SANDIA REPORT

SAND2025-11681R
Printed September 2025

Sandia National Laboratories

LDRD Final Report: Probabilistic Error Bounds for Low-Rank Tensor Decompositions Used in Large-Scale Data Analysis Applications

Daniel M. Dunlavy, Richard B. Lehoucq, Carlos Llosa-Vite, Oscar F. López, Arvind Prasadan

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185
Livermore, California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

National Nuclear Security Administration

ABSTRACT

This report documents a research project on analyzing low-rank tensor models for data analysis that took place at Sandia National Laboratories from October 2023–September 2025. The focus of this work was to extend theoretical frameworks from statistics and probability theory for use with models for scalar, vector, and matrix data to models with tensor, or general multi-dimensional array, data. Through this work, we have provided a new set of tools for bounding errors on low-rank tensor models of both complete and sampled data.

The remainder of this report is organized as follows. In Section 1, we describe the proposed work at the start of the project. Section 2 describes the research advances made as part of the project. Other research contributions in the form of conference presentations and software development is provided in Section 3. Workforce development at Sandia and Florida Atlantic University (via a subcontract on this project) is provided in Section 4.

This page intentionally left blank.

CONTENTS

This page intentionally left blank.

1. PROPOSED RESEARCH

Low-rank matrix and tensor models are a form of unsupervised machine learning that can aid in identifying and understanding latent behaviors in large, complicated data. Such models aim to represent the important signals and/or patterns in data using fewer variables than the amount of data available. To accomplish this goal, these models often leverage distributional assumptions on the data, structural forms of the relationships between data, and/or assumptions on the structure of noise that is part of the data. Principal Component Analysis (PCA) and its variants are examples of such low-rank matrix models that are widely used in many applications, including model reduction of hypersonic aerodynamics [3], cybersecurity analysis [2], and photovoltaic failure classification [17], among many others. Generalization of low-rank matrix models to more than two dimensions are referred to as low-rank tensor models, where data are represented in multi-dimensional arrays, or tensors. Examples of low-rank tensor models include Canonical Polyadic (CP) [4, 14, 15], Tucker [30], and Tensor Train [27] decompositions, which also have been used in many applications at Sandia, including cybersecurity analysis [1], combustion simulation data analysis [19], and climate modeling [29], respectively. When analyzing data with more than two dimensions—i.e., tensor data—it is crucial to use methods that work with all data dimensions rather than marginalizing data to lower dimensions where matrix or vector analysis can be performed [26]. At the beginning of this project, few results existed that provided theoretical guarantees on model errors for low-rank tensor models.

A major, important challenge in the area of low-rank tensor modeling of data is determining how trustworthy such models are when used in practice. Specifically, in this work we aimed to answer the following questions:

(Q1) Can we prove that an optimal low-rank tensor model (in terms of minimal variance of the model parameters) exists?

(Q2) Can we provide a bound on the estimation error for a given low-rank tensor model estimator?

Our proposed research is one of the first efforts to address these questions for low-rank tensor models. Recent work on randomized algorithms for computing low-rank matrix and tensor decompositions has provided efficient, scalable approaches to working with extremely large tensor data [5, 10, 18, 25, 28]. The majority of that work demonstrates that the iterative methods employed converge to an good approximate solution given sufficient data sampling. Unfortunately, conventional convergence analysis does not provide answers to questions (Q1) and (Q2). On this project, we proposed novel extensions to both classical statistical results and recent work in randomized linear and multilinear algebra to answer these questions and demonstrate the relationships between the amount of sampling and quality of low-rank tensor models of data.

In our proposed work, we focused on the analysis of methods for computing low-rank tensor decompositions via maximum likelihood estimation (MLE) of tensor entries (e.g., as described in [6]). Such methods are widely used in practice and are good starting points for attempting to answer questions (Q1) and (Q2) identified above due to the extensive body of existing research on MLE methods.

