

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

Near-Efficient and Non-Asymptotic Multiway Inference

Oscar López¹, Arvind Prasad², Carlos Llosa-Vite²,
Richard B. Lehoucq², Daniel M. Dunlavy²

¹*Harbor Branch Oceanographic Institute, Florida Atlantic University, Ft Pierce, FL*

²*Sandia National Laboratories, Albuquerque, NM and Livermore, CA*

Abstract

We establish non-asymptotic efficiency guarantees for tensor decomposition–based inference in count data models. Under a Poisson framework, we consider two related goals: (i) *parametric inference*, the estimation of the full distributional parameter tensor, and (ii) *multiway analysis*, the recovery of its canonical polyadic (CP) decomposition factors. Our main result shows that in the rank-one setting, a rank-constrained maximum-likelihood estimator achieves multiway analysis with variance matching the Cramér–Rao Lower Bound (CRLB) up to absolute constants and logarithmic factors. This provides a general framework for studying “near-efficient” multiway estimators in finite-sample settings. For higher ranks, we illustrate that our multiway estimator may not attain the CRLB; nevertheless, CP-based parametric inference remains nearly minimax optimal, with error bounds that improve on prior work by offering more favorable dependence on the CP rank. Numerical experiments corroborate near-efficiency in the rank-one case and highlight the efficiency gap in higher-rank scenarios.

1 Introduction

Parametric inference and multiway analysis play crucial roles in high-dimensional data analysis. Both perspectives are useful in practice: parametric inference estimates the tensor of distributional parameters as a whole, while multiway analysis yields its latent factors for interpretation [1]. Both tasks rely fundamentally on tensor decompositions to represent and exploit underlying structure. However, computing tensor decompositions is notoriously difficult. Degeneracy phenomena lead to non-unique or ill-conditioned factorizations [2] and many tensor problems are NP-hard [3], making even approximate computation intractable in general. These issues put into question the reliability of existing tensor-based inference methods. They are particularly pronounced for the canonical polyadic (CP) decomposition [2], which, despite its widespread use, lacks the theoretical guarantees enjoyed by other tensor formats. Computing CP factors, i.e., multiway analysis, with minimal variance across multiple sets of observations would enhance the reliability of multiway analysis and parametric inference, offering practitioners more confidence in their results while reducing the need for extensive data collection.



This work was supported by the Laboratory Directed Research and Development program (Project 229313) at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. SAND2025-14006R

To study the reliability of general inference procedures, the Cramér–Rao and minimax lower bounds provide fundamental benchmarks for variance-optimal estimators. In Cramér–Rao lower bound (CRLB) based analysis [4], the variance of an estimator is bounded from below by a term involving the target parameter’s inverse Fisher information matrix. Estimators that achieve the CRLB are known as *efficient* estimators and are particularly valuable because they guarantee optimal performance, such as maximizing prediction stability across experiments [5, 6, 7, 8]. Similarly, minimax lower bounds capture the best achievable performance of any estimation method for worst-case parameter choices. Notably, minimax analysis establishes worst-case variance guarantees over a class of parameters and all estimators, whereas the CRLB characterizes the smallest achievable variance at a specific parameter value for a given model. Thus, CRLB and minimax analyses offer complementary notions of optimality in structured statistical inference—one local and estimator-specific, the other global and distribution-agnostic—each with its own advantages and limitations.

However, deriving CRLB- and minimax-based guarantees for tensor methods presents significant challenges. On the CRLB side, modern estimators are often defined through constrained or nonconvex optimization procedures, which rarely yield the closed-form expressions required in classical efficiency analysis. Moreover, traditional CRLB results are typically asymptotic—they assume access to infinitely many independent measurements—whereas contemporary data science operates in the finite-sample regime, where only a limited number of observations are available and computational constraints matter. This motivates reformulating efficiency concepts to be meaningful at realistic sample sizes.

On the minimax side, informative lower bounds have been established for tensor estimation problems, and much of the current effort is devoted to matching these bounds with sharp upper bounds. This requires developing new analytical tools that yield sample-complexity guarantees with optimal—i.e., linear—dependence on the tensor rank. However, obtaining such results remains challenging, in large part because it depends on unresolved structural questions surrounding the CP decomposition. Unlike other tensor formats, CP imposes minimal structural assumptions, which makes it widely applicable but also more difficult to analyze rigorously in terms of identifiability and sample complexity.

In this work, we introduce a flexible framework for CRLB-based assessments, establishing several non-asymptotic guarantees for multiway analysis. Focusing on count data modeled by a Poisson distribution, we consider the setting where the array of Poisson rate parameters admits a rank-one representation. In this regime, we show that a rank-one maximum likelihood estimator of the decomposition factors attains an estimation error that matches the CRLB up to absolute constants and logarithmic factors. This result demonstrates that rank-one factor estimators can be *near-efficient*—approaching the theoretical lower variance bound. We propose a relaxed interpretation of the CRLB that is more appropriate for finite-sample settings and enables effective benchmarking of modern estimation methods.

In contrast, the general rank case exhibits significantly different behavior. We argue that estimators for higher-rank CP decompositions generally fail to attain the CRLB as in the rank-one setting. Through numerical experiments, we examine the “CRLB gap” and demonstrate that our factor estimation method does not achieve the CRLB. These results reflect a well-known issue of degeneracy in tensor decompositions [2]. Nevertheless, when estimating the full Poisson parameter tensor, we show that CP-based parametric inference remains nearly minimax optimal, up to a rank-dependent factor and a logarithmic term. Our results improve upon previous bounds by providing sharper characterizations of CP rank dependence, thereby reinforcing the utility of decomposition-based inference as a reliable approach for estimating structured distributional parameters.

The remainder of the paper is organized as follows: Subsection 1.1 discusses related work in the

literature. In Section 2, Subsections 2.1 and 2.2 provide background on the CRLB and minimax lower bounds, including a result that we will use to simplify CRLB-based analysis. Afterwards, Section 2.3 will discuss our multiway inference scenario and present a simplified version of our main results along with discussion in Section 2.4. Numerical experiments are presented in Section 3 that explore our theoretical findings. Sections 4-8 provide expanded versions of the simplified results presented in Section 2, along with required lemmas, and proofs. Concluding remarks and future work are discussed in Section 9.

1.1 Related Work

Our CRLB analysis is motivated by a non-asymptotic perspective, in which efficiency guarantees are established using a finite number of observations. Classical CRLB results—fundamental to statistical inference—are typically asymptotic, assuming the number of samples grows without bound [9, 10]. This approach has been especially prevalent in multiway inference, where efficiency is often justified by letting dimension sizes or the number of measurements tend to infinity. However, many applications in modern data science operate in finite-sample regimes, where the number of observations is limited relative to the model dimension. In such settings, estimators may not attain the CRLB [11], e.g., when Fisher information is concentrated in fine-scale model features that cannot be detected from a small number of samples [12]. These observations do not diminish the value of the classical CRLB; rather, they highlight the need to adapt its principles to finite-sample scenarios. Our work contributes toward this goal by developing a non-asymptotic framework for studying efficiency that preserves the spirit of CRLB analysis while making it applicable to modern high-dimensional tensor models.

Beyond adapting the CRLB to finite-sample settings, a growing body of work has examined statistical inference directly from a non-asymptotic perspective. In [12], the authors propose a framework for extending CRLB-style efficiency guarantees to finite samples and show, through concrete examples, that no estimator can achieve the CRLB in general. They address this by smoothing the underlying density and introducing a relaxed notion of efficiency that is attainable with finitely many observations. Other authors consider the accuracy of estimators under finite samples, but not in the context of the CRLB. In [13], under smoothness conditions on the log-likelihood function, the authors show a concentration inequality of the maximum likelihood estimator from its expectation. The authors in [14] consider finite-sample complexity needed to estimate a distribution’s mean to a certain accuracy and probability of success. They present an estimator that works for any distribution and matches the optimal sample complexity achieved in the Gaussian case. The work in [15] is very similar to [14], they show that a wide class of distributions allow for mean estimators that are optimal (near-Gaussian) under finite samples. The work in [16, 17] derive concentration properties of maximum likelihood and M -estimators, under finite samples and misspecified parametric assumptions.

2 Main Results

In this section we present simplified versions of our theoretical results. We begin with background on the CRLB and minimax lower bounds in Sections 2.1 and 2.2. Afterwards, Section 2.3 presents our tensor decomposition-based inference problem and results.

2.1 Cramér–Rao Lower Bound

This section presents a multivariate version of the CRLB, the definition of an efficient estimator, and a result that simplifies CRLB-based analysis. To elaborate, suppose $\mathbf{x} \in \mathbb{R}^n$ is a random vector containing the sample where $f(\mathbf{x}; \boldsymbol{\theta})$ denotes its probability density function (PDF) with $\boldsymbol{\theta} \in \mathbb{R}^p$ encompassing the distributional parameters we wish to estimate. A simplified version of our CRLB is the following:

Theorem 1. Consider an unbiased estimator $\hat{\boldsymbol{\theta}}(\mathbf{x})$ of $\boldsymbol{\theta}$ (i.e., $\mathbb{E}\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$), where the PDF satisfies the following interchanging of integral and derivative

$$\frac{\partial}{\partial \boldsymbol{\theta}} \int_{\mathbb{R}^n} [\hat{\boldsymbol{\theta}}(\mathbf{x}) - \boldsymbol{\theta}] f(\mathbf{x}; \boldsymbol{\theta}) d\mathbf{x} = \int_{\mathbb{R}^n} \frac{\partial}{\partial \boldsymbol{\theta}} ([\hat{\boldsymbol{\theta}}(\mathbf{x}) - \boldsymbol{\theta}] f(\mathbf{x}; \boldsymbol{\theta})) d\mathbf{x}.$$

Define the Fisher information matrix (FIM) as

$$\mathcal{I}_{\boldsymbol{\theta}} := \mathbb{E} \left[\left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{x}; \boldsymbol{\theta}) \right) \left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{x}; \boldsymbol{\theta}) \right)^\top \right].$$

Then

$$\mathcal{I}_{\boldsymbol{\theta}}^\dagger \preceq \text{Var}(\hat{\boldsymbol{\theta}}),$$

where

$$\text{Var}(\hat{\boldsymbol{\theta}}) := \mathbb{E} \left[(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})^\top \right]$$

is the covariance matrix, $A \preceq B$ means that $B - A$ is positive semi-definite (PSD), and the \dagger superscript denotes the pseudo-inverse.

Theorem 1 follows from a more general result that additionally considers biased estimators. For ease of exposition, we focus on this CRLB for now but refer the reader to Section 6 for the general result and its proof. This fundamental result provides a manner to gauge the variability of an estimator. The approach considers a PSD-based hierarchy and labels an estimator as *efficient* if it achieves the “lowest” possible covariance:

Definition 1. Under the conditions of Theorem 1, $\hat{\boldsymbol{\theta}}$ is said to be an *efficient* estimator if

$$\mathcal{I}_{\boldsymbol{\theta}}^\dagger = \text{Var}(\hat{\boldsymbol{\theta}}). \tag{1}$$

Establishing (1) for an estimator is a major concern in statistical inference, to assure an estimation procedure is optimal in the sense of error variance. This is a challenging task in multivariate problems where computations of the FIM and its pseudo-inverse (FIM-PI) are cumbersome. Furthermore, many modern inference procedures do not provide a straightforward manner to compute an estimator in closed form, which makes it difficult to compute the covariance matrix. Finally, as noted in Section 1.1, (1) may not hold in non-asymptotic scenarios [12, 11]. This last observation stresses that CRLB-based analysis may not be adequate for finite-sample settings, which are most informative in practice.

To alleviate non-asymptotic CRLB-based analysis, our following result simplifies the establishment of estimator efficiency:

Theorem 2. Under the conditions of Theorem 1, $\hat{\boldsymbol{\theta}}$ is efficient if and only if

$$\text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^\dagger) = \mathbb{E} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2, \tag{2}$$

where $\text{tr}(\circ)$ is the matrix trace.

The proof is straightforward, we quickly present it before continuing.

Proof of Theorem 2. Since Theorem 1 follows, we have that $\text{Var}(\hat{\boldsymbol{\theta}}) - \mathcal{I}_{\boldsymbol{\theta}}^{\dagger}$ is a PSD matrix. We now apply a well known trace lower bound for PSD matrices (e.g., see Section 2 in [18]):

$$\text{if } 0 \preceq A, \text{ then } \|A\| \leq \text{tr}(A),$$

where $\|\circ\|$ is the operator norm. Therefore

$$\left\| \text{Var}(\hat{\boldsymbol{\theta}}) - \mathcal{I}_{\boldsymbol{\theta}}^{\dagger} \right\| \leq \text{tr} \left(\text{Var}(\hat{\boldsymbol{\theta}}) - \mathcal{I}_{\boldsymbol{\theta}}^{\dagger} \right) = \mathbb{E} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2 - \text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^{\dagger}),$$

and the result follows. \square

While simple, Theorem 2 seems to be unnoticed in the literature. This result severely simplifies the study of efficiency since we must now only consider scalar terms, the FIM-PI trace and the mean squared error $\mathbb{E} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2$ (MSE). The main advantage is that the MSE is well studied for outputs of optimization programs under finite-samples, even when closed form solutions are not available. We propose to use (2) to gauge how close an estimator is to achieving the CRLB. Our main argument is that this approach alleviates non-asymptotic CRLB-based analysis. To this end, we introduce the more flexible concept of *near-efficient* estimators:

Definition 2. *Given a random observation $\mathbf{x} \sim f(\boldsymbol{\theta})$, let $\tilde{\boldsymbol{\theta}}$ be an estimator of $\boldsymbol{\theta}$. We say that $\tilde{\boldsymbol{\theta}}$ is a **near-efficient** estimator if the terms*

$$\text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^{\dagger}) \quad \text{and} \quad \text{Var}(\tilde{\boldsymbol{\theta}})$$

match up to absolute constants and logarithmic terms.

