
Back to Square Roots: An Optimal Bound on the Matrix Factorization Error for Multi-Epoch Differentially Private SGD

Nikita P. Kalinin

Institute of Science and Technology (ISTA)
Klosterneuburg, Austria
nikita.kalinin@ist.ac.at

Ryan McKenna

Google Research
mckennar@google.com

Jalaj Upadhyay

Rutgers University
New Jersey, USA
jalaj.upadhyay@rutgers.edu

Christoph H. Lampert

Institute of Science and Technology (ISTA)
Klosterneuburg, Austria
chl@ist.ac.at

Abstract

Matrix factorization mechanisms for differentially private training have emerged as a promising approach to improve model utility under privacy constraints. In practical settings, models are typically trained over multiple epochs, requiring matrix factorizations that account for repeated participation. Existing theoretical upper and lower bounds on multi-epoch factorization error leave a significant gap. In this work, we introduce a new explicit factorization method, Banded Inverse Square Root (BISR), which imposes a banded structure on the inverse correlation matrix. This factorization enables us to derive an explicit and tight characterization of the multi-epoch error. We further prove that BISR achieves asymptotically optimal error by matching the upper and lower bounds. Empirically, BISR performs on par with state-of-the-art factorization methods, while being simpler to implement, computationally efficient, and easier to analyze.

1 Introduction

Private machine learning has become increasingly important as the use of sensitive data in model training continues to grow. Ensuring privacy while maintaining model accuracy presents a critical challenge, particularly in fields like healthcare, finance, and personal data analysis. *Differential Privacy* (DP) has emerged as a fundamental framework for formalizing privacy guarantees in machine learning. It provides a mathematically rigorous way to limit the influence of any individual data point on the model’s output, thereby preserving privacy. One effective approach to achieving DP in iterative training is through the use of structured noise mechanisms that balance privacy guarantees with model utility.

In this work, we focus on the *Matrix Factorization Mechanism* for ensuring DP, a method extensively studied in recent years in the context of private learning (Kairouz et al., 2021; Denisov et al., 2022; Fichtenberger et al., 2023; Henzinger et al., 2023; Andersson & Pagh, 2023; Henzinger et al., 2024; Kalinin & Lampert, 2024; Andersson & Pagh, 2025; Henzinger & Upadhyay, 2025; Henzinger et al., 2025). The idea behind this approach is to add a correlated noise to the gradients to preserve accuracy of the model training, with correlation matrix $C \in \mathbb{R}^{n \times n}$.

The concept of *multi-epoch participation* with matrix factorization was first introduced by Choquette-Choo et al. (2023), where the problem was formulated as an optimization problem over banded

matrices. However, a key limitation of existing methods is the lack of precise theoretical guarantees on the *factorization error* in multi-epoch participation. While Kalinin & Lampert (2024) established a general lower bound and provided an upper bound for *Square Root Factorization*, the error bounds for *Banded Square Root Factorization*, where the correlation matrix C is made p -banded, remained imprecise.

In this work, we propose a novel approach to matrix factorization: rather than imposing a banded structure on the correlation matrix C , we introduce a *banded inverse square root*, enforcing the banded structure on C^{-1} . This shift¹ offers several key advantages. First, it allows for precise control over the resulting factorized matrices, enabling us to derive **explicit upper bounds** on the factorization error with clear dependence on the bandwidth. Second, the method is **computationally efficient**, as it requires one just to convolve the previous noise with a quickly computable fixed sequence of coefficients, which can be done for instance via Fast Fourier Transform (FFT), making it suitable for large-scale machine learning tasks. Most importantly, we prove that our method achieves **asymptotically optimal factorization error**: we establish a **new lower bound** that matches our upper bound, closing a significant theoretical gap in the literature.

By refining the theoretical understanding of banded factorization in multi-epoch settings, our work provides both theoretical insights and practical benefits for privacy-preserving ML training. Our main contributions are:

1. We introduce a new factorization method, the **Banded Inverse Square Root (BISR)**, which is scalable, efficient, and agnostic to the underlying training objective.
2. We prove that BISR is asymptotically **optimal**, by deriving tight upper and lower bounds on the multi-epoch factorization error, with explicit dependence on bandwidth and workload properties.
3. We conduct a thorough empirical evaluation, comparing BISR to existing techniques in multi participation training—including Banded Square Root (BSR), Buffered Linear Toeplitz (BLT), and Banded Matrix Factorization (Band-MF), showing that BISR achieves a higher or comparable accuracy for the large matrix sizes.
4. In the low-memory regime, we propose an optimization method, **BandInvMF**, which directly optimizes the coefficients of the matrix C^{-1} . This approach achieves error rates comparable to state-of-the-art factorization methods, while being easy and efficient to implement.

2 Background

Matrix Factorization (MF). MF mechanisms provide a promising approach to the private matrix multiplication problem, which has applications in continual counting and Stochastic Gradient Descent for machine learning. Specifically, we aim to estimate the product of a public matrix of coefficients $A \in \mathbb{R}^{n \times n}$ and a private matrix $X \in \mathbb{R}^{n \times d}$. Instead of doing so directly, we adopt a factorization $A = BC$, allowing us to estimate AX privately as $\widehat{AX} = B(CX + Z)$ or, equivalently, $\widehat{AX} = A(X + C^{-1}Z)$. Here, $Z \sim \mathcal{N}(0, sI)$ is appropriately scaled Gaussian noise, which ensures that $CX + Z$ is private; the multiplication by B preserves the privacy guarantees due to DP’s post-processing property.

The choice of factorization $A = BC$ can significantly impact the quality of the private estimation. We quantify the *approximation quality* by the expected Frobenius error of the estimated product,

$$\mathcal{E}(B, C)^2 = \frac{1}{n} \mathbb{E}_Z \|AX - \widehat{AX}\|_F^2, \quad (1)$$

where $\|\cdot\|_F$ is the Frobenius norm. An elementary analysis (Li et al., 2015) shows that

$$\mathcal{E}(B, C)^2 = \frac{s^2}{n} \|B\|_F^2, \quad (2)$$

and that the required noise strength, s , scales proportionally to the *sensitivity* of the matrix C . Let $X \sim X'$ indicates that the update vector sequences differ only in entries corresponding to a single data item. Then sensitivity of the matrix C is defined as

¹The inverse correlation matrix has been receiving more attention recently. In the concurrent work McMahan & Pillutla (2025), the authors consider the inverse correlation matrix of BLT.

$$\text{sens}(C) := \sup_{X \sim X'} \|CX - CX'\|_F \quad (3)$$

Private SGD. In this work, we consider the task of model training with SGD with (optional) weight decay and momentum. The corresponding update equations are $\theta_{i+1} = \alpha\theta_i - m_{i+1}$ and $m_{i+1} = \beta m_i + x_i$, where $\theta_1, \dots, \theta_n \in \mathbb{R}^D$ are the model parameters after each update step, x_1, \dots, x_n are the gradient vectors computed in each update step, $0 < \alpha \leq 1$ is the weight decay factor, and $0 \leq \beta < 1$ is the momentum strength².

Following Kalinin & Lampert (2024), we rewrite the dynamics in the matrix form as $\Theta = A_{\alpha,\beta}X$, with $\Theta = (\theta_1, \dots, \theta_n)^\top \in \mathbb{R}^{n \times D}$, $X = (x_1, \dots, x_n)^\top \in \mathbb{R}^{n \times D}$, and $A_{\alpha,\beta}$ is the *SGD workload matrix* defined as follows:

$$A_{\alpha,\beta} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \alpha + \beta & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \sum_{k=0}^{n-1} \alpha^k \beta^{n-1-k} & \sum_{k=0}^{n-2} \alpha^k \beta^{n-2-k} & \dots & 1 \end{pmatrix} \in \mathbb{R}^{n \times n}. \quad (4)$$

Note that, in contrast to the naive MF setting, in the SGD case any input data (gradient) x_i depends on the previously computed model parameters, θ_{i-1} , that is, we aim for *adaptive privacy*. However, Denisov et al. (2022) shows that for Gaussian noise, adaptive privacy follows from the non-adaptive one, i.e., it suffices for us to solve the case in which the X matrix is an arbitrary fixed data matrix. Consequently, we estimate $A_{\alpha,\beta}X$ privately using the form $\widehat{A_{\alpha,\beta}X} = A_{\alpha,\beta}(X + C^{-1}Z)$. This form corresponds to running SGD, but each individual gradient update is perturbed by a correlated noise vector. That has the advantage that we do not need to store any previous gradients, and we can rely on any existing implementation of the SGD procedure.

In multi-epoch SGD, each data sample might contribute to more than one gradient update vector. As a suitable notion of sensitivity, we adopt the setting of b -min-separated repeated participation (two participations of any data point occur at least b update steps apart). The resulting sensitivity can be bounded as Choquette-Choo et al. (2023):

$$\text{sens}_{k,b}(C) \leq \max_{\pi \in \Pi_{k,b}} \sqrt{\sum_{i,j \in \pi} |(C^\top C)_{[i,j]}|} \quad (5)$$

where $\Pi_{k,b} = \{\pi \subset \{1, \dots, n\} : |\pi| \leq k \wedge (\{i, j\} \subset \pi \Rightarrow i = j \vee |i - j| \geq b)\}$ represents the set of possible b -min-separated index sets with at most k participation. This bound becomes an equality if all entries of $C^\top C$ are non-negative.

Optimal factorization. Better choices of factorization matrices can achieve the same privacy levels with less added noise, potentially leading to higher utility. Therefore, various factorizations have been proposed and studied theoretically as well as empirically.

Choquette-Choo et al. (2023) defines the *optimal factorization* as the one that minimizes the expected approximation error (1), and proposed an optimization problem to (approximately) compute this factorization. A downside of this approach is that the optimization problem is computationally expensive and the numeric solution does not provide theoretical insights, such as the optimal (i.e. lowest) rate of growth of the approximation error.

On the other hand, a square root factorization introduced by Henzinger et al. (2024), is an explicit factorization, defined by $A_{\alpha,\beta} = C_{\alpha,\beta}^2$. Kalinin & Lampert (2024) showed that the factorization error of the square root factorization under multi-epoch participation is worse than that of the optimal factorization and they introduced *banded square root* (BSR) factorization, which is defined by making the matrix $C_{\alpha,\beta}$ banded. A limitation of BSR is that its guarantees are implicit in terms of the used bandwidth, which does not allow concluding how they relate to the optimal multi-epoch factorization at a theoretical level.

²For simplicity of exposition, we use an implicit learning rate of 1. Because of the linearity of the operations, the general case can be recovered by pre-scaling x_1, \dots, x_n accordingly.

3 Banded Inverse Square Root Factorization

In this section, we present our main theoretical results: we prove a new lower bound on the achievable approximation error (Theorem 1), we introduce the BISR factorization (Definition 1), and we prove that BISR achieves this (therefore optimal) rate (Theorem 2).

We first show an improved version of the lower bounds of the approximation error for general factorizations from Kalinin & Lampert (2024) in Section 4.

Theorem 1 (General Multi-Participation Lower Bound). *Let $A_{\alpha,\beta} \in \mathbb{R}^{n \times n}$ be the SGD workload matrix (4). In the multi-participation setting with separation $1 \leq b \leq n$ and $k = \lceil \frac{n}{b} \rceil$, for any factorization $A_{\alpha,\beta} = BC$, it holds*

$$\mathcal{E}(B, C) = \begin{cases} \Omega(\sqrt{k} \log n + k) & \text{if } \alpha = 1, \\ \Omega_{\alpha}(\sqrt{k}) & \text{if } \alpha < 1. \end{cases} \quad (6)$$

As our second main contribution, we now introduce the BISR factorization for multi-epoch SGD.

Definition 1 (Banded Inverse Square Root (BISR)). *For a given workload matrix A , let C be the matrix square root (i.e. $C^2 = A$) with positive values on the diagonal. Let C^p be the matrix obtained by: i) computing the inverse matrix C^{-1} , ii) imposing a banded structure with p bands by setting all elements below the p -th diagonal to zero, iii) inverting the resulting banded matrix back. Then, we denote by BISR the matrix factorization $A = B^p C^p$, with $B^p = A(C^p)^{-1}$.*

BISR can be seen as an alternative realization of the insights behind the BSR (Banded Square Root) factorization from Kalinin & Lampert (2024). There, the intuition was that making the matrix C p -banded reduces its sensitivity without increasing the Frobenius norm of the subsequent postprocessing matrix too much, thereby resulting in an overall reduction of the approximation error. The authors did not derive exact rates, though, because the dependence on p is not explicit.

For BISR, we instead make the matrix C^{-1} p -banded. This also leads to a reduction of the approximation error compared to the non-banded case, but with two additional advantages. First, the resulting algorithm is time- and memory-efficient because the product of $(C^p)^{-1}Z$ can be represented as a convolution with p elements. Therefore, the computation can be performed efficiently: in a streaming setting, only p rows of the matrix Z need to be stored at any time, while in an offline setting (which requires more storage), it can be accelerated further using the Fast Fourier Transform. Second, and mainly, it allows us to derive more explicit expressions of the approximation error with respect to the bandwidth p , as illustrated by the following theorem proved in Section 4.