To answer question (Q1) above, we planned to extend the Cramér-Rao lower bound on the variance of an estimator from the classical result for scalar model parameters [7] to the case of maximum likelihood estimators associated with low-rank tensor decomposition models. In practice, Cramér-Rao bounds can be used to identify how good an estimator we can expect to compute given a model and sample data. For our purposes, we planned to use the Cramér-Rao lower bound to identify the best low-rank tensor decomposition that can be computed for a given data tensor via maximum likelihood estimation. The benefit of our proposed Cramér-Rao bound approach is that we will no longer require ground truth model parameters—available in general only in controlled experiments involving synthetically generated data—to assess the quality of a low-rank tensor model, thus allowing us to analyze more readily models associated with real-world data. Our initial plan in extending the Cramér-Rao bound for maximum likelihood estimators associated with low-rank tensor decompositions was to focus on CP tensor models of count data using the Poisson probability density function. We note that answers to questions (Q1) and (Q2) will be the first of their kind for low-rank tensor models.

As many tensor analysis applications involve many different data types (e.g., count, continuous, binary, etc.) and distributions (e.g., Poisson, Gaussian, Gamma, etc.), we will extend the Cramér-Rao bounds from the low-rank CP Poisson case to general data distributions supported by Generalized CP (GCP) low-rank tensor models [16], which allow for arbitrary single-parameter probability density functions. Extending the computation of Cramér-Rao bounds to general models will thus answer question (Q1) for a large class of low-rank tensor models.

After answering question (Q1) for the various low-rank tensor models used in practice (CP, Tucker, and Tensor Train) on different data types, we will have estimates for the best low-rank models for a given tensor data set. We planned to then focus on answering question (Q2) for these different models and data types. We planned to leverage results from compressive sensing and theoretical results associated with concentration inequalities to address question (Q2).

2. COMPLETED RESEARCH

In this section, we present summaries of the research that was completed on this project. More details can be found in the publication(s) included in each section or in one of the appendices for research that has not yet been published.

2.1. Zero-Truncated Poisson CP Decomposition

We developed a novel statistical inference methodology for multiway count data corrupted by false zeros that are indistinguishable from true zero counts [24]. Our approach consisted of zero-truncating the Poisson distribution to neglect all zero values. This simple truncated approach dispenses with the need to distinguish between true and false zero counts and reduces the amount of data to be processed. We accomplished inference via tensor completion that imposes low-rank tensor structure on the Poisson parameter space. Our main result demonstrated that an N-way rank-R parametric tensor $\mathcal{M} \in (0, \infty)^{I \times \cdots \times I}$ generating Poisson observations can be accurately estimated by zero-truncated Poisson regression from approximately $IR^2 \log_2^2(I)$ non-zero counts under the nonnegative canonical polyadic decomposition. Our result also quantified the error made by zero-truncating the Poisson distribution when the parameter is uniformly bounded from below. Therefore, under a low-rank multiparameter model, our approach guaranteed accurate regression in under-determined scenarios with substantial corruption by false zeros. Several numerical experiments are presented in [24] demonstrating the theoretical results.

2.2. Poisson CP Decomposition as a Latent Variable Model

In this project we established parameter inference for the rank-one Poisson Canonical Polyadic (PCP) tensor model and extended our findings to the general-rank case through a latent-variable reformulation. Our approach exploited the observation that any random PCP tensor can be derived by marginalizing an unobservable random tensor of one dimension larger. The loglikelihood of this larger dimensional tensor, referred to as the "complete" loglikelihood, is comprised of multiple rank-one PCP loglikelihoods. Using this methodology, we first derived non-iterative maximum likelihood estimators for the rank-one PCP model and demonstrated that the traditional algorithms for fitting non-negative matrix and tensor factorizations are Expectation-Maximization algorithms. Next, we derived the Fisher information matrix for the rank-one PCP model and extended these results to the general-rank case using the missing information principle. The Fisher information provided us with crucial insights into the well-posedness of the tensor model, such as the role that tensor rank plays in identifiability and underdeterminacy. This work is to appear on arXiv as a pre-print and subsequently be submitted to a peer-reviewed journal.