Notice that Definition 2 does not require $\tilde{\boldsymbol{\theta}}$ to be unbiased, which differs from the unbiasedness assumption in Theorem 1. Intuitively, demonstrating that an estimator is near-efficient under this definition shows that its variance remains comparable to that of the best unbiased estimator. Although this introduces a mild mismatch in assumptions, the concept remains highly informative in practice, as it enables CRLB-based assessments in realistic finite-sample settings where unbiasedness is rarely attainable. Our approach will consist of deriving error bounds to compare the MSE to the FIM-PI trace, to show that Definition 2 is satisfied. This will be particularly useful in multiway analysis, where only asymptotic results seem to be available in the literature.

2.2 Minimax Lower Bounds

We now provide background on minimax lower bounds and explain their relation to the CRLB framework discussed previously. Both approaches characterize fundamental limits of statistical estimation, but they differ in perspective and scope.

Minimax theory studies the best possible MSE that any estimator can achieve over a class of parameters S . Formally, one considers

$$\inf_{\hat{\boldsymbol{\theta}}} \sup_{\boldsymbol{\theta} \in S} \mathbb{E} \|\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2. \quad (3)$$

A lower bound on (3) implies that no estimator can uniformly attain smaller MSE over the entire class S . Such bounds typically depend on the sample size, the ambient dimension, constraints such as sparsity or tensor rank, or properties of the noise model. Importantly, minimax bounds are *information-theoretic*: they apply to all estimators, not just a particular algorithm.

Both minimax theory and the CRLB quantify optimality, but in complementary ways:

- The CRLB is a *local* bound: it applies to a specific parameter θ and typically requires computing the FIM and the covariance of a chosen estimator.
- Minimax bounds are *global*: they measure the best achievable performance uniformly over a class S and do not depend on any particular estimator.
- The CRLB yields a matrix inequality, whereas minimax bounds typically yield scalar lower bounds such as (3).
- Classical CRLB results are primarily asymptotic, whereas minimax theory often provides non-asymptotic error rates.

Thus, while the CRLB characterizes the smallest variance achievable by an estimator at a fixed parameter, minimax bounds characterize the smallest possible worst-case error across all estimators and all parameters in S .

The choice between CRLB and minimax analysis depends on the application. If the goal is to evaluate a specific estimator and quantify how close it is to an information-theoretic limit at a fixed θ , the CRLB is the appropriate tool. However, CRLB analysis often requires detailed knowledge of the estimator and may become analytically intractable for large, structured models. Minimax bounds, in contrast, provide broad performance guarantees and are easier to generalize, although they may offer less detailed insight at specific parameter points. In this sense, CRLB and minimax approaches should be viewed as complementary rather than competing frameworks for understanding fundamental limits of structured statistical inference.

2.3 Near-Efficient Multiway Inference

We now describe our multiway inference problem, estimation method, and main results in the rank-one and general CP rank case. We consider N -way Poisson observations from an entry-wise independent parameter model:

$$\mathbf{X} \sim \text{Poisson}(\mathcal{M}).$$

Here, $\mathbf{X} \in \mathbb{N}^{I \times \dots \times I}$ and $\mathcal{M} \in \mathbb{R}_+^{I \times \dots \times I}$ are N -way tensors where \mathbb{R}_+ denotes the set of positive real numbers. The entry of \mathbf{X} at the multi-index \mathbf{i} is an independent random variable (denoted as $x_{\mathbf{i}}$) such that for any $k \in \mathbb{N}$

$$\mathbb{P}(x_{\mathbf{i}} = k) = \frac{m_{\mathbf{i}}^k e^{-m_{\mathbf{i}}}}{k!},$$

where $m_{\mathbf{i}}$ denotes the respective entry of \mathcal{M} . The array \mathbf{X} consists of observed count data and \mathcal{M} is the tensor of Poisson parameters we wish to estimate, either as an array (parametric inference) or in decomposed form (multiway analysis).

Before proceeding, it is useful to distinguish these two perspectives. Parametric inference refers to estimating the parameter tensor \mathcal{M} , i.e., the $I \times \dots \times I$ independent Poisson parameters, a well-defined statistical task for which minimax lower bounds provide sharp benchmarks. In contrast, multiway analysis aims to recover the canonical polyadic (CP) decomposition factors of \mathcal{M} , which is typically much more challenging and may be ill-posed in larger-rank settings due to degeneracy and identifiability issues. Our results highlight this distinction: in the rank-one case, multiway analysis achieves near-efficiency relative to the CRLB, while in the general rank case, CP factor estimation falls short of the CRLB even though parametric inference of \mathcal{M} remains nearly minimax optimal.

2.3.1 Rank-one Case: Near-Efficient Multiway Analysis

We first assume \mathcal{M} is a rank-one tensor

$$\mathcal{M} = \mathbf{u}^{(1)} \circ \dots \circ \mathbf{u}^{(N)},$$

where $\{\mathbf{u}^{(n)}\}_{n=1}^N \subset \mathbb{R}^I$ are the factors and \circ denotes the vector outer product. Entry-wise, this means that at the multi-index $\mathbf{i} = (i_1, \dots, i_N)$

$$m_{\mathbf{i}} = u_{i_1}^{(1)} u_{i_2}^{(2)} \dots u_{i_N}^{(N)}$$

where $u_{i_n}^{(n)}$ is the respective entry of $\mathbf{u}^{(n)}$. We remark that most tensor decomposition choices agree in the rank-one case so the result of this section applies beyond the CP decomposition.

Given \mathcal{X} , our goal is to estimate factors of the Poisson parameter tensor $\{\mathbf{u}^{(n)}\}_{n=1}^N$, i.e., multiway analysis. Throughout, we denote $\boldsymbol{\theta} \in \mathbb{R}^{NI}$ as the concatenation of the $\mathbf{u}^{(n)}$'s. We will produce an estimator of $\boldsymbol{\theta}$ whose variance roughly matches the respective CRLB. We will consider an estimator that exploits the rank-one Poisson model and that the factor entries lie in some interval $[\beta, \alpha]$ with $\beta > 0$, by searching for estimators in

$$\mathcal{S}_1(\beta, \alpha) := \left\{ \mathcal{J} = \mathbf{w}^{(1)} \circ \dots \circ \mathbf{w}^{(N)} \mid \{\mathbf{w}^{(n)}\}_{n=1}^N \subset [\beta, \alpha]^I \right\}. \quad (4)$$

Our approach will output an estimator of \mathcal{M} as

$$\widetilde{\mathcal{M}}_1 := \operatorname{argmax}_{\mathcal{J} \in \mathcal{S}_1} \mathcal{L}_{\mathcal{X}}(\mathcal{J}) := \operatorname{argmax}_{\mathcal{J} \in \mathcal{S}_1} \sum_{\mathbf{i}} (x_{\mathbf{i}} + 1) \log(t_{\mathbf{i}} + 1) - (t_{\mathbf{i}} + 1), \quad (5)$$

where we write $\mathcal{S}_1(\beta, \alpha)$ as \mathcal{S}_1 for compactness. This method is very similar to standard Poisson max-likelihood [19, 20], but differs slightly due to the “+1” terms that appear in the objective. We refer to this estimation method as *shifted-Poisson* regression. The +1 shifts advantageously avoid zero counts in order to simplify our analysis.

The estimator $\widetilde{\mathcal{M}}_1$ is output in rank-one format, with factors $\{\widetilde{\mathbf{u}}^{(n)}\}_{n=1}^N$ and we let $\widetilde{\boldsymbol{\theta}}$ denote the estimate factors concatenated as a vector. We now present a simplified theorem that summarizes our rank-one results. The following statement treats terms that do not involve the size of the ambient dimensions, I , as constants for ease of exposition.

Theorem 3 (rank-one multiway analysis). *Let $\beta > 0$ and assume that $\mathcal{M} \in \mathcal{S}_1$. Under our rank-one Poisson observation setting, the FIM-PI satisfies*

$$\operatorname{tr}(\mathcal{I}_{\boldsymbol{\theta}}^\dagger) = \mathcal{O}\left(\frac{1}{I^{N-2}}\right).$$

Furthermore, consider the factor estimation procedure based on (5) with the factors of \mathcal{M} and $\widetilde{\mathcal{M}}_1$ scaled such that $\|\mathbf{u}^{(1)}\|_2 = \dots = \|\mathbf{u}^{(N)}\|_2$ and $\|\widetilde{\mathbf{u}}^{(1)}\|_2 = \dots = \|\widetilde{\mathbf{u}}^{(N)}\|_2$. If I is large enough (see (11)), the factor estimator obeys

$$\mathbb{E}\|\widetilde{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2 = \mathcal{O}\left(\frac{\log^2(I)}{I^{N-2}}\right).$$

The proof combines Theorem 6 and Corollary 2 that provide explicit bounds for $\operatorname{tr}(\mathcal{I}_{\boldsymbol{\theta}}^\dagger)$ and the MSE respectively. We refer the reader to Sections 4 and 5 for the full statements and proofs.

Theorem 3 states that up to absolute constants and logarithmic terms, the trace of the FIM-PI and the MSE of our estimator match. According to Definition 2, we have shown that our estimator is near-efficient. The result shows that our method is nearly as efficient as the best possible unbiased estimation method. However, as a corollary, we can establish asymptotic efficiency and unbiasedness.

Corollary 1. *Under the setting of Theorem 3, $\tilde{\boldsymbol{\theta}}$ is asymptotically unbiased and efficient.*

The corollary follows by allowing I or N to be arbitrarily large, which shows that both $\text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^\dagger)$ and $\mathbb{E}\|\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2$ approach zero. In the limit, we must have $\mathbb{E}\tilde{\boldsymbol{\theta}} = \boldsymbol{\theta}$ and both terms match so that Theorem 2 is satisfied. This corollary demonstrates the flexibility of our approach, which further allows us to establish asymptotic results.

2.3.2 Larger Rank Case: Near-Minimax Optimal Parametric Inference

We further consider the general case where \mathcal{M} is allowed to have rank larger than one. In contrast to the rank-one case, the larger rank scenario presents many complications when considering the CP decomposition. In this setting, our analysis cannot guarantee near-efficient CP factor estimation. Section 3 will explore this situation numerically, illustrating that rank-2 and rank-5 CP decomposition estimation may not achieve the CRLB.

However, Poisson parameter inference can be shown to be near minimax optimal via our shifted estimator. In other words, we now consider estimating the entire array \mathcal{M} rather than targeting its CP factors as in the rank-one case. We will provide an estimator $\tilde{\mathcal{M}}$ and study its accuracy with an MSE bound, which we then gauge in optimality via minimax analysis rather than the CRLB.

To elaborate, in the rank- R scenario, we now consider the search space

$$\mathcal{S}_R(\tilde{\beta}, \tilde{\alpha}) := \left\{ \mathcal{J} = \sum_{r=1}^R \mathbf{w}_r^{(1)} \circ \dots \circ \mathbf{w}_r^{(N)} \mid \mathbf{w}_r^{(n)} \in \mathbb{R}^I, \tilde{\beta} \leq t_{\mathbf{i}} \leq \tilde{\alpha} \right\},$$

and produce the estimator

$$\tilde{\mathcal{M}} := \underset{\mathcal{J} \in \mathcal{S}_R}{\text{argmax}} \mathcal{L}_{\mathcal{X}}(\mathcal{J}), \quad (6)$$

where we write $\mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ as \mathcal{S}_R for compactness. Notice that (6) extends (5) to search for rank- R tensors, but imposes bounded Poisson parameters rather than factor entry bounds. Our main result in this setting is the following:

Theorem 4 (Rank- R parametric inference). *Let $\tilde{\beta} > 0$, assume that $\mathcal{M} \in \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ and let $\tilde{\mathcal{M}}$ be given as in (6). There exists absolute constants $C_1, C_2 > 0$ such that for any $\delta > 0$ we have*

$$\|\tilde{\mathcal{M}} - \mathcal{M}\|_F^2 \leq \frac{C_2(\tilde{\alpha} + 1)^4(N + \delta)^2 \log^2(I)NIR^2}{(\tilde{\beta} + 1)^2} \quad (7)$$

with probability exceeding $1 - I^{-\delta} - \exp(-C_1NI \log(N))$.

In contrast to the rank-one case, which bounds the CP factors $\boldsymbol{\theta}$, the rank greater than one case bounds the Poisson parameters. To gauge the optimality of this result we deviate from CRLB analysis and instead present a minimax lower bound.

