

Zero-Truncated Poisson Regression for Zero-Inflated Multiway Count Data*

Oscar F. López[†], Daniel M. Dunlavy[‡], and Richard B. Lehoucq[‡]

Abstract.

We propose a novel statistical inference paradigm for zero-inflated multiway count data that dispenses with the need to distinguish between true and false zero counts. Our approach ignores all zero entries and applies zero-truncated Poisson regression on the positive counts. Inference is accomplished via tensor completion that imposes low-rank structure on the Poisson parameter space. Our main result shows that an N -way rank- R parametric tensor $\mathcal{M} \in (0, \infty)^{I \times \dots \times I}$ generating Poisson observations can be accurately estimated from approximately $IR^2 \log_2^2(I)$ non-zero counts for a nonnegative canonical polyadic decomposition. Several numerical experiments are presented demonstrating that our zero-truncated paradigm is comparable to the ideal scenario where the locations of false zero counts are known *a priori*.

Key words. Zero-inflation, overdispersion, Poisson regression, canonical polyadic tensor decomposition, tensor completion, zero-truncated Poisson distribution

AMS subject classifications. 15A69, 62D10, 62F99, 62R07

1. Introduction. Count data arises in many data science applications including poll analysis [22], network communications [28, 12], single photon count imaging [34, 35], and ecology [5]. Statistical interpretation of count data typically involves estimating parametric distributions likely to generate the counts via regression and maximum likelihood estimation [29, 14, 3]. Though useful for analysis and decision making, in most practical settings the collected data are corrupted by false counts that mislead the inference procedure. In particular, such arrays are frequently congested by zeros, either false or true, in an indistinguishable manner [38, 13, 37, 19]. Such structured noise is a form of *overdispersion* [29, 14] known as *zero-inflation* [38, 13]. In our context we consider a portion of the zeros as false (whose locations are unknown)—i.e., erroneous counts, structural zeros denoting missing values, etc. The source of such zero-inflation is largely an artifact of the standard practice to initialize arrays with all zero entries prior to data collection paired with flawed counting procedures. However, many probability distributions that govern the observed counts are expected to generate a large amount of true zero counts, e.g., Poisson and Bernoulli distributions. This gives the set of zero values a central role in count data, where distinguishing and appropriately handling overdispersed zeros is crucial for accurate analysis and has long been a challenge in the field; see e.g. [5] for a discussion and many citations to this problem in the literature.

Further complicating the task of count data analysis is the inexorable growth in the volume and dimensionality of collected data—e.g., due to the expansion of global communication and social networks and the analysis of the corresponding data. In such large-dimensional settings,

*The authors are thankful for the following support for this work:

Funding: 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.

[†]Florida Atlantic University (lopezo@fau.edu)

[‡]Sandia National Laboratories (dmdunla@sandia.gov, rblehou@sandia.gov)

multiway data analysis and *tensor decompositions* extract insight to interpret the role of each independent data component [1]. When applied to tensors containing redundant and/or correlated information, such factorized representations provide a compressive manner by which to process data that are otherwise too large to handle efficiently. Due to the relative simplicity of many data generation processes, the underlying multiway distributions can be modeled accurately by parametric tensors with few components relative to the ambient dimensions (i.e., low-rank tensors [23]). For this reason, tensor decompositions are a numerically efficient tool by which to achieve multivariate statistical inference [31].

Our contribution is to propose a novel statistical inference paradigm that efficiently handles zero-inflation in multiway count data. Focusing on the commonly used multiparameter *Poisson model*, we condition this distribution on the positive integers to ignore all zero values and treat the respective data array entries as unobserved. Under a low-rank parametric tensor model, we achieve parameter estimation via *tensor completion* that imposes large *zero-truncated Poisson* likelihood. In this manner, we exploit the low-dimensional structure found in many multiway parametric models to accurately infer the underlying mean values of the entire volume in an underdetermined setting that only considers non-zero count values. In contrast to other methods that accommodate for an abundance of zeros, our approach does not introduce additional parameters to be tuned or solved [38, 19, 37] and does not require the zero values to be classified as true or false zeros [5]. Our contribution is a simple but accurate work flow that is numerically efficient while reducing the potential for tuning and misclassification errors.

We begin with a theorem that summarizes our two main results (see Appendix A). The summary theorem provides an error bound for parametric estimators with relatively large log-likelihood as a function of the data and factor dimensions along with the number of non-zero observations. The result implies that our proposed method achieves error bounds comparable to the ideal scenario in which an “oracle” identifies the false zeros and Poisson regression can be applied. These implications are validated in Section 2, where numerical experiments present several realistic situations in which the performance of our zero-truncated paradigm is comparable to the oracle.

In order to present the theoretical statement, we introduce a minimal amount of required notation and definitions. We mainly use the conventions in [11], but also rely on the results of [18, 16, 17, 9] and adopt similar terms for consistency. For any integer $N \in \mathbb{N}$, $[N]$ denotes the set $\{1, 2, \dots, N\}$. We focus on nonnegative tensors and their nonnegative *canonical polyadic decomposition* (NNCP). To elaborate, given $I_1, I_2, \dots, I_N \in \mathbb{N}$ and a tensor with nonnegative entries $\mathcal{J} \in \mathbb{R}_+^{I_1 \times \dots \times I_N}$, where \mathbb{R}_+ denotes the values in \mathbb{R} that are nonnegative, we define its NNCP rank as

$$\text{rank}_+(\mathcal{J}) := \min \left\{ R \in \mathbb{N} \mid \mathcal{J} = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(N)} \text{ with } \mathbf{a}_r^{(n)} \in \mathbb{R}_+^{I_n} \ \forall r \in [R], n \in [N] \right\}.$$

In other words, the NNCP rank is similar to the usual definition of CP rank [23] but only applies to nonnegative tensors and imposes nonnegative constraints on the factors. Such nonnegative matrix and tensor decompositions have received increasing amounts of attention due to their uniqueness properties [27], resulting in an enhanced ability to extract particularly

meaningful data components sought by practitioners [15, 11].

Henceforth, let $0 < \beta \leq \alpha$ be fixed but arbitrary bounds for our sought distribution parameters. Given a nonnegative rank upper bound R , our estimator search space will be

$$S_R^+(\beta, \alpha) := \left\{ \mathcal{T} \in \mathbb{R}^{I_1 \times \cdots \times I_N} \mid \beta \leq t_{\mathbf{i}} \leq \alpha \text{ and } \text{rank}_+(\mathcal{T}) \leq R \right\},$$

where $\mathbf{i} = (i_1, i_2, \dots, i_N) \in [I_1] \times [I_2] \times \cdots \times [I_N]$ denotes a multi-index and $t_{\mathbf{i}}$ is the respective entry of \mathcal{T} .

Provided with multiway count data $\mathcal{X} \in \mathbb{Z}_+^{I_1 \times \cdots \times I_N}$, where \mathbb{Z}_+ denotes the values in \mathbb{Z} that are nonnegative, we seek a low-rank factor model obeying the NNCP $\mathcal{M} \in S_R^+(\beta, \alpha)$ likely to generate the count data via the Poisson distribution, i.e.,

$$(1.1) \quad x_{\mathbf{i}} \sim \text{Poisson}(m_{\mathbf{i}}).$$

To this end, given a subset of multi-indices Ω , we consider the Poisson log-likelihood function

$$(1.2) \quad f_{\Omega}(\mathcal{M}, \mathcal{X}) := \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(m_{\mathbf{i}}) - m_{\mathbf{i}} - \log(x_{\mathbf{i}}!),$$

and the zero-truncated Poisson log-likelihood function

$$(1.3) \quad \tilde{f}_{\Omega}(\mathcal{M}, \mathcal{X}) := \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(m_{\mathbf{i}}) - \log(\exp(m_{\mathbf{i}}) - 1) - \log(x_{\mathbf{i}}!).$$

These and other likelihood functions are typically applied as metrics that quantify how likely a parameter model is to generate the observed data under an assumed distribution [30]. Intuitively, with \mathcal{X} fixed, \mathcal{M} providing larger likelihood function values in (1.2) implies an increasing confidence that the event in (1.1) is true when $\mathbf{i} \in \Omega$. In our setting, Ω indicates our trusted data entries where we wish to apply these metrics. In the oracle scenario, f_{Ω} will be applied with Ω identifying all non-zeros and true zeros. Our approach uses \tilde{f}_{Γ} where Γ specifies the multi-indices of all non-zero observations.

Theorem 1.1. *Let $I := \max_n \{I_n\}$, $\mathcal{M} \in S_R^+(\beta, \alpha)$, and Ω be a subset of multi-indices selected uniformly at random from all subsets of the same cardinality. Suppose $\mathcal{X} \in \mathbb{Z}_+^{I_1 \times \cdots \times I_N}$ is a random tensor with each entry in Ω generated independently as in (1.1) and let $\Gamma \subseteq \Omega$ contain the indices of the non-zero entries of \mathcal{X} restricted to Ω . Then the following statements each hold with probability no less than $1 - |\Omega|^{-1}$ when $\min_n \{I_n\} \geq (N-1) \log_2^2(\max_n \{I_n\}) + 1$:*

$$(1.4) \quad \text{If } \widehat{\mathcal{M}} \in S_R^+(\beta, \alpha) \text{ is such that } f_{\Omega}(\widehat{\mathcal{M}}, \mathcal{X}) \geq f_{\Omega}(\mathcal{M}, \mathcal{X}), \text{ then } \frac{\|\mathcal{M} - \widehat{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2} \leq \epsilon.$$

$$(1.5) \quad \text{If } \widetilde{\mathcal{M}} \in S_R^+(\beta, \alpha) \text{ is such that } \tilde{f}_{\Gamma}(\widetilde{\mathcal{M}}, \mathcal{X}) \geq \tilde{f}_{\Gamma}(\mathcal{M}, \mathcal{X}), \text{ then } \frac{\|\mathcal{M} - \widetilde{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2} \leq \kappa \epsilon,$$

where $\epsilon := O(RI^{\frac{1}{2}} \log_2(I) |\Omega|^{-\frac{1}{2}})$ and $\kappa := \frac{(4+\beta)e^{\beta}-4}{2(e^{\beta}-\beta-1)}$.

The result states that if the number of non-zeros adheres to the parametric complexity (in terms of the NNCP), then estimators with relatively large likelihoods (1.2) and (1.3) are accurate approximations of the true data model. The proof is postponed until Appendix A, where Theorem 1.1 results from combining Theorems A.1 and A.2. To further develop the implications of the result, we narrow down the context to specify our approach and compare it with the ideal oracle scenario mentioned before.

Suppose our given count data \mathcal{X} suffers from zero-inflation but otherwise possesses true non-zero counts. Let us further suppose that an oracle provides us with Ω specifying all true counts obeying (1.1). Notice that Ω contains all non-zeros along with true zero counts, which we assume are distributed in a random manner. Then Γ is simply the set of all non-zero entries of \mathcal{X} , which can always be identified in practice regardless of Ω . However, in this non-oracle scenario, Ω still plays an important role (albeit implicitly) since it determines the degree of zero-inflation in our observations.

To be concrete, let us produce our estimators via maximum likelihood

$$(1.6) \quad \widehat{\mathcal{M}} = \arg \max_{\mathcal{T} \in S_R^+(\beta, \alpha)} f_{\Omega}(\mathcal{T}, \mathcal{X}) \quad \text{and} \quad \widetilde{\mathcal{M}} = \arg \max_{\mathcal{T} \in S_R^+(\beta, \alpha)} \widetilde{f}_{\Gamma}(\mathcal{T}, \mathcal{X}),$$

which will satisfy the conditions in (1.4) and (1.5). The zero-truncated estimator $\widetilde{\mathcal{M}}$ corresponds to our approach, while $\widehat{\mathcal{M}}$ is the oracle estimator. Intuitively, our approach ignores all zeros, including false zeros, and compensates by truncating the distribution. The main novelty of our work is summarized in Theorem 1.1, which states that our approach provides parameter estimation (1.5) comparable in accuracy to the ideal scenario where Ω is known (1.4) while remaining implementable. Note that the error bound in (1.5) contains the term

$$(1.7) \quad \kappa := \frac{(4 + \beta)e^{\beta} - 4}{2(e^{\beta} - \beta - 1)},$$

which is always larger than one and thereby amplifies the error in our approach with respect to that of the ideal scenario. In particular, this term implies an increased amount of sensitivity when our approach is applied to near zero parameter models, i.e., when $\beta \approx 0$.

