appendix C.3.

Additionally, diagonal damping is a common feature of preconditioned methods due to their contribution to numerical stability. While we do not need this for numerical stability, we find empirically that diagonal damping still helps reduce the effect of gradient noise and smooth out the loss curve. We describe how this can be applied to KrADagrad* in further detail in Appendix C.4.

So far, our proposed algorithms and analyses apply to matrix parameters. As noted earlier, we can extend this to tensor parameters in a fairly straightforward manner as mainly a matter of additional notation and bookkeeping. Hence, we relegate the extension to Appendix C.5.

The baseline Pytorch Shampoo implementation we use is not optimized, so is not totally fair to the true capability of the baseline. Similarly, our KrADagrad variants, being derived from the mentioned implementation of Shampoo, did not have optimized implementations. Nonetheless, we have conducted some preliminary wall clock time comparisons between our implementation of KrADagrad and Shampoo, which we summarize in Table 1. For each data set, we report the time in seconds to run a single epoch for (KrADagrad, KrADagrad*, Shampoo) along with a 2 standard deviation interval, computed from epochs 5 through 10 on an NVidia A40 GPU.

Table 1: Comparison of wall clock times (in seconds) between our implementations of KrADagrad variants and Shampoo on a single epoch and a 2 standard deviation interval computed from 5 epochs for CIFAR-10/100 data sets on an NVidia A40 GPU.

DATA SET	KRADAGRAD	KRADAGRAD*	SHAMPOO
CIFAR-10	29.50± 0.88	45.76± 0.20	35.62 ± 0.17
CIFAR-100	54.71± 0.36	76.12± 1.72	66.45 ± 0.33

4.1 COMPUTE AND MEMORY COSTS

For KrADagrad, computing $\Delta \mathbf{L}_k$ has a computational cost of $(2N(2m + n) + 2m^2 + 2m^3)$. While we still require $O(m^3 + n^3)$ matrix square roots for the KrADagrad update, this is less difficult numerically and thus computationally than the -4 th root required by Shampoo.

KrADagrad* requires 4th root computation, which is comparable in cost to the inverse 4th root, but without the need

for high precision.

Storage involves tracking the two matrix factors, and so is $O(m^2 + n^2)$ for all methods.

We summarize these costs in Table 2.

Table 2: Comparison of complexity between our implementations of KrADagrad variants and Shampoo.

	KRADAGRAD	KRADAGRAD*	SHAMPOO
COMPUTE	$O(m^3 + n^3)$	$O(m^3 + n^3)$	$O(m^3 + n^3)$
MEMORY	$O(m^2 + n^2)$	$O(m^2 + n^2)$	$O(m^2 + n^2)$

5 EXPERIMENTS

We address the following: 1) How sensitive is our method to matrix conditioning compared to Shampoo? 2) How does convergence speed compare in the number of training steps? (as measured by task-specific validation metrics) 3) How do our methods compare in the achieved model quality at or near convergence? The goal of each is to compare optimizers in various challenging loss landscapes, rather than to achieve state of the art performance.

To answer 1, in a synthetic experiment we minimize a multi-dimensional quadratic function with a non-diagonal, poorly-conditioned Hessian that neatly factorizes into a Kronecker product of two individually poorly-conditioned PD matrices.

To answer 2 and 3, we compare KrADagrad and KrADagrad* to alternatives across a variety of tasks: image classification (IC), autoencoder problems (AE), recommendation (RecSys), continual learning (CL). For IC experiments we train ResNet-32/56 [He et al., 2016] without BatchNorm (BN) on CIFAR-10/100 [Krizhevsky, 2009]. For AE, we train simple autoencoders consistent with [Goldfarb et al., 2020] on MNIST [Lecun et al., 1998], as well as CURVE and FACES [Hinton and Salakhutdinov, 2006]. For recommendation, we train H+Vamp Gated [Kim and Suh, 2019] on MovieLens20M [Harper and Konstan, 2015]. In the continual learning setting we train on two benchmarks from [Lomonaco et al., 2021]: GEM [Lopez-Paz and Ranzato, 2017] on Permuted MNIST [Goodfellow et al., 2013] and LaMAML [Gupta et al., 2020] on Split CIFAR100 [Zenke et al., 2017].

We choose Shampoo as a baseline 2nd order optimizer as: a.) KrADagrad variants are most similar to Shampoo b.) we expect at best similar performance unless KrADagrad’s approximations and lower precision are in practice detrimental. To ground all comparisons, we always include SGD, Adam, plus any unique optimizer from an existing benchmark.

In all our experiments, we initialized from common seeds across optimizers. However, the first points on the training curves follows the number of training steps per evaluation interval, so they do not visually appear to start from the




Figure 1: Loss on synthetic quadratic in log scale, relative to SGD. Each curve is *divided* by that of SGD.

same point. Doing so would have required modifying the codebases separately for each experiment. In addition, ideally we would have had the resources to run multiple shared seeds for the selected HPs of every task we explored. For smaller tasks where convergence is achieved quickly, such as the autoencoder experiments, it was feasible for us to do this and we share this in Figures 3 and 8. In the tradeoff of how to use our compute and time resources, we opted to present fair evaluations of the optimizers (i.e. ensuring shared seeds and optimized HPs per optimizer per task) across a diversity of tasks, and repeatability statistics from different initial guesses as much as possible.

5.1 CONDITIONING EXPERIMENT

To see the effects of Hessian conditioning on optimizers of interest, we create a synthetic deterministic convex problem. For $\mathbf{X} \in \mathbb{R}^{128 \times 128}$, we minimize a quadratic loss function

$$\min_{\mathbf{X}} \text{tr}(\mathbf{X}^T \mathbf{A} \mathbf{X} \mathbf{B})$$

where $\mathbf{A}, \mathbf{B} \succ 0$ are non-diagonal and have condition numbers $\kappa(\mathbf{A}) = \kappa(\mathbf{B}) = 10^{10}$. We seed each optimizer with the same starting point and sweep learning rates and pick the one with the lowest loss for each optimizer. We provide further details in Appendix G. While this stationary problem without a validation or test dataset is different from the online setting assumed in Theorem 3.6, and even from a typical offline machine learning problem, it helps isolate differences in behavior between optimizers in a badly conditioned convex loss landscape.

In Figure 1, while Shampoo in double precision outperforms the others in terms of final loss achieved, single precision KrAdagrad* outperforms single precision Shampoo. Adam does not converge quite as quickly or to as good a loss, as the adaptivity it provides is aligned to the canonical axis, while \mathbf{A}, \mathbf{B} being non-diagonal act in a rotated basis.

5.2 REAL DATASETS

For each task and optimizer we sweep hyperparameters (HP), select those yielding the best validation metric at some epoch or iteration, and display the corresponding learning curves in Figures 2 through 5. We set the preconditioner update rate for Shampoo and KrAdagrad* to every 20 training

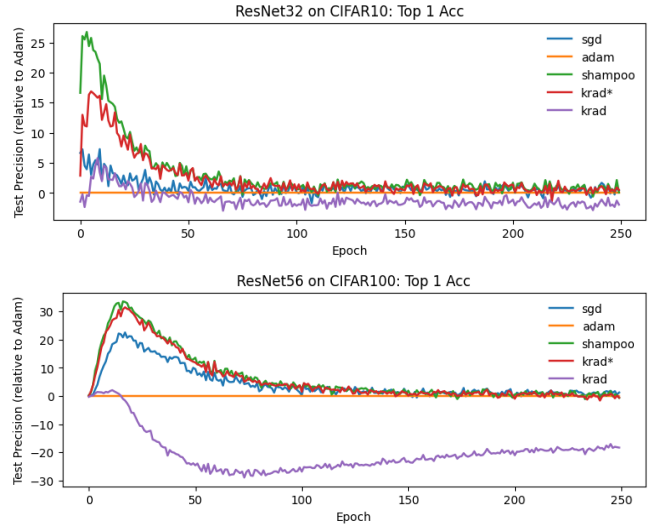


Figure 2: Vision experiments. Each curve is initialized from the same single seed and *subtracted* from that of Adam. Top: Top 1 Accuracy (in %) of ResNet-32 (without BN) on CIFAR-10, relative to Adam; Bottom: Top 1 Accuracy (in %) of ResNet-56 (without BN) on CIFAR-100, relative to Adam.

steps. For all experiments, Shampoo uses double precision for negative matrix roots, and SGD includes momentum, unless stated otherwise. We summarize our observations here and provide additional detail in Appendix G.

For the three autoencoder experiments, after HP tuning we retrain the best HPs with 5 unique seeds shared across optimizers, with results in Figures 3 and 8. We also find the auto encoder tasks to be relevant baselines for evaluating Shampoo alternatives because Shampoo consistently outperforms SGD and Adam. While KrAdagrad* also outperforms SGD and Adam, Shampoo reaches better solutions or similar solutions in fewer steps. Following the synthetic experiments, we test the hypothesis that KrAdagrad* might perform better than 32-bit Shampoo by running our best HPs for Shampoo with single precision. The 32- and 64-bit versions perform similarly on these tasks, falsifying the hypothesis in the general sense. KrAdagrad, slower in the number of steps and performing worse in general, eventually reaches similar performance to KrAdagrad* on CURVES, with some runs reaching that of 32 bit Shampoo (see Figure 8b). One hypothesis is that on these particular tasks, effectively traversing the loss landscape requires eigenvectors that the KrAdagrad approximation has more difficulty capturing than Shampoo.

While Gupta et al. [2018] performs experiments on “standard” machine learning tasks, their regret results pertain to an online setting. We include continual learning problems to test the limits of the theoretical setting. Figure 4a displays learning curves for GEM on Permuted MNIST relative to SGD to illuminate small differences (we include a plot of the actual accuracies in Appendix G). On average, KrAdagrad




Figure 3: Reconstruction validation error for fully connected auto encoder a) mean cross entropy on MNIST b) mean cross entropy on CURVES c) mean squared error on FACES. The learning curves shown are averaged across 5 unique seeds with CI95 error bars. We do this analysis for the autoencoder experiments since they’re relatively inexpensive to run.




Figure 4: a) Top 1 Accuracy of GEM on Permuted MNIST, relative to SGD. Each curve averages multiple seeds, subtracted from the corresponding SGD curve to illuminate otherwise unobservable differences. See Appendix G for absolute curves. b) Top 1 Accuracy of LaMAML on Split CIFAR-100, relative to Adam.

stays slightly ahead of others throughout training, unlike KrADagrad*.

Figure 4b shows similar accuracy curves for LaMAML on Split CIFAR100 relative to Adam (again, the absolute accuracies are in Appendix G). KraADagrad* closely follows Shampoo and the adaptive learning rate optimizer from [Gupta et al., 2020], while KrADagrad performs only better than Adam.

We also evaluate Shampoo and KrADagrad variants against the baseline for H+Vamp Gated³. Shampoo reaches a lower training loss compared to KrADagrad*, while both Adam variants converge faster and to lower loss values (Figure 5).