2.3. Near-Efficient and Non-Asymptotic Multiway Analysis

This project investigated the statistical efficiency of tensor decomposition-based inference, with a focus on the CP decomposition. Our objective was to establish rigorous guarantees that demonstrate the reliability of the CP decomposition as a tool for finite-sample multiway analysis. We developed a flexible framework for assessing estimators relative to the Cramér–Rao Lower Bound (CRLB) in non-asymptotic settings. In the rank-1 case, we proved that a constrained maximum likelihood estimator achieves variance matching the CRLB up to constants and logarithmic terms, establishing that near-efficient estimation is possible with finite samples. This provided one of the first rigorous guarantees that CP factorization methods can closely approach optimal inference performance in structured count-data problems. Extending to higher-rank tensors, we illustrated that CP estimators generally fail to achieve the CRLB, due in part to the degeneracy and instability inherent in tensor decompositions. Nonetheless, we demonstrated that CP-based inference remains nearly minimax optimal, and we provided sharper error bounds with improved dependence on tensor CP rank compared to previous work. Numerical experiments corroborated our theoretical results, highlighting both the efficiency of rank-1 estimation and the CRLB gap for higher ranks. This work is to appear on arXiv as a pre-print and subsequently be submitted to a peer-reviewed journal.

2.4. The Average Spectrum Norm and Near-Optimal Tensor Completion

In this work, we extended matrix analysis techniques to the tensor setting by introducing a novel family of norms designed for multidimensional arrays. Classical tools such as the nuclear and spectral matrix norms have been central in low-rank matrix recovery, but they fail to capture sampling requirements that scale appropriately with the degrees of freedom in tensor decompositions. To address this gap, we proposed the *average spectrum norms*, a family of tensor norms that generalize the spectral norm and provide a more flexible framework for multidimensional analysis.

In the tensor case, the spectral norm suffers from sensitivity to "spiky" elements—sparse or highly coherent tensors—since it is defined via the supremum over the rank-1 unit ball:

$$\|\mathcal{X}\| = \sup_{\|\mathbf{u}^{(1)}\|_2=\cdots=\|\mathbf{u}^{(N)}\|_2=1} \left|\langle \mathcal{X}, \mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \cdots \circ \mathbf{u}^{(N)} \rangle\right|. \tag{2.1}$$

This makes the norm overly influenced by pathological rank-1 tensors.

We instead defined the p-th order average spectrum norm:

$$\|\mathcal{X}\|_{\mu(p)} = \left(\mathbb{E}\left|\langle \mathcal{X}, \mathbf{u}^{(1)} \circ \cdots \circ \mathbf{u}^{(N)} \rangle\right|^p\right)^{1/p}, \tag{2.2}$$

where the vectors $\mathbf{u}^{(n)}$ are drawn uniformly from the unit sphere. As $p \to \infty$, these norms converge to the tensor spectral norm, but for finite p they capture the average rather than worst-case behavior. By leveraging concentration of measure in high dimensions, the average spectrum norms substantially diminish the influence of spiky tensors. With appropriate choice of p, this yields bounds that scale linearly with the ambient dimensions and rank, enabling sharper analysis of tensor recovery problems. See [23] for a report of this work.

2.5. The Poisson Tensor Completion Entropy Estimator

In this work, we introduced the *Poisson tensor completion (PTC) estimator*, a non-parametric differential entropy estimator [9]. The PTC estimator leverages inter-sample relationships to compute a low-rank Poisson tensor decomposition of the frequency histogram. Our crucial observation was that the histogram bins associated with samples from a multivariate random variable are an instance of a space partitioning of counts and thus can be identified with a spatial Poisson process. The Poisson tensor decomposition led to a completion of the intensity measure over all bins—including those containing few to no samples—and led to our PTC differential entropy estimator. A Poisson tensor decomposition models the underlying distribution of the count data and guarantees non-negative estimated values and so can be safely used directly in entropy estimation. Our estimator was the first tensor-based estimator that exploits the underlying spatial Poisson process related to the histogram explicitly when estimating the probability density with low-rank tensor decompositions for the purpose of tensor completion. Furthermore, we demonstrated that our PTC estimator was a substantial improvement over standard histogram-based estimators for sub-Gaussian probability distributions because of the concentration of norm phenomenon.