For low-rank tensors, the number of true counts required by the result for an accurate estimator is small relative to the ambient dimensions, i.e., $|\Omega| \sim IR^2 \log_2^2(I) \ll I_1 \cdots I_N$. This allows for statistical inference via multiway analysis under significantly underdetermined scenarios, which otherwise would require the entire volume to be observed in a setting free of false zeros. Theorem 1.1 is slightly pessimistic since the number of free variables that specifies an element of $S_R^+(\beta, \alpha)$ is at most $\sim IR$ with optimal sampling rate conjecture to be $|\Omega| \sim IR \log(I)$, where the logarithmic term is unavoidable in matrix and tensor completion under random sampling models [8]. Despite this, our derived sampling complexity is novel in that it improves upon current results in the literature, which involve super-quadratic dependence on R and I for N -way arrays with $N \geq 3$ [36, 25, 7]. However, it is important to notice that we consider the NNCP rank rather than the general CP rank so that this comparison is difficult to make fairly.

Theorem 1.1 does not provide a method for parameter estimation and instead assumes an estimator $\widehat{\mathcal{M}}$ is available. We state the result in this abstract manner in order to remain

flexible and practical. Indeed, outputs of the form (1.6) are NP-hard to compute [20], so that no tractable algorithm is guaranteed to achieve the global optimizer. For this reason we do not specify how $\widetilde{\mathcal{M}}$ should be produced and instead attempt to state minimal conditions that an accurate estimate should satisfy, in order to guide practitioners into developing appropriate methods. In fact, the result only requires for an estimator to have large likelihood relative to the true parameter tensor. Therefore, a global optimum of (1.6) is not needed and the result remains applicable to local optima and will be informative for other less greedy methods.

2. Numerical experiments. We present a series of experiments to illustrate the influence of several problem parameters on Theorem 1.1 in practice. Specifically, we demonstrate the errors associated with the estimators $\widehat{\mathcal{M}}$ and $\widetilde{\mathcal{M}}$ with respect to \mathcal{M} when these estimators are computed using the method of maximum likelihood estimation. These experiments illustrate some of the practical ramifications of Theorem 1.1.

2.1. Experimental Data. We generate synthetic data using the approach first described by Chi and Kolda in [11], which is implemented in the Matlab Tensor Toolbox [2] in the method `create_problem`. Specifically, we generate random instances of N -way tensors \mathcal{M} , with all dimensions of size I , having rank- R multilinear structure as represented in the CP model:

$$(2.1) \quad \mathcal{M} = \llbracket \boldsymbol{\lambda}; \mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)} \rrbracket = \sum_{r=1}^R \lambda_r \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(N)},$$

where $\mathbf{A}^{(n)} \in \mathbb{R}^{I \times R} \forall n \in [N]$.

We create the desired low-rank, multilinear structure such that all of the entries in \mathcal{M} lie in the interval $[\beta, \alpha]$, as prescribed in Theorem 1.1 via a sampling of the entries in the factor matrices, $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$, uniformly from $[(\beta/R)^{1/N}, (\alpha/R)^{1/N}]$, and set $\lambda_r = 1$. The result is that the entries in \mathcal{M} follow a truncated normal distribution in the interval $[\beta, \alpha]$. Figure 1 illustrates the distribution of entries of an instance of \mathcal{M} generated using $\beta = 1.5$ and $\alpha = 2.5$.

We generate instances of \mathcal{X} by first creating an instance of \mathcal{M} using the procedure above, and then use the Poisson random sampler, `poissrnd`, from Matlab's Statistics and Machine Learning Toolbox¹, to generate the entries of \mathcal{X} .

Instances of the index set Ω are constructed by uniformly sampling without replacement from the linearized index set of \mathcal{X} , given by $[I^N]$. Thus, when simulating false zeros in \mathcal{X} , the values at the indices in $[I^N] \setminus \Omega$ are set to 0.

2.2. Maximum Likelihood Estimation Methods. Given a data tensor \mathcal{X} whose entries are each assumed to be a draw from a Poisson distribution with parameters in \mathcal{M} , as defined in (1.1), we compute estimators for \mathcal{M} using the method of maximum likelihood estimation [30]. We solve the maximum likelihood estimation problem by minimizing the negative of the log-likelihood function associated with the distributions of interest. Specifically, in our experiments, we minimize $-f_{\Omega}(\mathcal{M}, \mathcal{X})$ from (1.2) and $-\tilde{f}_{\Gamma}(\mathcal{M}, \mathcal{X})$ from (1.3) to compute estimators $\widehat{\mathcal{M}}$ and $\widetilde{\mathcal{M}}$, respectively.

¹<https://www.mathworks.com/products/statistics.html>




Figure 1: Histograms of entries of example factor matrices $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}$ (left) and tensor \mathcal{M} (right) generated via `create_problem` with $\beta = 1$, $\alpha = 2.5$, $N = 3$, $I = 100$, and $R = 5$.

The Generalized Canonical Polyadic (GCP) method for computing low-rank CP decompositions [21, 24] provides a method for maximum likelihood estimation using general loss functions that we use here in our experiments. Specifically, we use the Matlab Tensor Toolbox implementation of GCP, provided in the method `gcp_opt`, to compute maximum likelihood estimators for \mathcal{M} . In `gcp_opt`, we use the limited-memory bound-constrained quasi-Newton optimization method [4, 6]; i.e., the input parameter `opt` is set to `'lbfgsb'`.

We compute three estimators denoted *Poisson*, *Oracle*, and *ZTP*:

- ***Poisson***. This approach was introduced in [11] for computing CP decompositions of data tensors with count values. It computes an estimate by minimizing $-f_{\Omega}(\mathcal{M}, \mathcal{X})$ over all values in \mathcal{X} , i.e., by setting $\Omega = [I^N]$. Thus, it treats both true and false zeros as zero values in the data. In `gcp_opt`, the input parameter `type` is set to `'count'` to specify this method.
- ***Oracle***. This approach is similar to the *Poisson* method except that the estimate uses only the true zeros and non-zeros in \mathcal{X} . Thus, the estimate ignores the zeros values in \mathcal{X} that correspond to false zeros by removing the indices of the false zeros from Ω . In general, this information about the specific types of zero values in a data tensor is unknown. However, since we generate Ω in our experiments, this information is known explicitly. Thus, we can use this estimate when zero values in data are known to be true or false *a priori*. In `gcp_opt`, the input parameter `mask` is set to be a tensor of the same size of \mathcal{X} whose values at indices in Ω are equal to 1 and all other entries are equal to 0. This provides the information to GCP to minimize only over the true zeros and non-zeros in \mathcal{X} when computing an estimator. All other input parameters are the same as those used for the *Poisson* method.
- ***ZTP***. This approach computes an estimate by minimizing $-f_{\Gamma}(\mathcal{M}, \mathcal{X})$, where $\Gamma \subseteq \Omega$ denotes the indices of the non-zeros of \mathcal{X} . Thus, no zero values are used in computing an estimator with this method, which is accounted for in the zero-truncated Poisson

log-likelihood function, defined in (1.3). In `gcp_opt`, the input parameters `func` and `grad` are set to anonymous function handles for code to compute $-\tilde{f}_\Gamma(\mathcal{M}, \mathcal{X})$ and $-\nabla \tilde{f}_\Gamma(\mathcal{M}, \mathcal{X})$, respectively. As for the *Oracle* method, the input parameter `mask` is set to be a tensor of the same size of \mathcal{X} whose values at indices in Γ are equal to 1 and all other entries are equal to 0.

To compute $f_\Omega(\mathcal{M}, \mathcal{X})$, $\nabla f_\Omega(\mathcal{M}, \mathcal{X})$, $\tilde{f}_\Gamma(\mathcal{M}, \mathcal{X})$ and $\nabla \tilde{f}_\Gamma(\mathcal{M}, \mathcal{X})$, we follow the existing conventions in `gcp_opt` of protecting against division by zero and computing $\log(0)$ in function and gradient calculations by shifting values by 10^{-10} in divisors and logarithm computations.

2.3. Average Relative Error. Estimator errors are computed as the relative difference between the estimators and Poisson parameter tensors, as in (1.4) and (1.5). For each instance pair $(\mathcal{M}, \mathcal{X})$, we report the averages relative error (denoted as *Average Relative Error* in the plots presented in §2.5) across k randomly selected instances of the index set Ω . Table 1 presents the maximum likelihood estimate (MLE) methods, the indices of entries in \mathcal{X} used for each MLE method, and the corresponding relative error expressions. Note that estimators $\widehat{\mathcal{M}}$ and $\widetilde{\mathcal{M}}$ are those computed using the Poisson log-likelihood (1.2) and zero-truncated Poisson log-likelihood (1.3) functions, respectively.

<i>MLE Method</i>	<i>Data Indices</i>	<i>Relative Error</i>
<i>Poisson</i>	$[I^N]$	$\ \mathcal{M} - \widehat{\mathcal{M}}\ /\ \mathcal{M}\ $
<i>Oracle</i>	Ω	$\ \mathcal{M} - \widehat{\mathcal{M}}\ /\ \mathcal{M}\ $
<i>ZTP</i>	Γ	$\ \mathcal{M} - \widetilde{\mathcal{M}}\ /\ \mathcal{M}\ $

Table 1: Data indices relative error expressions used for the MLE methods in experiments.

2.4. Experimental Setup. Our experiments illustrate the differences in computing maximum likelihood estimators for \mathcal{M} using the various methods described in §2.2. Specifically, in these experiments, we vary the size of the number of trusted data tensor entries, $|\Omega|$, and the ranges of the Poisson parameter tensor entries, $[\beta, \alpha]$.

We run several experiments by varying $|\Omega|$, β , and α . In all experiments, we use $N = 3$ and $R = 5$. Since the minimum requirement for each dimension of these experiments is $I \geq 82$, as specified in the setup of Theorem 1.1, we use values of $I \in \{50, 100, 200\}$ to illustrate the impact of dimension size on the results. For each experiment, we use β and α to generate instances of \mathcal{M} and \mathcal{X} as described in §2.1. For each instance pair of $(\mathcal{M}, \mathcal{X})$, we generate $k = 50$ instances of Ω .

Across the experiments, we vary the problem parameters $|\Omega|$, β , and α as follows.

- **Varying $|\Omega|$.** We vary the size of the set of true zero and non-zero values, $|\Omega|$, such that $|\Omega|/I^N$ falls in the range $[0, 1]$. Results for the different methods are reported as a function of $|\Omega|/I^N$, even though different amounts of data are used in computing the estimators with the different methods, as discussed in §2.2. We run experiments with $|\Omega|/I^N \in \{0.01, 0.02, 0.03, 0.04, 0.05, 0.10, 0.15, \dots, 0.95, 1\}$.
- **Varying β .** The probability of generating true zeros in \mathcal{X} increases as $\beta \rightarrow 0$. Since

the different estimator methods treat zeros differently, it is important to understand the impact of the number of true zeros in \mathcal{X} on the estimator errors. We run experiments with $\beta \in \{0.001, 0.01, 0.01, 0.1, 1\}$.

- **Varying α .** The probability of generating true zeros in \mathcal{X} decreases with increasing α . When there are no true zeros in \mathcal{X} , the Oracle and ZTP methods are equivalent. Moreover, when there are no true or false zeros in \mathcal{X} —i.e., when $|\Omega| = I^N$ —all three methods described in §2.2 are equivalent. We run experiments with $\alpha \in \{2.5, 5, 10, 25, 50\}$.

All experiments were conducted using Matlab Tensor Toolbox v3.2.1 [2] in Matlab R2021b.

2.5. Results. We present results for experiments involving the methods defined as *Poisson*, *Oracle*, and *ZTP* in §2.2 to demonstrate the results of Theorem 1.1 in practice.