6 CONCLUSION

We introduced KrADagrad*, a preconditioned gradient optimizer that avoids matrix inversion, and proved that it

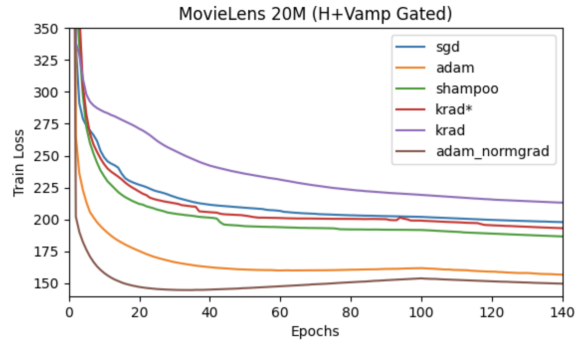


Figure 5: Training loss of H+Vamp Gated on ML-20M. We choose to compare the training loss for this experiment because it comprises a weighted combination of components, $RE + \beta \cdot KL$, where β changes according to a predefined schedule (causing loss growth in the 1st 100 epochs), which is not consistent in the validation loss.

has optimal regret properties. We showed that its performance exceeds that of Shampoo in single precision on an ill-conditioned synthetic convex optimization problem. Our experiments on real datasets show that KrADagrad* often performs similarly to Shampoo.

Future work includes tractable methods to compute low-rank updates to matrix roots, for instance with rational Krylov subspace methods, to improve the performance of KrADagrad* and other Kronecker-factored preconditioners. It is also worth considering the application of the optimizer to other application areas, such as decision making and controls, or scientific discovery.

Acknowledgements

We thank Luminous Computing for funding this work. In particular, we thank David Baker for encouraging this line of investigation and David Scott for fruitful discussions about numerical linear algebra.

³<https://github.com/psywaves/EVCF>

References

- Naman Agarwal, Brian Bullins, Xinyi Chen, Elad Hazan, Karan Singh, Cyril Zhang, and Yi Zhang. Efficient full-matrix adaptive regularization. In *International Conference on Machine Learning*, pages 102–110. PMLR, 2019.
- T. Ando, Chi Kwong Li, and Roy Mathias. Geometric means. *Linear Algebra and Its Applications*, 385:305–334, 7 2004. ISSN 00243795. doi: 10.1016/J.LAA.2003.11.019.
- Rohan Anil, Vineet Gupta, Tomer Koren, Kevin Regan, and Yoram Singer. Scalable second order optimization for deep learning. *arXiv preprint arXiv:2002.09018*, 2020.
- Rohan Anil, Sandra Gadoho, Da Huang, Nijith Jacob, Zhuoshu Li, Dong Lin, Todd Phillips, Cristina Pop, Kevin Regan, Gil I. Shamir, Rakesh Shivanna, and Qiqi Yan. On the factory floor: ML engineering for industrial-scale ads recommendation models. *ACM RecSys*, 9 2022. doi: 10.48550/arxiv.2209.05310. URL <https://orsum.inesctec.pt/orsum2022/assets/files/paper6.pdf>.
- Bernhard Baumgartner. An inequality for the trace of matrix products, using absolute values. 6 2011. doi: 10.48550/arxiv.1106.6189.
- Richard Bellman. Some inequalities for positive definite matrices. In *General Inequalities 2*, pages 89–90. Springer, 1980.
- Rajendra Bhatia. *Positive Definite Matrices*. Princeton University Press, 12 2009. ISBN 9781400827787. doi: 10.1515/9781400827787.
- Stephen P Boyd and Lieven Vandenbergh. *Convex optimization*. Cambridge university press, 2004.
- Jack Dongarra, Mark Gates, Azzam Haidar, Jakub Kurzak, Piotr Luszczek, Stanimire Tomov, and Ichitaro Yamazaki. Accelerating numerical dense linear algebra calculations with gpus. *Numerical Computations with GPUs*, pages 1–26, 2014.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive sub-gradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(7), 2011.
- Zachary Frangella, Pratik Rathore, Shipu Zhao, and Madeleine Udell. SketchySGD: Reliable stochastic optimization via robust curvature estimates. 11 2022. doi: 10.48550/arxiv.2211.08597. URL <https://arxiv.org/abs/2211.08597v2>.
- Donald Goldfarb, Yi Ren, and Achraf Bahamou. Practical quasi-newton methods for training deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2386–2396, 2020.
- Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks, 2013. URL <https://arxiv.org/abs/1312.6211>.
- Gunshi Gupta, Karmesh Yadav, and Liam Paull. Look-ahead meta learning for continual learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11588–11598. Curran Associates, Inc., 2020.
- Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pages 1842–1850. PMLR, 2018.
- F. Maxwell Harper and Joseph A. Konstan. The MovieLens datasets: History and context. *ACM Trans. Interact. Intell. Syst.*, 5(4), dec 2015. ISSN 2160-6455. doi: 10.1145/2827872.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. URL <http://image-net.org/challenges/LSVRC/2015/>.
- Nicholas J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, 2008. ISBN 978-0-898716-46-7.
- G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507, 2006. doi: 10.1126/science.1127647. URL <https://www.science.org/doi/abs/10.1126/science.1127647>.
- Yerlan Idelbayev. Proper ResNet implementation for CIFAR10/CIFAR100 in PyTorch. https://github.com/akamaster/pytorch_resnet_cifar10. Accessed: 2022-12-12.
- Norman P. Jouppi, Doe Hyun Yoon, Matthew Ashcraft, Mark Gottscho, Thomas B. Jablin, George Kurian, James Laudon, Sheng Li, Peter Ma, Xiaoyu Ma, Thomas Norrie, Nishant Patil, Sushma Prasad, Cliff Young, Zongwei Zhou, and David Patterson. Ten Lessons From Three Generations Shaped Google’s TPUv4i Industrial Product. 2021. doi: 10.1109/ISCA52012.2021.00010.
- Daeryong Kim and Bongwon Suh. Enhancing VAEs for collaborative filtering. In *Proceedings of the 13th ACM Conference on Recommender Systems*. ACM, sep 2019. doi: 10.1145/3298689.3347015.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015.

- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. doi: 10.1109/5.726791.
- Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L. Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido van de Ven, Martin Mundt, Qi She, Keiland Cooper, Jeremy Forest, Eden Belouadah, Simone Calderara, German I. Parisi, Fabio Cuzzolin, Andreas Toliias, Simone Scardapane, Luca Antiga, Subutai Amhad, Adrian Popescu, Christopher Kanan, Joost van de Weijer, Tinne Tuytelaars, Davide Bacciu, and Davide Maltoni. Avalanche: an end-to-end library for continual learning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2nd Continual Learning in Computer Vision Workshop, 2021. URL <https://github.com/ContinualAI/continual-learning-baselines>.
- David Lopez-Paz and Marc' Aurelio Ranzato. Gradient episodic memory for continual learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- Haiping Lu, Konstantinos N. Plataniotis, and Anastasios N. Venetsanopoulos. *Multilinear subspace learning: Dimensionality reduction of multidimensional data*. CRC Press, 2013. ISBN 9781439857298. doi: 10.1201/B16252.
- James Martens. New insights and perspectives on the natural gradient method. *Journal of Machine Learning Research*, 21:1–76, 2020. ISSN 1533-7928.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International Conference on Machine Learning*, pages 2408–2417, 2015.
- NVidia Team. cuSOLVER API Reference, January 2023. URL <https://docs.nvidia.com/cuda/cusolver>.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- PyTorch Team. The torch.linalg module: Accelerated linear algebra with autograd in PyTorch | PyTorch, June 2021. URL <https://pytorch.org/blog/torch-linalg-autograd/>.
- Sashank J Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of Adam and beyond. *International Conference on Learning Representations*, 2018.
- Yi Ren and Donald Goldfarb. Tensor normal training for deep learning models. *Advances in Neural Information Processing Systems*, 34:26040–26052, 12 2021.
- Lavanya Shukla. Understanding the Landscape of the latest Large Models. *Advances in neural information processing systems*, 35, 2022.
- Yue Song, Nicu Sebe, and Wei Wang. Why approximate matrix square root outperforms accurate SVD in global covariance pooling? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1115–1123, 2021.
- Nati Srebro, Karthik Sridharan, and Ambuj Tewari. On the universality of online mirror descent. *Advances in neural information processing systems*, 24, 2011.
- Paul Tseng. Convergence of a block coordinate descent method for nondifferentiable minimization. *Journal of Optimization Theory and Applications* 2001 109:3, 109: 475–494, 6 2001. ISSN 1573-2878. doi: 10.1023/A:1017501703105.
- Charles F. Van Loan. The ubiquitous kronecker product. *Journal of Computational and Applied Mathematics*, 123:85–100, 11 2000. ISSN 03770427. doi: 10.1016/S0377-0427(00)00393-9.
- Charles F. Van Loan and Nikos Pitsianis. Approximation with kronecker products. In *Linear Algebra for Large Scale and Real-Time Applications*, pages 293–314. Springer, Dordrecht, 1993. doi: 10.1007/978-94-015-8196-7_17.
- Max A Woodbury. *Inverting modified matrices*. Statistical Research Group, 1950.
- Minghan Yang, Dong Xu, Zaiwen Wen, Mengyun Chen, and Pengxiang Xu. Sketch-based empirical natural gradient methods for deep learning. *Journal of Scientific Computing*, 92:1–29, 9 2022. ISSN 15737691. doi: 10.1007/S10915-022-01911-X/METRICS.
- Zhewei Yao, Amir Gholami, Sheng Shen, Mustafa Mustafa, Kurt Keutzer, and Michael W Mahoney. AdaHessian: An adaptive second order optimizer for machine learning. 2021.
- Friedemann Zenke, Ben Poole, and Surya Ganguli. Improved multitask learning through synaptic intelligence. *CoRR*, abs/1703.04200, 2017. URL <http://arxiv.org/abs/1703.04200>.

KrADagrad: Kronecker Approximation-Domination Gradient Preconditioned Stochastic Optimization (Supplementary Material)

Jonathan Mei^{1†}

Alexander Moreno^{1†}

Luke Walters^{1†}

¹Independent Researcher

A KrADagrad	13
A.1 KrADagrad: direct Adagrad-style alternating updates	13
B Deferred proofs	14
B.1 Proof of Lemma 3.1	14
B.2 Proof of Proposition 3.2	15
B.3 Proof of Lemma 3.4	16
B.4 Proof of Theorem 3.6	18
C Implementation details	18
C.1 KrADagrad	18
C.2 Delayed updates	18
C.3 Matrix Roots	19
C.4 Diagonal damping	20
C.5 Generalization to tensors	20
D Properties and extended derivations	21
D.1 Linear algebra properties	21
D.2 Additional and extended derivations	22
E Memory and compute cost analysis	23
E.1 Kradagrad: statistic update	23
E.2 Kradagrad: matrix square root	24
E.3 Kradagrad: applying preconditioner	24
E.4 Kradagrad*: statistics updates	24
E.5 KrADagrad*: matrix p -th roots	24
E.6 Kradagrad*: applying preconditioners	24

F Other Kronecker approximations	24
G Experiment Details	25
G.1 Synthetic	25
G.2 Vision Experiments	26
G.3 Autoencoder Experiments	26
G.4 Recommender Experiment	27
G.5 Continual Learning Experiments	28

A KRADAGRAD

A.1 KRADAGRAD: DIRECT ADAGRAD-STYLE ALTERNATING UPDATES

We propose an alternative scheme for updating the statistics. We alternate between directly using KrAD to update left and right statistics in different iterations. This has the practical advantage of only requiring one matrix statistic update per iteration and only needing a matrix square root (even in the extension to tensors!). However, the scheme only attains optimal regret when we stop updating the right preconditioner after a fixed number of iterations, though we note that we describe how to choose an update schedule to obtain regret arbitrarily close to optimal.