Theorem 5 (Minimax lower bound). *Under the Poisson observation model, $\mathbf{X} \sim \text{Poisson}(\mathcal{M})$, let $\inf_{\widehat{\mathcal{M}}} \mathcal{M}$ denote the infimum over all estimators $\widehat{\mathcal{M}} \in \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ based on the count tensor observation \mathbf{X} . Assume that $R \leq I$, $IR > 16$, and that*

$$\left(\frac{IR}{16} - 1\right) \frac{\tilde{\beta} \log 2}{(\tilde{\alpha} - \tilde{\beta})^2} \leq I^N. \quad (8)$$

Then,

$$\inf_{\widehat{\mathcal{M}}} \sup_{\mathcal{M} \in \mathcal{S}_R} \mathbb{E} \|\widehat{\mathcal{M}} - \mathcal{M}\|_F^2 \geq \frac{\tilde{\beta} \log 2}{128} \left(\frac{IR}{16} - 1\right). \quad (9)$$

The bound (9) illustrates the best derivable MSE rate for parameter tensors in \mathcal{S}_R . The result shows that “worst-case” Poisson parameters (achieving the supremum) will suffer an MSE of the order IR , and therefore one cannot expect to improve upon this rate for general tensors. Comparing (9) to the MSE upper bound (7), $\sim IR^2 \log^2(I)$, we see that our shifted-Poisson estimator is optimal up to a rank factor and a logarithmic term. Theorem 4 improves on previous CP-based bounds from the literature in terms of rank dependence; see Section 2.4 for further discussion.

2.4 Discussion

As in [12, 11], our framework illustrates an interesting trade-off where non-asymptotic analysis comes at the loss of strict efficiency. Our approach introduces the more flexible concept of near-efficiency which can be achieved under limited observations. In contrast to other work, the methodology provides a simplified approach to CRLB-based analysis under finite samples while not deviating from the original definition of efficiency. The approach simplifies the study of estimators derived from constrained optimization programs, which typically do not involve straightforward and closed-form solutions and its utility goes beyond multiway analysis.

In the literature of tensor decompositions, our work is novel in its attempt to express error bounds in terms of the CP factors $\boldsymbol{\theta}$ (multiway analysis, as in Theorem 3). In contrast, most results focus on error bounds involving the parameter tensor \mathcal{M} (i.e., parametric inference). In this context, Theorem 4 improves upon the literature in terms of CP rank dependency. To the best of our knowledge, all results in the literature exhibit error terms with polynomial dependence on the CP rank or largest array in terms of N (e.g., [21] with $\text{MSE} \sim IR^{N-1}$). This gap is also apparent in cases with missing information [22, 23, 24, 25, 26, 27] (e.g., tensor completion). In contrast, Theorem 4 derives the rate $\text{MSE} \sim IR^2 \log^2(I)$ which removes the exponential dependence on the order.

We note that work in the literature has been able to provide optimal sample complexity results for other factorization choices: Tucker decomposition [28, 29], higher order singular value decomposition, tensor train, and hierarchical Tucker format [30]. This highlights that the CP decomposition lacks theoretical understanding and CP-based analysis is more challenging than alternative multiway array decomposition choices. One of our contributions is to improve nuclear norm-based bounds in terms of the CP rank (see Theorem 10). This leads to our improved error rates and may be of interest on its own.

3 Numerical Experiments

This section presents numerical experiments that validate Theorem 3 in the rank-one case and illustrate why its conclusions do not extend to rank larger than one settings. All experiments were

implemented in MATLAB using the Tensor Toolbox [31] and its generalized CP (GCP) optimization framework. The purpose of these experiments is to empirically evaluate the correspondence between the mean squared error (MSE) and the trace of the FIM-PI, $\text{tr}(\mathcal{I}_\theta^\dagger)$, under the shifted-Poisson regression model to achieve multiway analysis.

3.1 Rank-one Case

We construct a Poisson parameter tensor \mathcal{M} of order N and dimensions I with entries constrained to the interval $[\beta, \alpha]$. The tensor is formed by taking the outer product of N nonnegative factor vectors, each generated randomly with entries uniformly distributed in the interval $[\beta^{1/N}, \alpha^{1/N}]$. Thus, $\mathcal{M} = \mathbf{u}^{(1)} \circ \dots \circ \mathbf{u}^{(N)}$, ensuring that all tensor entries lie within the prescribed bounds.

Given \mathcal{M} , we generate Poisson-distributed observations $\mathcal{Y} \sim \text{Poisson}(\mathcal{M})$ using `poissrnd`, and estimate the tensor via shifted-Poisson regression:

$$\widetilde{\mathcal{M}} := \underset{\mathcal{J} \in \widetilde{\mathcal{S}}_1}{\text{argmax}} \sum_{\mathbf{i}} (x_{\mathbf{i}} + 1) \log(t_{\mathbf{i}} + 1) - (t_{\mathbf{i}} + 1). \quad (10)$$

where

$$\widetilde{\mathcal{S}}_1 := \left\{ \mathcal{J} = \mathbf{w}^{(1)} \circ \dots \circ \mathbf{w}^{(N)} \mid \{\mathbf{w}^{(n)}\}_{n=1}^N \subset (0, \infty)^I \right\}.$$

This program is very similar to (5), but only imposes a non-negativity constraint on the factor entries. We solve (10) using the `gcp_opt` solver, which outputs our estimator in rank-one format $\widetilde{\boldsymbol{\theta}} = (\widetilde{\mathbf{u}}^{(1)}, \dots, \widetilde{\mathbf{u}}^{(N)})$. The optimization is performed using the `lbfgsb` quasi-Newton solver with nonnegativity constraints. We scale the factors such that $\|\mathbf{u}^{(1)}\|_2 = \dots = \|\mathbf{u}^{(N)}\|_2$ and $\|\widetilde{\mathbf{u}}^{(1)}\|_2 = \dots = \|\widetilde{\mathbf{u}}^{(N)}\|_2$, and then compute the error $\|\boldsymbol{\theta} - \widetilde{\boldsymbol{\theta}}\|_2^2$.

To compute $\text{tr}(\mathcal{I}_\theta^\dagger)$, we use Lemma 4 in [32] that provides the structure of our Poisson Fisher Information matrix. The true factors of \mathcal{M} are further normalized to have equal ℓ_1 -norm (i.e., equal sums across modes). This normalization is to agree with the proof of Theorem 6 that simplifies calculations, but we note that this choice does not seem to severely affect the value of $\text{tr}(\mathcal{I}_\theta^\dagger)$ in contrast to other normalization options.

The quantities MSE and $\text{tr}(\mathcal{I}_\theta^\dagger)$ are averaged over 50 realizations of random parameter tensors with $\beta = 1$ and $\alpha = 2$. Each experiment is repeated for ambient dimension sizes $I = 50, 75$, and 100, and for tensor orders $N = 3$ and $N = 4$. Figure 1 reports the mean value with error bars corresponding to the minimum and maximum observed values across repetitions.

Figure 1 demonstrates that Theorem 3 holds in practice. In all scenarios, we observe near-efficiency since the estimated MSE and FIM-PI trace match. Furthermore, the order of these terms is in agreement with our main result as I and N vary. We note that the Poisson parameter bounds (β, α) do not seem to have a significant influence on the behavior of the plots.

3.2 Larger Rank Cases

The rank-2 and rank-5 cases follow the same structure, with the key modification that \mathcal{M} is generated as a sum of $R = 2$ or $R = 5$ rank-one components:

$$\mathcal{M} = \sum_{r=1}^R \mathbf{u}_r^{(1)} \circ \dots \circ \mathbf{u}_r^{(N)}$$

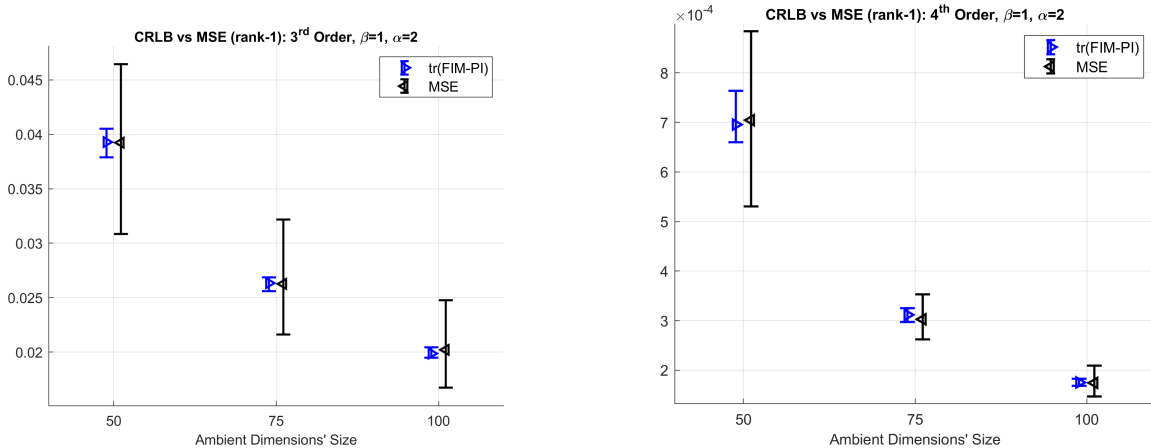


Figure 1: Rank-one experiments: dimension sizes (I) vs mean value of $\text{tr}(\mathcal{I}_{\theta}^{\dagger})$ and $\|\theta - \tilde{\theta}\|_2^2$, with error bars. Left plot shows the $N = 3$ dimensional case and right plot shows $N = 4$.

where each factor vector $\mathbf{u}_r^{(n)}$ has entries uniformly distributed in

$$[\beta^{1/N}/R^{1/N}, \alpha^{1/N}/R^{1/N}],$$

ensuring that the tensor entries remain within the prescribed range $[\beta, \alpha]$.

As in the rank-one setting, Poisson observations are generated using `poissrnd`, and the shifted-Poisson regression problem is solved via `gcp_opt` with a rank- R constraint and only imposing nonnegative factors. After estimation, the factor vectors of both the ground truth and the recovered tensor are normalized so that all vectors in a given rank-one component have equal ℓ_2 -norm, preserving the total component scale. The factors are stacked into a tall vector and the MSE is computed as before. For the Fisher information term, the factors are again normalized to have equal ℓ_1 -norm per component for consistency to the rank-one case. The FIM is computed using Theorem 3 in [32], pseudo-inverted, and its trace recorded.

Figures 2 and 3 report the empirical MSE and the quantity $\text{tr}(\mathcal{I}_{\theta}^{\dagger})$, averaged over 50 independent trials, as functions of the ambient dimension $I \in \{50, 75, 100, 125, 150\}$ for both $N = 3$ and $N = 4$. For each dimension, the mean, minimum, and maximum values across realizations are shown. Unlike the rank-one setting, these figures reveal a noticeable gap between the average MSE and the FIM trace term. This indicates that, even in low-rank tensor settings, the rank-constrained estimators do not generally achieve the CRLB. Nevertheless, the spread of the curves suggests that a substantial fraction of the estimators lie near the CRLB, particularly in the large scale $N = 4$ experiments.

To better understand this behavior and highlight the concentration of estimators near the CRLB, we include quantile-based visualizations. Figure 4 displays the median together with quantile bands: the rank-2 case uses a 0–90% band, while the rank-5 case uses a more restrictive 0–65% band due to the increased variability. These plots demonstrate that, although a formal proof of near-efficiency is not yet available, the numerical evidence strongly suggests that most realizations achieve performance close to the CRLB. Specifically, in the rank-2 experiments, only a small number of outliers fall outside the efficient range (beyond the 90th percentile), while in the rank-5 case a greater number of outliers are observed, yet approximately 32 out of 50 trials still lie within the 65% efficiency band. This indicates an interesting trade-off between tensor rank and the empirical attainability of CRLB-level efficiency.

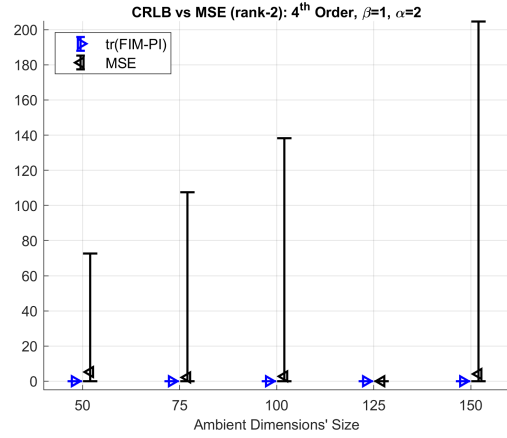



Figure 2: Rank-2 experiments: dimension sizes (I) vs mean value of $\text{tr}(\mathcal{I}_\theta^\dagger)$ and $\|\theta - \tilde{\theta}\|_2^2$, with error bars. Left plot shows the $N = 3$ dimensional case and right plot shows $N = 4$.