Varying $|\Omega|$. Figure 2 presents the average relative errors of estimators using the three methods as a function of $|\Omega|/I^N$, which is the fraction of the number true zeros and non-zeros to the total number of entries in the data tensors. In these experiments, we set $\beta = 1$, $\alpha = 2.5$, $N = 3$, $I = 100$, $R = 5$, and generate 50 replicates of Ω for each value of $|\Omega|/I^N$. As expected, the *Oracle* method, which only computes estimators using true zeros and non-zeros, leads to the best results for all values of $|\Omega|/I^N$. When $|\Omega|/I^N = 1$, the *Poisson* and *Oracle* methods are identical, since there are no false zeros, as illustrated in the right side of the plot. In such cases, though, the *ZTP* method ignores all zeros and thus incurs more error in the estimates. As predicted by Theorem 1.1, we see that the average errors of the *ZTP* estimators track those of the *Oracle* estimators, differing only by a small multiplicative value at each value of $|\Omega|/I^N$. In these experiments, the predicted difference in relative error in Theorem 1.1 should be bounded by a factor of $\sqrt{\kappa} \approx 2.6$ (with $\beta = 1$), which aligns well with these results.

Varying β . Figure 3 presents the average relative errors of estimators using the three methods as a function of β , which influences the number of true zeros in the data tensors. As expected, as $\beta \rightarrow 0$, κ increases, and thus there are greater differences in the average errors between the estimators computed with the *Oracle* and *ZTP* methods. Moreover, these




Figure 2: Results varying $|\Omega|$: $\beta = 1$, $\alpha = 2.5$, $N = 3$, $I = 100$, $R = 5$, and 50 replicates.

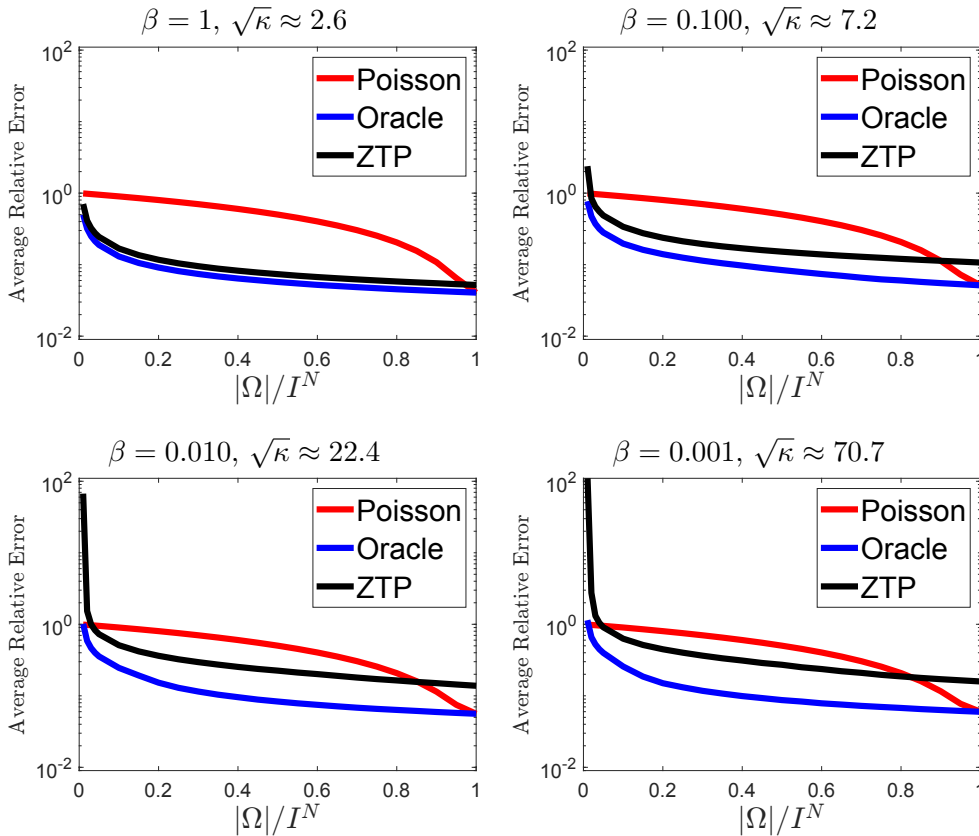


Figure 3: Results varying β : $\alpha = 2.5$, $N = 3$, $I = 100$, $R = 5$, and 50 replicates.

differences are much more extreme as $|\Omega|/I^N \rightarrow 0$ —i.e., as the numbers of false zeros in the data tensors increase. When β is close to 0, there are few observations used by the *ZTP* method to compute the estimator, and thus we see that the average relative errors can be large, whereas the average relative errors for the *Oracle* method are still bounded by the results of computing estimators using the *Poisson* method. Thus, we recommend that the *ZTP* method be used only when there are a sufficient number of non-zero entries in the data tensors; the specific fractions will be determined by the number of dimensions, sizes of those dimensions, and the distributions of values of the non-zero entries.

Varying α . Figure 4 presents the average relative errors of estimators using the three methods as a function of α , which also influences the number of true zeros in the data tensors. We see that for fixed values of β (in this case $\beta = 0.1$), as α increases, there is very little difference in average relative errors between estimators computing using the *Oracle* and *ZTP* methods, despite no change in the factor κ . These results are due to the fact that as α increases, the probability of generating true zeros in the data tensors decreases. Thus, with fewer true zeros, the differences between these methods are diminished.

Varying I . Figure 5 presents the average relative errors of estimators using the three

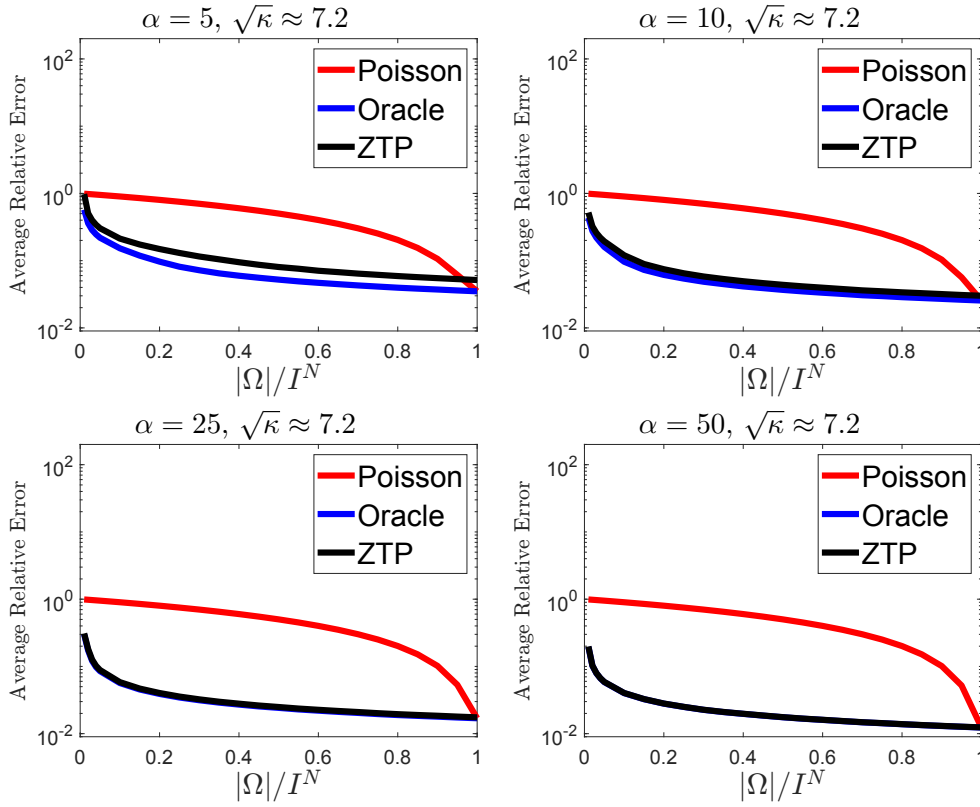


Figure 4: Results varying α : $\beta = 0.1$, $N = 3$, $I = 100$, $R = 5$, and 50 replicates.




Figure 5: Results varying I : $\beta = 1$, $\alpha = 2.5$, $N = 3$, $R = 5$, and 50 replicates.

methods for values of $I \in \{50, 200\}$, which represents smaller and much larger dimension sizes than those required for the results in Theorem 1.1. For the results presented here,

$\beta = 1$ and $\alpha = 2.5$. Recall that when $N = 3$ and $R = 5$, we require that $I \geq 82$ for the results in Theorem 1.1 to hold. We see that when this requirement is not satisfied—e.g., when $I = 50$ —the average relative errors are worse than expected, with rapid increases as $|\Omega|/I^N \rightarrow 0$. Alternatively, as I increases well above the minimum value required to support the conclusions of Theorem 1.1—e.g., when $I = 200$, we see that both the *Oracle* and *ZTP* methods produce even better results in terms of average relative errors for the estimators computed. Since the relative errors in Theorem 1.1 are functions of I for fixed values of β , α , N , R , and $|\Omega|$, these results indicate good agreement between theory and practice.

3. Conclusions. We propose a novel statistical inference paradigm for zero-inflated multi-way count data that does not require the user to distinguish between true and false zero counts. This work debuts the approach on the multi-parameter Poisson model, where we condition this distribution on the positive integers in order to appropriately ignore zero values and treat the respective array entries as unobserved. Under a low-rank parametric model, our approach applies zero-truncated Poisson regression only on the non-zeros. The low-dimensional parametric assumption allows us to achieve Poisson estimation on the entire volume in an underdetermined setting that only considers these positive count values. We show that the approach is efficient at approximating the mean values when the level of zero-inflation is not excessive relative to the parametric complexity. For an N -way parametric tensor $\mathcal{M} \in \mathbb{R}^{I \times \dots \times I}$ with NNCP rank R that generates Poisson observations, our main result states that $\sim IR^2 \log_2^2(I)$ non-zeros provide an accurate estimate via our methodology.

Our numerical experiments explore the implementation of the approach via maximum likelihood and its effectiveness by comparing it to ideal “oracle” scenario, in which the locations of false zeros are known. The presented cases show that in many situations our approach is comparable to the oracle while allowing for practical implementation. We further numerically explore the limitations of the method, including its sensitivity to the mean value bounds β and α . The experiments reveal that when β is relatively large (e.g., $\beta \geq .1$), our approach accurately achieves Poisson regression on the entire array. On the other hand, when the parametric values are small (e.g., $\beta \leq .01$ and $\alpha \leq 1$), the efficiency of our approach is degraded since such situations with sparse data generate an overwhelming amount of true and false zeros.

Several extensions remain to be explored as future work. The current work focuses on the multi-parameter Poisson distribution. However, the paradigm can be applied to any count data model, such as the negative binomial distribution, or even continuous counterparts for other applications, such as the normal distribution. Furthermore, we only consider the case of zero-inflation since it is the most common type of data corruption in the literature of count data. As an extension, any range of integers can be truncated to allow for other types of untrusted count values in data. In the case of continuous models, distributions can be conditioned to any interval of trusted observations. These types of generalizations, paired with more ample theoretical results, can help launch our proposed statistical inference paradigm to handle severe corruption in a wide range of applications that involve multi-dimensional data processing.

Appendix A. Main Theorems and Proofs. We now present our main results for Poisson tensor completion and the zero-truncated counterpart. The following statements are meant

to be general, but under specific circumstances we will combine these theorems to produce Theorem 1.1. For compactness, in this section we modify the log-likelihood functions to

$$f_{\Omega}(\mathcal{M}) := \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(m_{\mathbf{i}}) - m_{\mathbf{i}},$$

and

$$\tilde{f}_{\Omega}(\mathcal{M}) := \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(m_{\mathbf{i}}) - \log(\exp(m_{\mathbf{i}}) - 1),$$

so that their dependency on the count data \mathcal{X} is implicit and the terms $-\log(x_{\mathbf{i}}!)$ are removed. We note that any $\widehat{\mathcal{M}}, \widetilde{\mathcal{M}} \in S_R^+(\beta, \alpha)$ satisfying

$$f_{\Omega}(\widehat{\mathcal{M}}) \geq f_{\Omega}(\mathcal{M}) \quad \text{and} \quad \tilde{f}_{\Gamma}(\widetilde{\mathcal{M}}) \geq \tilde{f}_{\Gamma}(\mathcal{M})$$

will also satisfy the requirements in (1.4) and (1.5). Therefore, this modification does not change the statement and simply serves as a means to compress our proofs.