Arguably, the most intuitive way to update $(\mathbf{L}_k, \mathbf{R}_k)$ is to select an index set $\mathcal{I} \subset \mathbb{N}$ as a subset of the natural numbers, so that if $k \in \mathcal{I}$ we update \mathbf{L}_k and hold \mathbf{R}_k constant, and vice versa for $k \notin \mathcal{I}$.

It is straightforward to see that we simply need to compute a matrix square root on KrAD statistics to apply Adagrad-style preconditioning,

$$\mathbf{W}_{k+1} = \mathbf{W}_k - \eta_k \mathbf{L}_k^{1/2} \mathbf{G}_k \mathbf{R}_k^{1/2}. \quad (16)$$

In an iteration $k \in \mathcal{I}$, since we have held \mathbf{R}_k constant, we can store its square root from the previous iteration and only need to compute a square root for \mathbf{L}_k that we updated in this iteration. In this case, we refer to our intermediate matrices $(\mathbf{L}_k, \mathbf{R}_k)$ as our statistics and the pair $(\mathbf{L}_k^{1/2}, \mathbf{R}_k^{1/2})$ as the actual preconditioners. Thus, we can see that our storage requirement is now $(\mathbf{L}_k, \mathbf{L}_k^{1/2}, \mathbf{R}_k, \mathbf{R}_k^{1/2})$.

We call this update scheme KrADagrad and summarize in Appendix C.1.

A.1.1 Regret

We establish some results for these updates to provide the groundwork for the theoretical properties of the algorithm derived from the updates.

Lemma A.1. *The corresponding $\mathbf{Q}_k \succeq \frac{1}{nt} \mathbf{G}_k \mathbf{G}_k^\top$, where $t = \max_k t_k$.*

Proof. Taking $\Delta \mathbf{L}_k = \frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}$ and $\mathbf{L}_k = \mathbf{L}_{k-1} - \Delta \mathbf{L}_k$, we have from the chain of inequalities preceding Equation (13),

$$\begin{aligned} \mathbf{L}_k^{-1} &= (\mathbf{L}_{k-1} - \Delta \mathbf{L}_k)^{-1} \succeq \mathbf{B}_{k-1} + \Delta \mathbf{B}_k \\ \Rightarrow \mathbf{Q}_k &= \mathbf{R}_k^{-1} \otimes \mathbf{L}_k^{-1} \succeq \mathbf{C}_k \otimes \tilde{\mathbf{B}}_k \succeq \frac{1}{nt} \mathbf{G}_k \mathbf{G}_k^\top \end{aligned}$$

where \mathbf{L}_k is invertible due to it being PD as shown in Corollary 3.3, and the last domination inequality is due to Equation (11). We can see that $t = \max_k t_k$ suffices. \square

Lemma A.1 is half of what we need to show good regret guarantees, as per proof of Theorem 7 in Gupta et al. [2018]. The other half is the Lemma 3.4 from the main text.

Thus, we have the next result, which plugs in the trace bounds of Lemma 3.4 into Lemma 3.5.

Theorem A.2. *Over K iterations, assume we update \mathbf{R}_k in $K' = o(K)$ iterations and \mathbf{L}_k in the remaining $K - K'$ iterations. The regret from using KrADagrad with Adagrad-style $\alpha = 1/2$ exponent scales as*

$$R(K) = O(\sqrt{KK'})$$

Proof. With the modified form of the Hölder inequality for PSD matrices from (P6),

$$\begin{aligned} \text{tr}(\mathbf{A}^{1/2}\mathbf{B}^{1/2}) &\leq \text{tr}(\mathbf{A})^{1/2} \text{tr}(\mathbf{B})^{1/2} \\ \Rightarrow \text{tr}(\mathbf{B}_k^{1/2}) &\leq m^{1/2} \text{tr}(\mathbf{B}_k)^{1/2}. \end{aligned}$$

Then, from Equation (14),

$$\begin{aligned} \text{tr}(\mathbf{B}_k^{1/2}) &\leq m^{1/2} \text{tr}(\mathbf{B}_k)^{1/2} \\ &\leq \sqrt{m \text{tr}(\mathbf{B}_0) + k' m^2 s \|\mathbf{R}_0\|_2}. \end{aligned}$$

Analogously, from Equation (15)

$$\begin{aligned} \text{tr}(\mathbf{C}_k^{1/2}) &\leq n^{1/2} \text{tr}(\mathbf{C}_k)^{1/2} \\ &\leq \sqrt{n \text{tr}(\mathbf{C}_0) + (k - k') n^2 s \|\mathbf{L}_0\|_2}. \end{aligned}$$

Assuming w.l.o.g. that $k' = o(k)$, the regret bounded by $\text{tr}(\mathbf{B}_k^{1/2}) \text{tr}(\mathbf{C}_k^{1/2})$ is further bounded by

$$\sqrt{(am + m^2 b(k - k'))(cn + n^2 dk')} = O(\sqrt{kk'}),$$

where a, b, c, d are constants w.r.t. m, n, k, k' that can be expressed in terms of $\text{tr}(\mathbf{B}_0)$, $\text{tr}(\mathbf{C}_0)$, $\|\mathbf{L}_0\|_2$, and $\|\mathbf{R}_0\|_2$. Finally, $t_k = 1 + \|\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_2 \leq s$ for all $k > 0$, so $t = \max_k t_k \leq s$. \square

We see that this is greater than the optimal rate of $O(\sqrt{K})$ unless $K' = O(1)$. Taking an update scheme in which $K' \sim K^{2\varepsilon}$, we can achieve a rate of $O(K^{1/2+\varepsilon})$, which can be made arbitrarily close to the optimal. In one sense, keeping one Kronecker factor as identity (as Shampoo does) can be seen as optimal according to our regret result, as it does not introduce the additional factor of k' in the bound. This inspires an alternative way of updating \mathbf{R}_k , resulting in KrADagrad* from the main text.

B DEFERRED PROOFS

Here we provide proofs that were either too long for the main text or pertained to less consequential results.

B.1 PROOF OF LEMMA 3.1

Lemma B.1. *Let PD matrix $\mathbf{C} \in \mathbb{R}^{n \times n}$, $\mathcal{U} \in \mathbb{R}^{mn \times r}$, \mathbf{u}_i be the i -th column of \mathcal{U} , $\mathbf{U}_i = \text{vec}_{m,n}^{-1}(\mathbf{u}_i)$. Then*

$$\begin{aligned} \mathbf{B} &= \sum_{i=1}^r \mathbf{U}_i \mathbf{C}^{-1} \mathbf{U}_i^\top \succeq 0 \\ \Rightarrow \mathbf{C} \otimes \mathbf{B} &\succeq \frac{1}{n} \mathcal{U} \mathcal{U}^\top. \end{aligned}$$

Proof. We largely follow Gupta et al. [2018]. We consider the case where $r = 1$ and drop the subscript on \mathbf{u}_i for clarity (i.e. $\mathbf{U} = \mathbf{U}_1$) since summation holds across the inequality. Let the singular value decomposition (SVD) of $\mathbf{R} = \mathbf{V} \mathbf{\Lambda} \mathbf{V}^\top$, and \mathbf{v}_j be the j -th column of \mathbf{V} and $\lambda_j = \mathbf{\Lambda}_{j,j}$ the corresponding singular value. Let $\mathbf{Z} = \mathbf{U} \mathbf{V}$. Then

$$\begin{aligned} \mathbf{U} &= \mathbf{Z} \mathbf{V}^\top \\ \mathbf{L} &= \mathbf{U} \mathbf{R}^{-1} \mathbf{U}^\top = \mathbf{Z} \mathbf{\Lambda}^{-1} \mathbf{Z}^\top. \end{aligned}$$

Note that \mathbf{Z} does not necessarily have any special structure, since it was not chosen according to any factorization of \mathbf{U} . Let \mathbf{z}_j be the j -th column of \mathbf{Z} . Then

$$\begin{aligned}\mathbf{U} &= \sum_{i=1}^n \mathbf{z}_i \mathbf{v}_i^\top \\ \Rightarrow \mathbf{u} &= \sum_{i=1}^n \mathbf{v}_i \otimes \mathbf{z}_i. \quad \because \text{(P18)}\end{aligned}$$

Next, we note that for a set of vectors $\{\mathbf{y}_i\}$ and any arbitrary vector \mathbf{x} , using equivalence of norms

$$\begin{aligned}\mathbf{x}^\top \left(\sum_{i=1}^n \mathbf{y}_i \right) \left(\sum_{i=1}^n \mathbf{y}_i^\top \right) \mathbf{x} &= \left(\sum_{i=1}^n \mathbf{y}_i^\top \mathbf{x} \right)^2 \\ &\leq \left(\sum_{i=1}^n |\mathbf{y}_i^\top \mathbf{x}| \right)^2 \\ &\leq n \sum_{i=1}^n (\mathbf{y}_i^\top \mathbf{x})^2 \\ &= n \mathbf{x}^\top \left(\sum_{i=1}^n \mathbf{y}_i \mathbf{y}_i^\top \right) \mathbf{x} \\ \Rightarrow \left(\sum_{i=1}^n \mathbf{y}_i \right) \left(\sum_{i=1}^n \mathbf{y}_i^\top \right) &\preceq n \sum_{i=1}^n \mathbf{y}_i \mathbf{y}_i^\top.\end{aligned}$$

Then consider

$$\mathbf{u} \mathbf{u}^\top = \left(\sum_{i=1}^n \mathbf{v}_i \otimes \mathbf{z}_i \right) \left(\sum_{i=1}^n \mathbf{v}_i \otimes \mathbf{z}_i \right)^\top \quad (17)$$

$$\preceq n \sum_{i=1}^n (\mathbf{v}_i \otimes \mathbf{z}_i) (\mathbf{v}_i \otimes \mathbf{z}_i)^\top \quad (18)$$

$$= n \sum_{i=1}^n (\mathbf{v}_i \mathbf{v}_i^\top \otimes \mathbf{z}_i \mathbf{z}_i^\top) \quad \because \text{(P14)} \quad (19)$$

$$= n \sum_{i=1}^n (\lambda_i \mathbf{v}_i \mathbf{v}_i^\top \otimes \lambda_i^{-1} \mathbf{z}_i \mathbf{z}_i^\top) \quad (20)$$

$$\preceq n \sum_{i=1}^n (\mathbf{R} \otimes \lambda_i^{-1} \mathbf{z}_i \mathbf{z}_i^\top) \quad (21)$$

$$= n \mathbf{R} \otimes (\mathbf{Z} \mathbf{\Lambda}^{-1} \mathbf{Z}^\top) \quad (22)$$

$$= n \mathbf{R} \otimes \mathbf{L}. \quad (23)$$

Here Equation 21 holds since for any i ,

$$\mathbf{R} - \lambda_i \mathbf{v}_i \mathbf{v}_i^\top = \sum_{j \neq i} \lambda_j \mathbf{v}_j \mathbf{v}_j^\top \succeq 0.$$

□

B.2 PROOF OF PROPOSITION 3.2

First, we prove our intermediate result 3.2, which we repeat here for convenience.