2.6. Simple and Nearly-Optimal Sampling for Rank-1 Tensor Completion via Gauss-Jordan

We presented a novel perspective on the *rank-1 tensor completion problem*, where the objective is to recover a rank-1 tensor $\mathcal{X} \in \otimes_{i=1}^{N} \mathbb{R}^{d}$ given access to m uniformly sampled entries. In [11], we provided a new characterization that reduced the recovery task to solving a random linear system, leading to sharp bounds on both sample and computational complexity. In particular, our algorithm achieves exact recovery with $m \lesssim d^2 \log d$ samples and runtime $O(md^2)$, and we proved that $m \gtrsim d \log d$ samples are information-theoretically necessary. These bounds significantly tighten prior results, which were looser, dependent on tensor structure, and relied on more complicated algorithms. A key novelty of our analysis was a reduction to a matrix sketching problem, which enabled both algorithm design and lower-bound proofs. Overall, this work established near-optimal sample guarantees for rank-1 tensor completion and provided a promising approach to study higher-rank cases.

2.7. Spectral Gap-Based Deterministic Tensor Completion

This work advanced the theory of tensor completion under deterministic sampling patterns, a setting of practical importance in recommender systems and other applications with structured missing data. In [13], we analyzed two completion methods—Poisson loss minimization and atomic norm minimization—and established sharper generalization error bounds with improved deterministic sample complexity in terms of the canonical polyadic (CP) tensor rank R. For a tensor of order N, our analysis reduces rank dependence from $R^{2(N-1)(N^2-N-1)}$ of previous work to $R^{2(N-1)(3N-5)}$, with guarantees governed by the spectral gap of the sampling pattern. Our approach derived new and general properties of the atomic tensor norm, yielding improved theoretical tools that apply beyond the deterministic setting. Our numerical experiments illustrated that the spectral gap perspective accurately predicts error behavior across a variety of methods, including max-quasinorm, ridge-penalty, and Poisson loss minimization. Overall, this work established spectral gap analysis as a powerful and unifying tool for both the theory and practice of tensor completion.

2.8. On the Fisher Information Matrix for Arbitrary Tensor Decompositions

In this project we derived and studied the Fisher information matrix for arbitrary tensor factorization models. Our analysis demonstrated that the expected Fisher information matrix can be expressed in the form $\mathbf{J}\mathbf{\Lambda}\mathbf{J}^\top$, where $\mathbf{J}$ denotes a rectangular Jacobian matrix associated with the low-rank tensor factorization, and $\mathbf{\Lambda}$ represents a diagonal matrix associated with the distributional assumption of the random tensor. While our results are broadly applicable, we specifically focused on Jacobian matrices $\mathbf{J}$ associated to low-rank Tucker, canonical polyadic, and tensor train factorizations, and on diagonal matrices $\mathbf{\Lambda}$ associated to exponential dispersion family distributions with arbitrary link function, with particular emphasis on Gaussian, Poisson, and binomial distributions. This formulation of the Fisher information matrix facilitated a comprehensive and unifying examination of the unidentifiable and underdetermined nature of tensor decompositions, providing a robust framework for future research in statistical inference of low-rank tensor decompositions. This work is to appear on arXiv as a pre-print and subsequently be submitted to a peer-reviewed journal.

2.9. Parameter Inference for Poisson Tucker and Tensor Train Decompositions

In this study, we investigated likelihood-based parameter inference methodologies for the Poisson Tucker and Poisson tensor train tensor models by employing latent variable formulations. This approach enabled us to leverage the analytical tools developed for the Poisson canonical polyadic (PCP) tensor model. Specifically, we derived multiplicative algorithms for maximum likelihood estimation for both Poisson Tucker and Poisson tensor train models, drawing on foundational work by Lee and Seung [20, 21] and Chi and Kolda [6] pertaining to the PCP model. Furthermore, we formulated the expected complete loglikelihood function, which served as a basis for developing Expectation-Maximization (EM) algorithms and computing Fisher information matrices, as previously demonstrated in the work described in Section 2.2 above. Our findings contributed to the understanding of parameter inference in tensor models under Poisson assumptions, providing a robust framework for future research in this domain. This work is to appear on arXiv as a pre-print and subsequently be submitted to a peer-reviewed journal.