In the interest of generality, we will also state our results in terms of the CP rank [23] defined as

$$\text{rank}(\mathcal{J}) := \min \left\{ R \in \mathbb{N} \mid \mathcal{J} = \sum_{r=1}^R \mathbf{a}_r^{(1)} \circ \mathbf{a}_r^{(2)} \circ \dots \circ \mathbf{a}_r^{(N)} \quad \text{with} \quad \mathbf{a}_r^{(n)} \in \mathbb{R}^{I_n} \quad \forall r \in [R], n \in [N] \right\},$$

which simply removes the nonnegative constraints on the factors. We also define the respective the search space

$$S_R(\beta, \alpha) := \left\{ \mathcal{J} \in \mathbb{R}^{I_1 \times \dots \times I_N} \mid \beta \leq t_{\mathbf{i}} \leq \alpha \quad \text{and} \quad \text{rank}(\mathcal{J}) \leq R \right\}.$$

We note that we always have $\text{rank}(\mathcal{J}) \leq \text{rank}_+(\mathcal{J})$. We now proceed to the statement of our zero-truncated methodology.

Theorem A.1. *Suppose $\mathcal{M} \in S_R^+(\beta, \alpha)$ and let $\Omega \subseteq [I_1] \times \dots \times [I_N]$ be a subset of cardinality $|\Omega| \leq I_1 \dots I_N$, chosen uniformly at random from all subsets of the same cardinality. Let $\mathcal{X} \in \mathbb{Z}_+^{I_1 \times \dots \times I_N}$ be a random tensor, with each entry in Ω generated independently via (1.1) and let $\Gamma \subseteq \Omega$ be the set of nonzero entries of \mathcal{X} restricted to Ω . Further suppose that $\min_n \{I_n\} \geq (N-1) \log_2^2(\max_n \{I_n\}) + 1$. Fix $\tilde{R} \in \mathbb{N}$, then for any $\widetilde{\mathcal{M}} \in S_{\tilde{R}}^+(\beta, \alpha)$ such that*

$$(A.1) \quad \tilde{f}_{\Gamma}(\widetilde{\mathcal{M}}) \geq \tilde{f}_{\Gamma}(\mathcal{M}),$$

we have

$$(A.2) \quad \frac{\|\mathcal{M} - \widetilde{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2} \leq \frac{64\alpha(\alpha+1) \left((4+\beta)e^{\beta} - 4 \right)}{(e^{\beta} - \beta - 1)\beta^3} \left(\alpha(e^2 - 2) + 3 \log_2(|\Omega|) \right) \times \frac{(R + \tilde{R}) \sqrt{\sum_{n=1}^N I_n}}{\sqrt{|\Omega|}},$$

with probability exceeding $1 - \frac{2}{|\Omega|}$. Furthermore, in the general case where $\mathcal{M} \in S_R(\beta, \alpha)$ and $\widetilde{\mathcal{M}} \in S_{\tilde{R}}(\beta, \alpha)$ but otherwise under the same assumptions, we have

$$(A.3) \quad \frac{\|\mathcal{M} - \widetilde{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2} \leq \frac{64\alpha(\alpha+1)((4+\beta)e^\beta - 4)}{(e^\beta - \beta - 1)\beta^3} (\alpha(e^2 - 2) + 3\log_2(|\Omega|)) \times \\ \left((R\sqrt{R})^{N-1} + (\tilde{R}\sqrt{\tilde{R}})^{N-1} \right) \frac{\sqrt{\sum_{n=1}^N I_n}}{\sqrt{|\Omega|}}$$

with probability greater than $1 - \frac{2}{|\Omega|}$.

See Section A.1 for the proof. The result provides an explicit error bound of our methodology with respect to the CP rank and nonnegative CP rank, the proof is postponed until the next section. This statement is more general than what Theorem 1.1 permits, mainly since we may choose $\tilde{R} < R$, i.e., the rank of the estimate $\widetilde{\mathcal{M}}$ may be smaller than the rank of the tensor of interest \mathcal{M} . We stress that such a rank value for which (A.4) holds may not exist since, in general, this assumption is only guaranteed when $\tilde{R} \geq R$, e.g., by setting

$$\widetilde{\mathcal{M}} = \arg \max_{\mathcal{J} \in S_{\tilde{R}}^+(\beta, \alpha)} f_\Omega(\mathcal{J}),$$

a feasible problem since $\mathcal{M} \in S_{\tilde{R}}^+(\beta, \alpha)$ for $\tilde{R} \geq R$ whose output will satisfy (A.4).

Despite this, we state Theorem A.2 in this flexible manner since a practitioner is typically oblivious to the model's true structure, so \tilde{R} will likely be chosen smaller than R in practice. In such a scenario, the main result remains applicable and informative for practitioners. As a silver lining, tensors suffer from degeneracy [23], i.e., tensors may be approximated arbitrarily well by a factorization of lower rank. It is therefore conceivable that even when the true rank is known there may exist $\tilde{R} < R$ and $\widetilde{\mathcal{M}}$ satisfying (A.4), which will reduce the numerical complexity involved in producing such an estimate.

The statement for the oracle scenario is very similar, but does not consider the set of nonzero entries Γ . Though Theorem A.1 is this work's main contribution due to the novel methodology, the following result may be of independent interest to the reader since it generalizes the work in [9] to the tensor case with best sampling complexity to date.

Theorem A.2. *Under the setup of Theorem A.1, fix $\hat{R} \in \mathbb{N}$. Then for any $\widehat{\mathcal{M}} \in S_{\hat{R}}^+(\beta, \alpha)$ such that*

$$(A.4) \quad f_\Omega(\widehat{\mathcal{M}}) \geq f_\Omega(\mathcal{M}),$$

we have

$$(A.5) \quad \frac{\|\mathcal{M} - \widehat{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2} \leq \frac{128\alpha(\alpha+1)}{\beta^3} (\alpha(e^2 - 2) + 3\log_2(|\Omega|)) \frac{(R + \hat{R})\sqrt{\sum_{n=1}^N I_n}}{\sqrt{|\Omega|}}$$

with probability exceeding $1 - \frac{2}{|\Omega|}$. Furthermore, in the general case where $\mathcal{M} \in S_R(\beta, \alpha)$ and

$\widehat{\mathbf{M}} \in S_{\widehat{R}}(\beta, \alpha)$ but otherwise under the same assumptions, we have

$$(A.6) \quad \frac{\|\mathbf{M} - \widehat{\mathbf{M}}\|^2}{\|\mathbf{M}\|^2} \leq \frac{128\alpha(\alpha+1)}{\beta^3} (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \times \\ \left((R\sqrt{R})^{N-1} + (\widehat{R}\sqrt{\widehat{R}})^{N-1} \right) \frac{\sqrt{\sum_{n=1}^N I_n}}{\sqrt{|\Omega|}}$$

with probability greater than $1 - \frac{2}{|\Omega|}$.

The proof is postponed until Section A.2. From this result, we see a simplified error bound in contrast to (A.3) which contains the multiplicative term

$$\frac{(4 + \beta)e^\beta - 4}{2(e^\beta - \beta - 1)}.$$

This is the error amplification factor κ defined in (1.7) that we encounter in the error bound (1.5) of the introductory result. Thus, Theorem 1.1 is obtained by combining Theorems A.1 and A.2 with absolute constants omitted, disregarding the dependence on N, α and β (i.e., assuming $N, \alpha, \beta \sim O(1)$), using $\sum_n I_n \leq N \max_n I_n$ and $|\Omega| \leq \max_n I_n^N$. Furthermore, for Theorem 1.1 we choose $R = \widehat{R} = \widetilde{R}$, which improves the probability of success to $1 - |\Omega|^{-1}$ since the union bounds used in the proofs of Theorems A.1 and A.2 to apply two distinct ranks are no longer needed (see Sections A.1 and A.2).

Notice that in the CP rank bounds both results exhibit polynomial dependence $(R\sqrt{R})^{N-1}$ on the rank due to the novel work of [18, 16, 17]. While pessimistic, the approach improves on all tensor sampling complexity results to date, particularly on the dependence of the ambient dimensions $\sum_n I_n$ (see Section A.4.1 for further discussion). A minor contribution of this work is that the same proof strategy can be applied to the nonnegative CP rank with severely improved rank dependence.

Sections A.1 and A.2 prove Theorems A.1 and A.2 respectively. We note that the proof of both results is very similar, where the proof of Theorem A.1 requires several additional steps. For this reason, we prove the zero-truncated result first which allows an expedited proof of Theorem A.2.

A.1. Zero-Truncated Poisson Tensor Completion: Proof. In this section we prove Theorem A.1. For the proof, we will need two lemmas, which we state now and prove in Sections A.3.2 and A.4.2. For the first lemma, we define the KL divergence between two zero-truncated Poisson probability distributions $p, q > 0$ as

$$(A.7) \quad D_0(p||q) := \frac{p}{1 - e^{-p}} \log\left(\frac{p}{q}\right) - (\log(e^p - 1) - \log(e^q - 1)).$$

We will require the following lower bound for this KL divergence.

Lemma A.3. *For any $p, q \in [\beta, \alpha]$, we have*

$$(1 - e^{-p})D_0(p||q) \geq \frac{e^\beta - \beta - 1}{2\alpha(e^\beta - 1)}(p - q)^2 \geq 0.$$

We postpone the proof until Section A.3.2, but comment that as a consequence of the proof it can be shown that $D_0(p||q) = 0$ if and only if $p = q$. The second lemma is the main component in the proof of Theorem A.1. We note that this result holds for any deterministic set of observed entries, Ω .

Lemma A.4. *Let $\Omega \subseteq [I_1] \times \dots \times [I_N]$ be any subset of entries and $\mathbf{X} \in \mathbb{Z}_+^{I_1 \times \dots \times I_N}$ be generated as in Theorem A.1 with $\Gamma \subseteq \Omega$ indicating the set of nonzero entries of \mathbf{X} restricted to Ω . Then for any $R \in \mathbb{N}$*

$$(A.8) \quad \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{T}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{T})| \\ \leq 32 \left(\frac{(4 + \beta)e^\beta - 4}{(e^\beta - 1)\beta} \right) (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) (\alpha + 1) R \sqrt{|\Omega| \sum_{n=1}^N I_n},$$

with probability exceeding $1 - \frac{1}{|\Omega|}$, where the probability and expectation are both over the draw of \mathbf{X} . Furthermore, under the same assumptions we have

$$(A.9) \quad \sup_{\mathcal{T} \in S_R(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{T}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{T})| \\ \leq 32 \left(\frac{(4 + \beta)e^\beta - 4}{(e^\beta - 1)\beta} \right) (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) (\alpha + 1) (R\sqrt{R})^{N-1} \sqrt{|\Omega| \sum_{n=1}^N I_n},$$

with probability exceeding $1 - \frac{1}{|\Omega|}$.

The proof can be found in Section A.4.2. We may now proceed to the proof of Theorem A.1.