Proposition B.2. Taking $t_k = 1 + \|\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\|_2$ (or the looser but more computationally friendly $t_k = 1 + \|\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\|_F$), the PSD matrix

$$\mathbf{M}_k \triangleq \frac{1}{t_k}\mathbf{L}_{k-1}^{1/2}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\mathbf{L}_{k-1}^{1/2} \prec \mathbf{I}.$$

Proof. Since $\mathbf{L}_{k-1}^{1/2} \succ 0$, it is invertible. Thus the matrices $t_k\mathbf{M}_k = \mathbf{L}_{k-1}^{1/2}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\mathbf{L}_{k-1}^{1/2}$ and $\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top$ are similar as

$$\mathbf{L}_{k-1}^{1/2}t_k\mathbf{M}_k\mathbf{L}_{k-1}^{-1/2} = \mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top.$$

This implies that these matrices have the same matrix norms

$$\|t_k\mathbf{M}_k\|_2 = \|\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\|_2.$$

Next,

$$t_k \geq 1 + \|\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\|_2.$$

Finally,

$$\|\mathbf{M}_k\|_2 = \frac{1}{t_k}\|\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\|_2 < 1.$$

□

B.2.1 Proof of Corollary 3.3

Corollary B.3. If $\mathbf{L}_{k-1} \succ 0$, the updated $0 \prec \mathbf{L}_k \preceq \mathbf{L}_{k-1}$.

Proof. It is enough to consider the matrix

$$\begin{aligned} \mathbf{L}_k &= \mathbf{L}_{k-1} - \frac{1}{t_k}\mathbf{L}_{k-1}\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top\mathbf{L}_{k-1} \\ &= \mathbf{L}_{k-1}^{1/2}(\mathbf{I} - \mathbf{M}_k)\mathbf{L}_{k-1}^{1/2} \end{aligned}$$

so that by (P10) we only need to examine the middle factor

$$\mathbf{I} - \mathbf{M}_k.$$

If this matrix is positive definite, then $\mathbf{L}_k \succ 0$. Since we have already proven this in Proposition 3.2, we are done. □

B.3 PROOF OF LEMMA 3.4

Lemma B.4. Assume a 1-Lipschitz loss, implying $\|\mathbf{G}_k\|_2 \leq 1 \forall k$. Suppose by iteration k , \mathbf{L}_k is updated in $k - k'$ of those steps for $0 \leq k' \leq k$ (and thus \mathbf{R}_k is updated in the remaining k' steps). Letting $\mathbf{B}_k = \mathbf{L}_k^{-1}$, $\mathbf{C}_k = \mathbf{R}_k^{-1}$, and $s = 1 + \|\mathbf{L}_0\|_2\|\mathbf{R}_0\|_2$,

$$\text{tr}(\mathbf{B}_k) \leq \text{tr}(\mathbf{B}_0) + k' m s \|\mathbf{R}_0\|_2 \quad (14)$$

$$\text{tr}(\mathbf{C}_k) \leq \text{tr}(\mathbf{C}_0) + (k - k') n s \|\mathbf{L}_0\|_2. \quad (15)$$

Proof. Suppose w.l.o.g. that k is an iteration in which we have updated \mathbf{L}_k . Let $\tilde{\mathbf{N}}_k = (\mathbf{G}_k\mathbf{R}_{k-1}\mathbf{G}_k^\top)^{1/2}$, $\mathbf{N}_k = \frac{1}{t_k^{1/2}}\tilde{\mathbf{N}}_k$, $\mathbf{T}_k = \mathbf{N}_k\mathbf{L}_{k-1}\mathbf{N}_k$, and $\tilde{\mathbf{T}}_k = \tilde{\mathbf{N}}_k\mathbf{L}_{k-1}\tilde{\mathbf{N}}_k$. Now, $\mathbf{T}_k = \frac{1}{t_k}\tilde{\mathbf{T}}_k$ and $t_k \geq 1 + \|\tilde{\mathbf{T}}_k\|_2$. First note the following generalized binomial expansion (or equivalently, the infinite geometric series),

$$(\mathbf{I} - \mathbf{M}_k)^{-1} = \sum_{i=0}^{\infty} \mathbf{M}_k^i,$$

which converges if $\|\mathbf{M}_k\|_2 < 1$. Since we have shown this in our proof of Proposition 3.2, we may use the series expansion. Now note that

$$\begin{aligned}
\mathbf{B}_k &= [\mathbf{L}_{k-1}^{1/2}(\mathbf{I}-\mathbf{M}_k)\mathbf{L}_{k-1}^{1/2}]^{-1} \\
&= \mathbf{L}_{k-1}^{-1/2}(\mathbf{I}-\mathbf{M}_k)^{-1}\mathbf{L}_{k-1}^{-1/2} \\
&= \mathbf{L}_{k-1}^{-1/2} \sum_{i=0}^{\infty} \mathbf{M}_k^i \mathbf{L}_{k-1}^{-1/2} && \because \|\mathbf{M}_k\| < 1 \\
&= \mathbf{L}_{k-1}^{-1} + \frac{1}{t_k} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top + \frac{1}{t_k^2} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top + \dots \\
&= \mathbf{B}_{k-1} + \mathbf{N}_k \sum_{i=0}^{\infty} (\mathbf{N}_k \mathbf{L}_{k-1} \mathbf{N}_k)^i \mathbf{N}_k,
\end{aligned}$$

Then

$$\begin{aligned}
\text{tr}(\mathbf{B}_k) &= \text{tr}(\mathbf{B}_{k-1}) + \text{tr}(\mathbf{N}_k \sum_{i=0}^{\infty} \mathbf{T}_k^i \mathbf{N}_k) \\
\Rightarrow \text{tr}(\mathbf{B}_k) - \text{tr}(\mathbf{B}_{k-1}) &\leq \|\mathbf{N}_k^2\|_2 \sum_{i=0}^{\infty} \text{tr}(\mathbf{T}_k^i) \\
&\leq \|\mathbf{N}_k^2\|_2 \left(m + \sum_{i=1}^{\infty} \|\mathbf{T}_k\|_2^{i-1} \text{tr}(\mathbf{T}_k) \right) && \because (\text{P7}) \\
&= \|\mathbf{N}_k^2\|_2 \left(m + \frac{\text{tr}(\tilde{\mathbf{T}}_k)}{\|\tilde{\mathbf{T}}_k\|_2} \sum_{i=1}^{\infty} \|\mathbf{T}_k\|_2^i \right) \\
&= \|\mathbf{N}_k^2\|_2 \left(m + \frac{\text{tr}(\tilde{\mathbf{T}}_k)}{\|\tilde{\mathbf{T}}_k\|_2} \sum_{i=1}^{\infty} \left(\frac{\|\tilde{\mathbf{T}}_k\|_2}{t_k} \right)^i \right) \\
&\leq \|\mathbf{N}_k^2\|_2 \left(m + \frac{\text{tr}(\tilde{\mathbf{T}}_k)}{\|\tilde{\mathbf{T}}_k\|_2} \sum_{i=1}^{\infty} \left(\frac{\|\tilde{\mathbf{T}}_k\|_2}{1 + \|\tilde{\mathbf{T}}_k\|_2} \right)^i \right) && \because t_k \geq 1 + \|\tilde{\mathbf{T}}_k\|_2 \\
&= \|\mathbf{N}_k^2\|_2 \left(m + \frac{\text{tr}(\tilde{\mathbf{T}}_k)}{\|\tilde{\mathbf{T}}_k\|_2} \sum_{i=1}^{\infty} \left(\frac{\|\tilde{\mathbf{T}}_k\|_2}{1 + \|\tilde{\mathbf{T}}_k\|_2} \right)^i \right) \\
&= \|\mathbf{N}_k^2\|_2 \left(m + \text{tr}(\tilde{\mathbf{N}}_k \mathbf{L}_{k-1} \tilde{\mathbf{N}}_k) \right),
\end{aligned}$$

where the last line follows by evaluating the convergent geometric series. Applying 1-Lipschitz loss, we have

$$\begin{aligned}
\text{tr}(\tilde{\mathbf{N}}_k \mathbf{L}_{k-1} \tilde{\mathbf{N}}_k) &\leq \|\mathbf{L}_{k-1}\|_2 \text{tr}(\tilde{\mathbf{N}}_k^2) \because (\text{P2}) \& \mathcal{L}(\text{P5}) \\
&= \|\mathbf{L}_{k-1}\|_2 \text{tr}(\mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top) \\
&= \|\mathbf{L}_{k-1}\|_2 \text{tr}(\mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{G}_k) \\
&\leq \|\mathbf{L}_{k-1}\|_2 \|\mathbf{R}_{k-1}\|_2 \text{tr}(\mathbf{G}_k^\top \mathbf{G}_k) \\
&= \|\mathbf{L}_{k-1}\|_2 \|\mathbf{R}_{k-1}\|_2 \|\mathbf{G}_k\|_F^2 \\
&\leq m \|\mathbf{L}_{k-1}\|_2 \|\mathbf{R}_{k-1}\|_2
\end{aligned}$$

where the last line applies equivalence of matrix norms in addition to the 1-Lipschitz loss. Further,

$$\begin{aligned}
\|\mathbf{N}_k^2\|_2 &= \frac{\|\mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_2}{t_k} \\
&\leq \frac{\|\mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_2}{1 + \|\tilde{\mathbf{T}}_k\|_2} \\
&\leq \|\mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_2 \\
&\leq \|\mathbf{G}_k\|_2 \|\mathbf{R}_{k-1}\|_2 \|\mathbf{G}_k^\top\|_2 \\
&\leq \|\mathbf{R}_{k-1}\|_2,
\end{aligned}$$

where the last line again uses 1-Lipschitz loss. Putting this together,

$$\mathrm{tr}(\mathbf{B}_k) - \mathrm{tr}(\mathbf{B}_{k-1}) \leq m \|\mathbf{R}_{k-1}\|_2 (1 + \|\mathbf{L}_{k-1}\|_2 \|\mathbf{R}_{k-1}\|_2)$$