Theorem 2 (BISR Approximation Error). *For $1 \leq p \leq n$ and $1 \leq k \leq \frac{n}{b}$ the following upper bound holds for the matrix factorization error of the BISR $A_{\alpha,\beta} = B_{\alpha,\beta}^p C_{\alpha,\beta}^p$ (as in Definition 1):*

$$\mathcal{E}(B_{\alpha,\beta}^p, C_{\alpha,\beta}^p) = \begin{cases} O_{\beta} \left(\sqrt{k} \log p + \sqrt{\frac{nk}{b}} + \sqrt{\frac{nk \log p}{p}} + \sqrt{\frac{kp \log p}{b}} \right) & \text{for } \alpha = 1, \\ O_{\alpha,\beta}(\sqrt{k}) & \text{for } \alpha < 1. \end{cases} \quad (7)$$

For comparison, Kalinin & Lampert (2024) proved a bound $O\left(\sqrt{\frac{nk \log p}{p}}\right) + O_p(\sqrt{k})$ on the approximation error of the BSR in the case $\alpha = 1, \beta = 0$. While the first term also appears in (7), the second is non-informative about the effect of the bandwidth, p , and therefore does not allow making a statement about the optimality of the BSR.

In contrast, Theorem 2 is explicit about the role of p . Choosing its value such that the occurring terms in (7) are minimized, we obtain the following corollary.

Corollary 1 (Optimized BISR Approximation Error). *Let $A_{\alpha,\beta} = B_{\alpha,\beta}^p C_{\alpha,\beta}^p$ be the BISR factorization defined of Definition 1. For $1 \leq b \leq n$ let $k = \lceil \frac{n}{b} \rceil$. Then, for $p^* \sim b \log b$, the matrix factorization error admits the following optimized upper bound:*

$$\mathcal{E}(B_{\alpha,\beta}^{p^*}, C_{\alpha,\beta}^{p^*}) = \begin{cases} O_{\beta} \left(\sqrt{k} \log n + k \right), & \text{for } \alpha = 1, \\ O_{\alpha,\beta}(\sqrt{k}), & \text{for } \alpha < 1. \end{cases} \quad (8)$$

Comparing (8) with (6), we obtain as direct corollaries.

Corollary 2 (Tightness of Theorem 1). *The lower bounds of Theorem 1 are tight, in the sense that there exists a factorization method that achieves them.*

Corollary 3 (Rate-optimality of BISR). *No matrix factorization method can achieve better rates than BISR.*

Following the work of Andersson & Yehudayoff (2025), we show that the space complexity of the matrix $(C_{\alpha,\beta}^p)^{-1}$ is equal to p , meaning that exact multiplication with a random real vector z in a streaming setting, performing continuous operations, requires storing p real values (not including the memory needed to store the matrix coefficients). This implies that, for memory-efficient computation, one must either use a small bandwidth p or consider approximate multiplication. We formally state the result in the following lemma:

Lemma 1. *The space complexity—defined as the minimum buffer size required by a **streaming** algorithm to correctly process an input—for computing the product of the Toeplitz matrix $(C_{\alpha,\beta}^p)^{-1}$ with an arbitrary vector $z \in \mathbb{R}^n$, for $n \geq 2p - 1$, is exactly p , and at least $\frac{n-5}{2}$ for $C_{\alpha,\beta}^{-1}$.*

4 Proofs

Proof of Lemma 1. Throughout this proof, we use only results from Andersson & Yehudayoff (2025), and thus any reference to a theorem or lemma should be understood as coming from that work. We use their Lemma 7, which lower-bounds the space complexity of a Toeplitz matrix using the rank of a corresponding Hankel matrix of its coefficients. The matrix $C_{\alpha,\beta}^{-1}$ has a generating function $f = \sqrt{(1 - \alpha x)(1 - \beta x)}$. The proof of their Corollary 16 implies that the Hankel matrix $H[f]$ has corank at most 3. Thus, Lemma 7 implies that $C_{\alpha,\beta}^{-1}$ has space complexity at least $\frac{n+1}{2} - 3$. For the matrix $(C_{\alpha,\beta}^p)^{-1}$, the generating function is a rational function of degree p ; therefore, for $n \geq 2p - 1$, their Theorem 2 implies that the space complexity is exactly p , concluding the proof. \square

Proof of Theorem 1. We use the probabilistic method in Lemma 9 to obtain the bounds $\Omega_\alpha(\sqrt{k})$ for $\alpha < 1$ and $\mathcal{E}(B, C) = \Omega(\sqrt{k} \log n)$ for $\alpha = 1$. It remains to prove that for $\alpha = 1$, we also have the lower bound $\mathcal{E}(B, C) = \Omega(k)$.

We begin with the following observation: given a matrix C , we can compute an optimal participation scheme represented by a vector with ones at positions corresponding to participating columns, denoted by π_C^* . As a lower bound, we consider a specific participation vector π_1 , with ones in columns indexed by $1 + ib$ for $i \in [0, k - 1]$, such that $|\pi_1| = k$. By construction, we have $\text{sens}_{k,b}(C) := \|C\pi_C^*\|_2 \geq \|C\pi_1\|_2$. Therefore, the error can be bounded as follows:

$$\mathcal{E}(B, C) = \frac{1}{\sqrt{n}} \|B\|_F \text{sens}_{k,b}(C) \geq \frac{1}{\sqrt{n}} \|B\|_F \|C\pi_1\|_2 \geq \frac{1}{\sqrt{n}} \|BC\pi_1\|_2 = \frac{1}{\sqrt{n}} \|A_{1,\beta}\pi_1\|_2. \quad (9)$$

As a lower bound, we consider $\beta = 0$ as $\|A_{1,\beta}\pi_1\|_2 \geq \|A_{1,0}\pi_1\|_2$. The elements of the matrix $A_{1,0}$ are positive and non-increasing. Therefore, by Theorem 3, the (k, b) -sensitivity of $A_{1,0}$ is exactly $\|A_{1,0}\pi_1\|_2$. By Theorem 9 from Kalinin & Lampert (2024), this sensitivity is at least $\frac{k\sqrt{n}}{\sqrt{3}}$, resulting in the lower bound:

$$\mathcal{E}(B, C) \geq \frac{k}{\sqrt{3}} = \Omega(k), \quad (10)$$

which concludes the proof. \square

4.1 Proof of Theorem 2

Before proving the main theorem, we state several properties of BISR decomposition, except for the ones taken from prior work, we provide their proofs in the appendix.

We first re-state the analytic form of the coefficients of $C = A^{1/2}$.

Lemma 2 (Theorem 1 in Kalinin & Lampert (2024) – Square Root of the Matrix $A_{\alpha,\beta}$). *For $k \geq 0$, let $r_k = \left\lfloor \binom{-1/2}{k} \right\rfloor$. For $0 \leq \beta < \alpha \leq 1$, the square root matrix $C_{\alpha,\beta} = A_{\alpha,\beta}^{1/2}$ has the following*

explicit form:

$$C_{\alpha,\beta} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ c_1^{\alpha,\beta} & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ c_{n-1}^{\alpha,\beta} & c_{n-2}^{\alpha,\beta} & \cdots & 1 \end{pmatrix} \quad \text{with coefficients} \quad c_k^{\alpha,\beta} = \sum_{j=0}^k \alpha^j \beta^{k-j} r_j r_{k-j}, \quad (11)$$

The follow lemma provides analytic expressions for the elements of the inverse matrix $C_{\alpha,\beta}^{-1}$.

Lemma 3 (Inverse Square Root of the Matrix $A_{\alpha,\beta}$). *For $k \geq 0$, let $\tilde{r}_k = (-1)^k \binom{1/2}{k} = \frac{-r_k}{2k-1} = \frac{-1}{2k-1} \left| \binom{-1/2}{k} \right|$. The inverse of the matrix $C_{\alpha,\beta} = A_{\alpha,\beta}^{1/2}$ defined in (2), for $0 \leq \beta < \alpha \leq 1$, is*

$$C_{\alpha,\beta}^{-1} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ \tilde{c}_1^{\alpha,\beta} & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{c}_{n-1}^{\alpha,\beta} & \tilde{c}_{n-2}^{\alpha,\beta} & \cdots & 1 \end{pmatrix}, \quad \text{where} \quad \tilde{c}_k^{\alpha,\beta} = \sum_{j=0}^k \tilde{r}_j \beta^j \tilde{r}_{k-j} \alpha^{k-j}. \quad (12)$$

We can now compute the values of the matrices $B_{\alpha,\beta}^p$ and $C_{\alpha,\beta}^p$ using the following lemmas.

Lemma 4 (Bounds on diagonals of $B_{\alpha,\beta}^p$). *The matrix $B_{\alpha,\beta}^p$ in the BISR factorization is a lower triangular Toeplitz matrix. The values on its diagonals are*

$$(1, c_1^{\alpha,\beta}, c_2^{\alpha,\beta}, \dots, c_{p-1}^{\alpha,\beta}, b_p^{\alpha,\beta}, \dots, b_{n-1}^{\alpha,\beta}) \quad \text{where} \quad 0 \leq b_i^{\alpha,\beta} \leq \alpha^i c_{p-1}^{1,\beta/\alpha} \quad \text{for} \quad i \geq p \quad (13)$$

where $c_i^{1,\beta/\alpha}$ for $1 \leq i \leq p-1$ is as defined in equation (11).

Lemma 5 (Bounds on diagonals of $C_{\alpha,\beta}^p$). *The matrix $C_{\alpha,\beta}^p$ in the BISR factorization is a lower triangular Toeplitz matrix. The values on its diagonals are $(1, \tilde{c}_1^{\alpha,\beta}, \tilde{c}_2^{\alpha,\beta}, \dots, \tilde{c}_{p-1}^{\alpha,\beta}, g_p^{\alpha,\beta}, \dots, g_{n-1}^{\alpha,\beta})$, where $\tilde{c}_i^{1,\beta/\alpha}$ (for $1 \leq i \leq p-1$) is as defined in equation (11) and*

$$0 \leq g_i^{\alpha,\beta} \leq \alpha^i \min(c_i^{1,\beta/\alpha}, c_p^{1,\beta/\alpha} \gamma_{\beta/\alpha}^{i-p}) \quad \text{for} \quad \gamma_{\beta/\alpha} = \left(1 + \frac{(1-\beta/\alpha)^2}{4p(1+\beta/\alpha)}\right)^{-1} \quad \text{and} \quad i \geq p. \quad (14)$$

We then use the multi-participation sensitivity stated in the form of the following theorem.

Theorem 3 (Theorem 2 from Kalinin & Lampert (2024)). *Let M be a lower triangular Toeplitz matrix with decreasing non-negative entries $m_0 \geq m_1 \geq m_2 \geq \dots m_{n-1} \geq 0$ on the diagonals. Then the sensitivity of M in the setting of b -min-separation is*

$$\text{sens}_{k,b}(M) = \left\| \sum_{j=0}^{k-1} M_{[\cdot, 1+jb]} \right\|_2 \quad (15)$$

where $M_{[\cdot, 1+jb]}$ denotes the $(1+jb)$ -th column of M .

To apply Theorem 3 we need to prove that the values of the matrix $C_{\alpha,\beta}^p$ are decreasing, which we formulate by the following lemma

Lemma 6 (Decreasing values). *The values $(1, \tilde{c}_1^{\alpha,\beta}, \dots, \tilde{c}_{p-1}^{\alpha,\beta}, g_p^{\alpha,\beta}, \dots, g_{n-1}^{\alpha,\beta})$ of matrix $C_{\alpha,\beta}^p$ as defined in Lemma 5 are decreasing.*

To prove Theorem 2, we use Lemma 4 to bound the Frobenius norm $\|B_{\alpha,\beta}^p\|_F$. Then, Theorem 3, together with Lemma 6, provides an explicit way to express the sensitivity $\text{sens}_{k,b}(C_{\alpha,\beta})$. To bound the sensitivity, we apply Lemma 5 to bound individual values. The product of the bounds on $\|B_{\alpha,\beta}^p\|_F$ and $\text{sens}_{k,b}(C_{\alpha,\beta})$ yields the result. The full proof of Theorem 2 can be found in the appendix.

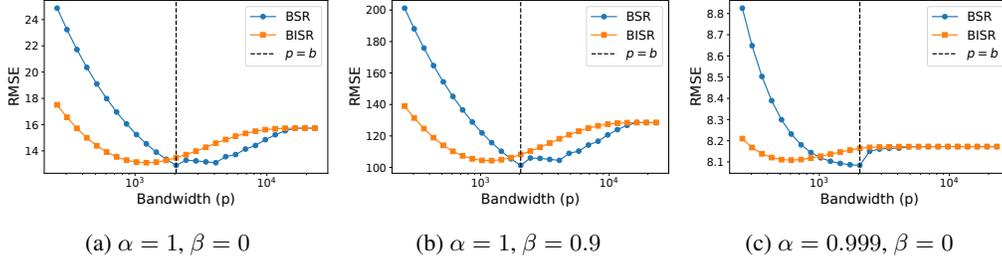


Figure 1: RMSE comparison for Banded Square Root (BSR) and Banded Inverse Square Root (BISR) methods across varying bandwidths (p). The results are shown for a fixed matrix size of $n = 16384$ and a participation number of $k = 8$.