Proof of Theorem A.1. We will first show (A.2). Afterward, establishing bound (A.3) only requires a minor modification. We begin by computing $\mathbb{E} \tilde{f}_\Gamma(\mathcal{T})$ for $\mathcal{T} \in \mathbb{R}_+^{I_1 \times \dots \times I_N}$, where the expectation is taken with respect to \mathbf{X} (since \tilde{f}_Γ depends on \mathbf{X}). Let \mathbf{u} be a random binary tensor with entries generated as

$$u_{\mathbf{i}} := \begin{cases} 0 & \text{if } x_{\mathbf{i}} = 0 \\ 1 & \text{if } x_{\mathbf{i}} \neq 0, \end{cases}$$

which allows us to write

$$\tilde{f}_\Gamma(\mathcal{T}) = \sum_{\mathbf{i} \in \Omega} u_{\mathbf{i}} \left[x_{\mathbf{i}} \log(t_{\mathbf{i}}) - \log(\exp(t_{\mathbf{i}}) - 1) \right] \\ = \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(t_{\mathbf{i}}) - u_{\mathbf{i}} \log(\exp(t_{\mathbf{i}}) - 1).$$

Notice that $\mathbb{E} u_{\mathbf{i}} = 1 - \mathbb{P}(x_{\mathbf{i}} = 0) = 1 - \exp(-m_{\mathbf{i}})$, and therefore

$$\mathbb{E} \tilde{f}_\Gamma(\mathcal{T}) = \sum_{\mathbf{i} \in \Omega} m_{\mathbf{i}} \log(t_{\mathbf{i}}) - (1 - \exp(-m_{\mathbf{i}})) \log(\exp(t_{\mathbf{i}}) - 1).$$

With this in mind, we apply our assumptions on $\mathcal{M} \in S_R^+(\beta, \alpha)$ and $\tilde{\mathcal{M}} \in S_{\tilde{R}}^+(\beta, \alpha)$ and insert terms that take the marginal expectation with respect to \mathcal{X} only to obtain

$$\begin{aligned}
0 &\leq \tilde{f}_\Gamma(\tilde{\mathcal{M}}) - \tilde{f}_\Gamma(\mathcal{M}) \\
&= \mathbb{E} \left[\tilde{f}_\Gamma(\tilde{\mathcal{M}}) - \tilde{f}_\Gamma(\mathcal{M}) \right] + \left(\tilde{f}_\Gamma(\tilde{\mathcal{M}}) - \mathbb{E} \tilde{f}_\Gamma(\tilde{\mathcal{M}}) \right) + \left(\mathbb{E} \tilde{f}_\Gamma(\mathcal{M}) - \tilde{f}_\Gamma(\mathcal{M}) \right) \\
&\leq \mathbb{E} \left[\tilde{f}_\Gamma(\tilde{\mathcal{M}}) - \tilde{f}_\Gamma(\mathcal{M}) \right] + \sup_{\mathcal{J} \in S_{\tilde{R}}^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| + \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| \\
&= - \sum_{\mathbf{i} \in \Omega} \left[m_{\mathbf{i}} \log \left(\frac{m_{\mathbf{i}}}{\tilde{m}_{\mathbf{i}}} \right) - (1 - \exp(-m_{\mathbf{i}})) (\log(\exp(m_{\mathbf{i}}) - 1) - \log(\exp(\tilde{m}_{\mathbf{i}}) - 1)) \right] \\
&\quad + \sup_{\mathcal{J} \in S_{\tilde{R}}^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| + \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| \\
&= - \sum_{\mathbf{i} \in \Omega} (1 - \exp(-m_{\mathbf{i}})) D_0(m_{\mathbf{i}} \| \tilde{m}_{\mathbf{i}}) \\
&\quad + \sup_{\mathcal{J} \in S_{\tilde{R}}^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| + \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right|.
\end{aligned}$$

In the last line we used the definition of the KL divergence between two zero-truncated Poisson probability distributions (A.7). Since $m_{\mathbf{i}}, \tilde{m}_{\mathbf{i}} \in [\beta, \alpha]$ for all $\mathbf{i} \in [I_1] \times \dots \times [I_N]$, using Lemma A.3, this term can be lower bounded as

$$\sum_{\mathbf{i} \in \Omega} (1 - \exp(-m_{\mathbf{i}})) D_0(m_{\mathbf{i}} \| \tilde{m}_{\mathbf{i}}) \geq \frac{e^\beta - \beta - 1}{2\alpha(e^\beta - 1)} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2.$$

Gathering our bounds and applying equation (A.8) from Lemma A.4 for both R and \tilde{R} , we have established that for any Ω

$$\begin{aligned}
(\text{A.10}) \quad &\frac{e^\beta - \beta - 1}{2\alpha(e^\beta - 1)} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2 \\
&\leq \sup_{\mathcal{J} \in S_{\tilde{R}}^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| + \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J}) \right| \\
&\leq 32 \left(\frac{(4 + \beta)e^\beta - 4}{(e^\beta - 1)\beta} \right) (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) (\alpha + 1)(R + \tilde{R}) \sqrt{|\Omega| \sum_{n=1}^N I_n},
\end{aligned}$$

with probability exceeding $1 - \frac{2}{|\Omega|}$ by a union bound. We now apply our assumption on Ω .

Notice that in terms of the distribution on Ω , the final term above is deterministic since its cardinality $|\Omega|$ is fixed for all outcomes. Therefore, given \mathcal{X} such that the bound holds, we have bounded the random variable $\sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2$. Since \mathcal{X} and Ω are independently

generated, the upper bound holds for the expected value over Ω as well, i.e.,

$$\begin{aligned} & \mathbb{E} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2 \\ & \leq 64 \left(\frac{(4 + \beta)e^\beta - 4}{(e^\beta - \beta - 1)\beta} \right) (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \alpha(\alpha + 1)(R + \tilde{R}) \sqrt{|\Omega| \sum_{n=1}^N I_n}. \end{aligned}$$

We finish the proof by computing the expected value above. Define $K := \binom{I_1 I_2 \cdots I_N}{|\Omega|}$, which is the number of subsets of $[I_1] \times \cdots \times [I_N]$ of size $|\Omega|$ and let $\{\Omega_k\}_{k=1}^K$ list all such subsets. Then

$$\begin{aligned} \mathbb{E} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2 &= \frac{1}{K} \sum_{k=1}^K \sum_{\mathbf{i} \in \Omega_k} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2 \\ &= \frac{1}{K} \sum_{\mathbf{i} \in [I_1] \times \cdots \times [I_N]} \binom{I_1 \cdots I_N - 1}{|\Omega| - 1} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2, \end{aligned}$$

where the last equality holds since for any tensor entry $\mathbf{i} \in [I_1] \times \cdots \times [I_N]$ there will be a total of $\binom{I_1 \cdots I_N - 1}{|\Omega| - 1}$ subsets of size $|\Omega|$ that contain \mathbf{i} . Therefore, in the sum over k each term $(m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2$ will appear exactly $\binom{I_1 \cdots I_N - 1}{|\Omega| - 1}$ times. The proof ends by noticing that

$$\frac{1}{K} \binom{I_1 \cdots I_N - 1}{|\Omega| - 1} = \binom{I_1 \cdots I_N}{|\Omega|}^{-1} \binom{I_1 \cdots I_N - 1}{|\Omega| - 1} = \frac{|\Omega|}{I_1 \cdots I_N}.$$

and

$$\frac{|\Omega|}{I_1 \cdots I_N} \sum_{\mathbf{i} \in [I_1] \times \cdots \times [I_N]} (m_{\mathbf{i}} - \tilde{m}_{\mathbf{i}})^2 = \frac{|\Omega| \|\mathcal{M} - \tilde{\mathcal{M}}\|^2}{I_1 \cdots I_N} \geq \frac{|\Omega| \beta^2 \|\mathcal{M} - \tilde{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2}.$$

The proof of (A.3) is analogous with respect to $S_R(\beta, \alpha)$ and $S_{\tilde{R}}(\beta, \alpha)$, where we instead apply equation (A.9) from Lemma A.4 in (A.10). This replaces the term $\tilde{R} + R$ in (A.10) with $(\tilde{R}\sqrt{\tilde{R}})^{N-1} + (R\sqrt{R})^{N-1}$. The remaining terms are unchanged and the result follows. ■

A.2. Poisson Tensor Completion Proof. The proof of Theorem A.2 is very similar to the proof of Theorem A.1. For brevity, we will refer the reader to the proof of Theorem A.1 when similar steps are applied. The main difference will be to consider instead the KL divergence between Poisson probability distributions, defined as

$$(A.11) \quad D(p||q) := p \log \left(\frac{p}{q} \right) - (p - q).$$

The first lemma establishes a lower bound for the KL divergence.

Lemma A.5. *For any $p, q \in (0, \alpha]$, we have*

$$D(p||q) \geq \frac{(p - q)^2}{2\alpha}.$$

The proof of this lemma is postponed until Section A.3.1. The second lemma is an analogous version of Lemma A.4 used for the zero-truncated result.

Lemma A.6. *Let $\Omega \subseteq [I_1] \times \cdots \times [I_N]$ be any subset of entries, $\mathbf{X} \in \mathbb{Z}_+^{I_1 \times \cdots \times I_N}$ be generated as in Theorem A.2, and the function f_Ω (which depends on \mathbf{X}) be defined as in (1.2). Then for any $R \in \mathbb{N}$*

$$(A.12) \quad \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} |f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T})| \leq (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \times \frac{64(\alpha + 1)R}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n},$$

with probability exceeding $1 - \frac{1}{|\Omega|}$, where the probability and expectation are both over the draw of \mathbf{X} . Furthermore, under the same assumptions we have

$$(A.13) \quad \sup_{\mathcal{T} \in S_R(\beta, \alpha)} |f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T})| \leq (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \times \frac{64(\alpha + 1) (R\sqrt{R})^{N-1}}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n},$$

with probability exceeding $1 - \frac{1}{|\Omega|}$.

See Section A.4 for the proof. We may now proceed to the proof of Theorem A.2.

Proof of Theorem A.2. We will first show (A.5). Afterward, establishing (A.6) only requires a minor modification. We begin by noting that for any $\mathcal{T} \in \mathbb{R}_+^{I_1 \times \cdots \times I_N}$

$$\mathbb{E}f_\Omega(\mathcal{T}) = \mathbb{E} \sum_{\mathbf{i} \in \Omega} x_{\mathbf{i}} \log(t_{\mathbf{i}}) - t_{\mathbf{i}} = \sum_{\mathbf{i} \in \Omega} m_{\mathbf{i}} \log(t_{\mathbf{i}}) - t_{\mathbf{i}},$$

where the expectation is taken with respect to \mathbf{X} . Applying our assumptions on $\mathcal{M} \in S_R^+(\beta, \alpha)$ and $\widehat{\mathcal{M}} \in S_R^+(\beta, \alpha)$, we insert terms that take the marginal expectation with respect to \mathbf{X} only and obtain

$$\begin{aligned} 0 &\leq f_\Omega(\widehat{\mathcal{M}}) - f_\Omega(\mathcal{M}) = \mathbb{E} \left[f_\Omega(\widehat{\mathcal{M}}) - f_\Omega(\mathcal{M}) \right] + \left(f_\Omega(\widehat{\mathcal{M}}) - \mathbb{E}f_\Omega(\widehat{\mathcal{M}}) \right) + \left(\mathbb{E}f_\Omega(\mathcal{M}) - f_\Omega(\mathcal{M}) \right) \\ &\leq \mathbb{E} \left[f_\Omega(\widehat{\mathcal{M}}) - f_\Omega(\mathcal{M}) \right] + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right| + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right| \\ &= - \sum_{\mathbf{i} \in \Omega} \left[m_{\mathbf{i}} \log \left(\frac{m_{\mathbf{i}}}{\widehat{m}_{\mathbf{i}}} \right) - (m_{\mathbf{i}} - \widehat{m}_{\mathbf{i}}) \right] \\ &\quad + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right| + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right| \\ &= - \sum_{\mathbf{i} \in \Omega} D(m_{\mathbf{i}} \| \widehat{m}_{\mathbf{i}}) + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right| + \sup_{\mathcal{T} \in S_R^+(\beta, \alpha)} \left| f_\Omega(\mathcal{T}) - \mathbb{E}f_\Omega(\mathcal{T}) \right|. \end{aligned}$$

In the last line we used the definition of the KL divergence between two Poisson probability distributions (A.11). Since $m_{\mathbf{i}}, \hat{m}_{\mathbf{i}} \in [\beta, \alpha]$ for all $\mathbf{i} \in [I_1] \times \cdots \times [I_N]$, using Lemma A.5, this term can be lower bounded as

$$\sum_{\mathbf{i} \in \Omega} D(m_{\mathbf{i}} \| \hat{m}_{\mathbf{i}}) \geq \frac{1}{2\alpha} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \hat{m}_{\mathbf{i}})^2.$$

Gathering our bounds and applying equation (A.12) from Lemma A.6 for both R and \hat{R} , we have established that for any Ω

$$\begin{aligned} \text{(A.14)} \quad \frac{1}{2\alpha} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \hat{m}_{\mathbf{i}})^2 &\leq \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| f_{\Omega}(\mathcal{J}) - \mathbb{E} f_{\Omega}(\mathcal{J}) \right| + \sup_{\mathcal{J} \in S_{\hat{R}}^+(\beta, \alpha)} \left| f_{\Omega}(\mathcal{J}) - \mathbb{E} f_{\Omega}(\mathcal{J}) \right| \\ &\leq (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \frac{64(\alpha + 1)(R + \hat{R})}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n}, \end{aligned}$$

with probability exceeding $1 - \frac{2}{|\Omega|}$ by a union bound. We now apply our assumption on Ω .