2.10. Improving Runtime Performance of Tensor Kernels using Rust From Python

We investigated improving the runtime performance of key computational kernels in the Python Tensor Toolbox (`pyttb`), a package for analyzing tensor data across a wide variety of applications. Recent runtime performance improvements have been demonstrated using Rust, a compiled language, from Python via extension modules leveraging the Python C API—e.g., web applications, data parsing, data validation, etc. Using this same approach, we studied the runtime performance of key tensor kernels of increasing complexity, from simple kernels involving sums of products over data accessed through single and nested loops to more advanced tensor multiplication kernels that are key in low-rank tensor decomposition and tensor regression algorithms. In numerical experiments involving synthetically generated tensor data of various sizes and these tensor kernels, we demonstrated consistent improvements in runtime performance when using Rust from Python over using Python alone and using Python combined with the `NumPy` Python package for scientific computing. This work is to appear on arXiv as a pre-print and subsequently be submitted to a peer-reviewed journal.

2.11. Constrained Maximum Likelihood Estimation

This work studied the problem of finding Cramér-Rao lower bounds under parameter constraints. One question that we might ask is how the Fisher Information and hence the Cramér-Rao lower bound changes under these constraints. As an example, consider estimating the factors of a CP-rank-1 3-way tensor, $\mathbf{u}^{(1)} \circ \mathbf{u}^{(2)} \circ \mathbf{u}^{(3)}$; we might rescale any pair of the vectors $\mathbf{u}^{(i)}$ and obtain the same estimate (e.g., multiply the first by a scalar c and divide the second by the same value c). That is, these parameters are not identifiable without some form of constraint, and as a consequence, the Fisher Information matrix associated with the likelihood is singular. This was early work and has been superceded by the PCP latent variable formulation described in Section 2.2. We present some fundamental results about the CRLB and interpretation of a non-singular Fisher Information matrix in Appendix A.

This page intentionally left blank.

3. OTHER RESEARCH CONTRIBUTIONS

3.1. Technical Presentations

The following technical presentations at conferences, workshops, and university colloquia reflected some of the research advances on this project that are described in Section 2.

- **Faster Tensor Kernels via Rust-From-Python Implementation**, Kimmie Harding, Daniel M. Dunlavy, *Sandia Intern Symposium & Graduate School Fair*, Albuquerque, NM, USA, July 2025.
- **On a Latent-Variable Formulation of the Poisson Canonical Polyadic Tensor Model**, Carlos Llosa, Daniel M. Dunlavy, Richard B. Lehoucq, Oscar Lopez, Arvind Prasadan, *Conference on Data Analysis (CoDA)*, Los Alamos, NM, USA, February 2025.
- **The Poisson Canonical Polyadic Tensor Model as a Latent-Variable Model**, Carlos Llosa, Daniel M. Dunlavy, Richard B. Lehoucq, Oscar Lopez, Arvind Prasadan, *SIAM Conference on Mathematics of Data Science*, Atlanta, GA, USA, October 2024.
- **Near Efficient Factor Estimation: Approaching the CRLB under Multiway Poisson Noise**, Oscar Lopez, Carlos Llosa, Arvind Prasadan, Daniel Dunlavy, Richard Lehoucq, *SIAM Conference on Mathematics of Data Science*, Atlanta, GA, USA, October 2024.
- **Zero-Truncated Poisson Regression for Sparse Multiway Count Data Data Corrupted by False Zeros**, Oscar Lopez, Danny Dunlavy, Rich Lehoucq, *SIAM Conference on Mathematics of Data Science*, Atlanta, GA, USA, October 2024.
- **Probabilistic Guarantees for Low-Rank Tensor Decompositions**, Carlos Llosa, Daniel M. Dunlavy, Richard B. Lehoucq, Oscar Lopez, Arvind Prasadan, *Sandia Machine Learning and Deep Learning (MLDL) Workshop*, Albuquerque, NM, USA, September 2024.
- **On a Latent-Variable Formulation of the Poisson Canonical Polyadic Tensor Model**, Carlos Llosa, Daniel M. Dunlavy, Richard B. Lehoucq, Oscar Lopez, Arvind Prasadan, 2024 *Joint Statistical Meetings*, Portland, OR, USA, July 2024.
- **Low-rank Tensor Decompositions for Count Data**, Danny Dunlavy, *College of William & Mary Computer Science Colloquium Series*, Williamsburg, VA, USA, March 2024.
- **The Poisson Canonical Polyadic Tensor Model as a Latent-Variable Model**, Carlos Llosa, Daniel M. Dunlavy, Richard B. Lehoucq, Oscar Lopez, Arvind Prasadan, *6th International Conference of the ERCIM WG on Computational and Methodological Statistics (CMStatistics 2023)*, Berlin, Germany, December 2023.