5 Experiments

In this section, we present numerical results from evaluating various factorization methods in the multi-participation regime.

We first study the effect of using different bandwidths for BSR and BISR, as shown in Figure 1. We found that, in most settings, the optimal bandwidth for BSR is equal to the separation parameter b , whereas for BISR, the optimal value is not only different from b but can also perform noticeably worse when b is used. Therefore, in future comparisons with BISR, we propose optimizing over the bandwidth to find an optimal value p^* whenever computationally feasible.

We compare BISR with several other methods, including Banded Square Root (BSR), Banded Matrix Factorization (BandMF), introduced by McKenna (2025), and Buffered Linear Toeplitz (BLT), introduced by Dvijotham et al. (2024) and adapted for multi-participation training by McMahan et al. (2024). We use a buffer size of 4, as recommended, and observe that the error saturates quickly as the buffer size increases. We use BandMF with bandwidth equal to b , as we did not observe any benefit from using a larger bandwidth. Moreover, we conjecture that optimal multi-epoch participation can always be achieved on a banded lower triangular matrix with bandwidth b .

We emphasize that BLT has only been described, analyzed, and implemented for prefix-sum matrices. Therefore, we do not show BLT results for momentum and weight decay plots. For all methods except BISR, we use efficient implementations from the jax-privacy library (Balle et al., 2025).

Our experiments (Figure 2) show that banded inverse square root factorization consistently matches or exceeds BSR in quality across all regimes and outperforms it in scenarios with a large number of participations. The improvement is particularly pronounced when the participation count is high ($k = 16$). BISR achieves RMSE comparable to that of BLT, but has the advantage of easier implementation for both factorization and training, as it only requires convolving previous noise with a fixed set of coefficients—an "embarrassingly parallel" operation (see McKenna (2025)). While BandMF achieves slightly better RMSE at $k = 16$, it requires solving a computationally expensive optimization problem, making it impractical for matrix sizes beyond $n = 4096$.

6 From BISR to BandInvMF

In the previous sections, we established that BISR has asymptotically optimal rate for large bandwidths $p \sim b \log b$. However, in practice, one might want to work with a smaller value of p to save memory and computational resources. In this section, we showcase a modification to BISR with improved practical properties in this regime. We propose to keep the construction of a banded inverse matrix with Toeplitz structure, but to set its values not by the closed form expressions (12) but by a numeric optimization. Specifically, we optimize an upper bound on b -min separation participation, given by Equation 5.

For the sake of numerical optimization, we assume that the optimum is achieved for indices of the form $i + kb$. This assumption can be justified, as we observed that the resulting solution for matrix C is positive and decreasing, which guarantees optimality. We use banded inverse square root

factorization as an initialization for the coefficients. We provide an efficient JAX implementation in the Appendix (see Algorithm 1) as well as the convergence plots in Figure 5.

The numerical results are presented in Figure 3 referred to as BandInvMF. We observe that the error decreases drastically even with the addition of a single band, compared to a trivial factorization. This observation is supported theoretically by the following lemma.

Lemma 7. *Let the matrix $C_\lambda^{-1} = \text{LTT}(1, -\lambda, 0, \dots, 0)$ be a lower triangular Toeplitz matrix with 1 on the main diagonal and $-\lambda$ on the subdiagonal. Then, for a single participation and the prefix sum matrix $A_{1,0}$, we can prove the following bound on the matrix factorization error:*

$$\inf_{\lambda \in (0,1)} \mathcal{E}(A_{1,0}C_\lambda^{-1}, C_\lambda) = O(n^{1/4}) \quad (16)$$

Proof. If the matrix C_λ^{-1} is given by $\text{LTT}(1, -\lambda, 0, \dots, 0)$, then its inverse is $C_\lambda = \text{LTT}(1, \lambda, \lambda^2, \dots, \lambda^{n-1})$. The product $A_{1,0}C_\lambda^{-1} = \text{LTT}(1, 1 - \lambda, \dots, 1 - \lambda)$, which leads to the following error:

$$\mathcal{E}(A_{1,0}C_\lambda^{-1}, C_\lambda)^2 = \frac{1}{n} (1 + (1 - \lambda)^2(n - 1)) \sum_{k=0}^{n-1} \lambda^{2k} = \frac{(1 + (1 - \lambda)^2(n - 1))(1 - \lambda^{2n})}{n(1 - \lambda^2)}. \quad (17)$$

Therefore,

$$\inf_{\lambda \in (0,1)} \mathcal{E}(A_{1,0}C_\lambda^{-1}, C_\lambda)^2 \leq \frac{(2 - \frac{1}{n}) \left(1 - \left(1 - \frac{1}{\sqrt{n}}\right)^n\right)}{\sqrt{n}} \leq \frac{2}{\sqrt{n}}, \quad (18)$$

when $\lambda = \sqrt{1 - \frac{1}{\sqrt{n}}}$ as $1 - \lambda \leq 1 - \left(1 - \frac{1}{\sqrt{n}}\right) = \frac{1}{\sqrt{n}}$. The bound follows. \square

This lemma shows that the optimized inverse banded matrix factorization can achieve an asymptotically better bound than a trivial factorization $A \times I$, which yields an error of $O(\sqrt{n})$. Moreover, from Theorem 2, for small p , the leading term for banded inverse square root factorization remains of order $O(\sqrt{n})$. Therefore, we advocate for optimizing the coefficients when the bandwidth is small. However, we do not see any benefit from using this optimization for a full matrix, as it is more straightforward and computationally efficient to optimize over the coefficients of the matrix C directly.

We compare Band-Inv-MF with other methods for training the model on CIFAR-10 (see Figure 4), both with and without amplification by subsampling. For a fairer comparison, we use a recently

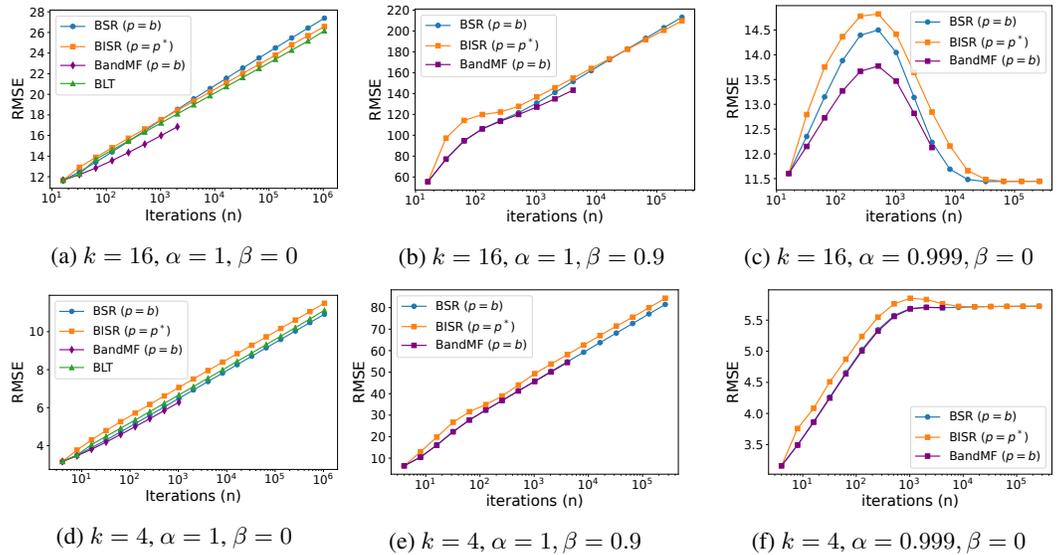


Figure 2: RMSE across varying matrix sizes for different factorization methods under multiple optimizer settings and participation levels.

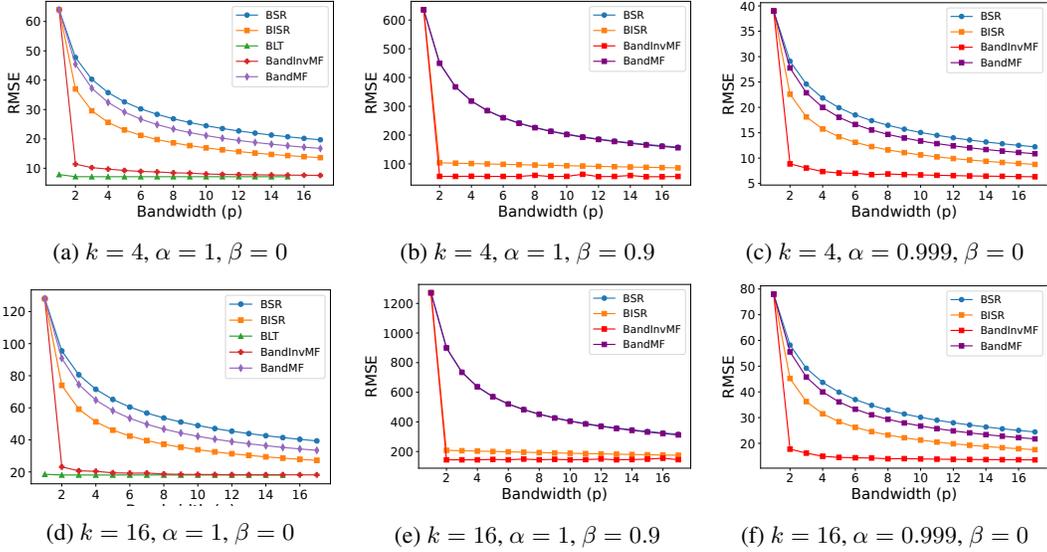


Figure 3: RMSE across different factorizations and optimization parameters α, β , with small bandwidth.

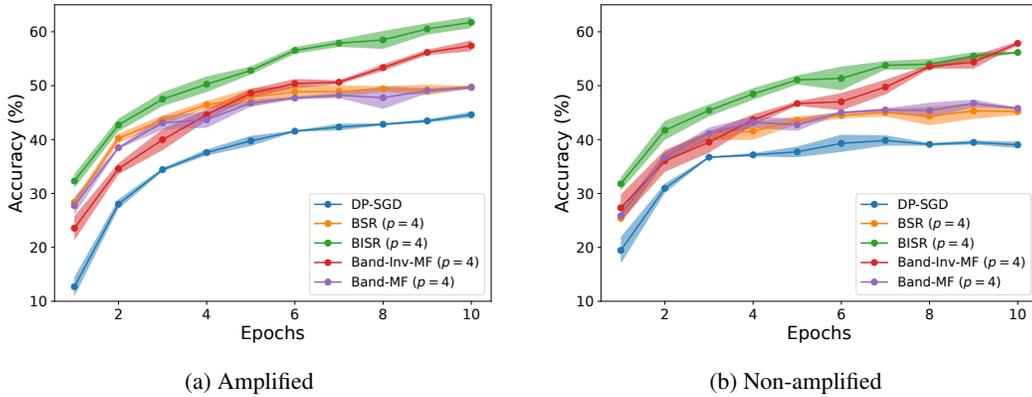


Figure 4: **CIFAR-10 accuracy for small bandwidth (low-memory regime)**. Both the amplified (left) and non-amplified (right) results show that inverse factorization methods, BISR and Band-Inv-MF, achieve significantly higher accuracy compared to Band-MF. Both plots correspond to $(9, 10^{-5})$ -DP, with training performed using momentum $\beta = 0.9$ and weight decay $\alpha = 0.9999$, which we found to be optimal. The tables with hyperparameters (Table 1) and accuracies (Table 2) can be found in the appendix.

proposed bins-and-balls subsampling mechanism (Chua et al., 2025), which combines the accuracy benefits of Poisson subsampling with improved implementation efficiency. More importantly, it supports the matrix mechanism via the MCMC accountant (Choquette-Choo et al., 2024a,b), even when the matrix C is not banded. Our results indicate that in a low-memory regime, inverse correlation matrix methods—BISR and Band-Inv-MF—achieve significantly higher accuracy than BSR and Band-MF, and consistently outperform DP-SGD, with and without amplification.

7 Discussion

This work demonstrates that imposing a banded structure on the inverse correlation matrix, rather than on the matrix itself, leads to both theoretical and practical benefits for differentially private training across multiple participations. Our Banded Inverse Square Root (BISR) method enables explicit factorization, supporting clean error analysis and efficient implementation.

We prove that BISR achieves asymptotically optimal factorization error by improving upon previously established lower bounds and showing that BISR matches the asymptotics precisely, thereby closing a fundamental gap in the theory.

In the low-memory regime, we find it beneficial to optimize directly over the coefficients of the inverse correlation matrix. Our Band-Inv-MF method achieves a lower matrix factorization error compared to BISR. However, these improvements do not yet translate to gains in model accuracy when training with the amplification by subsampling. Future research should focus on optimizing the matrix coefficients while explicitly accounting for amplification, in order to bridge this gap.