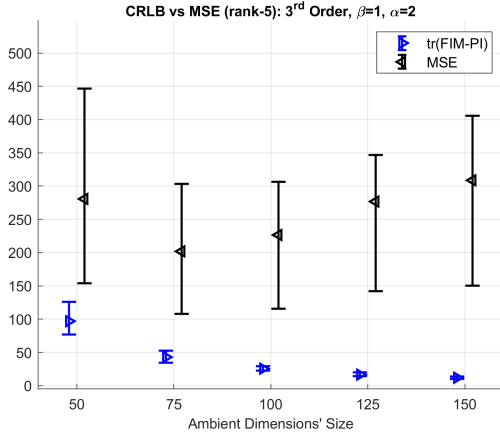


Figure 3: Rank-5 experiments: dimension sizes (I) vs mean value of $\text{tr}(\mathcal{I}_\theta^\dagger)$ and $\|\theta - \tilde{\theta}\|_2^2$, with error bars. Left plot shows the $N = 3$ dimensional case and right plot shows $N = 4$.

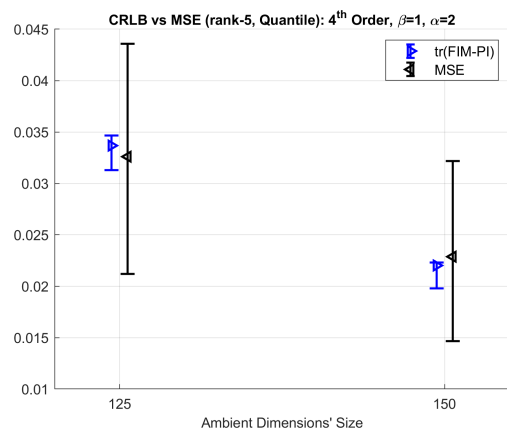
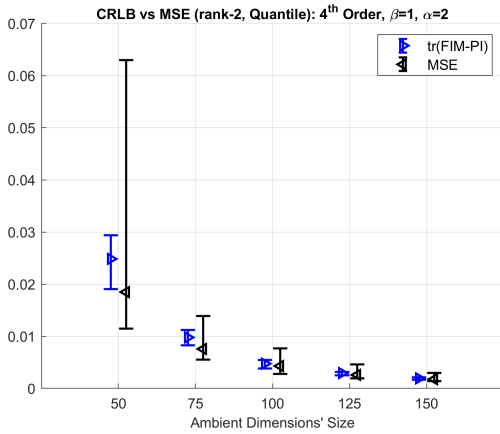


Figure 4: 4th order quantile plots: left plot shows the rank-2 case with median and a 0–90% quantile band and the right plot shows the rank-5 case with median and a 0–65% quantile band.

4 Fisher Information Bounds

In this section we focus on FIM computations, bounding the pseudo inverse's trace. Our main result in this section is the following:

Theorem 6. *Assume that $\mathbf{M} \in \mathcal{S}_1(\beta, \alpha)$ with $\beta > 0$, and let $\boldsymbol{\theta}$ denote its rank-one factors concatenated as a single vector (as in Section 2.3). Under the Poisson noise model*

$$\frac{\beta(I-1)N}{\alpha^{N-1}I^{N-1}} \leq \text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^{\dagger}) \leq \frac{\alpha(1+(I-1)N)}{\beta^{N-1}I^{N-1}}.$$

Ignoring terms that do not depend on I , this result establishes the showcased $\text{tr}(\mathcal{I}_{\boldsymbol{\theta}}^{\dagger}) = \mathcal{O}(I^{2-N})$ in Theorem 3.

We now dedicate the remainder of the section for the proof of Theorem 6. The proof will consist of bounding $\mathcal{F}_{\beta}^{\dagger} \preceq \mathcal{I}_{\boldsymbol{\theta}}^{\dagger} \preceq \mathcal{F}_{\alpha}^{\dagger}$, where \mathcal{F}_{β} and \mathcal{F}_{α} are PSD matrices that are easier to analyze. The result will then follow by computing the trace of these matrices.

Proof of Theorem 6. Without loss of generality, we scale the factors so that

$$\lambda := \|\mathbf{u}^{(1)}\|_1 = \|\mathbf{u}^{(2)}\|_1 = \dots = \|\mathbf{u}^{(N)}\|_1$$

which enforces equal factor sums (the factors are non-negative since $\beta > 0$). Using Lemma 4 in [32], the FIM of $\boldsymbol{\theta}$ is given in terms of $I \times I$ blocks as

$$\mathcal{I}_{\boldsymbol{\theta}} = \lambda^{N-1} \begin{bmatrix} \text{diag}(\mathbf{u}^{(1)})^{-1} & \lambda^{-1}\mathbf{1}\mathbf{1}^* & \dots & \lambda^{-1}\mathbf{1}\mathbf{1}^* \\ \lambda^{-1}\mathbf{1}\mathbf{1}^* & \text{diag}(\mathbf{u}^{(2)})^{-1} & \dots & \lambda^{-1}\mathbf{1}\mathbf{1}^* \\ \vdots & \vdots & \ddots & \vdots \\ \lambda^{-1}\mathbf{1}\mathbf{1}^* & \lambda^{-1}\mathbf{1}\mathbf{1}^* & \dots & \text{diag}(\mathbf{u}^{(N)})^{-1} \end{bmatrix} \in \mathbb{R}^{NI \times NI},$$

where $\mathbf{1} \in \{1\}^I$ is the all-ones vector and $\text{diag}(\mathbf{x}) \in \mathbb{R}^{I \times I}$ is a diagonal matrix with $\mathbf{x} \in \mathbb{R}^I$ in the diagonal. Furthermore, Theorem 4 in [32] states that $\mathcal{I}_{\boldsymbol{\theta}}$ has rank $(I-1)N+1$ for any $\boldsymbol{\theta}$ with positive entries.

Let $\mathbf{v} \in \mathbb{R}^{NI}$ be any vector. For $n \in \{1, \dots, N\}$, let $C_n = \{k \in \mathbb{N} \mid (n-1)I+1 \leq k \leq nI\}$. For fixed n , any entry $k \in C_n$ of $\mathcal{I}_{\boldsymbol{\theta}}\mathbf{v}$ is given as

$$\frac{(\mathcal{I}_{\boldsymbol{\theta}}\mathbf{v})_k}{\lambda^{N-1}} = \frac{v_k}{u_k^{(n)}} + \sum_{q \in [N]/\{n\}} \lambda^{-1} \sum_{\ell \in C_q} v_{\ell},$$

where $\bar{k} = k - (n-1)I$. Therefore

$$\begin{aligned} \frac{\mathbf{v}^* \mathcal{I}_{\boldsymbol{\theta}} \mathbf{v}}{\lambda^{N-1}} &= \sum_{n=1}^N \sum_{k \in C_n} \frac{v_k (\mathcal{I}_{\boldsymbol{\theta}} \mathbf{v})_k}{\lambda^{N-1}} = \sum_{n=1}^N \sum_{k \in C_n} \left(\frac{v_k^2}{u_k^{(n)}} + v_k \sum_{q \in [N]/\{n\}} \lambda^{-1} \sum_{\ell \in C_q} v_{\ell} \right) \\ &\leq \sum_{n=1}^N \sum_{k \in C_n} \left(\frac{v_k^2}{\beta} + v_k \sum_{q \in [N]/\{n\}} \lambda^{-1} \sum_{\ell \in C_q} v_{\ell} \right) := \frac{\mathbf{v}^* \mathcal{F}_{\beta} \mathbf{v}}{\lambda^{N-1}}, \end{aligned}$$

where we define

$$\mathcal{F}_{\beta} = \lambda^{N-1} \begin{bmatrix} \beta^{-1}\mathbf{I} & \lambda^{-1}\mathbf{1}\mathbf{1}^* & \dots & \lambda^{-1}\mathbf{1}\mathbf{1}^* \\ \lambda^{-1}\mathbf{1}\mathbf{1}^* & \beta^{-1}\mathbf{I} & \dots & \lambda^{-1}\mathbf{1}\mathbf{1}^* \\ \vdots & \vdots & \ddots & \vdots \\ \lambda^{-1}\mathbf{1}\mathbf{1}^* & \lambda^{-1}\mathbf{1}\mathbf{1}^* & \dots & \beta^{-1}\mathbf{I} \end{bmatrix}$$

and \mathbf{I} denotes the $I \times I$ identity matrix. We note that \mathcal{F}_β is also of rank $(I-1)N+1$ according to Theorem 4 in [32].

So far, we have shown that

$$\mathcal{I}_\theta \preceq \mathcal{F}_\beta.$$

Since these are of equal rank, Theorem 1 in [33] gives that

$$\mathcal{F}_\beta^\dagger \preceq \mathcal{I}_\theta^\dagger$$

and therefore $\text{tr}(\mathcal{F}_\beta^\dagger) \leq \text{tr}(\mathcal{I}_\theta^\dagger)$ can be used as a lower bound. The matrix \mathcal{F}_β is much simpler to study. Its eigenvectors are given as the columns of the following matrix (defined in terms of $I \times 1$ blocks):

$$\Gamma = \begin{bmatrix} \mathbf{1} & [\mathbf{e}_1 - \mathbf{e}_2] & [\mathbf{e}_1 + \mathbf{e}_2 - 2\mathbf{e}_3] & \dots & [\sum_{k=1}^{I-1} \mathbf{e}_k - (I-1)\mathbf{e}_I] & \mathbf{0} & \dots \\ \mathbf{1} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & [\mathbf{e}_1 - \mathbf{e}_2] & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \dots \\ \mathbf{1} & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{0} & \dots \end{bmatrix} \in \mathbb{R}^{NI \times (I-1)N+1},$$

where $\{\mathbf{e}_j\}_{j=1}^I \subset \mathbb{R}^I$ denote the canonical basis vectors and the pattern above repeats N times. The respective eigenvalues are

$$\Lambda_1 = \lambda^{N-1} \left(\frac{1}{\beta} + \frac{(N-1)I}{\lambda} \right) \quad \text{and} \quad \Lambda_r = \frac{\lambda^{N-1}}{\beta} \quad \text{for } r \geq 2.$$

Using the factor entry bound α gives

$$\Lambda_r \leq \frac{\alpha^{N-1} I^{N-1}}{\beta} \quad \text{for } r \geq 2$$

and we obtain

$$\text{tr}(\mathcal{I}_\theta^\dagger) \geq \text{tr}(\mathcal{F}_\beta^\dagger) = \sum_{r=1}^{(I-1)N+1} \frac{1}{\Lambda_r} > \sum_{r=2}^{(I-1)N+1} \frac{1}{\Lambda_r} = \frac{\beta(I-1)N}{\alpha^{N-1} I^{N-1}}.$$

The upperbound is very similar. The matrix \mathcal{F}_α , analogous to \mathcal{F}_β (replacing β 's with α 's), can be shown to satisfy $\text{tr}(\mathcal{F}_\alpha^\dagger) \geq \text{tr}(\mathcal{I}_\theta^\dagger)$. This matrix has the same eigenvectors, and similar eigenvalues that replace β 's with α 's. We can show that

$$\text{tr}(\mathcal{I}_\theta^\dagger) \leq \text{tr}(\mathcal{F}_\alpha^\dagger) = \lambda^{1-N} \left(\frac{1}{\alpha} + \frac{(N-1)I}{\lambda} \right)^{-1} + \frac{\alpha(I-1)N}{\lambda^{N-1}} \leq \frac{\alpha(1+(I-1)N)}{\beta^{N-1} I^{N-1}},$$

which finishes the proof. □

5 Mean Squared Error Bounds

We now analyze the accuracy of the estimator $\widetilde{\mathcal{M}}$ for parametric inference. This section will prove Theorem 4, and then modify the proof in the rank-one case to obtain the following corollary:

Corollary 2. Assume $\mathbf{M} \in \mathcal{S}_1(\beta, \alpha)$ and let $\widetilde{\mathbf{M}}_1$ be given as in (5) with rank-one factors $\{\widetilde{\mathbf{u}}^{(n)}\}_{n=1}^N$. Scale the factors such that $\|\mathbf{u}^{(1)}\|_2 = \dots = \|\mathbf{u}^{(N)}\|_2$ and $\|\widetilde{\mathbf{u}}^{(1)}\|_2 = \dots = \|\widetilde{\mathbf{u}}^{(N)}\|_2$, and let $\boldsymbol{\theta}, \widetilde{\boldsymbol{\theta}}$ denote the factors concatenated as a vectors (resp.). There exists an absolute constant $C_3 > 0$ such that if

$$IN^3 \log^2(I) \geq \frac{(\beta^N + 1)^2 \beta^{2(N-1)} (\alpha - \beta)^2}{C_3 (\alpha^N + 1)^4} \quad \text{and} \quad C_1 I \log(N) \geq \log(I), \quad (11)$$

where C_1 is as in Theorem 4, then

$$\mathbb{E} \|\boldsymbol{\theta} - \widetilde{\boldsymbol{\theta}}\|_2^2 \leq \frac{C_3 (\alpha^N + 1)^4 (N + \delta)^2 N^2 \log^2(I)}{(\beta^N + 1)^2 \beta^{2(N-1)} I^{N-2}}.$$

This corollary is used to obtain the simplified statement $\mathbb{E} \|\widetilde{\boldsymbol{\theta}} - \boldsymbol{\theta}\|_2^2 = \mathcal{O}\left(\frac{\log^2(I)}{I^{N-2}}\right)$ in Theorem 3, which treats terms that do not depend on I as a constant. We first provide the proof of Theorem 4, and then specify the changes needed to obtain Corollary 2.