Notice that in terms of the distribution on Ω , the final term above is deterministic since the cardinality $|\Omega|$ is fixed for all outcomes. Therefore, given \mathcal{X} such that the bound holds, we have bounded the random variable $\sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \hat{m}_{\mathbf{i}})^2$. Since \mathcal{X} and Ω are independently generated, the upper bound holds for the expected value over Ω as well, i.e.,

$$\begin{aligned} \mathbb{E} \sum_{\mathbf{i} \in \Omega} (m_{\mathbf{i}} - \hat{m}_{\mathbf{i}})^2 &= \frac{|\Omega| \|\mathcal{M} - \widehat{\mathcal{M}}\|^2}{I_1 \cdots I_N} \\ &\leq (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \frac{128\alpha(\alpha + 1)(R + \hat{R})}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n}. \end{aligned}$$

The proof ends by noting that

$$\frac{|\Omega| \|\mathcal{M} - \widehat{\mathcal{M}}\|^2}{I_1 \cdots I_N} \geq \frac{|\Omega| \beta^2 \|\mathcal{M} - \widehat{\mathcal{M}}\|^2}{\|\mathcal{M}\|^2}.$$

The proof of (A.6) is analogous with respect to $S_R(\beta, \alpha)$ and $S_{\hat{R}}(\beta, \alpha)$, where we instead apply equation (A.13) from Lemma A.6 in (A.14). This replaces the term $\hat{R} + R$ in (A.14) with $(\hat{R}\sqrt{\hat{R}})^{N-1} + (R\sqrt{R})^{N-1}$. The remaining terms are unchanged and the result follows. \blacksquare

A.3. Lower Bounds for KL Divergence. We dedicate this section solely to the proofs Lemmas A.5 and A.3. We will first produce the lower bound for the KL-divergence between two Poisson probability distributions, this in turn will be used to obtain the lower bound for the divergence between two zero-truncated Poisson distributions.

A.3.1. Proof of Lemma A.5. Using the work in [32], the authors in [9] produce a lower bound for the KL divergence between two Poisson probability distributions. In this work, using the work in [32] we are able to obtain a tighter bound.

Proof of Lemma A.5. In [32], the author establishes in equation 11 of Chapter 3 that

$$(1+x)\log(1+x) = x + \frac{x^2}{2(1+x^*)}$$

holds for $x > -1$ and some x^* between 0 and x . With the choice $x = (p-q)/q > -1$, if we multiply through by q we obtain

$$p \log\left(\frac{p}{q}\right) - (p-q) = \frac{(p-q)^2}{2q(1+x^*)}.$$

We now lower bound the right hand side by upper bounding the term $1+x^*$, which we note is always strictly positive. Consider the two possible cases $p \geq q$ and $p < q$. When $p \geq q$, we have $x \geq 0$ so that $x^* \in [0, (p-q)/q]$ and therefore

$$1+x^* \leq 1 + \frac{p-q}{q}.$$

Otherwise, if $p < q$ then $x^* \in [(p-q)/q, 0)$ and

$$1+x^* < 1.$$

Using both of these upper bounds, our assumption $p, q \leq \alpha$ gives that

$$\frac{1}{q(1+x^*)} \geq \frac{1}{q} \min\left\{1, \frac{1}{1+\frac{p-q}{q}}\right\} = \min\left\{\frac{1}{q}, \frac{1}{p}\right\} \geq \frac{1}{\alpha}$$

and therefore

$$\frac{(p-q)^2}{2q(1+x^*)} \geq \frac{(p-q)^2}{2\alpha}.$$

In terms of the KL divergence between two Poisson probability distributions (A.11), we have shown that for $p, q \in (0, \alpha]$

$$D(p\|q) \geq \frac{(p-q)^2}{2\alpha}. \quad \blacksquare$$

A.3.2. Proof of Lemma A.3. We now prove Lemma A.3, which applies the lower bound established in Lemma A.5.

Proof of Lemma A.3. Using basic calculus, we will show that for some term $c_\beta > 0$ depending only on β , we have

$$(1 - e^{-p})D_0(p\|q) \geq c_\beta D(p\|q)$$

for all $p, q \geq \beta > 0$ where $D(p\|q)$ is defined in (A.11). Using Lemma A.5 will then establish the claim.

To this end, let $c_\beta > 0$ be an arbitrary constant (independent of p and q) and consider $p \geq \beta$ fixed, so that we only vary q in $(1 - e^{-p})D_0(p\|q)$ and $c_\beta D(p\|q)$. Notice that these univariate functions intersect at $q = p$ since $D_0(p\|p) = 0 = D(p\|p)$. We compute c_β so

that $(1 - e^{-p})D_0(p||q)$ has a greater rate of change than $c_\beta D(p||q)$ for $q > p$. Taking partial derivatives we obtain

$$\partial q \left[(1 - e^{-p})D_0(p||q) \right] = \frac{e^q(e^p - 1)}{e^p(e^q - 1)} - \frac{p}{q}$$

and

$$\partial q \left[c_\beta D(p||q) \right] = c_\beta \left(1 - \frac{p}{q} \right).$$

Notice that for $q > p$ we have $\partial q(1 - e^{-p})D_0(p||q) > 0$ and $(1 - p/q) > 0$, and we therefore achieve our greater rate of change if

$$c_\beta \leq \frac{\frac{e^q(e^p-1)}{e^p(e^q-1)} - \frac{p}{q}}{\left(1 - \frac{p}{q}\right)} = \frac{qe^q(e^p - 1)}{(q - p)e^p(e^q - 1)} - \frac{p}{q - p} := f(q)$$

holds for all $p \geq \beta$ and $q > p$.

Examining $f(q)$, we see that $f'(q) > 0$ for all $q > p$ and therefore $f(q) \geq f(p)$ where

$$f(p) = \lim_{q \rightarrow p} \left(\frac{qe^q(e^p - 1)}{(q - p)e^p(e^q - 1)} - \frac{p}{q - p} \right) = \frac{e^p - p - 1}{e^p - 1}.$$

This allows us to choose

$$c_\beta := \frac{e^\beta - \beta - 1}{e^\beta - 1} \leq \frac{e^p - p - 1}{e^p - 1},$$

where the inequality holds for all $p \geq \beta$ since $f(p)$ is a monotonically increasing function with respect to p .

We have chosen $c_\beta > 0$ such that $(1 - e^{-p})D_0(p||q)$ and $c_\beta D(p||q)$ agree at $q = p$ and $\partial_q(1 - e^{-p})D_0(p||q) \geq \partial_q c_\beta D(p||q)$ when $q > p$. Therefore $(1 - e^{-p})D_0(p||q) \geq c_\beta D(p||q)$ when $q > p$. The same argument can be applied when $q < p$ (but now with negative rates of change), where the same choice for c_β will give $(1 - e^{-p})D_0(p||q) \geq c_\beta D(p||q)$ when $p > q$. Using Lemma A.5, we have shown for all $p, q \in [\beta, \alpha]$

$$(1 - e^{-p})D_0(p||q) \geq c_\beta D(p||q) \geq \frac{c_\beta(p - q)^2}{2\alpha}. \quad \blacksquare$$

A.4. Proof of the Main Lemmas. The main bulk of our work will be to prove Lemmas A.4 and A.6, the main components in the proofs of Theorems A.1 and A.2. We note that both proofs are very similar, requiring only different terms but applying the same proof strategy. The proof of Lemma A.4 requires more terms to be bounded, aside from analogous terms found in the proof of Lemma A.6. For this reason we will focus on a detailed proof of Lemma A.4 and as a consequence the proof of Lemma A.6 can be achieved in a condensed manner.

To this end, we collect several additional lemmas that will be used in both proofs.

A.4.1. Required Lemmas. We begin by gathering some standard tools from probability in Banach spaces [26]. The following is the symmetrization inequality in diluted form, simplified to be directly applicable to our context (see [26] for the full result).

Lemma A.7 (Symmetrization Inequality, Lemma 6.3 in [26]). *Let $F : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be convex. Let $\{y_\ell\}_{\ell=1}^L \subset \mathbb{R}$ be a finite sequence of independent random variables with $\mathbb{E}|y_\ell| < \infty$ and $\epsilon_1, \epsilon_2, \dots, \epsilon_L$ be i.i.d. Rademacher random variables. Then for any bounded $U \subset \mathbb{R}$*

$$\mathbb{E}F \left(\sup_{(u_1, \dots, u_L) \in U^L} \left| \sum_{\ell=1}^L u_\ell (y_\ell - \mathbb{E}y_\ell) \right| \right) \leq \mathbb{E}F \left(2 \sup_{(u_1, \dots, u_L) \in U^L} \left| \sum_{\ell=1}^L \epsilon_\ell u_\ell y_\ell \right| \right),$$

where the expected value on the right hand side is taken over y_ℓ and ϵ_ℓ .

The symmetrization technique is by now standard, allowing simplified computations by translating these with respect to well studied Rademacher random variables. Subsequently, introducing a Rademacher sequence will pair well with the next result.

Lemma A.8 (Contraction Inequality, Theorem 4.12 in [26]). *Let $F : \mathbb{R}_+ \mapsto \mathbb{R}_+$ be convex and increasing. For $\ell \in [L]$, let ϵ_ℓ be i.i.d. Rademacher random variables and $\varphi_\ell : \mathbb{R} \mapsto \mathbb{R}$ be contractions such that $\varphi_\ell(0) = 0$. Then for any bounded $U \subset \mathbb{R}^L$*

$$\mathbb{E}F \left(\frac{1}{2} \sup_{(u_1, \dots, u_L) \in U^L} \left| \sum_{\ell=1}^L \epsilon_\ell \varphi_\ell(u_\ell) \right| \right) \leq \mathbb{E}F \left(\sup_{(u_1, \dots, u_L) \in U^L} \left| \sum_{\ell=1}^L \epsilon_\ell u_\ell \right| \right),$$

where the expected value is taken with respect to the ϵ_ℓ .

In our proof, the contraction inequality will help us deal with the logarithmic terms introduced by the log-likelihood of the Poisson distribution.

We now consider the atomic M -norm for tensors [18, 16, 17], an approach that will allow our optimal sampling complexity dependence in terms of the tensor dimensions $\{I_n\}_{n=1}^N$. First, define

$$\mathcal{T}_\pm := \left\{ \mathcal{J} \in \{-1, 1\}^{I_1 \times \dots \times I_N} \mid \text{rank}(\mathcal{J}) = 1 \right\}.$$

The atomic M -norm of a tensor $\mathcal{J} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is defined as the gauge (see [10, 33]) of \mathcal{T}_\pm , i.e.,

$$\|\mathcal{J}\|_M := \inf\{t > 0 \mid \mathcal{J} \in t \text{conv}(\mathcal{T}_\pm)\},$$

where $\text{conv}(\mathcal{T}_\pm)$ is the convex envelope of \mathcal{T}_\pm . The M -norm is a convex norm [18, 16, 10] and we will require the following bounds when acting on bounded rank- R tensors.

Lemma A.9. *Assume $\mathcal{J} \in \mathbb{R}^{I_1 \times \dots \times I_N}$ is a rank- R tensor with $\|\mathcal{J}\|_\infty \leq \alpha$. Then*

$$(A.15) \quad \|\mathcal{J}\|_M \leq \alpha \left(R\sqrt{R} \right)^{N-1}.$$

Furthermore, if $\mathcal{J} \in \mathbb{R}_+^{I_1 \times \dots \times I_N}$ with $\text{rank}_+(\mathcal{J}) \leq R_+$ then

$$(A.16) \quad \|\mathcal{J}\|_M \leq \alpha R_+.$$

This result is essentially Theorem 7 in [18], where (A.15) is established. The bound (A.16) is a simple corollary, which we prove briefly before continuing.

Proof of Lemma A.9. As discussed, we only need to show (A.16) using (A.15). By assumption

$$\mathcal{J} = \sum_{r=1}^{R_+} \mathbf{a}_r^{(1)} \circ \dots \circ \mathbf{a}_r^{(N)} := \sum_{r=1}^{R_+} \mathcal{J}_r,$$

where each rank one component \mathcal{J}_r is nonnegative. Since $\|\mathcal{J}\|_\infty \leq \alpha$, by nonnegativity it is easy to see that $\|\mathcal{J}_r\|_\infty \leq \alpha$ for all $r \in [R_+]$. Due to the fact that the M -norm is a norm [18, 16, 10], the triangle inequality gives

$$\|\mathcal{J}\|_M \leq \sum_{r=1}^{R_+} \|\mathcal{J}_r\|_M \leq \sum_{r=1}^{R_+} \alpha = \alpha R_+$$

where the second inequality holds by (A.15) since each \mathcal{J}_r is rank one with $\|\mathcal{J}_r\|_\infty \leq \alpha$. \blacksquare

We also consider the M -norm's dual norm

$$\|\mathcal{J}\|_M^* := \max_{\|\mathbf{u}\|_M \leq 1} \langle \mathcal{J}, \mathbf{u} \rangle = \max_{\mathbf{u} \in \mathcal{T}_\pm} \langle \mathcal{J}, \mathbf{u} \rangle,$$

where the second equality is established in [18]. We will require a bound on the expectation of this dual norm when acting on random tensors of a certain structure.