Acknowledgments

We thank Arun Ganesh for providing the code for the MCMC accountant. We thank Joel Andersson and Monika Henzinger for valuable comments on the early version of the draft. We thank Christian Lebeda for a fruitful discussion on the lower bound theorem. Jalaj Upadhyay’s research was funded by the Rutgers Decanal Grant no. 302918 and an unrestricted gift from Google. This work is supported in part by the Austrian Science Fund (FWF) [10.55776/COE12] and the Scientific Service Units (SSU) of ISTA through resources provided by Scientific Computing (SciComp).

References

- Andersson, J. D. and Pagh, R. A smooth binary mechanism for efficient private continual observation. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Andersson, J. D. and Pagh, R. Streaming private continual counting via binning. In *IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 2025.
- Andersson, J. D. and Yehudayoff, A. On the space complexity of online convolution, 2025. arXiv:2505.00181.
- Balle, B., Berrada, L., Charles, Z., Choquette-Choo, C. A., De, S., Doroshenko, V., Dvijotham, D., Galen, A., Ganesh, A., Ghalebikesabi, S., Hayes, J., Kairouz, P., McKenna, R., McMahan, B., Pappu, A., Ponomareva, N., Privilov, M., Rush, K., Smith, S. L., and Stanforth, R. JAX-Privacy: Algorithms for privacy-preserving machine learning in JAX, 2025. URL http://github.com/google-deepmind/jax_privacy.
- Böttcher, A. and Grudsky, S. M. *Toeplitz Matrices, Asymptotic Linear Algebra, and Functional Analysis*. Springer, 2000.
- Choquette-Choo, C. A., Ganesh, A., McKenna, R., McMahan, H. B., Rush, J. K., Thakurta, A. G., and Zheng, X. (Amplified) banded matrix factorization: A unified approach to private training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.
- Choquette-Choo, C. A., Ganesh, A., Haque, S., Steinke, T., and Thakurta, A. Near exact privacy amplification for matrix mechanisms. In *International Conference on Learning Representations (ICLR)*, 2024a.
- Choquette-Choo, C. A., Ganesh, A., Steinke, T., and Thakurta, A. Privacy amplification for matrix mechanisms. In *International Conference on Learning Representations (ICLR)*, 2024b.
- Chua, L., Ghazi, B., Harrison, C., Leeman, E., Kamath, P., Kumar, R., Manurangsi, P., Sinha, A., and Zhang, C. Balls-and-Bins sampling for DP-SGD. In *Conference on Uncertainty in Artificial Intelligence (AISTATS)*, 2025.
- Denisov, S., McMahan, H. B., Rush, J., Smith, A., and Thakurta, G. A. Improved Differential Privacy for SGD via optimal private linear operators on adaptive streams. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2022.
- Dvijotham, K., McMahan, H. B., Pillutla, K., Steinke, T., and Thakurta, A. Efficient and near-optimal noise generation for streaming differential privacy. In *Symposium on Foundations of Computer Science (FOCS)*, 2024.

- Fichtenberger, H., Henzinger, M., and Upadhyay, J. Constant matters: Fine-grained complexity of differentially private continual observation using completely bounded norms. In *International Conference on Machine Learning (ICML)*, 2023.
- Henzinger, M. and Upadhyay, J. Improved differentially private continual observation using group algebra. In *Symposium on Discrete Algorithms (SODA)*, 2025.
- Henzinger, M., Upadhyay, J., and Upadhyay, S. Almost tight error bounds on differentially private continual counting. In *Symposium on Discrete Algorithms (SODA)*, 2023.
- Henzinger, M., Upadhyay, J., and Upadhyay, S. A unifying framework for differentially private sums under continual observation. In *Symposium on Discrete Algorithms (SODA)*, 2024.
- Henzinger, M., Kalinin, N. P., and Upadhyay, J. Binned group algebra factorization for differentially private continual counting, 2025. arXiv:2504.04398.
- Kairouz, P., McMahan, B., Song, S., Thakkar, O., Thakurta, A., and Xu, Z. Practical and private (deep) learning without sampling or shuffling. In *International Conference on Machine Learning (ICML)*, 2021.
- Kalinin, N. and Lampert, C. H. Banded square root matrix factorization for differentially private model training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2024.
- Li, C., Miklau, G., Hay, M., McGregor, A., and Rastogi, V. The matrix mechanism: Optimizing linear counting queries under Differential Privacy. *International Conference on Very Large Data Bases (VLDB)*, 2015.
- McKenna, R. Scaling up the banded matrix factorization mechanism for differentially private ML. In *International Conference on Learning Representations (ICLR)*, 2025.
- McMahan, H. B. and Pillutla, K. An inversion theorem for Buffered Linear Toeplitz (BLT) matrices and applications to streaming differential privacy, 2025. arXiv:2504.21413.
- McMahan, H. B., Xu, Z., and Zhang, Y. A hassle-free algorithm for private learning in practice: Don't use tree aggregation, use BLTs. In *Conference on Empirical Methods on Natural Language Processing (EMNLP)*, 2024.
- Strang, G. A proposal for Toeplitz matrix calculations. *Studies in Applied Mathematics*, 1986.

Appendix

Key Inequalities and Relationships

This section compiles the fundamental inequalities and key relationships employed throughout this study. Each entry is presented with a concise explanation of its origin or the context in which it arises.

1. $r_k = \left| \binom{-1/2}{k} \right| = \frac{1}{4^k} \binom{2k}{k}$ *Henzinger et al. (2024)*
2. $\frac{1}{\sqrt{\pi(k+1)}} \leq r_k \leq \frac{1}{\sqrt{\pi k}}$ *Lemma 2.1 from Dvijotham et al. (2024)*
3. $\sum_{k=0}^{p-1} r_k \leq 1 + \frac{1}{\sqrt{\pi}} \sum_{k=1}^{p-1} \frac{1}{\sqrt{k}} \leq 1 + \frac{2\sqrt{p}}{\sqrt{\pi}}$ *Integral inequality.*
4. $\sum_{k=0}^{p-1} r_k^2 \leq 1 + \log p$ *Lemma 8*
5. $c_k^{\alpha, \beta} = \sum_{j=0}^k \alpha^j \beta^{k-j} r_j r_{k-j}$ *Theorem 1 from Kalinin & Lampert (2024)*
6. $c_k^{1, \beta} \leq c_{k-1}^{1, \beta} \left[1 - \frac{(1-\beta)^2}{2k} \right]$ for $k \geq 1$. *Lemma 12*
7. $\sum_{j=0}^k c_j^{1, \beta} = \sum_{j=0}^k r_j \beta^j \tilde{r}_{k-j}$ *In the proof of Lemma 11*
8. $\sum_{j=0}^k c_j^{1, \beta} \beta^{k-j} = \sum_{j=0}^k r_j \beta^j \tilde{r}_{k-j}$ *In the proof of Lemma 11*
9. $\sum_{j=0}^k \frac{\tilde{c}_j^{1, \beta} (1 - \beta^{k-j+1})}{1 - \beta} = c_k^{1, \beta}$ *In the proof of Lemma 11*
10. $r_k (1 - \beta) \leq \sum_{j=0}^k \tilde{c}_j^{1, \beta} \leq c_k^{1, \beta} (1 - \beta)$ *Lemma 11*
11. $\frac{\log(k+1)}{4} \leq \sum_{j=0}^{k-1} (c_j^{1, \beta})^2 \leq \frac{1 + \log k}{(1 - \beta)^2}$ *Lemma 8*
12. $\frac{\alpha^k}{2\sqrt{k+1}} \leq c_k^{\alpha, \beta} \leq \frac{\alpha^k}{(1 - \beta/\alpha)\sqrt{k+1}}$ *Lemma 8*
13. $1 \leq \sum_{j=0}^{k-1} (c_j^{\alpha, \beta})^2 \leq \frac{1}{(\alpha - \beta)^2} \log \left(\frac{1}{1 - \alpha^2} \right)$ *Lemma 8*
14. $\tilde{r}_k = (-1)^k \binom{1/2}{k} = \frac{-1}{2k-1} r_k$ *Lemma 3*
15. $\tilde{c}_k^{\alpha, \beta} = \sum_{j=0}^k \tilde{r}_j \beta^j \tilde{r}_{k-j} \alpha^{k-j}$ *Lemma 3*
16. $\tilde{r}_k (1 + \beta) \leq \tilde{c}_k^{1, \beta} \leq 0$ *Lemma 10*

A Proofs

Lemma 8 (Lemma 7 from Kalinin & Lampert (2024)). *For $k \in \{1, \dots, n\}$ it holds for $c_i^{\alpha, \beta}$ as defined in equation (11):*

$$\frac{\log(k+1)}{4} \leq \sum_{i=0}^{k-1} (c_i^{1, \beta})^2 \leq \frac{1 + \log k}{(1 - \beta)^2} \quad (19)$$

for $\alpha = 1$, and otherwise

$$1 \leq \sum_{i=0}^{k-1} (c_i^{\alpha, \beta})^2 \leq \frac{1}{(\alpha - \beta)^2} \log \left(\frac{1}{1 - \alpha^2} \right). \quad (20)$$

Lemma 9. *Let $A_{\alpha, \beta} \in \mathbb{R}^{n \times n}$ be the SGD workload matrix (4). In the multi-participation setting with separation $1 \leq b \leq n$ and $k = \lceil \frac{n}{b} \rceil$, for any factorization $A_{\alpha, \beta} = BC$, it holds that*

$$\mathcal{E}(B, C) = \begin{cases} \Omega(\sqrt{k} \log n), & \alpha = 1 \\ \Omega(\sqrt{k}), & \alpha < 1 \end{cases} \quad (21)$$

Proof. Here, we refine Theorem 8 from Kalinin & Lampert (2024) by removing the assumption that the scalar products between the columns of the matrix C are non-negative, i.e., $C^\top C \geq 0$. We first prove that

$$\text{sens}_{k,b}(C)^2 \geq \frac{1}{4b} \|C\|_F^2. \quad (22)$$

To do so, we lower bound the b -min separation participation by the (k, b) -participation, where we have a fixed b separation between vectors but are allowed to include only a subset of them. This splits the set of all column indices into b disjoint subsets \mathcal{S}_j for $j \in [1, b]$ with $|\mathcal{S}_j| \leq k$. Then, the following inequality holds:

$$\text{sens}_{k,b}(C)^2 \geq \max_{j \in [1, b]} \sup_{S \subseteq \mathcal{S}_j} \left\| \sum_{i \in S} C_{[:,i]} \right\|_2^2, \quad (23)$$

where $C_{[:,i]}$ denotes the i -th column of the matrix C .

To prove a lower bound, we use the probabilistic method. Consider i.i.d. random variables $\epsilon_i \sim \text{Bernoulli}(\frac{1}{2})$. Then:

$$\begin{aligned} \sup_{S \subseteq \mathcal{S}_j} \left\| \sum_{i \in S} C_{[:,i]} \right\|_2^2 &\geq \mathbb{E} \left\| \sum_{i \in \mathcal{S}_j} C_{[:,i]} \epsilon_i \right\|_2^2 = \frac{1}{2} \sum_{i \in \mathcal{S}_j} \|C_{[:,i]}\|_2^2 + \frac{1}{4} \sum_{\substack{i \neq i' \\ i, i' \in \mathcal{S}_j}} \langle C_{[:,i]}, C_{[:,i']} \rangle \\ &= \frac{1}{4} \sum_{i \in \mathcal{S}_j} \|C_{[:,i]}\|_2^2 + \frac{1}{4} \left\| \sum_{i \in \mathcal{S}_j} C_{[:,i]} \right\|_2^2 \geq \frac{1}{4} \sum_{i \in \mathcal{S}_j} \|C_{[:,i]}\|_2^2. \end{aligned} \quad (24)$$

Thus,

$$\text{sens}_{k,b}(C)^2 \geq \max_{j \in [1, b]} \frac{1}{4} \sum_{i \in \mathcal{S}_j} \|C_{[:,i]}\|_2^2 \geq \frac{1}{4b} \sum_{i=1}^n \|C_{[:,i]}\|_2^2 = \frac{1}{4b} \|C\|_F^2. \quad (25)$$

Therefore,

$$\mathcal{E}(B, C) = \frac{1}{\sqrt{n}} \|B\|_F \text{sens}_{k,b}(C) \geq \frac{1}{2\sqrt{nb}} \|B\|_F \|C\|_F \geq \frac{1}{2\sqrt{nb}} \|BC\|_2 = \frac{1}{2\sqrt{nb}} \|A_{\alpha, \beta}\|_2. \quad (26)$$

The spectral norm of the matrix $A_{\alpha, \beta}$ has been lower bounded in Lemma 8 of Kalinin & Lampert (2024) by $\Omega(n \log n)$ for $\alpha = 1$, and by $\Omega(n)$ for $\alpha < 1$. Substituting $k = \lceil \frac{n}{b} \rceil$ concludes the proof. \square