Proof of Theorem 4. We will follow the proof of Theorem 1 in [21]. By second-order Taylor's theorem, for any $\mathcal{J} \in \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ there exists some $\widehat{\mathbf{M}} = \gamma \mathbf{M} + (1 - \gamma) \mathcal{J}$ with $\gamma \in [0, 1]$ such that

$$\mathcal{L}_{\mathcal{X}}(\mathcal{J}) = \mathcal{L}_{\mathcal{X}}(\mathbf{M}) + \langle \mathbf{S}_{\mathcal{X}}(\mathbf{M}), \mathcal{J} - \mathbf{M} \rangle + \frac{1}{2} \text{vec}(\mathcal{J} - \mathbf{M})^* H_{\mathcal{X}}(\widehat{\mathbf{M}}) \text{vec}(\mathcal{J} - \mathbf{M}). \quad (12)$$

Above, $\mathbf{S}_{\mathcal{X}}(\mathbf{M}) \in \mathbb{R}^{I \times \dots \times I}$ is the collection of the score functions evaluated at \mathbf{M} whose entry in the multi-index \mathbf{i} is given as

$$\frac{\partial \mathcal{L}_{\mathcal{X}}(\mathbf{M})}{\partial m_{\mathbf{i}}} = \frac{x_{\mathbf{i}} + 1}{m_{\mathbf{i}} + 1} - 1.$$

$H_{\mathcal{X}}(\widehat{\mathbf{M}}) \in \mathbb{R}^{I^N \times I^N}$ is the collection of the Hessian functions evaluated at $\widehat{\mathbf{M}}$. Identifying each $k \in \{1, \dots, I^N\}$ with a unique multi-index $\mathbf{i}_k \in \{1, \dots, I\} \times \dots \times \{1, \dots, I\}$, the entry (k, ℓ) of $H_{\mathcal{X}}(\widehat{\mathbf{M}})$ is given as

$$\frac{\partial^2 \mathcal{L}_{\mathcal{X}}(\widehat{\mathbf{M}})}{\partial \hat{m}_{\mathbf{i}_k} \partial \hat{m}_{\mathbf{i}_\ell}} = \begin{cases} -\frac{x_{\mathbf{i}} + 1}{(\hat{m}_{\mathbf{i}} + 1)^2} & \text{if } k = \ell \quad (\mathbf{i} = \mathbf{i}_k = \mathbf{i}_\ell) \\ 0 & \text{if } k \neq \ell. \end{cases}$$

We will bound certain terms in (12) and rearrange to obtain our result. Our shifted version of the Poisson loglikelihood gives that

$$\text{vec}(\mathcal{J} - \mathbf{M})^* H_{\mathcal{X}}(\widehat{\mathbf{M}}) \text{vec}(\mathcal{J} - \mathbf{M}) = - \sum_{\mathbf{i}} \frac{x_{\mathbf{i}} + 1}{(\hat{m}_{\mathbf{i}} + 1)^2} (t_{\mathbf{i}} - m_{\mathbf{i}})^2 \leq \frac{-1}{(\tilde{\alpha} + 1)^2} \|\mathcal{J} - \mathbf{M}\|_F^2. \quad (13)$$

To bound the linear term in (12), we use

$$\langle \mathbf{S}_{\mathcal{X}}(\mathbf{M}), \mathcal{J} - \mathbf{M} \rangle \leq \|\mathbf{S}_{\mathcal{X}}(\mathbf{M})\| \|\mathcal{J} - \mathbf{M}\|_*$$

where $\|\circ\|$ is the tensor spectral norm and $\|\circ\|_*$ is its dual norm (the tensor nuclear norm), see Lemma 1 in [21]. Using Theorem 10 in the next section we obtain

$$\|\mathcal{J} - \mathbf{M}\|_* \leq 2R \|\mathcal{J} - \mathbf{M}\|_F, \quad (14)$$

since $\mathcal{J} - \mathbf{M}$ is at most rank $2R$. To bound $\|\mathbf{S}_{\mathcal{X}}(\mathbf{M})\|$ we use Lemma 6 in [21]. Consider any $\delta > 0$, then the entries of $\mathbf{S}_{\mathcal{X}}(\mathbf{M})$ are independent, centered, and satisfy

$$\max_{\mathbf{i}} \left| \frac{\partial \mathcal{L}_{\mathcal{X}}(\mathbf{M})}{\partial m_{\mathbf{i}}} \right| \leq \max_{\mathbf{i}} \frac{(N + \delta) \log(I) + m_{\mathbf{i}} + 1}{m_{\mathbf{i}} + 1} - 1 \leq \frac{(N + \delta) \log(I)}{\tilde{\beta} + 1}.$$

The first inequality above holds with probability exceeding $1 - I^{-\delta}$ and can be shown using Lemma 1 in [34], which gives that $x_{\mathbf{i}} \geq m_{\mathbf{i}} + t$ with probability less than e^{-t} for $t \geq \tilde{\alpha}(e^2 - 3)$. If $(N + \delta) \log(I) \geq \tilde{\alpha}(e^2 - 3)$, we may choose $t = (N + \delta) \log(I)$ and obtain that $x_{\mathbf{i}} \geq m_{\mathbf{i}} + t$ for all \mathbf{i} with probability less than $I^N e^{-(N+\delta) \log(I)} = I^{-\delta}$ by a union bound.

Combining this bound with Lemma 6 in [21], there exists some absolute constants $c_1, c_2 > 0$ such that with probability at least $1 - I^{-\delta} - \exp(-c_1 NI \log(N))$

$$\|\mathcal{S}\mathbf{x}(\mathcal{M})\| \leq c_2 \frac{(N + \delta) \log(I)}{\tilde{\beta} + 1} \sqrt{NI}. \quad (15)$$

Finally, choosing $\mathcal{J} = \tilde{\mathcal{M}}$ in (12), since $\mathcal{L}_{\mathcal{X}}(\tilde{\mathcal{M}}) \geq \mathcal{L}_{\mathcal{X}}(\mathcal{M})$ if we apply (13), (14), (15), and rearrange we have shown that

$$\|\tilde{\mathcal{M}} - \mathcal{M}\|_F \leq \frac{4(\tilde{\alpha} + 1)^2 c_2 (N + \delta) \log(I) R \sqrt{NI}}{\tilde{\beta} + 1}, \quad (16)$$

holds with the prescribed probability. The main statement defines C_1, C_2 to absorb c_1, c_2 along with other absolute constants. \square

We now proceed to the proof of Corollary 2. The beginning of the proof is identical to the proof of Theorem 4, with a slight modification to obtain a bound analogous to (16). Lemma 1 from the next section will then translate the derived parameter tensor error (16) to an error on the CP factors.

Proof of Corollary 2. Mimicking the proof of Theorem 4, the main difference in the rank-one setting are the bounds imposed on the factors in the search space (4). If $\mathcal{J} \in \mathcal{S}_1(\beta, \alpha)$, notice that $\beta^N \leq t_{\mathbf{i}} \leq \alpha^N$. Setting $R = 1$, $\tilde{\beta} = \beta^N$, and $\tilde{\alpha} = \alpha^N$ in the definition of $\mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$, we can follow the proof of Theorem 4 to conclude

$$\|\tilde{\mathcal{M}}_1 - \mathcal{M}\|_F \leq \frac{C_2(\alpha^N + 1)^2 (N + \delta) \log(I) \sqrt{NI}}{\beta^N + 1} \quad (17)$$

with probability exceeding $1 - I^{-\delta} - \exp(-C_1 NI \log(N))$.

By Lemma 1, if we set $\|\mathbf{u}^{(1)}\|_2 = \dots = \|\mathbf{u}^{(N)}\|_2$ and $\|\tilde{\mathbf{u}}^{(1)}\|_2 = \dots = \|\tilde{\mathbf{u}}^{(N)}\|_2$ we can conclude

$$\|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}\|_2^2 = \sum_{n=1}^N \|\mathbf{u}^{(n)} - \tilde{\mathbf{u}}^{(n)}\|_2^2 \leq N \frac{C_2^2 (\alpha^N + 1)^4 (N + \delta)^2 N \log^2(I)}{(\beta^N + 1)^2 \beta^{2(N-1)} I^{N-2}} := \kappa_{\delta}$$

with probability exceeding $1 - I^{-\delta} - \exp(-C_1 NI \log(N))$. To bound the expected value, we use Lemma 2. Treat $y = \|\boldsymbol{\theta} - \tilde{\boldsymbol{\theta}}\|_2^2$ as a random variable, and notice that $y \in [0, NI(\alpha - \beta)^2]$. We have shown

$$\mathbb{P}(y > \kappa_{\delta}) \leq I^{-\delta} + \exp(-C_1 NI \log(N))$$

so that by Lemma 2 we will obtain $\mathbb{E}y \leq 2\kappa_{\delta}$ if

$$I^{-\delta} + \exp(-C_1 NI \log(N)) \leq \frac{\kappa_{\delta}}{NI(\alpha - \beta)^2}. \quad (18)$$

Under our assumptions, $C_1 NI \log(N) \geq \delta \log(I)$ holds with $\delta = N$ and we have that $\exp(-C_1 NI \log(N)) \leq I^{-\delta}$. Therefore (18) holds with $\delta = N$ if

$$I \geq \frac{2(\beta^N + 1)^2 \beta^{2(N-1)} (\alpha - \beta)^2}{C_2^2 (\alpha^N + 1)^4 (2N)^2 N \log^2(I)}.$$

The final statement defines $C_3 = 2C_2^2$. \square

6 CRLB Proof

In this section we state and prove a general version of Theorem 1, which allows for biased estimators. For the following result, we refer the reader to the setting and definitions of Section 2.1.

Theorem 7. *Consider an estimator $\hat{\boldsymbol{\theta}}(\mathbf{x})$ of $\boldsymbol{\theta}$ (not necessarily unbiased). Otherwise, under the conditions of Theorem 1,*

$$\left(\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \hat{\boldsymbol{\theta}}\right) \mathcal{I}_{\boldsymbol{\theta}}^{\dagger} \left(\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \hat{\boldsymbol{\theta}}\right)^{\top} \preceq \text{Var}(\hat{\boldsymbol{\theta}}),$$

where

$$\text{Var}(\hat{\boldsymbol{\theta}}) := \mathbb{E} \left[(\hat{\boldsymbol{\theta}} - \mathbb{E} \hat{\boldsymbol{\theta}}) (\hat{\boldsymbol{\theta}} - \mathbb{E} \hat{\boldsymbol{\theta}})^{\top} \right]$$

is the covariance matrix.

In contrast to Theorem 1, this general CRLB includes terms that consider the bias of $\hat{\boldsymbol{\theta}}$. In the unbiased case, notice that $\mathbb{E} \hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$ so that $\frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \hat{\boldsymbol{\theta}}$ is the $p \times p$ identity matrix and Theorem 1 follows. The main novelty of Theorem 7 is our the proof essentially shows that the CRLB is a consequence of the generalized Cauchy-Schwarz inequality.

Theorem 8 (Theorem 1 in [35], article C467). *Let $\mathbf{a} \in \mathbb{R}^{p_1}$ and $\mathbf{b} \in \mathbb{R}^{p_2}$ be random vectors satisfying $\mathbb{E} \|\mathbf{a}\|_2 < \infty$ and $\mathbb{E} \|\mathbf{b}\|_2 < \infty$. Then:*

$$\mathbb{E} \mathbf{a} \mathbf{b}^{\top} \left(\mathbb{E} \mathbf{b} \mathbf{b}^{\top} \right)^{\dagger} \mathbb{E} \mathbf{b} \mathbf{a}^{\top} \preceq \mathbb{E} (\mathbf{a} \mathbf{a}^{\top}).$$

The original matrix Cauchy-Schwarz inequality was shown in [36] and Theorem 8 extends the result to include the generalized inverse. We now proceed to the proof of our CRLB. The proof of Theorem 7 will use the matrix Cauchy-Schwarz inequality, followed by assumptions on the PDF.