By telescoping we get

$$\mathrm{tr}(\mathbf{B}_k) - \mathrm{tr}(\mathbf{B}_0) \leq \sum_{i \in \mathcal{I}} m \|\mathbf{R}_i\|_2 (1 + \|\mathbf{L}_i\|_2 \|\mathbf{R}_i\|_2),$$

where \mathcal{I} is the set of indices for which \mathbf{L}_k was updated. Then $|\mathcal{I}| = k'$. Finally, since we know that $\|\mathbf{L}_k\|_2 \leq \|\mathbf{L}_0\|_2$ and $\|\mathbf{R}_k\|_2 \leq \|\mathbf{R}_0\|_2$,

$$\mathrm{tr}(\mathbf{B}_k) - \mathrm{tr}(\mathbf{B}_0) \leq k' m \|\mathbf{R}_0\|_2 (1 + \|\mathbf{L}_0\|_2 \|\mathbf{R}_0\|_2)$$

□

B.4 PROOF OF THEOREM 3.6

Proof. Plugging in $(\mathbf{R}_{k,1}, \mathbf{L}_{k,2}) = (\mathbf{I}_m, \mathbf{I}_n)$ into (14)-(15) and again using the trace Hölder inequality (P6) yields

$$\begin{aligned}
\mathrm{tr}(\mathbf{B}_k^{1/4}) &\leq (m^3 \mathrm{tr}(\mathbf{B}_0) + km^4 (1 + \|\mathbf{L}_0\|_2))^{1/4} \\
\mathrm{tr}(\mathbf{C}_k^{1/4}) &\leq (n^3 \mathrm{tr}(\mathbf{C}_0) + kn^4 (1 + \|\mathbf{R}_0\|_2))^{1/4}.
\end{aligned}$$

The term in the regret $\mathrm{tr}(\mathbf{B}_k^{1/4}) \mathrm{tr}(\mathbf{C}_k^{1/4})$ is bounded by

$$(am^3 + m^4 bk)^{1/4} (cn^3 + n^4 dk)^{1/4} = O(mn\sqrt{k}),$$

where a, b, c, d are constants w.r.t. m, n, k, k' that can be expressed in terms of $\mathrm{tr}(\mathbf{B}_0)$, $\mathrm{tr}(\mathbf{C}_0)$, $\|\mathbf{L}_0\|_2$, and $\|\mathbf{R}_0\|_2$. □

C IMPLEMENTATION DETAILS

We provide a few additional implementation details here.

C.1 KRADAGRAD

We describe the algorithm for the alternate update scheme, which we dub KrADagrad.

C.2 DELAYED UPDATES

For KrADagrad*, we may choose an interval at which to update the preconditioners from the statistics. In this case, we may also use additional memory to batch the updates and reduce the FLOPs needed to compute the statistics to be equal to those needed for Shampoo. To derive the delayed update, consider the modification

$$\begin{aligned}
\tilde{\mathbf{Q}}_{k+\kappa} &= \mathbf{Q}_k + \frac{1}{nt_{k+\kappa}} \sum_{i=1}^{\kappa} \mathbf{g}_{k+i} \mathbf{g}_{k+i}^\top \\
\Rightarrow \Delta \mathbf{B}_{k+\kappa} &= \frac{1}{nt_{k+\kappa}} \sum_{i=1}^{\kappa} \mathbf{G}_{k+i} \mathbf{R}_k \mathbf{G}_{k+i}^\top.
\end{aligned}$$

Algorithm 2 KrADagrad

Input: Parameters $\mathbf{W}_0 \in \mathbb{R}^{m \times n}$, iterations K , update set \mathcal{I} , step size η , exponent $\alpha = 1/2$

Initialize: $(\mathbf{L}_0, \mathbf{R}_0) = (\mathbf{I}_m, \mathbf{I}_n)$

for $k = 1, \dots, K$ **do**

 Obtain gradient \mathbf{g}_k

if $k \in \mathcal{I}$ **then**

 Compute $\Delta \mathbf{L}_k = \frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{L}_{k-1}$

 Update $\mathbf{L}_k \leftarrow \mathbf{L}_{k-1} - \Delta \mathbf{L}_k$

else

 Compute $\Delta \mathbf{R}_k = \frac{1}{t_k} \mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1}$

 Update $\mathbf{R}_k \leftarrow \mathbf{R}_{k-1} - \Delta \mathbf{R}_k$

end if

 Approximate $(\mathbf{L}_k^\alpha, \mathbf{R}_k^\alpha)$ from $(\mathbf{L}_k, \mathbf{L}_{k-1}^\alpha, \mathbf{R}_k, \mathbf{R}_{k-1}^\alpha)$

 Apply preconditioned gradient step

$$\mathbf{W}_k = \mathbf{W}_{k-1} - \eta \mathbf{L}_k^\alpha \mathbf{G}_k \mathbf{R}_k^\alpha$$

end for

Then, using our extra memory to accumulate $\Delta \mathbf{B}_{k+\kappa}$, we can take

$$t_{k+\kappa} = 1 + \text{tr} \left(\mathbf{L}_k \sum_{i=1}^{\kappa} \mathbf{G}_{k+i} \mathbf{R}_k \mathbf{G}_{k+i}^\top \right)$$

and

$$\Delta \mathbf{L}_{k+\kappa} = \frac{1}{nt_{k+\kappa}} \mathbf{L}_k \sum_{i=1}^{\kappa} \mathbf{G}_{k+i} \mathbf{R}_k \mathbf{G}_{k+i}^\top \mathbf{L}_k.$$

C.3 MATRIX ROOTS

Matrix roots can be computed using numerical linear algebra methods (e.g. SVD) or iteratively. We covered the use of SVD in the main text, so we discuss iterative methods here. Using parenthesized superscript to denote the iterations, the simplest of these takes the form

$$\mathbf{X}^{(k+1)} = f(\mathbf{X}^{(k)}),$$

where f is a polynomial. More sophisticated methods introduce additional helper matrices. For example, one well-known method for iteratively computing the matrix square root and its inverse without matrix inversion is the Newton-Schulz (NS) method. The NS iteration is given as starting from $\mathbf{Y}^{(0)} = \mathbf{A}$, $\mathbf{Z}^{(0)} = \mathbf{I}$ and iterating

$$\begin{cases} \mathbf{X}^{(k)} &= \frac{1}{2} (3\mathbf{I} - \mathbf{Z}^{(k)} \mathbf{Y}^{(k)}) \\ \mathbf{Y}^{(k+1)} &= \mathbf{Y}^{(k)} \mathbf{X}^{(k)} \\ \mathbf{Z}^{(k+1)} &= \mathbf{X}^{(k)} \mathbf{Z}^{(k)} \end{cases} \quad (24)$$

until convergence. The values at convergence are $\mathbf{Y}^{(k)} \rightarrow \mathbf{A}^{1/2}$, $\mathbf{Z}^{(k)} \rightarrow \mathbf{A}^{-1/2}$. The iterations are derived from the so-called “matrix sign function” Higham [2008].

While the NS method converges quadratically, information about the matrix to be rooted \mathbf{A} is encoded not in the iterations themselves or in any invariants, but only in the initialization, making accurate warm starting of the estimate impossible (as far as we are aware). We could also use the numerically stable coupled Newton method for p -th roots of \mathbf{A} . However, these iterations involve matrix inversions, which are also fairly expensive, even using Newton iterations. Instead, we modify the also numerically stable coupled Newton’s method for $-p$ -th roots and compute $\mathbf{A}^{1/p} = \mathbf{A}(\mathbf{A}^{-1/p})^{p-1}$ Higham [2008]. While it may seem that this requires high precision to compute (as $\mathbf{A}^{-1/p}$ is exactly what Shampoo computes),

the difference is that in multiplying by \mathbf{A} , the impact of inaccuracies in small eigenvalues remains small. Initializing $(\mathbf{X}^{(0)}, \mathbf{M}^{(0)}) = (\mathbf{I}, \mathbf{A})$, we arrive at

$$\begin{cases} \mathbf{N}^{(k)} &= \frac{1}{p} ((p+1)\mathbf{I} - \mathbf{M}^{(k)}) \\ \mathbf{Y}^{(k)} &= \mathbf{N}^{(k)} \mathbf{X}^{(k)} \\ \mathbf{X}^{(k+1)} &= \frac{1}{2} (\mathbf{Y}^{(k)} + \mathbf{Y}^{(k)\top}) \\ \mathbf{M}^{(k+1)} &= (\mathbf{N}^{(k+1)})^p \mathbf{M}^{(k)}, \end{cases} \quad (25)$$

where $\mathbf{X}^{(k)} \rightarrow \mathbf{A}^{-1/p}$, $\mathbf{M}^{(k)} \rightarrow \mathbf{I}$. These iterations also appear not to encode information about \mathbf{A} ; however, for $p = 2\phi$, invariant in the iterations is $\mathbf{M}^{(k)} = (\mathbf{X}^{(k)})^\phi \mathbf{A} (\mathbf{X}^{(k)})^\phi$. For even integers p , this gives us exact guidance on how to initialize the iterations with a good guess $\widehat{\mathbf{X}} \approx \mathbf{A}^{-1/p}$ that leads to the desired fixed point: $(\mathbf{X}^{(0)}, \mathbf{M}^{(0)}) = (\widehat{\mathbf{X}}, \widehat{\mathbf{X}}^\phi \mathbf{A} \widehat{\mathbf{X}}^\phi)$. We note that if $\mathbf{N}^{(k)}$ and $\mathbf{X}^{(k)}$ commute, the step of symmetrizing $\mathbf{Y}^{(k)}$ is unnecessary; in practice, due to numerics and especially if we use a warm start, this step ensures that the sequence $\mathbf{X}^{(k)}$ stays symmetric. These iterations also require commutativity of \mathbf{A} and $\mathbf{X}^{(k)}$, and thus non-commutativity may still be a source of divergence. Of course, this initialization procedure now requires additional storage of $\mathbf{A}^{-1/p}$ in between preconditioning computation steps, so we may still opt to use the prescribed initialization depending on the compute-memory requirements of our hardware and software system.

C.4 DIAGONAL DAMPING

For numerical reasons, Adagrad and Shampoo add a small scaled identity $\epsilon \mathbf{I}$ prior to matrix inversion (i.e. to avoid singularity). For KrADagrad, while not necessary for numerical stability, we find that adding damping in some cases helps reduce the effect of noise in the gradient estimate on the preconditioner. We hypothesize that it reduces the effect of noise in the direction of eigenvectors with small eigenvalues. Assuming $\epsilon \mathbf{L}_k \prec \mathbf{I}$, we have

$$\begin{aligned} (\epsilon \mathbf{I} + \mathbf{B}_k)^{-1} &= (\epsilon \mathbf{L}_k + \mathbf{I})^{-1} \mathbf{L}_k \\ &\succeq \mathbf{L}_k (\mathbf{I} - \epsilon \mathbf{L}_k) \\ \Rightarrow (\mathbf{L}_k (\mathbf{I} - \epsilon \mathbf{L}_k))^{-1} &\succeq \epsilon \mathbf{I} + \mathbf{B}_k. \end{aligned}$$

Then letting $\mathbf{L}'_k = \mathbf{L}_k (\mathbf{I} - \epsilon \mathbf{L}_k)$, this looks like

$$\left(\mathbf{L}'_k - \frac{1}{t'_k} \mathbf{L}'_k \mathbf{G} \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}'_k \right)^{-1} \succeq \epsilon \mathbf{I} + \mathbf{B}_k + \frac{1}{t} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top.$$

This implies that the update should be computed with \mathbf{L}_k while the damped preconditioner should be computed with the modified \mathbf{L}'_k . However, as a computational savings, we can incorporate \mathbf{L}'_k into both the update and the preconditioner computations, and suffer negligible error as long as ϵ is set small enough.