Lemma 3 (Inverse Square Root of the Matrix $A_{\alpha,\beta}$). For $k \geq 0$, let $\tilde{r}_k = (-1)^k \binom{1/2}{k} = \frac{-r_k}{2k-1} = \frac{-1}{2k-1} \left| \binom{-1/2}{k} \right|$. The inverse of the matrix $C_{\alpha,\beta} = A_{\alpha,\beta}^{1/2}$ defined in (2), for $0 \leq \beta < \alpha \leq 1$, is

$$C_{\alpha,\beta}^{-1} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \tilde{c}_1^{\alpha,\beta} & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{c}_{n-1}^{\alpha,\beta} & \tilde{c}_{n-2}^{\alpha,\beta} & \dots & 1 \end{pmatrix}, \quad \text{where} \quad \tilde{c}_k^{\alpha,\beta} = \sum_{j=0}^k \tilde{r}_j \beta^j \tilde{r}_{k-j} \alpha^{k-j}. \quad (12)$$

Proof. The matrix for the momentum matrix is given by:

$$A_{\alpha,\beta} = A_{\alpha,0} \times A_{\beta,0} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \alpha & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \alpha^{n-1} & \alpha^{n-2} & \dots & 1 \end{pmatrix} \times \begin{pmatrix} 1 & 0 & \dots & 0 \\ \beta & 1 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \beta^{n-1} & \beta^{n-2} & \dots & 1 \end{pmatrix}. \quad (27)$$

The inverse square root then takes the form:

$$C_{\alpha,\beta}^{-1} = C_{\alpha,0}^{-1} \times C_{\beta,0}^{-1}, \quad (28)$$

since all lower triangular Toeplitz (LTT) matrices commute (see Strang (1986) or Böttcher & Grudsky (2000)). Therefore, it suffices to prove that the inverse square root of the matrix $C_{\alpha,0}^{-1}$ is a lower triangular Toeplitz matrix with elements $\tilde{r}_i \alpha^i$, which would imply the stated formula for $\tilde{c}_k^{\alpha,\beta}$, since the product of LTT matrices is given by the convolution of their elements.

The proof for the matrix $C_{\alpha,0}^{-1}$ is based on the identities of the generating functions of the sequences r_k and \tilde{r}_k , derived simultaneously using the binomial formula:

$$\begin{aligned} (1 - \alpha x)^{-1/2} &= \sum_{k=0}^{\infty} \binom{-1/2}{k} (-1)^k \alpha^k x^k = \sum_{k=0}^{\infty} r_k \alpha^k x^k, \\ (1 - \alpha x)^{1/2} &= \sum_{k=0}^{\infty} \binom{1/2}{k} (-1)^k \alpha^k x^k = \sum_{k=0}^{\infty} \tilde{r}_k \alpha^k x^k. \end{aligned} \quad (29)$$

Then the generating function of the product of the matrices C_α and the proposed C_α^{-1} is given by:

$$\sum_{n=0}^{\infty} x^n \left[\sum_{k=0}^n r_k \alpha^k \tilde{r}_{n-k} \alpha^{n-k} \right] = (1 - \alpha x)^{1/2} \times (1 - \alpha x)^{-1/2} = 1, \quad (30)$$

implying that $\tilde{r}_i \alpha^i$ are indeed the coefficients of C_α^{-1} , which concludes the proof. \square

Lemma 10 (Bounds on diagonal entries of $C_{1,\beta}^{-1}$). The diagonal elements of the inverse square root of the momentum matrix $C_{1,\beta}^{-1}$ defined in equation (12) with parameter $0 \leq \beta < 1$, denoted as $(1, \tilde{c}_1^{1,\beta}, \tilde{c}_2^{1,\beta}, \dots, \tilde{c}_{n-1}^{1,\beta})$, satisfy the following inequality:

$$\tilde{r}_k (1 + \beta) \leq \tilde{c}_k^{1,\beta} \leq 0, \quad \text{for } k \geq 1. \quad (31)$$

Proof. The values $\tilde{c}_k^{1,\beta}$ are given by the convolution of \tilde{r}_k and $\beta^k \tilde{r}_k$:

$$\tilde{c}_k^{1,\beta} = \sum_{j=0}^k \tilde{r}_j \tilde{r}_{k-j} \beta^j = (1 + \beta^k) \tilde{r}_k + \sum_{j=1}^{k-1} \tilde{r}_j \tilde{r}_{k-j} \beta^j. \quad (32)$$

Since \tilde{r}_j is negative for $j \geq 1$, the summation term is positive. Furthermore, $1 + \beta^k \leq 1 + \beta$, and since \tilde{r}_k is negative, we obtain the lower bound:

$$\tilde{c}_k^{1,\beta} \geq \tilde{r}_k (1 + \beta). \quad (33)$$

This bound is tight for $k = 1$ as $\tilde{c}_1^{1,\beta} = -\frac{1+\beta}{2}$.

For the upper bound, we first consider the special cases. When $\beta = 0$, we have $\tilde{c}_k^{1,0} = \tilde{r}_k < 0$. For $\beta = 1$, we formally obtain:

$$\tilde{c}_k^{1,1} = \sum_{j=0}^k \tilde{r}_j \tilde{r}_{k-j} = \begin{cases} 1, & k = 0, \\ -1, & k = 1, \\ 0, & \text{otherwise.} \end{cases} \quad (34)$$

This follows from the observation that $C_1^{-1} \times C_1^{-1} = A_1^{-1}$, which has the structure described in the equation.

Since the inequality holds for $k = 1$, we now consider $k \geq 2$, where $\tilde{c}_k^1 = 0$. We show the following, which establishes the upper bound:

Proposition 1 (Monotonicity of diagonal elements of $C_{1,\beta}^{-1}$). *Let $\tilde{c}_k^{1,\beta}$ be the diagonal elements of $C_{1,\beta}^{-1}$ defined in equation (12). Then $\tilde{c}_k^{1,\beta}$ is an increasing function of β , varying from \tilde{r}_k at $\beta = 0$ to 0 at $\beta = 1$.*

Proof. To do so, we differentiate $\tilde{c}_k^{1,\beta}$ with respect to β :

$$\frac{d\tilde{c}_k^{1,\beta}}{d\beta} = k\tilde{r}_k\beta^{k-1} + \sum_{j=1}^{k-1} \tilde{r}_j \tilde{r}_{k-j} j \beta^{j-1} = \beta^{k-1} \left(k\tilde{r}_k + \sum_{j=1}^{k-1} \tilde{r}_j \tilde{r}_{k-j} j \beta^{j-k} \right). \quad (35)$$

To prove that this expression is positive, we analyze the term in brackets. As $\beta \rightarrow 0$, the term tends to positive infinity since $\tilde{r}_j \tilde{r}_{k-j}$ are positive and $j - k$ is negative. Moreover, this term is decreasing as $\beta \rightarrow 1$, so it suffices to check its non-negativity at $\beta = 1$, i.e.,

$$\left. \frac{d\tilde{c}_k^{1,\beta}}{d\beta} \right|_{\beta=1} \geq 0 \quad (36)$$

Setting $\beta = 1$ in equation (35), we have

$$\left. \frac{d\tilde{c}_k^{1,\beta}}{d\beta} \right|_{\beta=1} = k\tilde{r}_k + \sum_{j=1}^{k-1} \tilde{r}_j \tilde{r}_{k-j} j. \quad (37)$$

To show this, we use an auxiliary identity for the values $\tilde{r}_j j$:

$$\tilde{r}_j j = -\frac{r_j j}{2j-1} = \frac{-1}{2} \left(r_j + \frac{r_j}{2j-1} \right) = \frac{-r_j}{2} + \frac{\tilde{r}_j}{2}. \quad (38)$$

Using the identity (38) in equation (37), we obtain:

$$\begin{aligned} \left. \frac{d(\tilde{c}_k^{1,\beta})}{d\beta} \right|_{\beta=1} &= \frac{1}{2}\tilde{r}_k - \frac{1}{2}r_k + \frac{1}{2} \sum_{j=1}^{k-1} \tilde{r}_j \tilde{r}_{k-j} - \frac{1}{2} \sum_{j=1}^{k-1} r_j \tilde{r}_{k-j} \\ &= \frac{1}{2} \sum_{j=0}^k \tilde{r}_j \tilde{r}_{k-j} - \frac{1}{2} \sum_{j=0}^k r_j \tilde{r}_{k-j} = 0. \end{aligned} \quad (39)$$

Since both sums vanish for $k \geq 2$, this concludes the proof of Proposition 1. \square

This completes the proof of the lemma. \square

Lemma 6 (Decreasing values). *The values $(1, c_1^{\alpha,\beta}, \dots, c_{p-1}^{\alpha,\beta}, g_p^{\alpha,\beta}, \dots, g_{n-1}^{\alpha,\beta})$ of matrix $C_{\alpha,\beta}^p$ as defined in Lemma 5 are decreasing.*

Proof. The first p values are decreasing, as shown in Kalinin & Lampert (2024). For the remaining values, we prove that

$$g_{p+k}^{\alpha,\beta} - g_{p+k+1}^{\alpha,\beta} = \sum_{j=1}^{p-1} (-\tilde{c}_j^{\alpha,\beta})(g_{p+k-j}^{\alpha,\beta} - g_{p+k+1-j}^{\alpha,\beta}) \geq 0. \quad (40)$$

In Lemma 10, we prove that $(-\tilde{c}_j^{\alpha,\beta}) \geq 0$, so each term in the summation is non-negative. Since the differences $(g_i^{\alpha,\beta} - g_{i+1}^{\alpha,\beta})$ are also non-negative by the induction step, the inequality follows, completing the proof. \square

Lemma 11 (Bound on the matrix diagonal sum of $C_{1,\beta}^{-1}$). *The diagonal elements of the inverse square root of the momentum matrix $C_{1,\beta}^{-1}$ defined in equation (12) with parameter $0 \leq \beta < 1$, denoted as $(1, \tilde{c}_1^{1,\beta}, \tilde{c}_2^{1,\beta}, \dots, \tilde{c}_{n-1}^{1,\beta})$, satisfy the following inequality:*

$$r_k(1 - \beta) \leq \sum_{j=0}^k \tilde{c}_j^{1,\beta} \leq c_k^{1,\beta}(1 - \beta), \quad \text{for } k \geq 1. \quad (41)$$

Here $\tilde{c}_i^{1,\beta}$ is as defined by equation (12).

Proof. We first state several properties of the sums of $\tilde{c}_j^{1,\beta}$:

$$(1) \sum_{j=0}^k \tilde{c}_j^{1,\beta} = \sum_{j=0}^k \tilde{r}_j \beta^j r_{k-j}, \quad (2) \sum_{j=0}^k \tilde{c}_j^{1,\beta} \beta^{k-j} = \sum_{j=0}^k r_j \beta^j \tilde{r}_{k-j}, \quad (3) \sum_{j=0}^k \frac{\tilde{c}_j^{1,\beta} (1 - \beta^{k-j+1})}{1 - \beta} = c_k^{1,\beta} \quad (42)$$

which can be derived from equating the coefficients of the following generating function identities, respectively:

$$(1) \quad \left[\sqrt{1-x} \sqrt{1-\beta x} \right] \times \frac{1}{1-x} = \frac{\sqrt{1-\beta x}}{\sqrt{1-x}}$$

$$(2) \quad \left[\sqrt{1-x} \sqrt{1-\beta x} \right] \times \frac{1}{1-\beta x} = \frac{\sqrt{1-x}}{\sqrt{1-\beta x}}$$

$$(3) \quad \left[\sqrt{1-x} \sqrt{1-\beta x} \right] \times \left[\frac{1}{1-\beta x} \frac{1}{1-x} \right] = \frac{1}{\sqrt{1-x} \sqrt{1-\beta x}}$$
(43)

Upper Bound. First, we rewrite the expression as follows by multiplying and dividing by $1 - \beta$ the terms $\tilde{c}_j^{1,\beta}$:

$$\begin{aligned} \sum_{j=0}^k \tilde{c}_j^{1,\beta} - c_k^{1,\beta}(1 - \beta) &= (1 - \beta) \sum_{j=0}^k \frac{\tilde{c}_j^{1,\beta} (1 - \beta^{k-j+1} + \beta^{k-j+1})}{1 - \beta} - c_k^{1,\beta}(1 - \beta) \\ &= \beta \sum_{j=0}^k \tilde{c}_j^{1,\beta} \beta^{k-j} = \beta \sum_{j=0}^k r_j \beta^j \tilde{r}_{k-j} \\ &= \beta^{k+1} \sum_{j=0}^k \tilde{r}_j \beta^{-j} r_{k-j}, \end{aligned} \quad (44)$$

where the third equality follows from equation 42 (2). For $\beta = 0$, the expression is identically 0. So, now consider when $\beta > 0$. We want to show that

$$\sum_{j=0}^k \tilde{r}_j \beta^{-j} r_{k-j} \geq 0 \quad (45)$$

for all $\beta \in (0, 1]$.