- **Zero-truncated Poisson Regression for Sparse Multiway Count Data Corrupted by False Zeros**, Oscar Lopez, Danny Dunlavy, Rich Lehoucq. *Workshop on Model Reduction and Numerical Linear Algebra*, Virginia Tech, Blacksburg, VA, USA, November 2023.
- **Spectral gap-based deterministic tensor completion**, Kameron Decker Harris, Oscar López, Angus Read, Yizhe Zhu, *Sampling Theory and Applications (SampTA) Conference*, New Haven, CT, USA, July 2023.
- **Low-Rank Models for Tensor Count Data Corrupted by False Zeros**, Daniel M. Dunlavy, Oscar Lopez, Richard B. Lehoucq, *Conference on Data Analysis (CoDA)*, Santa Fe, NM, USA, March 2023.
- **Zero-Truncated Poisson Regression for Sparse Multiway Count Data Corrupted by False Zeros**, Oscar Lopez, *Joint Mathematics Meeting*, Boston, MA, USA, January 2023.

3.2. Software Development

Although the majority of the work on this project focused on the development of theory for low-rank tensor model error analysis, implementation of the research was partially completed in the Python Tensor Toolbox `pyttb` as part of this work as well [8].

4. WORKFORCE DEVELOPMENT

As part of this project, several early career researchers were trained in the use of low-rank matrix and tensor decompositions and contributed to the research include above in Section 2.

Sandia National Laboratories:

- Jeremy Myers (graduate intern), *College of William & Mary*: low-rank tensor decompositions, spectral analysis of Poisson tensor factorization algorithms, Python tensor software development
- Jordan Medina (undergraduate intern), *New Mexico Institute of Mining and Technology*: low-rank tensor decompositions, Python tensor software development
- Kimmie Harding (undergraduate intern), *New Jersey Institute of Technology*: low-rank tensor decompositions, Python and Rust tensor software development

Florida Atlantic University, Harbor Branch Oceanographic Institute:

- Nick Fularczyk (postdoctoral fellow): low-rank tensor decompositions, image processing.
- Alejandro Gomez-Leos (graduate intern), *University of Texas at Austin*: low-rank tensor decompositions, sample complexity analysis, concentration inequality-based proofs, publication drafting and submission.
- Chandra Kundu (graduate intern), *University of Central Florida*: matrix completion analysis, hyperspectral image processing, machine learning techniques for hyperparameter selection, MATLAB and Python software development.

This page intentionally left blank.

REFERENCES

[1] E. Acar, D.M. Dunlavy, T.G. Kolda, and M. Mørup. Scalable tensor factorizations for incomplete data. *Chemometrics and Intelligent Laboratory Systems*, 106(1):41–56, 2011. URL: https://doi.org/10.1016/j.chemolab.2010.08.004.

[2] E. Acquesta, L.P. Swiler, and A. Pinar. Time series dimension reduction for surrogate models of port scanning cyber emulations. Technical report, Sandia National Laboratories, 2020. Technical Report, SAND-2020-10617. URL: https://www.osti.gov/servlets/purl/1673455.

[3] P.J. Blonigan, K. Carlberg, F. Rizzi, M. Howard, and J.A. Fike. Model reduction for hypersonic aerodynamics via conservative lspg projection and hyper-reduction. In *Proc. AIAA Scitech Forum*, 2020. URL: https://doi.org/10.2514/6.2020-0104.

[4] J. D. Carroll and J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of 'eckart–young' decomposition. *Psychometrika*, 35(3):283–319, 1970. URL: https://doi.org/10.1007/BF02310791.

[5] M. Che and Y. Wei. Randomized algorithms for the approximations of tucker and the tensor train decompositions. *Adv. Comput. Math.*, 45(1):395–428, 2019. URL: https://doi.org/10.1007/s10444-018-9622-8.