Proof of Theorem 7. Apply Theorem 8 with $\mathbf{b} = \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{x}; \boldsymbol{\theta})$ and $\mathbf{a} = \hat{\boldsymbol{\theta}} - \mathbb{E} \hat{\boldsymbol{\theta}}$. Then:

1. $\mathbb{E} \mathbf{a} \mathbf{a}^{\top} = \text{Var}(\hat{\boldsymbol{\theta}}),$
2. $\mathbb{E} \mathbf{b} \mathbf{b}^{\top} = \mathcal{I}_{\boldsymbol{\theta}},$
3. $\mathbb{E} \mathbf{a} \mathbf{b}^{\top} = \mathbb{E} \left[(\hat{\boldsymbol{\theta}} - \mathbb{E} \hat{\boldsymbol{\theta}}) \left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{x}; \boldsymbol{\theta}) \right)^{\top} \right] = \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \hat{\boldsymbol{\theta}},$

where the last equality can be shown using the PDF assumptions from Theorem 1:

$$\begin{aligned}
0 &= \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E}[\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}}] \\
&= \frac{\partial}{\partial \boldsymbol{\theta}} \int_{\mathbb{R}^n} [\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] f(\boldsymbol{x}; \boldsymbol{\theta}) d\boldsymbol{x} \\
&= \int_{\mathbb{R}^n} \frac{\partial}{\partial \boldsymbol{\theta}} \left([\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] f(\boldsymbol{x}; \boldsymbol{\theta}) \right) d\boldsymbol{x} \\
&= \int_{\mathbb{R}^n} \left(\frac{\partial}{\partial \boldsymbol{\theta}} [\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] \right) f(\boldsymbol{x}; \boldsymbol{\theta}) d\boldsymbol{x} + \int_{\mathbb{R}^n} [\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] \left(\frac{\partial}{\partial \boldsymbol{\theta}} f(\boldsymbol{x}; \boldsymbol{\theta}) \right)^\top d\boldsymbol{x} \\
&= \int_{\mathbb{R}^n} \left(\frac{\partial}{\partial \boldsymbol{\theta}} [\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] \right) f(\boldsymbol{x}; \boldsymbol{\theta}) d\boldsymbol{x} + \int_{\mathbb{R}^n} [\hat{\boldsymbol{\theta}}(\boldsymbol{x}) - \mathbb{E}\hat{\boldsymbol{\theta}}] \left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}; \boldsymbol{\theta}) \right)^\top f(\boldsymbol{x}; \boldsymbol{\theta}) d\boldsymbol{x} \\
&= \mathbb{E} \left[\frac{\partial}{\partial \boldsymbol{\theta}} [\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}}] \right] + \mathbb{E} \left[(\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}}) \left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}; \boldsymbol{\theta}) \right)^\top \right] \\
&= \mathbb{E} \left(\frac{\partial}{\partial \boldsymbol{\theta}} \hat{\boldsymbol{\theta}} \right) - \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E}\hat{\boldsymbol{\theta}} + \mathbb{E} \left[(\hat{\boldsymbol{\theta}} - \mathbb{E}\hat{\boldsymbol{\theta}}) \left(\frac{\partial}{\partial \boldsymbol{\theta}} \log f(\boldsymbol{x}; \boldsymbol{\theta}) \right)^\top \right].
\end{aligned}$$

The final term above appears in bullet point 3. We note that, under the assumption of estimability, $\frac{\partial}{\partial \boldsymbol{\theta}} \hat{\boldsymbol{\theta}} = 0$ since $\hat{\boldsymbol{\theta}}$ cannot depend on $\boldsymbol{\theta}$. This proves the last equality in bullet 3 and finishes the proof. \square

7 Minimax Lower Bound

In this section, we prove Theorem 5. A standard technique for bounding the minimax risk is the Generalized Fano's method [37, Prop. 5.12]:

Theorem 9 (Generalized Fano's Method Minimax Lower Bound). *Given a finite subset of the parameter space $\mathcal{F} \subseteq \mathcal{S}_R$ and a semi-metric loss function L , assume that $|\mathcal{F}| = M_\delta$ and that there exists $\delta > 0$ such that*

$$\delta \leq \min_{\mathcal{M}^{(i)}, \mathcal{M}^{(j)} \in \mathcal{F}, i \neq j} L(\mathcal{M}^{(i)}, \mathcal{M}^{(j)}).$$

Let J denote a random subset of \mathcal{F} and let W denote the observed data. If $I(W, J)$ is the mutual information between W and a random variable given by a mixture of the distributions in \mathcal{F} with indices in J , we have that

$$\inf_{\widehat{\mathcal{M}}} \sup_{\mathcal{M} \in \mathcal{S}_R} \mathbb{E} \Phi \left(L(\widehat{\mathcal{M}}, \mathcal{M}) \right) \geq \Phi \left(\frac{\delta}{2} \right) \left(1 - \frac{I(W, J) + \log 2}{\log M_\delta} \right),$$

for any strictly increasing function $\Phi : [0, \infty) \mapsto [0, \infty)$.

We associate each parameter tensor $\mathcal{M}^{(i)} \in \mathcal{F}$ with the corresponding probability distribution $P^{(i)} \in \mathcal{P}_{\mathcal{F}}$, in this case, a joint Poisson distribution over the entries of the element. We omit the definition of mutual information here, but note that by [38, Eq. (13)]

$$I(W, J) \leq \frac{1}{\binom{M_\delta}{2}} \sum_{P^{(i)}, P^{(j)} \in \mathcal{P}_{\mathcal{F}}, i \neq j} D_{KL} \left(P^{(i)} \| P^{(j)} \right) \leq \max_{P^{(i)}, P^{(j)} \in \mathcal{P}_{\mathcal{F}}, i \neq j} D_{KL} \left(P^{(i)} \| P^{(j)} \right),$$

where D_{KL} denotes the Kullback-Leibler (KL) divergence. For tensors with independent Poisson entries we may write:

$$D_{KL}\left(P^{(i)}\|P^{(j)}\right) = \sum_{\mathbf{k}\in\{1,\dots,I\}^N} \left[m_{\mathbf{k}}^{(i)} \log \frac{m_{\mathbf{k}}^{(i)}}{m_{\mathbf{k}}^{(j)}} - m_{\mathbf{k}}^{(i)} + m_{\mathbf{k}}^{(j)} \right].$$

This observation enables us to write the minimax lower bound in terms of the KL divergence:

Corollary 3. *Under the assumptions of Theorem 9, if we further have $0 < \gamma < \infty$ such that*

$$\gamma \geq \max_{i \neq j; P^{(i)}, P^{(j)} \in \mathcal{P}_{\mathcal{F}}} D_{KL}\left(P^{(i)}\|P^{(j)}\right),$$

we may write

$$\inf_{\widehat{\mathcal{M}}} \sup_{\mathcal{M} \in \mathcal{S}_R} \mathbb{E} \Phi\left(L\left(\widehat{\mathcal{M}}, \mathcal{M}\right)\right) \geq \Phi\left(\frac{\delta}{2}\right) \left(1 - \frac{\gamma + \log 2}{\log M_{\delta}}\right).$$

As a point of interest, if we were to observe $n > 1$ realizations (samples) of the tensor, the decoupling property of the KL divergence would scale the KL divergence by n (see [37, Eq. 15.11(a)]), and we would replace γ by $n\gamma$ in Corollary 3.

To obtain Theorem 5, we will choose $\Phi(x) = x^2$ and the Frobenius norm of $\widehat{\mathcal{M}} - \mathcal{M}$ as the loss function. To establish our minimax bound, we will control δ so that

$$\log M_{\delta} \geq 2(\gamma + \log 2),$$

leading to

$$\left(1 - \frac{\gamma + \log 2}{\log M_{\delta}}\right) \geq \frac{1}{2},$$

and a minimax lower bound of $\frac{1}{2}\Phi(\delta/2) = \frac{1}{8}\delta^2$. We have three steps in our analysis:

1. Define a packing set \mathcal{F} with separation δ .
2. Compute the upper bound γ required in Corollary 3.
3. Select δ such that $\log M_{\delta} \geq 2(\gamma + \log 2)$.

The proof applies Lemmas 4 and 5, which we state afterwards. We refer the reader to Section 8.2 for the required lemmas and their proofs.

Proof of Theorem 5. We begin by applying Lemma 4 to create a packing set in $\mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$. We then have the following: for some constant $\epsilon \in (0, 1]$ (to be specified later), we have a packing set \mathcal{F} with $M_{\delta} = |\mathcal{F}| \geq 2^{IR/8}$ and spacing $\delta = \frac{\epsilon}{4}(\tilde{\alpha} - \tilde{\beta})\sqrt{I^N}$.

We next apply Lemma 5 to upper bound the KL divergence between the elements of the packing set. We conclude that since the spacing of elements in the packing set is bounded above by $\epsilon(\tilde{\alpha} - \tilde{\beta})\sqrt{I^N}$ (by Lemma 4), the KL divergence is bounded above by

$$\gamma = \epsilon^2 \frac{(\tilde{\alpha} - \tilde{\beta})^2}{\tilde{\beta}} I^N.$$

We now need to tune δ so that $\log M_{\delta} \geq 2(\gamma + \log 2)$, which we accomplish by adjusting the parameter ϵ . That is, we choose $\epsilon \in (0, 1]$ so that

$$\frac{IR \log 2}{8} \geq 2 \left(\epsilon^2 \frac{(\tilde{\alpha} - \tilde{\beta})^2}{\tilde{\beta}} I^N + \log 2 \right).$$

Equivalently,

$$\epsilon^2 \leq \frac{\left(\frac{IR}{16} - 1\right) \tilde{\beta} \log 2}{I^N (\tilde{\alpha} - \tilde{\beta})^2}.$$

For $\epsilon > 0$, we require $IR/16 - 1 > 0$, or $IR > 16$ —a slightly stronger condition than the $IR > 8$ required by the construction of the packing set in Lemma 4. We also require that $\epsilon^2 \leq 1$, which imposes condition (8) in the statement of the theorem.

It remains to compute the final minimax bound, $\delta^2/8$:

$$\begin{aligned} \frac{1}{2}\Phi(\delta/2) &= \frac{1}{8}\delta^2 \\ &= \frac{1}{8} \left(\frac{\epsilon^2}{16} (\tilde{\alpha} - \tilde{\beta})^2 I^N \right) \\ &= \frac{\log 2 \left(\frac{IR}{16} - 1\right)}{128} \frac{\tilde{\beta}}{I^N (\tilde{\alpha} - \tilde{\beta})^2} (\tilde{\alpha} - \tilde{\beta})^2 I^N \\ &= \frac{\tilde{\beta} \log 2}{128} \left(\frac{IR}{16} - 1 \right). \end{aligned}$$

□

8 Required Lemmas

In this section we provide some necessary lemmas. Subsection 8.1 states results needed for the MSE bounds in Section 5, while Subsection 8.2 focuses on lemmas needed for the minimax bound in Section 7. The proofs are provided in Subsection 8.3.

8.1 MSE Lemmas

The first lemma translates a bound on the tensor error to a bound on rank-one factor error.

Lemma 1. *Let \mathcal{M} and $\tilde{\mathcal{M}}$ be N -way rank-one tensors with*

$$\mathcal{M} = \mathbf{u}^{(1)} \circ \dots \circ \mathbf{u}^{(N)} \quad \text{and} \quad \tilde{\mathcal{M}} = \tilde{\mathbf{u}}^{(1)} \circ \dots \circ \tilde{\mathbf{u}}^{(N)},$$

where $\mathbf{u}^{(n)}, \tilde{\mathbf{u}}^{(n)} \in [\beta, \infty)^I$ for all $n \in [N]$. Assume $\|\mathbf{u}^{(1)}\|_2 = \|\mathbf{u}^{(2)}\|_2 = \dots = \|\mathbf{u}^{(N)}\|_2$ and $\|\tilde{\mathbf{u}}^{(1)}\|_2 = \|\tilde{\mathbf{u}}^{(2)}\|_2 = \dots = \|\tilde{\mathbf{u}}^{(N)}\|_2$.

If

$$\left\| \mathcal{M} - \tilde{\mathcal{M}} \right\|_F^2 \leq \epsilon$$

for some $\epsilon > 0$, then

$$\|\mathbf{u}^{(n)} - \tilde{\mathbf{u}}^{(n)}\|_2^2 \leq \frac{\epsilon}{\beta^{2(N-1)} I^{N-1}}$$

for all $n \in [N]$.

We next present a result bounding the nuclear norm of a CP rank R tensor in terms of the Frobenius norm. The following result improves upon previous tensor nuclear norm bounds in the literature [21, 39]. We present it as a theorem since it may be of interest on its own.

Theorem 10. *If $\mathcal{X} \in \mathbb{C}^{I \times \dots \times I}$ is a CP rank R tensor, then*

$$\|\mathcal{X}\|_* \leq R \|\mathcal{X}\|_F. \tag{19}$$

We end this section with a lemma that bounds the expected value based on a tail distribution bound. This lemma helps translate high probability error bounds to MSE bounds.

Lemma 2. *Let $y \in [0, L]$ be a continuous random variable. If for some δ , it holds that $\mathbb{P}(y \geq \delta) \leq \frac{\delta}{L}$ then $\mathbb{E}y \leq 2\delta$.*

8.2 Minimax Lemmas

First, we recall the classical Varshamov-Gilbert bound, taken from [21, Lemma 7].

Lemma 3 (Varshamov-Gilbert bound). *Let $\Omega = \{(w_1, w_2, \dots, w_m) : w_i \in \{0, 1\}\} \subseteq \mathbb{R}^m$, and take $m > 8$. Then, there exists a subset $\{\mathbf{w}^{(0)}, \mathbf{w}^{(1)}, \dots, \mathbf{w}^{(M)}\}$ of Ω such that $\mathbf{w}^{(0)}$ is the zero vector and for $0 \leq j < k \leq M$,*

$$\left\| \mathbf{w}^{(j)} - \mathbf{w}^{(k)} \right\|_0 \geq \frac{m}{8}.$$

We next use the Varshamov-Gilbert bound to construct a packing set $\mathcal{F} \subseteq \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ with $\tilde{\beta} = 0$ (no lower bound on the entries of the tensor). Our approach is a modification of [40, Lemma 9] and [21, Lemma 8].