As β increases from 0 to 1, the sum is clearly increasing, since the only positive term does not have a β multiplier. For $\beta = 1$, the sum equals zero, as the sequences \tilde{r}_j and r_j have inverse generating functions. Therefore, the sum remains negative, concluding the proof of the upper bound.

Lower Bound.

For the lower bound, using equation (42) and the recurrence relation of \tilde{r}_j stated in Lemma 3, we get

$$\begin{aligned} \sum_{j=0}^k \tilde{c}_j^{1,\beta} - r_k(1-\beta) &= \sum_{j=0}^k \tilde{r}_j \beta^j r_{k-j} - r_k(1-\beta) = \sum_{j=1}^k \tilde{r}_j \beta^j r_{k-j} + \beta r_k \\ &= \beta \left[r_k - \sum_{j=1}^k \frac{r_j}{2j-1} r_{k-j} \beta^{j-1} \right] \geq \beta \left[r_k - \sum_{j=1}^k \frac{r_j}{2j-1} r_{k-j} \right] \\ &\geq \beta \sum_{j=0}^k \tilde{r}_j r_{k-j} = 0, \end{aligned} \quad (46)$$

concluding the proof. In the above, the first inequality follows from the fact that $0 < \beta \leq 1$. The fact is trivially true for $\beta = 0$. \square

Lemma 12 (Bound on diagonal values of the matrix $C_{1,\beta}$). *The diagonal values of the matrix $C_{1,\beta}$ (see equation (11)) with parameter $0 \leq \beta < 1$, denoted as $(1, c_1^{1,\beta}, c_2^{1,\beta}, \dots, c_{n-1}^{1,\beta})$, satisfy the inequality:*

$$c_k^{1,\beta} \leq c_{k-1}^{1,\beta} \left[1 - \frac{(1-\beta)^2}{2k} \right] \quad \text{for } k \geq 1. \quad (47)$$

Proof. We first prove that

$$c_{k-1}^{1,\beta} - c_k^{1,\beta} \geq (r_{k-1} - r_k)(1-\beta). \quad (48)$$

Using the expression of $c_k^{1,\beta}$, we have

$$\begin{aligned} c_{k-1}^{1,\beta} - c_k^{1,\beta} - (r_{k-1} - r_k)(1-\beta) &= \sum_{j=0}^{k-1} r_j r_{k-1-j} \beta^j - \sum_{j=0}^k r_j r_{k-j} \beta^j - (r_{k-1} - r_k)(1-\beta) \\ &= \beta(r_{k-1} - r_k) + \sum_{j=1}^{k-1} r_j (r_{k-j-1} - r_{k-j}) \beta^j - r_k \beta^k \\ &= \beta^k \left[\beta^{1-k} (r_{k-1} - r_k) + \sum_{j=1}^{k-1} r_j (r_{k-j-1} - r_{k-j}) \beta^{j-k} - r_k \right] \end{aligned} \quad (49)$$

We note that r_k is a decreasing sequence; therefore, the first two terms inside the brackets are positive, and the powers of β in front of them are non-positive. Therefore, as a lower bound, we can substitute $\beta = 1$ inside the sum:

$$\begin{aligned} c_{k-1}^{1,\beta} - c_k^{1,\beta} - (r_{k-1} - r_k)(1-\beta) &\geq \beta^k \left[r_{k-1} - r_k + \sum_{j=1}^{k-1} r_j (r_{k-j-1} - r_{k-j}) - r_k \right] \\ &= \beta^k [r_{k-1} - 2r_k + (1 - r_{k-1}) - (1 - 2r_k)] = 0 \end{aligned} \quad (50)$$

Using this inequality, we obtain:

$$\begin{aligned} \frac{c_k^{1,\beta}}{c_{k-1}^{1,\beta}} &= \frac{c_k^{1,\beta} - (c_{k-1}^{1,\beta} - c_k^{1,\beta})}{c_{k-1}^{1,\beta}} \leq 1 - \frac{r_{k-1} - r_k}{c_{k-1}^{1,\beta}} (1-\beta) \\ &= 1 - \frac{r_{k-1}}{2k} \cdot \frac{1-\beta}{c_{k-1}^{1,\beta}} \leq 1 - \frac{(1-\beta)^2}{2k}, \end{aligned} \quad (51)$$

concluding the proof. \square

Lemma 4 (Bounds on diagonals of $B_{\alpha,\beta}^p$). *The matrix $B_{\alpha,\beta}^p$ in the BISR factorization is a lower triangular Toeplitz matrix. The values on its diagonals are*

$$(1, c_1^{\alpha,\beta}, c_2^{\alpha,\beta}, \dots, c_{p-1}^{\alpha,\beta}, b_p^{\alpha,\beta}, \dots, b_{n-1}^{\alpha,\beta}) \quad \text{where } 0 \leq b_i^{\alpha,\beta} \leq \alpha^i c_{p-1}^{1,\beta/\alpha} \quad \text{for } i \geq p \quad (13)$$

where $c_i^{1,\beta/\alpha}$ for $1 \leq i \leq p-1$ is as defined in equation (11).

Proof. The first p values are identical to the square root factorization $c_i^{\alpha,\beta}$ due to the uniqueness of the inverse. The remaining values satisfy the following recurrence:

$$b_i^{\alpha,\beta} = \sum_{j=0}^{p-1} \tilde{c}_j^{\alpha,\beta} \frac{\alpha^{i-j+1} - \beta^{i-j+1}}{\alpha - \beta} = \alpha^i \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta/\alpha} \frac{1 - \beta^{i-j+1}}{1 - \beta/\alpha} = \alpha^i b_i^{1,\beta/\alpha}. \quad (52)$$

Therefore, it suffices to prove that $b_i^{1,\beta} \leq c_{p-1}^{1,\beta}$, since we can then substitute β with β/α .

$$\begin{aligned} b_i^{1,\beta} &= \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} \frac{1 - \beta^{i-j+1}}{1 - \beta} = \frac{1}{1 - \beta} \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} - \beta^{i+1-p} \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} \frac{\beta^{p-j}}{1 - \beta} \\ &= \frac{1}{1 - \beta} \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} + \beta^{i+1-p} \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} \frac{(-\beta^{p-j} + 1 - 1)}{1 - \beta} \\ &= \frac{1 - \beta^{i+1-p}}{1 - \beta} \sum_{j=0}^{p-1} \tilde{c}_j^{1,\beta} + c_{p-1}^{1,\beta} \beta^{i+1-p}. \end{aligned} \quad (53)$$

We now use Lemma 11 to first show that $b_i^{1,\beta} \geq 0$, since the sum of $\tilde{c}_j^{1,\beta}$ is non-negative and all other terms are positive. Second, we establish that:

$$b_i^{1,\beta} \leq \frac{1 - \beta^{i+1-p}}{1 - \beta} (1 - \beta) c_{p-1}^{1,\beta} + c_{p-1}^{1,\beta} \beta^{i+1-p} = c_{p-1}^{1,\beta}, \quad (54)$$

which completes the proof. \square

Lemma 5 (Bounds on diagonals of $C_{\alpha,\beta}^p$). *The matrix $C_{\alpha,\beta}^p$ in the BISR factorization is a lower triangular Toeplitz matrix. The values on its diagonals are $(1, c_1^{\alpha,\beta}, c_2^{\alpha,\beta}, \dots, c_{p-1}^{\alpha,\beta}, g_p^{\alpha,\beta}, \dots, g_{n-1}^{\alpha,\beta})$, where $c_i^{1,\beta/\alpha}$ (for $1 \leq i \leq p-1$) is as defined in equation (11) and*

$$0 \leq g_i^{\alpha,\beta} \leq \alpha^i \min(c_i^{1,\beta/\alpha}, c_p^{1,\beta/\alpha} \gamma_{\beta/\alpha}^{i-p}) \quad \text{for } \gamma_{\beta/\alpha} = \left(1 + \frac{(1 - \beta/\alpha)^2}{4p(1 + \beta/\alpha)}\right)^{-1} \quad \text{and } i \geq p. \quad (14)$$

Proof. The first p values of $C_{\alpha,\beta}^p$ are the same as those of $C_{\alpha,\beta}$ since the matrix is Lower Triangular Toeplitz (LTT). For the subsequent values, we first prove the following inequality by induction:

$$g_i^{\alpha,\beta} = \sum_{j=1}^{p-1} (-\tilde{c}_j^{\alpha,\beta}) g_{i-j}^{\alpha,\beta} \leq \sum_{j=1}^{p-1} (-\tilde{c}_j^{\alpha,\beta}) c_{i-j}^{\alpha,\beta} \leq \sum_{j=1}^i (-\tilde{c}_j^{\alpha,\beta}) c_{i-j}^{\alpha,\beta} = c_i^{\alpha,\beta} = \alpha^i c_i^{1,\beta/\alpha}. \quad (55)$$

We observe that for all sequences $c_i^{\alpha,\beta}$, $\tilde{c}_i^{\alpha,\beta}$, and $g_i^{\alpha,\beta}$, we can factor out α^i by replacing β with β/α . Therefore, it suffices to prove the inequality $g_i^{1,\beta} \leq c_p^{1,\beta} \gamma_{\beta}^{i-p}$, after which we may substitute β with β/α . For the subsequent p values, we establish the stated bound $g_i^{1,\beta} \leq c_p^{1,\beta} \gamma_{\beta}^{i-p}$ using Lemma 12.

$$\frac{g_{p+k}^{1,\beta}}{c_p^{1,\beta} \gamma_{\beta}^k} \leq \frac{c_{p+k}^{1,\beta}}{c_p^{1,\beta}} \left(1 + \frac{(1 - \beta)^2}{4p(1 + \beta)}\right)^k = \prod_{j=1}^k \left(1 - \frac{(1 - \beta)^2}{2(p+j)}\right) \left(1 + \frac{(1 - \beta)^2}{4p(1 + \beta)}\right) \leq 1. \quad (56)$$

Since each term in the product is less than 1 for $2p + 2j \leq 4p$, the inequality holds. For the induction step, we show:

$$\frac{g_{p+k}^{1,\beta}}{c_p^{1,\beta} \gamma_\beta^k} = \frac{1}{c_p^{1,\beta} \gamma_\beta^k} \sum_{j=1}^{p-1} (-\tilde{c}_j^{1,\beta}) g_{p+k-j}^{1,\beta} \leq \sum_{j=1}^{p-1} (-\tilde{c}_j^{1,\beta}) \gamma_\beta^{-j} = \sum_{j=1}^{p-1} (-\tilde{c}_j^{1,\beta}) \left(1 + \frac{(1-\beta)^2}{4p(1+\beta)}\right)^j. \quad (57)$$

For convenience, we denote $\phi_\beta = \frac{(1-\beta)^2}{1+\beta} < 1$. To proceed, we use the following auxiliary inequality for $j \leq p-1$:

$$\left(1 + \frac{\phi_\beta}{4p}\right)^j \leq e^{\frac{j\phi_\beta}{4p}} \leq 1 + \frac{5j\phi_\beta}{16p}, \quad (58)$$

since $e^x \leq 1 + 1.25x$ for $x \leq \frac{1}{4}$. Combining this inequality with Lemma 11, we obtain:

$$\frac{g_{p+k}^{1,\beta}}{c_p^{1,\beta} \gamma_\beta^k} \leq \sum_{j=1}^{p-1} (-\tilde{c}_j^{1,\beta}) \left(1 + \frac{5j\phi_\beta}{16p}\right) \leq 1 - r_{p-1}(1-\beta) + \frac{5\phi_\beta}{16p} \sum_{j=1}^{p-1} (-\tilde{c}_j^{1,\beta})j. \quad (59)$$

By Lemma 10, we can upper bound:

$$(-\tilde{c}_j^{1,\beta})j \leq (-\tilde{r}_j)j(1+\beta) = \frac{j r_j}{2j-1}(1+\beta) \leq r_j(1+\beta). \quad (60)$$

Using the known bounds $\frac{1}{\sqrt{\pi(j+1)}} \leq r_j \leq \frac{1}{\sqrt{\pi j}}$, we conclude:

$$\begin{aligned} \frac{g_{p+k}^{1,\beta}}{c_p^{1,\beta} \gamma_\beta^k} &\leq 1 - \frac{1-\beta}{\sqrt{\pi p}} + \frac{5(1-\beta)^2}{16p\sqrt{\pi}} \sum_{j=1}^{p-1} \frac{1}{\sqrt{j}} \leq 1 - \frac{1-\beta}{\sqrt{\pi p}} + \frac{5(1-\beta)^2}{16p\sqrt{\pi}} \cdot 2\sqrt{p} \\ &\leq 1 - \frac{1-\beta}{\sqrt{\pi p}} \left(1 - \frac{5}{8}(1-\beta)\right) < 1, \end{aligned} \quad (61)$$

where for the second inequality we used the integral estimate $\sum_{j=1}^{k-1} j^{-1/2} \leq \int_0^k x^{-1/2} dx = 2\sqrt{k}$.