[6] E.C. Chi and T.G. Kolda. On tensors, sparsity, and nonnegative factorizations. *SIAM Journal on Matrix Analysis and Applications*, 33:1272–1299, 2012. URL: https://doi.org/10.1137/110859063.

[7] A.C. Davison. *Statistical Models*. Cambridge University Press, 2003. Vol. 11. URL: https://doi.org/10.1017/CBO9780511815850.

[8] Daniel M. Dunlavy, Nicholas T. Johnson, et al. pyttb: Python Tensor Toolbox, 2025. URL: https://github.com/sandialabs/pyttb.

[9] Daniel M. Dunlavy, Richard B. Lehoucq, Carolyn D. Mayer, and Arvind Prasadan. The poisson tensor completion non-parametric differential entropy estimator, 2025. URL: https://arxiv.org/abs/2505.04957, arXiv:2505.04957.

[10] X. Fu, C. Gao, H. Wai, and K. Huang. Block-randomized stochastic proximal gradient for constrained low-rank tensor factorization. In *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 7485–7489, 2019. URL: https://doi.org/10.1109/ICASSP.2019.8682465.

[11] Alejandro Gomez-Leos and Oscar López. Simple and nearly-optimal sampling for rank-1 tensor completion via gauss-jordan, 2024. URL: https://arxiv.org/abs/2408.05431, arXiv:2408.05431.

[12] John D Gorman and Alfred O Hero. Lower bounds for parametric estimation with constraints. *IEEE Transactions on Information Theory*, 36(6):1285–1301, 1990. URL: https://doi.org/10.1109/18.59929.

[13] 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. doi:10.1109/SampTA59647.2023.10301415.

[14] R.A. Harshman. Foundations of the parafac procedure: Models and conditions for an explanatory multi-modal factor analysis. *UCLA Working Papers in Phonetics*, 16(84), 1970. No. 10,085. URL: https://web.archive.org/web/20041010092429/http://publish.uwo.ca/~harshman/wpppfac0.pdf.

[15] F.L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1–4):164–189, 1927. URL: https://doi.org/10.1002/sapm192761164.

[16] D. Hong, T.G. Kolda, and J.A. Duersch. Generalized canonical polyadic tensor decomposition. *SIAM Review*, 62:133–163, 2020. URL: https://doi.org/10.1137/18M1203626.

[17] M. Hopwood, T. Gunda, H. Seigneur, and J. Walters. An assessment of the value of principal component analysis for photovoltaic iv trace classification of physically-induced failures. In *Proc. IEEE Photovoltaic Specialists Conference*, pages 0798–0802, 2020. URL: https://doi.org/10.1109/PVSC45281.2020.9300601.

[18] T.G. Kolda and D. Hong. Stochastic gradients for large-scale tensor decomposition. *SIAM Journal on Mathematics of Data Science*, 2:1066–1095, 2020. URL: https://doi.org/10.1137/19m1266265.

[19] H. Kolla, K. Aditya, and J.H. Chen. *Higher Order Tensors for DNS Data Analysis and Compression*, pages 109–134. Springer International Publishing, 2020. URL: https://doi.org/10.1007/978-3-030-44718-2_6.

[20] D.D. Lee and H.S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999. URL: http://dx.doi.org/10.1038/44565.

[21] D.D. Lee and H.S. Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems*, volume 13, pages 556–562. MIT Press, 2001. URL: https://proceedings.neurips.cc/paper/1861-algorithms-for-non-negative-matrix-factorization.pdf.

[22] Yen-Huan Li and Ping-Cheng Yeh. An interpretation of the moore-penrose generalized inverse of a singular fisher information matrix. *IEEE Transactions on Signal Processing*, 60(10):5532–5536, 2012. URL: https://doi.org/10.1109/TSP.2012.2208105.

[23] Oscar López, Richard Lehoucq, Carlos Llosa-Vite, Arvind Prasadan, and Daniel M. Dunlavy. The average spectrum norm and near-optimal tensor completion, 2024. URL: https://arxiv.org/abs/2404.10085, arXiv:2404.10085.