C.5 GENERALIZATION TO TENSORS

Extending the algorithm to tensors is an exercise in bookkeeping. We first introduce some additional terminology and notation that is standard in the literature for analysis of tensors as well as practical implementation of tensor algebra Lu et al. [2013], Gupta et al. [2018], Ren and Goldfarb [2021], Paszke et al. [2017]. For the following, consider two tensors \mathbf{A}_1 and \mathbf{A}_2 with D_1 and D_2 dimensions, respectively,

$$\mathbf{A}^{(i)} \in \mathbb{R}^{n_1 \times \dots \times n_{D_i}}.$$

1. Let $\mathcal{D}_i \subset \{1, \dots, D_i\}$ be ordered sets $(d_{i,1}, \dots, d_{i,p})$ with $p = |\mathcal{D}_1| = |\mathcal{D}_2|$. Define the *tensordot* to be such that the element with index $i_1, \dots, i_{D_1+D_2-2p}$ is

$$\begin{aligned} &[\text{TD}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, \mathcal{D}_1, \mathcal{D}_2)]_{i_1, \dots, i_{D_1+D_2-2p}} \\ &= \sum_{s_p=1}^{n_{d_{1,p}}} \dots \sum_{s_1=1}^{n_{d_{1,1}}} \mathbf{A}_{i_1, \dots, i_{d_{1,1}-1}, s_1, i_{d_{1,1}+1}, \dots, i_{d_{2,p}-1}, s_p, i_{d_{2,p}+1}, \dots}^{(1)} \mathbf{A}_{i_1, \dots, i_{d_{2,1}-1}, s_1, i_{d_{2,1}+1}, \dots, i_{d_{2,p}-1}, s_p, i_{d_{2,p}+1}, \dots}^{(2)}, \end{aligned}$$

which is a summation over the indices in the sets \mathcal{D}_1 and \mathcal{D}_2 of the broadcasted product aligned over the indices specified by \mathcal{D}_1 and \mathcal{D}_2 . Note this is only well defined if $n_{d_{1,k}} = n_{d_{2,k}}$ for $k = 1, \dots, p$. The tensordot is a multidimensional

Algorithm 3 KrADagrad* for tensors

Input: Parameter tensor $\mathbf{W}_0 \in \mathbb{R}^{n_1 \times \dots \times n_D}$ of order D , iterations K , step size η , damping ϵ , exponent $\alpha = 1/2$

Initialize: $\mathbf{L}_{0,d} = \mathbf{I}_{n_d}$

for $k = 1, \dots, K$ **do**

 Obtain gradient \mathbf{G}_k

for $d = 1, \dots, D$ **do**

 Compute

$$\Delta \mathbf{L}_{k,d} = \frac{1}{t_{k,d}} \mathbf{L}_{k-1,d} \text{contr}_d(\mathbf{G}_k) \mathbf{L}_{k-1,d}$$

 Update $\mathbf{L}_{k,d} \leftarrow \mathbf{L}_{k-1,d} - \Delta \mathbf{L}_{k,d}$

 Compute $\mathbf{L}_{k,d}^{\alpha/D}$ from $\mathbf{L}_{k,d}$

end for

 Apply preconditioned gradient step

$$\mathbf{W}_k = \mathbf{W}_{k-1} - \eta \mathbf{G}_k \times_1 \mathbf{L}_{k,1}^{\alpha/D} \cdots \times_D \mathbf{L}_{k,D}^{\alpha/D}$$

end for

generalization of matrix multiplication and traces. To see this, suppose $\mathbf{A}^{(1)}, \mathbf{A}^{(2)}$ are matrices and \mathbf{x}, \mathbf{y} are column vectors,

$$\text{TD}(\mathbf{x}, \mathbf{y}, [1], [1]) = \mathbf{x}^\top \mathbf{y}$$

$$\text{TD}(\mathbf{A}^{(1)}, \mathbf{x}, [2], [1]) = \mathbf{A}^{(1)} \mathbf{x}$$

$$\text{TD}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, [2], [1]) = \mathbf{A}^{(1)} \mathbf{A}^{(2)}$$

$$\text{TD}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, [1], [1]) = \mathbf{A}^{(1)\top} \mathbf{A}^{(2)}$$

$$\text{TD}(\mathbf{A}^{(1)}, \mathbf{A}^{(2)}, [1, 2], [1, 2]) = \text{tr}(\mathbf{A}^{(1)\top} \mathbf{A}^{(2)}) = \mathbf{A}^{(1)} \cdot \mathbf{A}^{(2)}$$

2. Take $\mathcal{D}_{\setminus d} = [1, \dots, D] \setminus d$. Let the *contraction* of a tensor be the matrix of size $n_d \times n_d$ resulting from

$$\text{contr}_d(\mathbf{A}) = \text{TD}(\mathbf{A}, \mathbf{A}, \mathcal{D}_{\setminus d}, \mathcal{D}_{\setminus d})$$

3. For $1 \leq d \leq D$ and matrix $\mathbf{L}_d \in \mathbb{R}^{n_d \times m}$, let the d -th tensor-matrix product be

$$\mathbf{A} \times_d \mathbf{L}_d = \text{TD}(\mathbf{L}_d, \mathbf{A}, [2], [d]).$$

D PROPERTIES AND EXTENDED DERIVATIONS

Here we list the main linear algebra properties utilized in derivations and proofs. We then extend the derivations as a mini-tutorial in matrix manipulation for maximum clarity.

D.1 LINEAR ALGEBRA PROPERTIES

Matrix shapes are assumed to be conforming when left unspecified.

P1. $\mathbf{A} \cdot \mathbf{B} = \mathbf{B} \cdot \mathbf{A}$.

P2. (P1) implies $\text{tr}(\mathbf{ABC}) = \text{tr}(\mathbf{BCA})$.

P3. If $\mathbf{A}, \mathbf{B} \succ 0$, then $\mathbf{A} \cdot \mathbf{B} \succ 0$.

P4. If $\mathbf{A} \succeq 0$, then $\|\mathbf{A}\|_2 \leq \text{tr}(\mathbf{A})$.

P5. If $\mathbf{A}, \mathbf{B} \succeq 0$, then $\mathbf{A} \cdot \mathbf{B} \leq \|\mathbf{A}\|_2 \text{tr}(\mathbf{B})$.

- P6.** If $\mathbf{A}, \mathbf{B} \succ 0$, then for p, q such that $1/p + 1/q = 1$, $\mathbf{A}^{1/p} \cdot \mathbf{B}^{1/q} \leq (\text{tr}(\mathbf{A}))^{1/p} (\text{tr}(\mathbf{B}))^{1/q}$.
- P7.** (P4) and (P5) together imply that for $\mathbf{A} \succeq 0$, $\text{tr}(\mathbf{A}^i) \leq \|\mathbf{A}\|_2^{i-1} \text{tr}(\mathbf{A}) \leq \text{tr}(\mathbf{A})^i$ for $i \geq 1$.
- P8.** (P5) implies that for $\mathbf{A}, \mathbf{B} \succeq 0$, $\|\mathbf{AB}\|_F \leq \|\mathbf{A}\|_2 \|\mathbf{B}\|_F$.
- P9.** If $\mathbf{A} \succeq \mathbf{B} \succ 0$, then for $0 \leq p \leq 1$, $\mathbf{A}^p \succeq \mathbf{B}^p$.
- P10.** If $\mathbf{A}, \mathbf{B}, \mathbf{C} \succeq 0$, then $\mathbf{A} \succeq \mathbf{B} \Rightarrow \mathbf{CAC}^\top \succeq \mathbf{CBC}^\top$.
- P11.** If diagonal matrices $\mathbf{D} \succeq \mathbf{C} \succ 0$, then $\mathbf{C}^{-1} \succeq \mathbf{D}^{-1}$.
- P12.** (P11) implies that if $\mathbf{B} \succeq 0, \mathbf{C} \succ 0$, then $\mathbf{C}^{-1} \succeq (\mathbf{C} + \mathbf{B})^{-1}$.
- P13.** $\text{tr}(\mathbf{A} \otimes \mathbf{B}) = \text{tr}(\mathbf{A}) \text{tr}(\mathbf{B})$.
- P14.** $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD})$.
- P15.** $\mathbf{A} \otimes (\mathbf{B} + \mathbf{C}) = \mathbf{A} \otimes \mathbf{B} + \mathbf{A} \otimes \mathbf{C}$.
- P16.** If $\mathbf{A}, \mathbf{B} \succeq 0$, then $(\mathbf{A} \otimes \mathbf{B})^r = \mathbf{A}^r \otimes \mathbf{B}^r \forall r \in \mathbb{R}$.
- P17.** $(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top$.
- P18.** $(\mathbf{C}^\top \otimes \mathbf{A}) \text{vec}(\mathbf{B}) = \text{vec}(\mathbf{ABC})$.

D.2 ADDITIONAL AND EXTENDED DERIVATIONS

D.2.1 Proof of (P10)

For arbitrary \mathbf{x} , let $\mathbf{y} = \mathbf{C}^\top \mathbf{x}$. We have,

$$\begin{aligned} \mathbf{x}^\top \mathbf{CAC}^\top \mathbf{x} &= \mathbf{y}^\top \mathbf{A} \mathbf{y} \\ &\geq \mathbf{y}^\top \mathbf{B} \mathbf{y} && \because \mathbf{A} \succeq \mathbf{B} \\ &= \mathbf{x}^\top \mathbf{CBC}^\top \mathbf{x} \end{aligned}$$

D.2.2 Proof of (P12)

Let the SVD of $\mathbf{C}^{-1/2} \mathbf{B} \mathbf{C}^{-1/2} = \mathbf{V} \mathbf{\Sigma} \mathbf{V}^\top$. Then,

$$\begin{aligned} \mathbf{C}^{-1} &= \mathbf{C}^{-1/2} \mathbf{I} \mathbf{C}^{-1/2} \\ &= \mathbf{C}^{-1/2} \mathbf{V} \mathbf{I} \mathbf{V}^\top \mathbf{C}^{-1/2} && \because \mathbf{V}^{-1} = \mathbf{V}^\top \\ &\succeq \mathbf{C}^{-1/2} \mathbf{V} (\mathbf{I} + \mathbf{\Sigma})^{-1} \mathbf{V}^\top \mathbf{C}^{-1/2} && \because \text{(P10)\&(P11)} \\ &= \mathbf{C}^{-1/2} (\mathbf{V} \mathbf{V}^\top + \mathbf{V} \mathbf{\Sigma} \mathbf{V}^\top)^{-1} \mathbf{C}^{-1/2} \\ &= \mathbf{C}^{-1/2} (\mathbf{I} + \mathbf{C}^{-1/2} \mathbf{B} \mathbf{C}^{-1/2})^{-1} \mathbf{C}^{-1/2} \\ &= (\mathbf{C} + \mathbf{B})^{-1} \end{aligned}$$