Lemma A.10. *Assume $\mathcal{V} \in [-1, 1]^{I_1 \times \dots \times I_N}$ is a random tensor with p non-zero entries, which are independent mean zero discrete random variables. Define $\bar{I} := \sum_n I_n$ and $\tilde{I} := I_1 I_2 \dots I_N$. Then, for any $h > 0$ such that $\bar{I} - 1 \geq h \log_2 \left(\frac{\tilde{I}}{4\bar{I}} \right)$ we have*

$$\mathbb{E} (\|\mathcal{V}\|_M^*)^h \leq 2 \left(2\sqrt{p\bar{I}} \right)^h.$$

We postpone the proof of Lemma A.10 until Section A.5. Lemmas A.9 and A.10 produce our sampling complexity in terms of I and R , where $I = \max_{n \in [N]} I_n$. In contrast to previous approaches that try to generalize results for matrix norms, considering the M -norm reduces our sampling complexity from $\mathcal{O}(I^{N/2} \sqrt{R} \log^{3/2}(I))$ [36] to $\mathcal{O}(I(R\sqrt{R})^{2N-2} \log(I))$ in the general case and $\mathcal{O}(IR^2 \log(I))$ in the nonnegative case. Since $R \leq I_1 \dots I_N / I$, this results in a great improvement in many cases. However, the results are still sub-optimal in terms of its rank dependence which is an open problem conjectured to be linear $\mathcal{O}(IR \log(I))$.

A.4.2. Proof of Lemma A.4. We may now proceed to the proof of the main lemma for the zero-truncated case.

Proof of Lemma A.4. We first show (A.8). Afterward, establishing bound (A.9) will only require a slight modification. In what follows, recall that Ω is fixed and let \mathbf{u} be the random tensor with entries u_i defined as in the proof of Theorem A.1. Then, for any $\mathcal{J} \in \mathbb{R}_+^{I_1 \times \dots \times I_N}$ we can write

$$\tilde{f}_\Gamma(\mathcal{J}) = \sum_{i \in \Omega} x_i \log(t_i) - u_i \log(\exp(t_i) - 1),$$

which is a sum of independent random variables. We begin by bounding

$$\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})|^h$$

for arbitrary $h \geq 1$. Afterward, we will apply Markov's inequality for a specified value of h to obtain the statement with the prescribed probability. To this end, we symmetrize (Lemma A.7) by introducing a tensor $\mathcal{V} \in \{-1, 1\}^{I_1 \times \dots \times I_N}$ whose entries are i.i.d. Rademacher random variables to obtain

$$\begin{aligned} & \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})|^h \\ & \leq 2^h \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} \left[x_{\mathbf{i}} \log(t_{\mathbf{i}}) - u_{\mathbf{i}} \log(\exp(t_{\mathbf{i}}) - 1) \right] \right|^h \\ & \leq 2^{2h-1} \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} x_{\mathbf{i}} \log(t_{\mathbf{i}}) \right|^h + 2^{2h-1} \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \log(\exp(t_{\mathbf{i}}) - 1) \right|^h \end{aligned}$$

where the expectations are now over the draw of \mathcal{X} and \mathcal{V} and the last inequality holds since $(a+b)^h \leq 2^{h-1}(a^h + b^h)$ when $a, b > 0$ and $h \geq 1$. Both terms resulting from the last inequality can be bounded by applying Lemma A.8. For the first term, define $\varphi(t) := \beta \log(t+1)$, which is a contraction for $t \geq \beta - 1$ that vanishes at the origin (see [9]). We see that

$$\begin{aligned} 2^{2h-1} \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} \log(t_{\mathbf{i}}) x_{\mathbf{i}} v_{\mathbf{i}} \right|^h &= \frac{1}{2} \left(\frac{4}{\beta} \right)^h \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} \varphi(t_{\mathbf{i}} - 1) x_{\mathbf{i}} v_{\mathbf{i}} \right|^h \\ &\leq \frac{1}{2} \left(\frac{4}{\beta} \right)^h \mathbb{E} \left[\max_{\mathbf{i} \in \Omega} x_{\mathbf{i}}^h \right] \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} \varphi(t_{\mathbf{i}} - 1) v_{\mathbf{i}} \right|^h \\ &\leq \frac{1}{2} \left(\frac{8}{\beta} \right)^h \mathbb{E} \left[\max_{\mathbf{i} \in \Omega} x_{\mathbf{i}}^h \right] \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} (t_{\mathbf{i}} - 1) v_{\mathbf{i}} \right|^h, \end{aligned}$$

where the last inequality holds by Lemma A.8 since with $\mathcal{J} \in S_R^+(\beta, \alpha)$ we have $t_{\mathbf{i}} - 1 \geq \beta - 1$ for all $\mathbf{i} \in \Omega$. We now bound the two expectations in the last line.

For the term $\mathbb{E} \left[\max_{\mathbf{i} \in \Omega} x_{\mathbf{i}}^h \right]$, we argue as in [9] in the proof of Lemma 4. An analogous version of equation (65) therein gives in our context

$$(A.17) \quad \mathbb{E} \left[\max_{\mathbf{i} \in \Omega} x_{\mathbf{i}}^h \right] \leq 2^{2h-1} \left(\alpha^h + \alpha^h (e^2 - 3)^h + 2h! + \log^h(|\Omega|) \right).$$

For the remaining term, let $\Delta_\Omega \in \{0, 1\}^{I_1 \times \dots \times I_N}$ be the indicator tensor for Ω and $\mathbb{1} \in$

$\{1\}^{I_1 \times \dots \times I_N}$ be the all ones tensor so that

$$\begin{aligned} & \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} (t_{\mathbf{i}} - 1) v_{\mathbf{i}} \right|^h \\ &= \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} |\langle \mathcal{T} - \mathbb{1}, \mathcal{V} \circ \Delta_{\Omega} \rangle|^h \leq \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \|\mathcal{T} - \mathbb{1}\|_M^h \mathbb{E} (\|\mathcal{V} \circ \Delta_{\Omega}\|_M^*)^h, \end{aligned}$$

where the inequality holds by the definition of the dual norm. Applying equation (A.16) from Lemma A.9 and the fact that the M -norm is a norm [18, 16, 10], we have by the triangle inequality

$$(A.18) \quad \|\mathcal{T} - \mathbb{1}\|_M \leq \|\mathcal{T}\|_M + \|\mathbb{1}\|_M \leq \alpha R + 1 \leq (\alpha + 1)R,$$

where the second inequality holds since $\mathcal{T} \in \mathcal{S}_R^+$ and $\text{rank}_+(\mathbb{1}) = 1$ with $\|\mathbb{1}\|_{\infty} = 1$. Furthermore, $\mathcal{V} \circ \Delta_{\Omega}$ satisfies the conditions of Lemma A.10, so assuming h will be chosen such that

$$(A.19) \quad \sum_{n=1}^N I_n \geq h \log_2 \left(\frac{I_1 \cdots I_N}{4 \sum_{n=1}^N I_n} \right) + 1$$

we have

$$\mathbb{E} (\|\mathcal{V} \circ \Delta_{\Omega}\|_M^*)^h \leq 2 \left(2 \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h.$$

Thus far, we have shown

$$\begin{aligned} & 2^{2h-1} \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} x_{\mathbf{i}} \log(t_{\mathbf{i}}) \right|^h \\ & \leq \frac{1}{2} \left(\alpha^h + \alpha^h (e^2 - 3)^h + 2h! + \log^h(|\Omega|) \right) \left(\frac{64(\alpha + 1)R}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h. \end{aligned}$$

The remaining term can be bounded in a similar manner, considering $\phi(t) := (1 - e^{-\beta}) \log(\exp(t + \log(2)) - 1)$ which is a contraction for $t \geq \beta - \log(2)$ that vanishes at the origin. Using Lemma A.8 again we obtain

$$\begin{aligned} & 2^{2h-1} \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \log(\exp(t_{\mathbf{i}}) - 1) \right|^h \\ &= \frac{2^{2h-1}}{(1 - e^{-\beta})^h} \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \phi(t_{\mathbf{i}} - \log(2)) \right|^h \\ & \leq \frac{2^{3h-1}}{(1 - e^{-\beta})^h} \mathbb{E} \sup_{\mathcal{T} \in \mathcal{S}_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} (t_{\mathbf{i}} - \log(2)) \right|^h \leq \left(\frac{16(\alpha + 1)R}{1 - e^{-\beta}} \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h, \end{aligned}$$

where the last inequality holds as in the bound of the first term by considering the M -norm, its dual, and applying Lemma A.9 to $\mathcal{J} - \log(2)$ and Lemma A.10 to $\mathbf{u} \circ \mathbf{v} \circ \Delta_\Omega$.

In conclusion, we have shown

$$\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})|^h \leq \delta_0,$$

where

$$\begin{aligned} \delta_0 := & \frac{1}{2} \left(\alpha^h + \alpha^h (e^2 - 3)^h + 2h! + \log^h(|\Omega|) \right) \left(\frac{64(\alpha + 1)R}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h \\ & + \left(\frac{16(\alpha + 1)R}{1 - e^{-\beta}} \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h. \end{aligned}$$

Applying Markov's inequality, we have for any $\delta > 0$

$$\begin{aligned} \mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})| \geq \delta \right) &= \mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})|^h \geq \delta^h \right) \\ &\leq \frac{\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})|^h}{\delta^h} \leq \frac{\delta_0}{\delta^h}. \end{aligned}$$

Pick $\delta = 2\delta_0^{1/h}$ and $h = \log_2(|\Omega|)$, so that

$$\mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |\tilde{f}_\Gamma(\mathcal{J}) - \mathbb{E} \tilde{f}_\Gamma(\mathcal{J})| \geq 2\delta_0^{1/h} \right) \leq 2^{-h} = \frac{1}{|\Omega|}.$$

Using $(a^h + b^h)^{1/h} \leq a + b$, $h!^{1/h} \leq h$, and $(a^h + b^h + c^h + d^h)^{1/h} \leq a + b + c + d$ if $a, b, c, d > 0$, we can simplify the bound as

$$\begin{aligned} 2\delta_0^{1/h} &\leq (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \frac{128(\alpha + 1)R}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n} + \frac{32(\alpha + 1)R}{1 - e^{-\beta}} \sqrt{|\Omega| \sum_{n=1}^N I_n} \\ &= 32 (\alpha(e^2 - 2) + 3 \log_2(|\Omega|)) \left(\frac{(4 + \beta)e^\beta - 4}{(e^\beta - 1)\beta} \right) (\alpha + 1)R \sqrt{|\Omega| \sum_{n=1}^N I_n}. \end{aligned}$$

To finish, we note that (A.19) with our choice $h = \log_2(|\Omega|)$ is satisfied if

$$\min_n I_n \geq (N - 1) \log_2^2 \left(\max_n I_n \right) + \frac{1}{N}.$$

which holds under our assumed contexts defined in Theorems 1.1, A.1, and A.2.

To obtain (A.9), we use an analogous argument with respect to $S_R(\beta, \alpha)$. The only difference is that we apply equation (A.15) from Lemma A.9. This replaces the term $(\alpha + 1)R$ with $(\alpha + 1)(R\sqrt{R})^{N-1}$, and otherwise leaves all other terms unchanged thereby establishing (A.9) with the same probability. \blacksquare

A.4.3. Proof of Lemma A.6. Here we prove the main lemma of the Poisson tensor completion result. The proof is very similar to strategy used in the last section and for brevity we will apply bounds therein.