Lemma 4 (Minimax Packing Set). *Assume that $R \leq I$ and $IR > 8$. For a constant $\epsilon \in (0, 1]$ and bounds $0 < \tilde{\beta} < \tilde{\alpha}$, there is a finite set $\mathcal{F} \subset \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$ such that:*

1. *The cardinality satisfies $|\mathcal{F}| \geq 2^{IR/8}$;*
2. *For any pair of distinct elements $\mathcal{M}^{(i)}, \mathcal{M}^{(j)} \in \mathcal{F}$,*

$$\left\| \mathcal{M}^{(i)} - \mathcal{M}^{(j)} \right\|_F \geq \frac{\epsilon}{4} (\tilde{\alpha} - \tilde{\beta}) \sqrt{I^N};$$

3. *For any pair of distinct elements $\mathcal{M}^{(i)}, \mathcal{M}^{(j)} \in \mathcal{F}$,*

$$\left\| \mathcal{M}^{(i)} - \mathcal{M}^{(j)} \right\|_F \leq \epsilon (\tilde{\alpha} - \tilde{\beta}) \sqrt{I^N}.$$

We next state a useful lemma about the multivariate Poisson distribution, taken from [41, Lemma 19] but rephrased to fit our context and notation.

Lemma 5. *As in Section 7, associate each Poisson parameter tensor $\mathcal{M}^{(i)}$ with a corresponding multivariate Poisson distribution $P^{(i)}$. Then the Kullback-Leibler (KL) divergence is bounded as:*

$$D_{KL} \left(P^{(i)} \| P^{(j)} \right) \leq \frac{1}{\min_{\mathbf{i}} m_{\mathbf{i}}^{(j)}} \left\| \mathcal{M}^{(i)} - \mathcal{M}^{(j)} \right\|_F^2 \leq \frac{1}{\tilde{\beta}} \left\| \mathcal{M}^{(i)} - \mathcal{M}^{(j)} \right\|_F^2.$$

This lemma will allow us to upper bound the KL divergence between elements of the packing set by an upper bound on the distance between elements.

8.3 Proof of Lemmas

8.3.1 Proof of Lemma 1

Proof of Lemma 1. Under our assumptions, notice that

$$\left\| \mathcal{M} - \tilde{\mathcal{M}} \right\|_F^2 = \prod_{n=1}^N \left\| \mathbf{u}^{(n)} \right\|_2^2 + \prod_{n=1}^N \left\| \tilde{\mathbf{u}}^{(n)} \right\|_2^2 - 2 \prod_{n=1}^N \langle \mathbf{u}^{(n)}, \tilde{\mathbf{u}}^{(n)} \rangle \leq \epsilon.$$

Since $\langle \mathbf{u}^{(n)}, \tilde{\mathbf{u}}^{(n)} \rangle \geq 0$ for all n , we obtain

$$\begin{aligned}
& \|\mathbf{u}^{(n)} - \tilde{\mathbf{u}}^{(n)}\|_2^2 = \|\mathbf{u}^{(n)}\|_2^2 + \|\tilde{\mathbf{u}}^{(n)}\|_2^2 - 2\langle \mathbf{u}^{(n)}, \tilde{\mathbf{u}}^{(n)} \rangle \\
& \leq \|\mathbf{u}^{(n)}\|_2^2 + \|\tilde{\mathbf{u}}^{(n)}\|_2^2 + \prod_{k \neq n} \frac{\epsilon}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} - \|\mathbf{u}^{(n)}\|_2^2 \prod_{k \neq n} \frac{\|\mathbf{u}^{(k)}\|_2^2}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} - \|\tilde{\mathbf{u}}^{(n)}\|_2^2 \prod_{k \neq n} \frac{\|\tilde{\mathbf{u}}^{(k)}\|_2^2}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} \\
& = \prod_{k \neq n} \frac{\epsilon}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} + \|\mathbf{u}^{(n)}\|_2^2 \left(1 - \prod_{k \neq n} \frac{\|\mathbf{u}^{(k)}\|_2^2}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} \right) + \|\tilde{\mathbf{u}}^{(n)}\|_2^2 \left(1 - \prod_{k \neq n} \frac{\|\tilde{\mathbf{u}}^{(k)}\|_2^2}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} \right) \\
& \leq \prod_{k \neq n} \frac{\epsilon}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} + \|\mathbf{u}^{(n)}\|_2^2 \left(1 - \prod_{k \neq n} \frac{\|\mathbf{u}^{(k)}\|_2}{\|\tilde{\mathbf{u}}^{(k)}\|_2} \right) + \|\tilde{\mathbf{u}}^{(n)}\|_2^2 \left(1 - \prod_{k \neq n} \frac{\|\tilde{\mathbf{u}}^{(k)}\|_2}{\|\mathbf{u}^{(k)}\|_2} \right),
\end{aligned}$$

where the last line applies the Cauchy–Schwarz inequality.

Assume WLOG that $\|\tilde{\mathcal{M}}\|_F \leq \|\mathcal{M}\|_F$, and note this implies that $\|\tilde{\mathbf{u}}^{(k)}\|_2 \leq \|\mathbf{u}^{(k)}\|_2$ for all k . Then the last term in the last line above is non-negative. Using $\|\tilde{\mathbf{u}}^{(n)}\|_2 \leq \|\mathbf{u}^{(n)}\|_2$ we obtain

$$\|\mathbf{u}^{(n)} - \tilde{\mathbf{u}}^{(n)}\|_2^2 \leq \prod_{k \neq n} \frac{\epsilon}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle} + \|\mathbf{u}^{(n)}\|_2^2 \left(2 - \prod_{k \neq n} \frac{\|\mathbf{u}^{(k)}\|_2}{\|\tilde{\mathbf{u}}^{(k)}\|_2} - \prod_{k \neq n} \frac{\|\tilde{\mathbf{u}}^{(k)}\|_2}{\|\mathbf{u}^{(k)}\|_2} \right) \leq \prod_{k \neq n} \frac{\epsilon}{\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle}$$

where the last inequality follows from the fact that the function $f(x, y) = \frac{x}{y} + \frac{y}{x}$ is greater than or equal to 2 for all $x, y > 0$. If $\|\mathcal{M}\|_F \leq \|\tilde{\mathcal{M}}\|_F$, an analogous argument can be applied to obtain the same conclusion.

To finish, since $\mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \in [\beta, \infty)^I$ we have $\langle \mathbf{u}^{(k)}, \tilde{\mathbf{u}}^{(k)} \rangle \geq \beta^2 I$ and therefore

$$\|\mathbf{u}^{(n)} - \tilde{\mathbf{u}}^{(n)}\|_2^2 \leq \frac{\epsilon}{\beta^{2(N-1)} I^{N-1}}.$$

□

8.3.2 Proof of Theorem 10

Proof of Theorem 10. Consider a rank R tensor $\mathcal{X} \in \mathbb{C}^{I_1 \times \dots \times I_N}$ and write its CPD as

$$\mathcal{X} = \sum_{r=1}^R a_r \mathcal{X}_r \tag{20}$$

where $\{a_r\}_{r=1}^R \subset \mathbb{C}$ and each \mathcal{X}_r is a rank-one unit norm tensor. It is important to note that we may assume the $\{\mathcal{X}_r\}_{r=1}^R$ are linearly independent, otherwise we can obtain a CPD of lower rank. By the triangle inequality

$$\|\mathcal{X}\|_* \leq \sum_{r=1}^R |a_r| \|\mathcal{X}_r\|_* = \sum_{r=1}^R |a_r|.$$

The main component of the proof will show that $|a_r| \leq \|\mathcal{X}\|_F$ for each $r \in \{1, \dots, R\}$ and the result will follow.

Considering the CPD of \mathcal{X} above, WLOG assume that

$$|a_R| = \max_r |a_r|.$$

We now follow the Gram-Schmidt process to produce an orthogonal basis that spans the linear subspace spanned by $\{\mathbf{x}_r\}_{r=1}^R$. The new orthogonal basis is defined as

$$\begin{aligned}\tilde{\mathbf{x}}_1 &= \mathbf{x}_1, \\ \tilde{\mathbf{x}}_2 &= \mathbf{x}_2 - \frac{\langle \mathbf{x}_2, \tilde{\mathbf{x}}_1 \rangle}{\|\tilde{\mathbf{x}}_1\|_F} \tilde{\mathbf{x}}_1, \\ &\vdots \\ \tilde{\mathbf{x}}_R &= \mathbf{x}_R - \sum_{k=1}^{R-1} \frac{\langle \mathbf{x}_R, \tilde{\mathbf{x}}_k \rangle}{\|\tilde{\mathbf{x}}_k\|_F} \tilde{\mathbf{x}}_k.\end{aligned}$$

Then, there exists some $\{b_r\}_{r=1}^R \subset \mathbb{C}$ such that

$$\begin{aligned}\mathbf{x} &= \sum_{r=1}^R b_r \tilde{\mathbf{x}}_r \\ &= \sum_{r=1}^{R-1} b_r \tilde{\mathbf{x}}_r + b_R \left(\mathbf{x}_R - \sum_{k=1}^{R-1} \frac{\langle \mathbf{x}_R, \tilde{\mathbf{x}}_k \rangle}{\|\tilde{\mathbf{x}}_k\|_F} \tilde{\mathbf{x}}_k \right) \\ &= \sum_{r=1}^{R-1} c_r \tilde{\mathbf{x}}_r + b_R \mathbf{x}_R.\end{aligned}$$

The second equality holds by using the definition of $\tilde{\mathbf{x}}_R$ in the Gram-Schmidt process and the last equality defines $\{c_r\}_{r=1}^{R-1} \subset \mathbb{C}$ by absorbing all the respective coefficients of $\{\tilde{\mathbf{x}}_r\}_{r=1}^{R-1}$.

By the Gram-Schmidt process, the $\text{span}\{\tilde{\mathbf{x}}_r\}_{r=1}^{R-1} = \text{span}\{\mathbf{x}_r\}_{r=1}^{R-1}$. Therefore, there exists coefficients $\{d_r\}_{r=1}^{R-1} \subset \mathbb{C}$ such that

$$\mathbf{x} = \sum_{r=1}^{R-1} d_r \mathbf{x}_r + b_R \mathbf{x}_R.$$

Comparing this expansion to (20), by linear independence we obtain that $b_R = a_R$ and $d_r = a_r$ for each $r \in \{1, \dots, R-1\}$. We have shown

$$\max_r |a_r| = |a_R| = |b_R| = \left| \left\langle \mathbf{x}, \frac{\tilde{\mathbf{x}}_R}{\|\tilde{\mathbf{x}}_R\|_F} \right\rangle \right| \leq \|\mathbf{x}\|_F.$$

Our previous observations give that

$$\|\mathbf{x}\|_* \leq \sum_{r=1}^R |a_r| \leq R \|\mathbf{x}\|_F.$$

□

8.3.3 Proof of Lemma 4

Proof of Lemma 4. We adapt the proof of [40, Lemma 9] and [21, Lemma 8]. We first define the following set of matrices:

$$\mathcal{C} = \left\{ \mathbf{M} = (m_{ij}) \in \mathbb{R}^{I \times R} : m_{ij} \in \left\{ \tilde{\beta}, \tilde{\beta} + \epsilon(\tilde{\alpha} - \tilde{\beta}) \right\}, 1 \leq i \leq I, 1 \leq j \leq R \right\}.$$

We next construct a set of block tensors. First, for fixed $\mathbf{M} \in \mathcal{C}$, define the matrix

$$\mathbf{A}(\mathbf{M}) = [\mathbf{M} \ \mathbf{M} \ \dots \ \mathbf{M} \ \mathbf{M}[:, 1 : c_M]] \in \mathbb{R}^{I \times I},$$

where the matrix \mathbf{M} repeats $\lfloor I/R \rfloor$ times and then add the first $c_M = (I - R\lfloor I/R \rfloor)$ columns of \mathbf{M} . Here, $\lfloor \cdot \rfloor$ denotes the floor function (integer part). Then, consider the set

$$\mathcal{B} = \mathcal{B}(\mathcal{C}) = \{\mathbf{A}(\mathbf{M}) \circ \mathbf{1} \circ \mathbf{1} \circ \dots \circ \mathbf{1} : \mathbf{M} \in \mathcal{C}\},$$

where $\mathbf{1} \in \mathbb{R}^I$ is a vector of all 1s.