D.2.3 Extended derivation of KrADagrad update

We can use Woodbury directly on the inverse before applying KrAD,

$$\begin{aligned} \tilde{\mathbf{Q}}_k^{-1} &= (\mathbf{Q}_{k-1} + \frac{1}{nt_k} \mathbf{g}_k \mathbf{g}_k^\top)^{-1} \\ &= \mathbf{Q}_{k-1}^{-1} - \frac{1}{nt_k (1 + \mathbf{g}_k^\top \mathbf{Q}_{k-1}^{-1} \mathbf{g}_k)} \mathbf{Q}_{k-1}^{-1} \mathbf{g}_k \mathbf{g}_k^\top \mathbf{Q}_{k-1}^{-1} \\ \Rightarrow \mathbf{Q}_{k-1}^{-1} - \tilde{\mathbf{Q}}_k^{-1} &= \frac{1}{nt_k (1 + \text{tr}(\mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k))} \mathbf{Q}_{k-1}^{-1} \mathbf{g}_k \mathbf{g}_k^\top \mathbf{Q}_{k-1}^{-1}. \end{aligned}$$

Now using KrAD on the RHS above,

$$\begin{aligned} \frac{1}{n} \mathbf{Q}_{k-1}^{-1} \mathbf{g}_k \mathbf{g}_k^\top \mathbf{Q}_{k-1}^{-1} &= \frac{1}{n} \text{vec}(\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k) \text{vec}(\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k)^\top \\ &\succeq \mathbf{R}_k \otimes (\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{R}_k^{-1} \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}) \\ &= \mathbf{R}_k \otimes (\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}) \end{aligned}$$

so that

$$\tilde{\mathbf{Q}}_k^{-1} \succeq \mathbf{R}_k \otimes \mathbf{L}_{k-1} - \frac{1}{t_k(1 + \text{tr}(\mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k))} \mathbf{R}_k \otimes (\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1}).$$

Then letting

$$\Delta \mathbf{L}_k = \frac{1}{1 + \text{tr}(\mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k)} \mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_k \mathbf{G}_k^\top \mathbf{L}_{k-1},$$

take $t_k = 1$ and $\mathbf{L}_k = \mathbf{L}_{k-1} - \Delta \mathbf{L}_{k-1}$, and we have

$$\mathbf{Q}_k^{-1} = \mathbf{R}_k \otimes \mathbf{L}_k \preceq \tilde{\mathbf{Q}}_k^{-1} \preceq \left(\frac{1}{n} \mathbf{G}_k \mathbf{G}_k^\top \right)^{-1}.$$

D.2.4 Alternative derivation of KrADagrad update

We can also start with

$$\begin{aligned} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k &\succeq 0 \\ \Rightarrow \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k + \mathbf{C}_k &\succeq \mathbf{C}_k \\ \Rightarrow \mathbf{C}_k^{-1} &\succeq (\mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k + \mathbf{C}_k)^{-1}, \quad \because \text{(P12)} \end{aligned}$$

which implies that

$$\frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k (\mathbf{C}_k)^{-1} \mathbf{G}_k^\top \mathbf{L}_{k-1} \succeq \frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k (\mathbf{C}_k + \frac{1}{t_k} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k)^{-1} \mathbf{G}_k^\top \mathbf{L}_{k-1} \quad \because \text{(P10)}$$

and thus

$$-\frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k (\mathbf{C}_k + \frac{1}{t_k} \mathbf{G}_k^\top \mathbf{L}_{k-1} \mathbf{G}_k)^{-1} \mathbf{G}_k^\top \mathbf{L}_{k-1} \succeq -\frac{1}{t_k} \mathbf{L}_{k-1} \mathbf{G}_k (\mathbf{C}_k)^{-1} \mathbf{G}_k^\top \mathbf{L}_{k-1}$$

E MEMORY AND COMPUTE COST ANALYSIS

E.1 KRADAGRAD: STATISTIC UPDATE

For KrADagrad, we can actually cache some of the computations from the statistic update for later use in applying the preconditioner. Computing $\Delta \mathbf{L}_k$ has a computational cost of $2N(m+n) + 4m^2 + 4m^3 \leq 4m(2n^2 + n)$ from the following:

1. $\mathbf{G}_k \mathbf{R}_{k-1}^{1/2} \rightarrow 2mn^2 = 2Nn$
2. $\mathbf{G}_k \mathbf{R}_{k-1}^{1/2} (\mathbf{G}_k \mathbf{R}_{k-1}^{1/2})^\top \rightarrow 2m^2n = 2Nm$
3. $\mathbf{L}_{k-1} (\mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top) \rightarrow 2m^3$
4. $t_k = 1 + \|\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top\|_F \rightarrow 2m^2$
5. $\Delta \mathbf{L}_k = (\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{R}_{k-1} \mathbf{G}_k^\top) \mathbf{L}_{k-1} \rightarrow 2m^3$
6. $\frac{1}{t_k} \Delta \mathbf{L}_k \rightarrow m^2$
7. $\mathbf{L}_{k-1} - \frac{1}{t_k} \Delta \mathbf{L}_k \rightarrow m^2$

Similarly, computing $\Delta \mathbf{R}_k$ has a computational cost upper bounded by $2N(m+n) + 4n^2 + 4n^3 \leq 4n(2n^2 + n)$. Thus, using a simple alternating scheme, the per update computation cost is bounded by the average, $2(m+n)(2n^2 + n) \rightarrow O(n^3)$.

The memory cost is $m^2 + n^2 \rightarrow O(n^2)$ beyond gradient descent. We may cache $\mathbf{G}_k \mathbf{R}_{k-1}^{1/2}$ replacing \mathbf{G}_k in memory, which incurs no additional storage cost.

E.2 KRADAGRAD: MATRIX SQUARE ROOT

The cost comes from NS iterations or SVD. For each NS iteration, the cost is $O(n^3)$, and the SVD also has cost $O(n^3)$.

E.3 KRADAGRAD: APPLYING PRECONDITIONER

In iterations where we update \mathbf{L}_k , using the cached $\mathbf{G}_k \mathbf{R}_{k-1}^{1/2}$, we merely need to compute

$$\mathbf{L}_{k-1}^{1/2} (\mathbf{G}_k \mathbf{R}_{k-1}^{1/2}) \rightarrow 2m^2 n = 2Nm.$$

For iterations in which we update \mathbf{R}_k , we similarly have a cost of $2Nn$. Again using a simple alternating scheme, the average cost is then $(m+n)N \rightarrow O(n^3)$.

E.4 KRADAGRAD*: STATISTICS UPDATES

For Kradagrad*, computing $\Delta \mathbf{L}_k$ has a computational cost of $(2Nm + 2m^2 + 4m^3) = 2m^2(n + 2m + 1)$ from the following:

1. $\mathbf{G}_k \mathbf{G}_k^\top \rightarrow 2m^2 n = 2Nm$
2. $\mathbf{L}_{k-1} (\mathbf{G}_k \mathbf{G}_k^\top) \rightarrow 2m^3$
3. $\|\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{G}_k^\top\|_F \rightarrow 2m^2$
4. $(\mathbf{L}_{k-1} \mathbf{G}_k \mathbf{G}_k^\top) \mathbf{L}_{k-1} \rightarrow 2m^3$

Similarly, computing $\Delta \mathbf{R}_k$ has a computational cost of $(4Nn + 2n^2 + 2n^3) = 2n^2(n + 2m + 1)$ from the following:

1. $\mathbf{R}_k \mathbf{G}_k^\top \rightarrow 2mn^2 = 2Nn$
2. $(\mathbf{R}_{k-1} \mathbf{G}_k^\top) \mathbf{G}_k \rightarrow 2mn^2 = 2Nn$
3. $\|\mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{G}_k\|_F \rightarrow 2n^2$
4. $(\mathbf{R}_{k-1} \mathbf{G}_k^\top \mathbf{G}_k) \mathbf{R}_{k-1} \rightarrow 2n^3$

Thus the total per update compute cost is $2(m^2 + n^2)(n + 2m + 1) \rightarrow O(n^3)$.

Again the memory cost is straightforward at $m^2 + n^2 \rightarrow O(n^2)$.

E.5 KRADAGRAD*: MATRIX p -TH ROOTS

The cost comes from the coupled Newton iterations. The cost of a couple Newton iteration is $O(n^3)$, and the SVD also has cost $O(n^3)$.

E.6 KRADAGRAD*: APPLYING PRECONDITIONERS

The cost is

1. $\mathbf{L}_k^{1/2p} \mathbf{G}_k \rightarrow 2m^2 n = 2Nm$
2. $(\mathbf{L}_k^{1/2p} \mathbf{G}_k) \mathbf{R}_k^{1/2p} \rightarrow 2mn^2 = 2Nn$

for a total of $2(m+n)N \rightarrow O(n^3)$.

F OTHER KRONECKER APPROXIMATIONS

Here we describe a traditional Kronecker approximation that may point in a direction for further research. Given matrix $\mathbf{U} \in \mathbb{R}^{m \times r}$, we wish to minimize the Frobenius norm

$$(\mathbf{L}, \mathbf{R}) = \operatorname{argmin}_{\mathbf{L}', \mathbf{R}'} \|\mathbf{R}' \otimes \mathbf{L}' - \mathbf{U} \mathbf{U}^\top\|_F^2. \quad (26)$$

Note that there is an equivalence class for solutions, i.e. if (\mathbf{B}, \mathbf{C}) is a solution, then so is $(\alpha\mathbf{B}, \frac{1}{\alpha}\mathbf{C})$ for $\alpha \neq 0$. One solution is related to taking the singular vectors corresponding to the largest singular value of a permuted and reshaped matrix Van Loan and Pitsianis [1993]. Alternatively, for large matrices, we can apply alternating minimization to obtain iterative updates for $\mathbf{L}^{(t+1)}$ at iteration $t + 1$ given $\mathbf{R}^{(t)}$ at iteration t and vice versa,

$$\mathbf{L}^{(t+1)} = \underset{\mathbf{L}}{\operatorname{argmin}} \left\| \mathbf{R}^{(t)} \otimes \mathbf{L} - \mathcal{U}\mathcal{U}^\top \right\|_F^2 \quad (27)$$

$$\mathbf{R}^{(t+1)} = \underset{\mathbf{R}}{\operatorname{argmin}} \left\| \mathbf{R} \otimes \mathbf{L}^{(t+1)} - \mathcal{U}\mathcal{U}^\top \right\|_F^2. \quad (28)$$

Under mild conditions, this process converges to a local minimum Tseng [2001]. We can choose to perform half an iteration, iterate until convergence, or anywhere in between.