Thus, we have shown that $\frac{g_{p+k}^{1,\beta}}{c_p^{1,\beta} \gamma_\beta^k} \leq 1$ for all k , which completes the proof. \square

Theorem 2 (BISR Approximation Error). *For $1 \leq p \leq n$ and $1 \leq k \leq \frac{n}{b}$ the following upper bound holds for the matrix factorization error of the BISR $A_{\alpha,\beta} = B_{\alpha,\beta}^p C_{\alpha,\beta}^p$ (as in Definition 1):*

$$\mathcal{E}(B_{\alpha,\beta}^p, C_{\alpha,\beta}^p) = \begin{cases} O_\beta \left(\sqrt{k} \log p + \sqrt{\frac{nk}{b}} + \sqrt{\frac{nk \log p}{p}} + \sqrt{\frac{kp \log p}{b}} \right) & \text{for } \alpha = 1, \\ O_{\alpha,\beta}(\sqrt{k}) & \text{for } \alpha < 1. \end{cases} \quad (7)$$

Proof. We begin with the case $\alpha < 1$. To analyze this, we first consider the Frobenius norm:

$$\begin{aligned} \frac{\|B_{\alpha,\beta}^p\|_{\text{Fr}}^2}{n} &\leq \sum_{i=0}^{p-1} (c_i^{\alpha,\beta})^2 + \sum_{i=p}^{n-1} (b_i^{\alpha,\beta})^2 \leq \sum_{i=0}^{p-1} (c_i^{\alpha,\beta})^2 + (c_{p-1}^{1,\beta/\alpha})^2 \sum_{i=p}^{n-1} \alpha^{2i} \\ &\leq \frac{1}{(\alpha-\beta)^2} \log \left(\frac{1}{1-\alpha^2} \right) + \frac{\alpha^{2p}}{1-\alpha^2} = O_{\alpha,\beta}(1), \end{aligned} \quad (62)$$

where for the second inequality we used Lemma 4, and for the third inequality Lemma 7 from Kalinin & Lampert (2024).

For the (k, b) -sensitivity of the matrix $C_{\alpha,\beta}^p$, we use the fact that it is element-wise bounded by the full matrix $C_{\alpha,\beta}$ (see Lemma 5). For $C_{\alpha,\beta}$, we apply a bound from Theorem 7 of Kalinin & Lampert (2024), yielding $\text{sens}_{k,b}(C_{\alpha,\beta}) = O_{\alpha,\beta}(\sqrt{k})$, which concludes the case $\alpha < 1$.

For $\alpha = 1$, we use Lemma 4 to get:

$$\begin{aligned} \frac{\|B_{1,\beta}^p\|_{\text{Fr}}^2}{n} &\leq \sum_{i=0}^{n-1} (b_i^{1,\beta})^2 = \sum_{i=0}^{p-1} (c_i^{1,\beta})^2 + \sum_{i=p}^{n-1} (c_{p-1}^{1,\beta})^2 = \frac{1}{(1-\beta)^2} \sum_{i=0}^{p-1} r_i^2 + \frac{1}{(1-\beta)^2} \sum_{i=p}^{n-1} r_{p-1}^2 \\ &\leq \frac{1}{(1-\beta)^2} \left[1 + \log p + \frac{n-p}{p\pi} \right] = O_\beta \left(\log p + \frac{n}{p} \right). \end{aligned} \quad (63)$$

Next, we bound the sensitivity under k, b participation. Using Theorem 3, combined with Lemma 6 we obtain:

$$\text{sens}_{k,b}^2(C_{1,\beta}^p) = \sum_{j=0}^{k-1} \sum_{i=0}^{k-1} \langle (C_{1,\beta}^p)_{:,ib}, (C_{1,\beta}^p)_{:,jb} \rangle. \quad (64)$$

We split the sum into the following four terms:

$$\begin{aligned} \text{sens}_{k,b}^2(C_{1,\beta}^p) &= \underbrace{\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p+ib,n)-1-jb} c_t^{1,\beta} c_{jb-ib+t}^{1,\beta}}_{S_1} + \underbrace{\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p-1,n-1-jb)} c_t^{1,\beta} g_{jb-ib+t}^{1,\beta}}_{S_2} \\ &\quad + \underbrace{\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=p}^{n-1-jb} g_t^{1,\beta} g_{jb-ib+t}^{1,\beta}}_{S_3} + \underbrace{\sum_{i=0}^{k-1} \left[\sum_{t=0}^{\min(p-1,n-1-ib)} (c_t^{1,\beta})^2 + \sum_{t=p}^{n-1-ib} (g_t^{1,\beta})^2 \right]}_{S_4} \end{aligned} \quad (65)$$

Step 1 (S_1 Bound) We note that the case $b < p < n$ has not been considered in Kalinin & Lampert (2024) and is technically more challenging. Consider the half of the sum where $j > i$. The sum requires $jb - ib \leq p - 1$; otherwise, the upper limit would be negative. We bound the sum as follows:

$$\begin{aligned} \sum_{t=0}^{\min(p+ib,n)-1-jb} c_t^{1,\beta} c_{jb-ib+t}^{1,\beta} &\leq \frac{r_{jb-ib}}{1-\beta} + \frac{1}{\pi(1-\beta)^2} \sum_{t=1}^{p-1+ib-jb} \frac{1}{\sqrt{t(jb-ib+t)}} \\ &\leq \frac{1}{(1-\beta)^2} \left[1 + \frac{1}{\pi} \int_0^{p-1+ib-jb} \frac{dx}{\sqrt{x(jb-ib+x)}} \right] \\ &= \frac{1}{(1-\beta)^2} \left[1 + \frac{1}{\pi} f \left(\frac{jb-ib}{p-1+ib-jb} \right) \right], \end{aligned} \quad (66)$$

where $f(a) = 2 \log \left(\sqrt{\frac{1}{a} + 1} + \sqrt{\frac{1}{a}} \right)$. We then use the following auxiliary inequality for the function $f(a)$:

$$f(a) = \log \left(\frac{1}{a} + 1 \right) + 2 \log \left(1 + \frac{1}{\sqrt{a+1}} \right) \leq \log \left(\frac{1}{a} + 1 \right) + 2 \log 2. \quad (67)$$

This results in the following inequality:

$$\sum_{t=0}^{\min(p+ib,n)-1-jb} c_t^{1,\beta} c_{jb-ib+t}^{1,\beta} \leq \frac{1}{(1-\beta)^2} \left[1 + \frac{2 \log 2}{\pi} + \frac{1}{\pi} \log \left(\frac{p-1+ib-jb}{jb-ib} \right) \right] \mathbb{1}_{jb-ib \leq p-1}. \quad (68)$$

We can now upper bound the double sum:

$$\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p+ib,n)-1-jb} c_t^{1,\beta} c_{jb-ib+t}^{1,\beta} \leq \frac{2}{(1-\beta)^2} \sum_{i=0}^{k-1} \sum_{j=i+1}^{\min(k-1, i + \lfloor \frac{p-1}{b} \rfloor)} \left[\frac{3}{2} + \frac{1}{\pi} \log \left(\frac{p-1}{jb-ib} \right) \right]. \quad (69)$$

The first term gives us:

$$\frac{2}{(1-\beta)^2} \sum_{i=0}^{k-1} \sum_{j=i+1}^{\min(k-1, i + \lfloor \frac{p-1}{b} \rfloor)} \frac{3}{2} \leq \frac{3k}{(1-\beta)^2} \left\lfloor \frac{p-1}{b} \right\rfloor. \quad (70)$$

The second term is more involved. First, we upper bound the upper limit of the sum $\min(k-1, i + \lfloor \frac{p-1}{b} \rfloor)$ by $i + \lfloor \frac{p-1}{b} \rfloor$, since the summands are positive. We can then upper bound the expression by:

$$\sum_{i=0}^{k-1} \sum_{j=i+1}^{i + \lfloor \frac{p-1}{b} \rfloor} \log \left(\frac{\frac{p-1}{b}}{j-i} \right) = \log \prod_{i=0}^{k-1} \frac{(\frac{p-1}{b})^{\lfloor \frac{p-1}{b} \rfloor}}{(\lfloor \frac{p-1}{b} \rfloor)!} = k \log \frac{(\frac{p-1}{b})^{\lfloor \frac{p-1}{b} \rfloor}}{(\lfloor \frac{p-1}{b} \rfloor)!}. \quad (71)$$

Using the auxiliary inequality $k! \geq (\frac{k}{e})^k$, we show that:

$$\begin{aligned} \log \frac{(\frac{p-1}{b})^{\lfloor \frac{p-1}{b} \rfloor}}{(\lfloor \frac{p-1}{b} \rfloor)!} &\leq \left\lfloor \frac{p-1}{b} \right\rfloor \log \frac{p-1}{b} - \left\lfloor \frac{p-1}{b} \right\rfloor \log \left\lfloor \frac{p-1}{b} \right\rfloor + \left\lfloor \frac{p-1}{b} \right\rfloor \\ &= \left\lfloor \frac{p-1}{b} \right\rfloor \log \left\{ \frac{p-1}{b} \right\} + \left\lfloor \frac{p-1}{b} \right\rfloor \leq \left\lfloor \frac{p-1}{b} \right\rfloor. \end{aligned} \quad (72)$$

Resulting in:

$$\begin{aligned} \sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p+ib, n)-1-jb} c_t^{1, \beta} c_{jb-ib+t}^{1, \beta} &\leq \frac{1}{(1-\beta)^2} \left(3k \left\lfloor \frac{p-1}{b} \right\rfloor + \frac{2}{\pi} k \left\lfloor \frac{p-1}{b} \right\rfloor \right) \\ &\leq \frac{4k}{(1-\beta)^2} \left\lfloor \frac{p-1}{b} \right\rfloor, \end{aligned} \quad (73)$$

which concludes this part of the calculations.

Step 2 (Bound S_2). We can bound the inner sum as follows, assuming that $jb - ib \geq p$:

$$\begin{aligned} \sum_{t=0}^{\min(p-1, n-1-jb)} c_t^{1, \beta} g_{jb-ib+t}^{1, \beta} &\leq c_p^{1, \beta} \sum_{t=0}^{p-1} c_t^{1, \beta} \gamma_\beta^{jb-ib+t-p} = c_p^{1, \beta} \gamma_\beta^{jb-ib-p} \sum_{t=0}^{p-1} c_t^{1, \beta} \gamma_\beta^t \\ &\leq c_p^{1, \beta} \gamma_\beta^{jb-ib-p} \sum_{t=0}^{p-1} c_t^{1, \beta} \leq \frac{r_p \gamma_\beta^{jb-ib-p}}{(1-\beta)^2} \left(1 + \frac{1}{\sqrt{\pi}} \sum_{t=1}^{p-1} \frac{1}{\sqrt{t}} \right) \\ &\leq \frac{r_p \gamma_\beta^{jb-ib-p}}{(1-\beta)^2} \left(1 + \frac{2\sqrt{p}}{\sqrt{\pi}} \right) \leq \frac{3\gamma_\beta^{jb-ib-p}}{(1-\beta)^2} \end{aligned} \quad (74)$$

For our specific choice of $\gamma_\beta = 1 - \frac{\phi_\beta}{4p + \phi_\beta} = \left(1 + \frac{\phi_\beta}{4p} \right)^{-1}$, where $\phi_\beta = \frac{(1-\beta)^2}{1+\beta}$. We rewrite the bound using the following auxiliary inequality:

$$\gamma_\beta^{-p} = \left(1 + \frac{\phi_\beta}{4p} \right)^p \leq e^{\phi_\beta/4} \leq e^{1/4} \leq \frac{4}{3}, \quad (75)$$

This yields the upper bound for the whole sum:

$$\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p-1, n-1-jb)} c_t^{1, \beta} g_{jb-ib+t}^{1, \beta} \leq \frac{8}{(1-\beta)^2} \sum_{i=0}^{k-1} \sum_{j=i+1}^{k-1} \gamma_\beta^{jb-ib} \leq \frac{8k\gamma_\beta^b}{(1-\beta)^2(1-\gamma_\beta^b)} \quad (76)$$

We bound γ_β^b in the following way:

$$\gamma_\beta^b = \left(1 - \frac{\phi_\beta}{4p + \phi_\beta} \right)^b \leq e^{-\frac{b\phi_\beta}{4p + \phi_\beta}} = e^{-\frac{b\phi_\beta}{p} \frac{p}{4p+1}} \leq e^{-\frac{b\phi_\beta}{5p}} \quad (77)$$

Thus,

$$\begin{aligned} \sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=0}^{\min(p-1, n-1-jb)} c_t^{1,\beta} g_{jb-ib+t}^{1,\beta} &\leq \frac{8k}{(1-\beta)^2(\gamma_\beta^{-b}-1)} \\ &\leq \frac{8k}{(1-\beta)^2(e^{\frac{b\phi_\beta}{5p}}-1)} \leq \frac{40kp(1+\beta)}{b(1-\beta)^4}. \end{aligned} \quad (78)$$

Step 3 (Bound \mathcal{S}_3) We first bound the inner sum, assuming $j > i$:

$$\begin{aligned} \sum_{t=p}^{n-1-jb} g_t^{1,\beta} g_{jb-ib+t}^{1,\beta} &\leq (c_p^{1,\beta})^2 \sum_{t=p}^{n-1-jb} \gamma_\beta^{t-p} \gamma_\beta^{jb-ib+t-p} \\ &\leq \frac{(c_p^{1,\beta})^2 \gamma_\beta^{jb-ib}}{1-\gamma_\beta^2} \leq \frac{(c_p^{1,\beta})^2 \gamma_\beta^{jb-ib} (4p+\phi_\beta)}{2\phi_\beta} \\ &\leq \frac{r_p^2 (4p+1) \gamma_\beta^{jb-ib}}{2\phi_\beta (1-\beta)^2} \leq \frac{5r_p^2 \gamma_\beta^{jb-ib} p}{2\phi_\beta (1-\beta)^2} \\ &\leq \frac{5\gamma_\beta^{jb-ib}}{2\pi\phi_\beta (1-\beta)^2} \leq \frac{\gamma_\beta^{jb-ib} (1+\beta)}{(1-\beta)^4}. \end{aligned} \quad (79)$$

This yields the upper bound:

$$\sum_{i=0}^{k-1} \sum_{j \neq i}^{k-1} \sum_{t=p}^{n-1-jb} g_t^{1,\beta} g_{jb-ib+t}^{1,\beta} \leq \frac{2(1+\beta)}{(1-\beta)^4} \sum_{i=0}^{k-1} \sum_{j=i+1}^{k-1} \gamma_\beta^{jb-ib} \leq \frac{10kp(1+\beta)^2}{b(1-\beta)^6}. \quad (80)$$

analogously to the previous step.