By construction, every element in \mathcal{C} has rank at most R and hence the matrices \mathbf{A} also have rank at most R . It follows that every tensor in \mathcal{B} has CP rank at most R . Moreover, the entries of tensors in \mathcal{B} are either $\tilde{\beta}$ or $\tilde{\beta} + \epsilon(\tilde{\alpha} - \tilde{\beta})$. Noting that $0 < \tilde{\beta} < \tilde{\alpha}$, we have that

$$\tilde{\beta} + \epsilon(\tilde{\alpha} - \tilde{\beta}) \geq \tilde{\beta}.$$

We also see that

$$\epsilon\tilde{\alpha} + (1 - \epsilon)\tilde{\beta} < \epsilon\tilde{\alpha} + (1 - \epsilon)\tilde{\alpha} = \tilde{\alpha},$$

so that the entries of elements of \mathcal{B} are bounded below by $\tilde{\beta}$ and above by $\tilde{\alpha}$ and $\mathcal{B} \subset \mathcal{S}_R(\tilde{\beta}, \tilde{\alpha})$.

We now note that the set \mathcal{C} can be identified with the set of binary vectors $\{0, 1\}^{IR}$ after vectorization, subtracting $\tilde{\beta}$, and dividing by $\epsilon(\tilde{\alpha} - \tilde{\beta})$. Since the transformed \mathcal{C} contains the zero matrix, we may then apply Lemma 3 and conclude that there is a subset $\tilde{\mathcal{C}}$ of \mathcal{C} such that after transformation $\tilde{\mathcal{C}}$ contains the zero matrix, $|\tilde{\mathcal{C}}| \geq 2^{IR/8}$, and for distinct elements $\mathbf{M}^{(i)}, \mathbf{M}^{(j)} \in \tilde{\mathcal{C}}$,

$$\|\mathbf{M}^{(i)} - \mathbf{M}^{(j)}\|_0 \geq \frac{IR}{8}.$$

That is, the two matrices differ in at least $\frac{IR}{8}$ elements since $\|\cdot\|_0$ provides the number of non-zero entries. Noting that by construction, if two corresponding elements have a difference, their difference must be equal to $\epsilon(\tilde{\alpha} - \tilde{\beta})$, we conclude that

$$\|\mathbf{M}^{(i)} - \mathbf{M}^{(j)}\|_F \geq \epsilon(\tilde{\alpha} - \tilde{\beta}) \sqrt{\frac{I_1 R}{8}}. \quad (21)$$

As there is a one-to-one correspondence between \mathcal{B} and \mathcal{C} , we may identify the subset $\mathcal{F} \subseteq \mathcal{B}$ that corresponds to $\tilde{\mathcal{C}} \subseteq \mathcal{C}$. This subset \mathcal{F} has the same cardinality as $\tilde{\mathcal{C}}$.

It remains to bound the difference between distinct elements of \mathcal{F} . By construction, the elements of \mathcal{F} have constant values when fixing all but the first two indices. It follows that for distinct elements $\mathfrak{M}^{(i)}, \mathfrak{M}^{(j)} \in \mathcal{F}$,

$$\begin{aligned} \|\mathfrak{M}^{(i)} - \mathfrak{M}^{(j)}\|_F^2 &\geq \frac{IR}{8} \epsilon^2 (\tilde{\alpha} - \tilde{\beta})^2 \left[\frac{I}{R} \right] \prod_{i=3}^N I_i \\ &\geq \frac{IR}{8} \epsilon^2 (\tilde{\alpha} - \tilde{\beta})^2 \frac{I}{2R} \prod_{i=3}^N I_i \\ &= \frac{I^N}{16} \epsilon^2 (\tilde{\alpha} - \tilde{\beta})^2, \end{aligned}$$

where the first inequality follows by combining (21) with the structure of the tensors, the second by noting that $\lfloor I/R \rfloor \geq I/(2R)$ since $\lfloor x \rfloor/x \geq 1/2$ for $x \geq 1$ (and $I/R \geq 1$), and the third by simplification.

To bound the difference from above, we note that the difference in any corresponding entries of $\mathcal{M}^{(i)}$ and $\mathcal{M}^{(j)}$ is either 0 or $\epsilon(\tilde{\alpha} - \tilde{\beta})$. There are I_T total entries in the difference tensor, so that we immediately conclude that

$$\left\| \mathcal{M}^{(i)} - \mathcal{M}^{(j)} \right\|_F^2 \leq \epsilon^2 (\tilde{\alpha} - \tilde{\beta})^2 I^N.$$

□

9 Conclusions and Future Work

This study provides a theoretical foundation for assessing the reliability of tensor CP-based multi-way analysis through a non-asymptotic interpretation of the Cramér–Rao lower bound. By focusing on a rank-one Poisson tensor model, we demonstrate that constrained maximum likelihood estimators can achieve near-optimal decomposition factor estimation, approaching the CRLB up to constant and logarithmic factors. Our relaxed CRLB framework thus serves as a valuable benchmarking tool for evaluating estimator performance in structured, high-dimensional settings.

For higher-rank CP decompositions, the limitations of current estimation methods become more apparent. Our numerical experiments reveal a growing discrepancy between estimator variance and the CRLB as tensor rank increases. This divergence underscores the impact of degeneracy and non-identifiability that plagues high-rank tensor decompositions. Nonetheless, we find that CP-based inference retains near-minimax optimality in estimating the full parameter tensor, with only modest rank-dependent suboptimality.

Future work could extend this framework to broader classes of tensor models and noise distributions—such as Gaussian or overdispersed count data—to evaluate the generality of the near-efficiency results demonstrated for the Poisson rank-one case. Another important direction is the development of computationally tractable algorithms that approach the CRLB in higher-rank settings, where achieving near-efficiency remains a critical open challenge. In this work, we applied our recent results from [32], which provide closed-form expressions for the Fisher Information Matrix (FIM) for arbitrary tensor ranks. Further exploration of these explicit FIMs could reveal how to structure CP estimators to yield unique minimum variance solutions and ensure the FIM’s non-singularity. Understanding the null-space properties of these FIMs may ultimately offer a principled path toward designing more efficient tensor decomposition estimators.

References

- [1] Age K. Smilde, Rasmus Bro, and Paul Geladi. *Multi-Way Analysis with Applications in the Chemical Sciences*. J. Wiley, 2004.
- [2] Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.
- [3] Christopher J. Hillar and Lek-Heng Lim. Most tensor problems are NP-hard. *J. ACM*, 60(6), November 2013.
- [4] Yonina C Eldar et al. *Rethinking Biased Estimation: Improving Maximum Likelihood and the Cramér–Rao Bound*, volume 1. Now Publishers, Inc., 2008.
- [5] Vladimir Sovljanski and Mario Paolone. On the use of Cramér–Rao lower bound for least-variance circuit parameters identification of li-ion cells. *Journal of Energy Storage*, 94:112223, 2024.

- [6] Daniel E. Clark and Špela Ivekovič. The Cramér–Rao lower bound for 3-d state estimation from rectified stereo cameras. In *2010 13th International Conference on Information Fusion*, pages 1–8, 2010.
- [7] Dorian Bouchet, Valentina Krachmalnicoff, and Ignacio Izeddin. Cramér–Rao analysis of lifetime estimations in time-resolved fluorescence microscopy. *Opt. Express*, 27(15):21239–21252, 2019.
- [8] Dirk Robinson and Peyman Milanfar. Fundamental performance limits in image registration. *IEEE Transactions on Image Processing*, 13(9):1185–99, 2004.
- [9] Masafumi Akahira and Kei Takeuchi. *Asymptotic Efficiency*, pages 54–80. Springer New York, New York, NY, 1981.
- [10] Zhongwei Li, Qingfu Zeng, Huanen Xu, P.M. Grant, C.F.N. Cowan, and B. Mulgrew. The asymptotic Cramér–Rao lower bound and Maximum Likelihood estimation of parameters of an exponentially damped sinusoid from noisy measurements. In *[1992] Proceedings of the IEEE International Symposium on Industrial Electronics*, pages 64–68 vol.1, 1992.
- [11] Arie Yeredor, Amir Weiss, and Anthony J. Weiss. High-order analysis of the efficiency gap for Maximum Likelihood estimation in nonlinear gaussian models. *IEEE Transactions on Signal Processing*, 66(18):4782–4795, 2018.
- [12] Shivam Gupta, Jasper Lee, Eric Price, and Paul Valiant. Finite-sample Maximum Likelihood estimation of location. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 30139–30149. Curran Associates, Inc., 2022.
- [13] Yu Miao. Concentration inequality of the Maximum Likelihood Estimator. *Applied Mathematics Letters*, 23(10):1305–1309, 2010.
- [14] Jasper C.H. Lee and Paul Valiant. Optimal Sub-Gaussian mean estimation in \mathbb{R} . In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 672–683, 2022.
- [15] Luc Devroye, Matthieu Lerasle, Gabor Lugosi, and Roberto I. Oliveira. Sub-Gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.
- [16] Vladimir G. Spokoiny. Parametric estimation: Finite-sample theory. *Annals of Statistics*, 40:2877–2909, 2011.
- [17] Olivier Catoni. Challenging the empirical mean and empirical variance: A deviation study. *Annales de l’I.H.P. Probabilités et Statistiques*, 48(4):1148–1185, 2012.
- [18] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.
- [19] Yang Cao and Yao Xie. Poisson matrix recovery and completion. *IEEE Transactions on Signal Processing*, 64(6):1609–1620, 2016.
- [20] Oscar F López, Daniel M Dunlavy, and Richard B Lehoucq. Zero-truncated Poisson regression for sparse multiway count data corrupted by false zeros. *Information and Inference: A Journal of the IMA*, 12(3):1573–1611, 05 2023.

- [21] Miaoyan Wang and Lexin Li. Learning from binary multiway data: Probabilistic tensor decomposition and its statistical optimality. *Journal of Machine Learning Research*, 21(154):1–38, 2020.
- [22] Andrea Montanari and Nike Sun. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71(11):2381–2425, 2018.
- [23] Ming Yuan and Cun-Hui Zhang. Incoherent tensor norms and their applications in higher order tensor completion. *IEEE Transactions on Information Theory*, 63(10):6753–6766, 2017.
- [24] Navid Ghadermarzy, Yaniv Plan, and Özgür Yilmaz. Near-optimal sample complexity for convex tensor completion. *Information and Inference: A Journal of the IMA*, 8(3):577–619, 11 2018.
- [25] N. Ghadermarzy. *Near-Optimal Sample Complexity for Noisy or 1-Bit Tensor Completion*. PhD thesis, University of British Columbia, 2018.
- [26] N. Ghadermarzy, Y. Plan, and Ö. Yilmaz. Learning tensors from partial binary measurements. *IEEE Transactions on Signal Processing*, 67(1):29–40, 2019.
- [27] Kameron Decker Harris, Oscar López, Angus Read, and Yizhe Zhu. Spectral gap-based deterministic tensor completion. In *2023 International Conference on Sampling Theory and Applications (SampTA)*, pages 1–6, 2023.
- [28] Rungang Han, Rebecca Willett, and Anru R. Zhang. An optimal statistical and computational framework for generalized tensor estimation. *The Annals of Statistics*, 50(1):1–29, 2022.
- [29] Anru Zhang. Cross: Efficient low-rank tensor completion. *The Annals of Statistics*, 47(2):936–964, 2019.
- [30] Holger Rauhut, Reinhold Schneider, and Željka Stojanac. Low-rank tensor recovery via iterative hard thresholding. *Linear Algebra and its Applications*, 523:220–262, 2017.
- [31] Brett W. Bader, Tamara G. Kolda, Daniel M. Dunlavy, et al. Tensor toolbox for MATLAB, version 3.2.1. https://gitlab.com/tensors/tensor_toolbox/-/tree/v3.2.1, April 2021. Available online.
- [32] Carlos Llosa-Vite, Daniel M. Dunlavy, Richard Lehoucq, Oscar López, and Arvind Prasad. A latent-variable formulation of the Poisson canonical polyadic tensor model: Maximum Likelihood estimation and Fisher Information, 2025.
- [33] Hans-Joachim Werner. On the matrix monotonicity of generalized inversion. *Linear Algebra and its Applications*, 27:141–145, 1979.
- [34] Yang Cao and Yao Xie. Poisson matrix recovery and completion. *IEEE Transactions on Signal Processing*, 64(6):1609–1620, 2016.
- [35] J Gross and G Trenkler. C467. complementing remarks on a matrix extension of the Cauchy–Schwarz inequality. *Journal of Statistical Computation and Simulation*, 72(1):26–29, 2002.
- [36] Gautam Tripathi. A matrix extension of the Cauchy–Schwarz inequality. *Economics Letters*, 63(1):1–3, April 1999.

- [37] Martin J Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48. Cambridge University Press, 2019.
- [38] Xin Jiang, Garvesh Raskutti, and Rebecca Willett. Minimax optimal rates for Poisson inverse problems with physical constraints. *IEEE Transactions on Information Theory*, 61(8):4458–4474, 2015.
- [39] Xu Kong. New estimations on the upper bounds for the nuclear norm of a tensor. *Journal of Inequalities and Applications*, 282, 2018.
- [40] Chanwoo Lee and Miaoyan Wang. Tensor denoising and completion based on ordinal observations. In *International Conference on Machine Learning*, pages 5778–5788. PMLR, 2020.
- [41] Xin Jiang, Garvesh Raskutti, and Rebecca Willett. Minimax optimal rates for Poisson inverse problems with physical constraints. *IEEE Transactions on Information Theory*, 61(8):4458–4474, 2015.