Letting \mathcal{U}_i be the i -th column of \mathcal{U} , and $\mathbf{U}_i = \operatorname{vec}_{m,n}^{-1}(\mathcal{U}_i)$, note that

$$\begin{aligned} & \operatorname{tr} \left(\mathcal{U}^\top (\mathbf{R}^{(t)} \otimes \mathbf{L}) \mathcal{U} \right) \\ &= \sum_{i=1}^r \mathcal{U}_i^\top (\mathbf{R}^{(t)} \otimes \mathbf{L}) \mathcal{U}_i \\ &= \sum_{i=1}^r \mathcal{U}_i^\top \operatorname{vec}(\mathbf{L} \mathbf{U}_i \mathbf{R}^{(t)}) \\ &= \sum_{i=1}^r \operatorname{tr}(\mathbf{L} \mathbf{U}_i \mathbf{R}^{(t)} \mathbf{U}_i^\top) \end{aligned}$$

Then, we can rewrite and take the derivative of the objective function (dropping irrelevant constants) and set it to zero to solve for the partial optimum,

$$\begin{aligned} J(\mathbf{B}) &= -2 \operatorname{tr} \left(\mathcal{U}\mathcal{U}^\top (\mathbf{R}^{(t)} \otimes \mathbf{L}) \right) + \operatorname{tr}((\mathbf{R}^{(t)\top} \mathbf{R}^{(t)}) \otimes (\mathbf{L}^\top \mathbf{L})) \\ &= -2 \sum_{i=1}^r \operatorname{tr}(\mathbf{L} \mathbf{U}_i \mathbf{R}^{(t)} \mathbf{U}_i^\top) + \operatorname{tr}(\mathbf{L}^\top \mathbf{L}) \operatorname{tr}(\mathbf{R}^{(t)\top} \mathbf{R}^{(t)}) \\ \Rightarrow \frac{\partial J}{\partial \mathbf{B}} &= 2 \left\| \mathbf{R}^{(t)} \right\|_F^2 \mathbf{L} - 2 \sum_{i=1}^r \mathbf{U}_i \mathbf{R}^{(t)} \mathbf{U}_i^\top = 0 \\ \Rightarrow \mathbf{L}^{(t+1)} &= \frac{1}{\left\| \mathbf{R}^{(t)} \right\|_F^2} \sum_{i=1}^r \mathbf{U}_i \mathbf{R}^{(t)} \mathbf{U}_i^\top. \end{aligned}$$

Similarly,

$$\mathbf{R}^{(t+1)} = \frac{1}{\left\| \mathbf{L}^{(t+1)} \right\|_F^2} \sum_{i=1}^r \mathbf{U}_i^\top \mathbf{L}^{(t+1)} \mathbf{U}_i.$$

If we use traditional approximation to form the basis of a preconditioner, the resulting updates, while simple in form, would not satisfy the analogous properties in Equations (4)-(5) that would guarantee optimal regret. Thus, the traditional Kronecker approximation requires another set of properties that lead to optimal regret.

G EXPERIMENT DETAILS

G.1 SYNTHETIC

We form $\mathbf{A} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$, where \mathbf{U}^\top is a random orthonormal matrix (generated by taking the Q from a QR decomposition of a square matrix with i.i.d. $\mathcal{N}(0, 1)$ entries), and $\mathbf{\Lambda}$ is a diagonal matrix with entries $\lambda_i = 10^{(\sqrt{10} - \sqrt{10(i-1)/n})^2}$ for $i = 1, \dots, n$ (where the square/square root nonlinearities in the exponent make the spread in the top eigenvalues less disparate while not decreasing the overall condition number). We repeat a similar process for \mathbf{B} . We can see that the




Figure 6: Loss on synthetic quadratic (we do not observe pathologically bad behavior)




Figure 7: Left: Top 1 Accuracy of ResNet-32 (without BN) on CIFAR-10; Right: Top 1 Accuracy of ResNet-56 (without BN) on CIFAR-100. We do not observe pathologically bad behavior.

optimal is $\mathbf{X} = \mathbf{0}$ and that the Hessian of the loss function w.r.t. $\mathbf{x} = \text{vec}(\mathbf{X})$ is $2\mathbf{B} \otimes \mathbf{A}$ with a condition number of $\kappa(\mathbf{B} \otimes \mathbf{A}) = \kappa(\mathbf{A})\kappa(\mathbf{B}) = 10^{20}$, which is exceptionally badly conditioned. We found that for the preconditioned methods an interval of 3 steps in between updating the preconditioner performed best.

Table 3: Selected learning rate (epsilon) for synthetic quadratic.

Synthetic				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
1e-3, -	1e0, 1e-8	5e-3, 1e-8	5e-3, 1e-8	1e-2, 1e-8

The learning rate was swept over the range [1e-3, 2e-3, 5e-3, 1e-2, 2e-2, 5e-2, 1e-1, 2e-1, 5e-1, 1e0] for all methods, while epsilon was swept log uniformly with 6 steps over the range [1e-3, 1e-8] for all methods.

G.2 VISION EXPERIMENTS

We train a ResNet32 model on CIFAR-10 and a ResNet56 model on CIFAR-100, following an implementation by [Idelbayev], modified to be without (BN).

Figure 7 shows the absolute performance of these two models on the respective datasets they were trained on. The learning rate was swept over the range [1e-4, 2.5e-4, 1e-3, 2.5e-3, 1e-2, 2.5e-2, 1e-1] for preconditioned methods and [1e-3, 2e-3, 5e-3, 1e-2, 2e-2, 5e-2, 1e-1] for SGD and Adam, while epsilon was swept log uniformly with 6 steps over the range [1e-1, 1e-6] for all methods, and the batch size was set to 128 for all methods.

G.3 AUTOENCODER EXPERIMENTS

We closely follow [Goldfarb et al., 2020] Appendix D.2 Tables 3 and 4 in defining the datasets and hyperparameters for the autoencoder benchmarks. We first perform a hyperparameter search for each optimizer/dataset pair over the learning rate

Table 4: Selected learning rate (epsilon) for ResNet-32/56 on CIFAR-10/100 comparisons.

ResNet-32 on CIFAR-10				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
5e-3, -	5e-3, 1e-1	1e-2, 1e-4	1e-2, 1e-4	2.5e-2, 1e-4
ResNet-56 on CIFAR-100				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
1e-2, -	5e-3, 1e-3	1e-2, 1e-4	1e-2, 1e-4	1e-2, 1e-4




Figure 8: Reconstruction validation error curves for the best HPs for fully connected auto encoder a) mean cross entropy on MNIST b) mean cross entropy on CURVES c) mean squared error on FACES. All 5 seeds are shown per optimizer.

and epsilon parameters. For all optimizers and datasets we sample the learning rate log uniformly between 10^{-5} and 2, and sample epsilon log uniformly on the intervals $[10^{-10}, 10^{-1}]$ and $[10^{-4}, 1]$ for Adam and 2nd order optimizers, respectively. For each dataset we run 30 trials for SGD and ~ 90 trials for the other optimizers. We use a batch size of 100, momentum of 0.9 for SGD, and Shampoo and KrAdagrad* use block sizes of 100, 250, and 500 for the CURVES, MNIST, and FACES datasets, respectively. The best performing hyperparameters are provided in Table 5. After finding the best HPs from a single seed, we retrain with the selected HPs on five new, unique seeds shared across optimizers. See the raw learning curves in Figure 8 and averages in Figure 3.

Table 5: Selected learning rate (epsilon) for autoencoder comparisons.

MNIST				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
5.2e-3, -	3e-4, 1e-4	1e-3, 1e-2	1.4e-3, 3e-4	6.3e-3, 8.5e-4
CURVES				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
1e-3, -	1e-3, 1e-3	1.9e-3, 5.8e-4	3e-3, 3.2e-2	1.3e-2, 5e-3
FACES				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
1e-1, -	3e-3, 1e-3	3.1e-1, 3.6e-3	1.9e-1, 2e-4	3.2e-1, 1e-4

G.4 RECOMMENDER EXPERIMENT

For the MovieLens-20M experiment we start from the existing H+Vamp Gated baseline¹ and modify it to compare the optimizers of interest. Using the default hyperparameters otherwise, we find that learning rates greater than or equal to $4e-4$ cause the training loss to diverge for most of the optimizers. As a result we use the following learning rate grid search, and report the best performing rates in 6:

- SGD, Adam (norm grad)[
 - learning rates 1.25e-5, 2.5e-5, 5e-5, 1e-4, 2e-4, 4e-4]

¹<https://github.com/psywaves/EVCF>

- epsilon 1e-6 (Adam only)
- Adam, Shampoo, KrAdagrad*, KrAdagrad
 - learning rates [6.25e-6, 1.25e-5, 2.5e-5, 5e-5, 1e-4 2e-4]
 - epsilon 1e-4 (1e-3 for KrAdagrad)

Table 6: Selected learning rate (epsilon) for MovieLens 20M experiment.

SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad	Adam (norm grad)
2e-4	2e-4	2e-4	2e-4	2e-4	4e-4

For this experiment, we instead tune against the training loss (that which is optimized directly) after discovering that the equivalent validation loss function weights the loss components differently, as described in Figure 5.

G.5 CONTINUAL LEARNING EXPERIMENTS

In the Permuted MNIST benchmark [Goodfellow et al., 2013], the learner trains for 2000 steps, seeing a new pixel permutation of the images every 100 steps, resulting in 20 sequential non-overlapping experiences. Every 100 steps the Top 1 Accuracy is evaluated against all 20 experiences, including those not yet seen, yielding an aggregate measure of model accuracy in both current and prior domains. For the Split CIFAR-100 benchmark, the learner trains on exactly 5 unseen classes at a time every interval of 2500 training steps, evaluating Top 1 Accuracy against all 20 experiences after each interval, until reaching 50000 training steps. Starting with the GEM and LaMAML baselines ² from Lomonaco et al. [2021], we train with 3 unique seeds shared across optimizers, with learning curves of Top 1 Accuracy averaged across seeds for the best HPs in Figure 9. We share the same curves relative to SGD in Figure 4a, and relative to Adam in Figure 4b. We perform the following grid searches to select a best learning rate per optimizer (see Table 7 for the selected learning rates):

- SGD, Adam
 - learning rates [1e-1, 5e-2, 2e-2, 1e-2, 5e-3, 2e-3, 1e-3, 5e-4, 2e-4, 1e-4]³
 - epsilon 1e-6 (Adam only)
- Shampoo, KrAdagrad, KrAdagrad*
 - learning rates [1, 2.5e-1, 1e-1, 2.5e-2, 1e-2]
 - epsilon 1e-4

Table 7: Selected learning rates for continual learning comparisons.

GEM on Permuted MNIST				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
5e-3	1e-3	1e-2	1e-2	1e-2
LaMAML on Split CIFAR100				
SGD	Adam	Shampoo	KrAdagrad*	KrAdagrad
2e-2	1e-4	1e-2	1e-2	1e-2

Additionally, for both experiments we use a batch size of 10, a block size of 128 for the 2nd order optimizers, and beta=0.9 for SGD with momentum.

²<https://github.com/ContinualAI/continual-learning-baselines>

³the three smallest rates were included only for tuning Adam for LaMAML





Figure 9: Left: Top 1 Accuracy of GEM on Permutated MNIST; Right: Top 1 Accuracy of LaMAML on Split CIFAR-100. We do not observe pathologically bad behavior. The learning curves are extremely similar; hence the relative curves presented in the main body.