[24] 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. arXiv:https://academic.oup.com/imaiai/article-pdf/12/3/1573/51602600/iaad016.pdf, doi:10.1093/imaiai/iaad016.

[25] O.A. Malik and S. Becker. Low-rank tucker decomposition of large tensors using tensorsketch. In *Proc. Adv. Neural Inf. Process. Syst.*, pages 10096–10106, 2018. URL: https://proceedings.neurips.cc/paper/2018/file/45a766fa266ea2ebeb6680fa139d2a3d-Paper.pdf.

[26] C. Mu, B. Huang, J. Wright, and D. Goldfarb. Square deal: Lower bounds and improved relaxations for tensor recovery. In *Proc. 31st International Conference on Machine Learning*, volume 32, pages 73–81, 2014. URL: http://proceedings.mlr.press/v32/mu14.pdf.

[27] I.V. Oseledets. Tensor-train decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011. URL: https://doi.org/10.1137/090752286.

[28] W. Pu, S. Ibrahim, X. Fu, and M. Hong. Stochastic mirror descent for low-rank tensor decomposition under non-euclidean losses, 2021. arXiv:2104.14562v1. URL: https://doi.org/10.48550/arXiv.2104.14562.

[29] C. Safta, A. Gorodetsky, J. Jakeman, K. Sargsyan, and D. Ricciuto. Low-rank tensor network approximations for earth system model. In *Proc. SIAM Conference on Computational Science and Engineering*, 2021. URL: https://www.osti.gov/servlets/purl/1854317.

[30] L.R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966. URL: https://doi.org/10.1007/BF02289464.

This page intentionally left blank.

APPENDIX A. Constrained Maximum Likelihood Estimation

We provide some classical results on the Cramér-Rao lower bound under constraints. We refer the reader to [22] and [12] for a longer discussion and more details.

For a parameter vector θ, let J_θ denote the Fisher information matrix, that is,

$$J_\theta = \mathbb{E}\left[-\nabla_\theta^2 l(\theta)\right],$$

for a log likelihood function l.

Let θ be a parameter valued in a set Θ. Let Θ be such that for the function $G_\theta : \mathbb{R}^p \rightarrow \mathbb{R}^k$, $G_\theta = \mathbf{0}_k$ and for $H_\theta : \mathbb{R}^p \rightarrow \mathbb{R}^m$, $H_\theta \leq \mathbf{0}_m$. Assume that both G_θ and H_θ are continuously differentiable.

If at θ, we have $H_\theta < \mathbf{0}_m$ and the equality constraints are satisfied, we have:

$$\mathrm{Cov}(\widehat{\theta}) \geq (\nabla m_\theta) Q_\theta J_\theta^{-1} (\nabla m_\theta)^\top,$$

where the idempotent matrix Q_θ is defined as

$$Q_\theta = \mathcal{I}_p - J_\theta^{-1}(\nabla G_\theta)^\top \left[(\nabla G_\theta) J_\theta^{-1} (\nabla G_\theta)^\top\right]^+ (\nabla G_\theta).$$

If there are no equality constraints, $Q_\theta = \mathcal{I}_p$ and we recover the original Cramér-Rao bound. If (∇G_θ) has rank r_g, we have that Q_θ has rank $p - r_g$. Hence, equality constraints effectively reduce the rank of J_θ^{-1}; the inequality constraints do not affect the result.

What if J_θ is already singular? Assume that J_θ has rank r and singular value decomposition

$$J_\theta = \begin{bmatrix} U_{1:r} & U_{(r+1):p} \end{bmatrix} \begin{bmatrix} \Sigma_J & \mathbf{0}_{r\times(p-r)} \\ \mathbf{0}_{(p-r)\times r} & \mathbf{0}_{(p-r)\times(p-r)} \end{bmatrix} \begin{bmatrix} U_{1:r} & U_{(r+1):p} \end{bmatrix}^\top.$$

Then, J_θ^+ can be interpreted as a constrained Cramér-Rao lower bound for a constraint

$$G_\theta = U_{(r+1):p}^\top \theta + C = \mathbf{0}_r,$$

for some constant matrix C [22].

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Technical Library	1911	sanddocs@sandia.gov

This page intentionally left blank.

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.