Step 4 (Bound \mathcal{S}_4) We bound the sum of squared column norms as follows:

$$\begin{aligned} \sum_{i=0}^{k-1} \left[\sum_{t=0}^{\min(p-1, n-1-ib)} (c_t^{1,\beta})^2 + \sum_{t=p}^{n-1-ib} (g_t^{1,\beta})^2 \right] &\leq \frac{k}{(1-\beta)^2} \left[\sum_{t=0}^{p-1} r_t^2 + \frac{r_p^2}{1-\gamma_\beta^2} \right] \\ &\leq \frac{k}{(1-\beta)^2} \left[1 + \log p + \frac{5(1+\beta)}{2\pi(1-\beta)^2} \right]. \end{aligned} \quad (81)$$

Step 5 (Combination) Combining all steps together, we bound the k, b sensitivity as follows:

$$\begin{aligned} \text{sens}_{k,b}^2(C_{1,\beta}^p) &= \sum_{j=0}^{k-1} \sum_{i=0}^{k-1} \langle (C_{1,\beta}^p)_{:,ib}, (C_{1,\beta}^p)_{:,jb} \rangle \\ &\leq \frac{k}{(1-\beta)^2} \left(1 + \log p + \frac{5(1+\beta)}{2\pi} + 10 \frac{p(1+\beta)}{b(1-\beta)^2} \left(4 + \frac{1+\beta}{(1-\beta)^2} \right) + 4 \left\lfloor \frac{p-1}{b} \right\rfloor \right) \\ &\leq \frac{k(1+\beta)^2}{(1-\beta)^6} \left(2 + \log p + 54 \frac{p}{b} \right) \end{aligned} \quad (82)$$

Thus,

$$\mathcal{E}(B_{1,\beta}^p, C_{1,\beta}^p)^2 \leq \frac{k(1+\beta)^2}{(1-\beta)^8} \left(1 + \log p + \frac{n-p}{p\pi} \right) \left(2 + \log p + 54 \frac{p}{b} \right) \quad (83)$$

And

$$\mathcal{E}(B_{1,\beta}^p, C_{1,\beta}^p) = O_\beta \left(\sqrt{k} \log p + \sqrt{\frac{nk}{b}} + \sqrt{\frac{nk \log p}{p}} + \sqrt{\frac{kp \log p}{b}} \right). \quad (84)$$

□

B Additional Materials

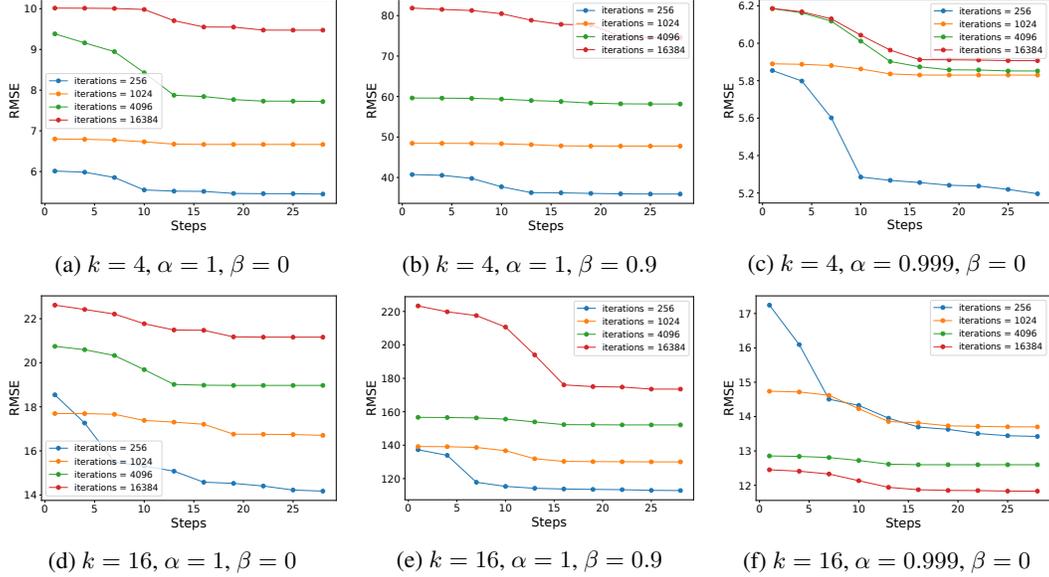


Figure 5: Convergence of Band-Inv-MF under different settings: for participation numbers $k = 4, 16$, with and without momentum (β) and weight decay (α), across various matrix sizes (iterations). In general, we observe that 20 steps are sufficient for the procedure to converge.

Table 1: Hyperparameters for CIFAR-10 Experiments. We train all methods with and without amplification to achieve $(9, 10^{-5})$ -differential privacy. Training uses a weight decay of 0.9999, momentum of 0.9, and batch size 512. Noise multipliers are computed via an MCMC accountant for the amplified case, and as $\sigma_{\epsilon, \delta} \times \text{sens}_{k, b}(C)$ for the non-amplified case, assuming 10 training epochs.

	Method	Noise Multiplier	Learning Rate	bandwidth	Clip Norm
Amplified	DP-SGD	1.2	0.1	1	10
	BSR	2.3	0.3	4	10
	BISR	4.4	0.7	4	10
	Band-MF	2.4	0.3	4	10
	Band-Inv-MF	8.2	0.4	4	10
Non-amplified	DP-SGD	1.8	0.1	1	10
	BSR	3.3	0.2	4	10
	BISR	5.8	0.7	4	10
	Band-MF	3.5	0.2	4	10
	Band-Inv-MF	9.1	0.5	4	10

Table 2: CIFAR-10 experiments with and without amplification, for $\epsilon = 9, \delta = 10^{-5}$ showing test accuracy (%) over 10 epochs. Mean \pm standard error computed over 3 runs.

	Method	Epoch 1	Epoch 2	Epoch 3	Epoch 4	Epoch 5	Epoch 6	Epoch 7	Epoch 8	Epoch 9	Epoch 10
Amp.	DP-SGD	12.7 \pm 2.2	28.0 \pm 1.1	34.4 \pm 0.4	37.6 \pm 0.7	39.8 \pm 1.2	41.6 \pm 0.2	42.3 \pm 0.8	42.8 \pm 0.3	43.5 \pm 0.4	44.6 \pm 0.7
	BSR	28.3 \pm 0.7	40.2 \pm 1.1	43.6 \pm 1.1	46.5 \pm 0.9	48.0 \pm 2.0	48.8 \pm 1.4	48.9 \pm 1.4	49.4 \pm 0.7	49.2 \pm 1.2	49.8 \pm 0.3
	BISR	32.3 \pm 0.7	42.7 \pm 1.1	47.5 \pm 1.1	50.3 \pm 0.9	52.8 \pm 2.0	56.5 \pm 1.4	57.9 \pm 1.4	58.5 \pm 0.7	60.5 \pm 1.2	61.8 \pm 0.3
	Band-MF	27.7 \pm 2.0	38.5 \pm 0.3	43.1 \pm 1.6	43.7 \pm 1.8	46.8 \pm 0.8	47.7 \pm 0.3	48.2 \pm 0.6	47.8 \pm 2.6	49.1 \pm 0.6	50.0 \pm 0.4
	Band-Inv-MF	23.6 \pm 2.8	34.6 \pm 1.3	40.0 \pm 2.4	44.6 \pm 1.3	48.6 \pm 1.0	50.4 \pm 1.0	50.6 \pm 0.5	53.4 \pm 0.8	56.2 \pm 0.6	57.4 \pm 1.2
Non-Amp.	DP-SGD	19.5 \pm 3.0	31.0 \pm 1.1	36.7 \pm 0.2	37.2 \pm 0.4	37.7 \pm 1.2	39.3 \pm 2.0	39.8 \pm 1.2	39.1 \pm 0.3	39.5 \pm 0.5	39.0 \pm 0.7
	BSR	25.4 \pm 1.2	36.7 \pm 1.2	40.8 \pm 1.1	41.6 \pm 2.0	43.6 \pm 0.9	44.5 \pm 0.7	45.0 \pm 0.9	44.4 \pm 2.1	45.3 \pm 1.8	45.2 \pm 0.8
	BISR	31.8 \pm 1.5	41.7 \pm 2.2	45.4 \pm 1.4	48.5 \pm 1.3	51.1 \pm 1.0	51.4 \pm 2.7	53.8 \pm 1.0	54.0 \pm 1.2	55.5 \pm 0.8	56.2 \pm 0.2
	Band-MF	25.9 \pm 1.5	36.7 \pm 0.9	41.1 \pm 1.4	43.2 \pm 1.3	42.8 \pm 1.4	45.0 \pm 0.2	45.5 \pm 0.4	45.4 \pm 1.8	46.7 \pm 0.9	45.8 \pm 0.2
	Band-Inv-MF	27.4 \pm 3.0	36.0 \pm 2.5	39.5 \pm 2.4	43.7 \pm 1.2	46.7 \pm 0.5	47.0 \pm 2.0	49.7 \pm 1.7	53.5 \pm 0.5	54.4 \pm 1.5	57.9 \pm 0.4

```

1 import jax_privacy
2 from jax_privacy.dpftrl_mechanisms import toeplitz
3 import jax.numpy as jnp
4 import functools
5 import numpy as np
6
7 def expected_mean_error(inv_coef, n, k, workload_coef) -> float:
8     inv_coef = jnp.pad(inv_coef, (0, n - inv_coef.size))
9     B_norm_squared = toeplitz.mean_error(noising_coef=inv_coef, n=n,
10     workload_coef=workload_coef, skip_checks=True)
11
12     coef = toeplitz.inverse_coef(inv_coef)
13     min_sep = n // k # assume divisible
14
15     sensitivity_squared = toeplitz.minsep_sensitivity_squared(coef,
16     min_sep, k, n, skip_checks=True)
17
18     return sensitivity_squared * B_norm_squared
19
20 def compute_square_root(x, n) -> np.ndarray:
21     y = np.zeros(n)
22     y[0] = np.sqrt(x[0])
23     for k in range(1, n):
24         y[k] = (x[k] - np.dot(y[1:k], y[1:k][::-1])) / (2 * y[0])
25     return y
26
27 def init(n, p, alpha = 1.0, beta = 0.0) -> jnp.ndarray:
28     x = jnp.array([1, -alpha - beta, alpha * beta] + [0]*(n-3))
29     return jnp.array(compute_square_root(x, n)[:p])
30
31 def Band_Inv_MF(n, b, k, p, alpha, beta, steps = 20):
32     # compute workload matrix
33     M = jnp.array([(alpha ** (k + 1) - beta ** (k + 1)) / (beta -
34     alpha) for k in range(n)])
35
36     # initialize with BISR coefficients
37     C_inv_init = init(n, p, alpha, beta)
38
39     # optimize!
40     C_inv_opt = toeplitz.optimize_banded_toeplitz(
41         n=n,
42         bands=p,
43         strategy_coef=C_inv_init,
44         loss_fn=functools.partial(expected_mean_error, k=k,
45         workload_coef=M),
46         max_optimizer_steps=steps,
47     )
48     return C_inv_opt

```

Listing 1: Python code for Band-Inv-MF factorization. The function "Band_Inv_MF" takes the matrix size (n), the minimum separation (b), the number of participations (k), the bandwidth (p), the weight decay (α), the momentum (β), and the number of optimization steps.