Proof of Lemma A.6. We first show (A.12). Afterward, establishing bound (A.13) will only require a slight modification. Notice that

$$f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J}) = \sum_{\mathbf{i} \in \Omega} \log(t_{\mathbf{i}}) (x_{\mathbf{i}} - \mathbb{E}x_{\mathbf{i}}),$$

where, with Ω fixed, we take expected value with respect to \mathcal{X} . To bound

$$\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})|^h$$

for arbitrary $h \geq 1$, we apply Lemma A.7 so that

$$\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})|^h \leq 2^h \mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} \left| \sum_{\mathbf{i} \in \Omega} \log(t_{\mathbf{i}}) x_{\mathbf{i}} v_{\mathbf{i}} \right|^h,$$

where $\mathbf{V} \in \{-1, 1\}^{I_1 \times \dots \times I_N}$ is a random tensor whose entries are i.i.d. Rademacher random variables and the expectation is now over the draw of \mathcal{X} and \mathbf{V} . This last term can be bounded exactly as in the proof of Lemma A.4, to obtain

$$\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})|^h \leq \delta_0,$$

where

$$\delta_0 := \left(\alpha^h + \alpha^h (e^2 - 3)^h + 2h! + \log^h(|\Omega|) \right) \left(\frac{32(\alpha + 1)R}{\beta} \sqrt{|\Omega| \sum_{n=1}^N I_n} \right)^h.$$

Applying Markov's inequality, we have for any $\delta > 0$

$$\begin{aligned} \mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})| \geq \delta \right) &= \mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})|^h \geq \delta^h \right) \\ &\leq \frac{\mathbb{E} \sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})|^h}{\delta^h} \leq \frac{\delta_0}{\delta^h}. \end{aligned}$$

Pick $\delta = 2\delta_0^{1/h}$ and $h = \log_2(|\Omega|)$, so that

$$\mathbb{P} \left(\sup_{\mathcal{J} \in S_R^+(\beta, \alpha)} |f_{\Omega}(\mathcal{J}) - \mathbb{E}f_{\Omega}(\mathcal{J})| \geq 2\delta_0^{1/h} \right) \leq 2^{-h} = \frac{1}{|\Omega|}.$$

For the advertised result, we further bound

$$\left(\alpha^h + \alpha^h (e^2 - 3)^h + 2h! + \log^h(|\Omega|) \right)^{1/h} \leq \alpha(e^2 - 2) + 3 \log_2(|\Omega|).$$

To obtain (A.13), we use an analogous argument with respect to $S_R(\beta, \alpha)$ and apply equation (A.15) from Lemma A.9 in (A.18). \blacksquare

A.5. Proof of Required Lemmas. From Section A.4.1, we need only to prove Lemma A.10 since the remaining lemmas are established in the respective citations. To obtain the lemma, we will use the following result for bounded discrete random variables.

Theorem A.11. *Let $y \in [0, L]$ be a discrete random variable. If for some $\delta \in (0, \infty)$ we have*

$$\mathbb{P}(y \geq \delta) \leq \frac{\delta}{L},$$

then

$$\mathbb{E}y \leq 2\delta.$$

The proof of Theorem A.11 is rather simple, we quickly provide the proof before continuing.

Proof of Theorem A.11. If $\delta \geq L$, then the conclusion is trivial. Otherwise, let $y_1 < y_2 < y_3 < \dots \leq L$ be the possible outcomes of y and let $k_0 \in \mathbb{N}$ be such that $y_{k_0} \leq \delta < y_{k_0+1}$. Then

$$\begin{aligned} \mathbb{E}y &= \sum_{k=1}^{\infty} y_k \mathbb{P}(y = y_k) = \sum_{k=1}^{k_0} y_k \mathbb{P}(y = y_k) + \sum_{k=k_0+1}^{\infty} y_k \mathbb{P}(y = y_k) \\ &\leq \delta \sum_{k=1}^{k_0} \mathbb{P}(y = y_k) + L \sum_{k=k_0+1}^{\infty} \mathbb{P}(y = y_k) = \delta \mathbb{P}(y \leq y_{k_0}) + L \mathbb{P}(y \geq y_{k_0+1}) \\ &\leq \delta + L \frac{\delta}{L} = 2\delta. \end{aligned}$$

With this in mind, we now proceed to the proof of Lemma A.10.

Proof of Lemma A.10. Recall that we have defined $\bar{I} := \sum_{n=1}^N I_n$, $\tilde{I} := I_1 I_2 \dots I_N$ and let $\Omega \subset [I_1] \times \dots \times [I_N]$ be the set of p non-zero entries of \mathbf{V} . Using equation (4.41) in [16] we have

$$\|\mathbf{V}\|_M^* \leq \sup_{\mathbf{u} \in \mathcal{T}_{\pm}} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \right|.$$

Notice that the term on the right hand side is a discrete random variable, taking values in $[0, p]$. For fixed $\mathbf{u} \in \mathcal{T}_{\pm}$, a standard Hoeffding's inequality for bounded random variables gives for $t > 0$

$$\mathbb{P} \left(\left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \right| \geq t \right) \leq 2 \exp \left(-\frac{t^2}{2p} \right).$$

Since $|\mathcal{T}_{\pm}| \leq 2^{\bar{I}}$ (see [18, 16]), a union bound and choosing $t = 2\sqrt{p\bar{I}}$ provides

$$\mathbb{P} \left(\sup_{\mathbf{u} \in \mathcal{T}_{\pm}} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \right| \geq 2\sqrt{p\bar{I}} \right) \leq 2^{\bar{I}+1} e^{-2\bar{I}} \leq e^{\bar{I}+1} e^{-2\bar{I}} = e^{-\bar{I}+1}.$$

Equally, for any $h > 0$ we have shown

$$\mathbb{P} \left(\sup_{\mathbf{u} \in \mathcal{T}_{\pm}} \left| \sum_{\mathbf{i} \in \Omega} v_{\mathbf{i}} u_{\mathbf{i}} \right|^h \geq \left(2\sqrt{p\bar{I}} \right)^h \right) \leq e^{-\bar{I}+1},$$

and by Theorem A.11 we end the proof if $e^{-\bar{I}+1} < \frac{(2\sqrt{p\bar{I}})^h}{p^h}$. To this end, using the Maclaurin series of the exponential function, we note that for any $\ell \in \mathbb{N}$

$$\exp(2(\bar{I}-1)/h) \geq \frac{2^\ell(\bar{I}-1)^\ell}{h^\ell \ell!} \geq \frac{2^\ell(\bar{I}-1)^\ell}{(h\ell)^\ell}$$

which in particular holds for the non-integer choice $\ell := \log_2(\tilde{I}/(4\bar{I}))$ in the last term. Recall our assumption $\bar{I}-1 \geq h \log_2(\tilde{I}/(4\bar{I})) = h\ell$, so that

$$p \leq \tilde{I} = 2^\ell 4\bar{I} \leq \frac{2^\ell 4\bar{I}(\bar{I}-1)^\ell}{(h\ell)^\ell} \leq 4\bar{I} \exp(2(\bar{I}-1)/h). \quad \blacksquare$$

We have shown $p \leq 4\bar{I} \exp(2(\bar{I}-1)/h)$, raising both sides to the power of $h/2$ and rearranging gives the desired inequality. We conclude $\mathbb{E}(\|\mathbf{V}\|_M^*)^h \leq 2 \left(2\sqrt{p\bar{I}}\right)^h$.

REFERENCES

- [1] E. ACAR AND B. YENER, *Unsupervised multiway data analysis: A literature survey*, IEEE Transactions on Knowledge and Data Engineering, 21 (2009), pp. 6–20.
- [2] B. W. BADER, T. G. KOLDA, ET AL., *Matlab tensor toolbox version 3.0-dev*. Available online, August 2017.
- [3] D. N. BARRON, *The analysis of count data: Overdispersion and autocorrelation*, Sociological Methodology, 22 (1992), pp. 179–220.
- [4] S. BECKER, *L-BFGS-B-C*, 2019.
- [5] A. BLASCO-MORENO, M. PÉREZ-CASANY, P. PUIG, M. MORANTE, AND E. CASTELLS, *What does a zero mean? understanding false, random and structural zeros in ecology*, Methods in Ecology and Evolution, 10 (2019), pp. 949–959.
- [6] R. H. BYRD, P. LU, AND J. NOCEDAL, *A limited memory algorithm for bound constrained optimization*, SIAM Journal on Scientific and Statistical Computing, 16 (1995), pp. 1190–1208.
- [7] C. CAI, G. LI, H. V. POOR, AND Y. CHEN, *Nonconvex low-rank tensor completion from noisy data*, in Advances in Neural Information Processing Systems, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, eds., vol. 32, Curran Associates, Inc., 2019.
- [8] E. J. CANDÉS AND T. TAO, *The power of convex relaxation: Near-optimal matrix completion*, IEEE Transactions on Information Theory, 56 (2010), pp. 2053–2080.
- [9] Y. CAO AND Y. XIE, *Poisson matrix recovery and completion*, IEEE Transactions on Signal Processing, 64 (2016), pp. 1609–1620.
- [10] V. CHANDRASEKARAN, B. RECHT, P. A. PARRILO, AND A. S. WILLSKY, *The convex geometry of linear inverse problems*, Foundations of Computational Mathematics, 12 (2012), pp. 805–849.
- [11] E. C. CHI AND T. G. KOLDA, *On tensors, sparsity, and nonnegative factorizations*, SIAM Journal on Matrix Analysis and Applications, 33 (2012), pp. 1272–1299.
- [12] J. CHIQUET, S. ROBIN, AND M. MARIADASSOU, *Variational inference for sparse network reconstruction from count data*, in Proceedings of the 36th International Conference on Machine Learning, K. Chaudhuri and R. Salakhutdinov, eds., vol. 97 of Proceedings of Machine Learning Research, PMLR, 09–15 Jun 2019, pp. 1162–1171.
- [13] H. CHOI, J. GIM, AND S. WON, *Network analysis for count data with excess zeros*, BMC Genet, 18 (2017).
- [14] S. COXE, S. G. WEST, AND L. S. AIKEN, *The analysis of count data: A gentle introduction to poisson regression and its alternatives*, Journal of Personality Assessment, 91 (2009), pp. 121–136. PMID: 19205933.

- [15] M. P. FRIEDLANDER AND K. HATZ, *Computing non-negative tensor factorizations*, Optimization Methods and Software, 23 (2008), pp. 631–647.
- [16] N. GHADERMARZY, *Near-Optimal Sample Complexity for Noisy or 1-bit Tensor Completion*, PhD thesis, University of British Columbia, 2018.
- [17] N. GHADERMARZY, Y. PLAN, AND Ö. YILMAZ, *Learning tensors from partial binary measurements*, IEEE Transactions on Signal Processing, 67 (2019), pp. 29–40.
- [18] N. GHADERMARZY, Y. PLAN, AND Ö. YILMAZ, *Near-optimal sample complexity for convex tensor completion*, Information and Inference: A Journal of the IMA, 8 (2018), pp. 577–619.
- [19] M. S. GILTHORPE, M. FRYDENBERG, Y. CHENG, AND V. BAELUM, *Modelling count data with excessive zeros: The need for class prediction in zero-inflated models and the issue of data generation in choosing between zero-inflated and generic mixture models for dental caries data*, Statistics in Medicine, 28 (2009), pp. 3539–3553.
- [20] C. J. HILLAR AND L.-H. LIM, *Most tensor problems are np-hard*, Journal of the AMC, 60 (2013).
- [21] D. HONG, T. G. KOLDA, AND J. A. DUERSCH, *Generalized canonical polyadic tensor decomposition*, SIAM Review, 62 (2020), pp. 133–163.
- [22] P. KLIMEK, Y. YEGOROV, R. HANEL, AND S. THURNER, *Statistical detection of systematic election irregularities*, Proceedings of the National Academy of Sciences, 109 (2012), pp. 16469–16473.
- [23] T. G. KOLDA AND B. W. BADER, *Tensor decompositions and applications*, SIAM Review, 51 (2009), pp. 455–500.
- [24] T. G. KOLDA AND D. HONG, *Stochastic gradients for large-scale tensor decomposition*, SIAM Journal on Mathematics of Data Science, 2 (2020), pp. 1066–1095.
- [25] A. KRISHNAMURTHY AND A. SINGH, *Low-rank matrix and tensor completion via adaptive sampling*, in Advances in Neural Information Processing Systems, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, eds., vol. 26, Curran Associates, Inc., 2013.
- [26] M. LEDOUX AND M. TALAGRAND, *Probability in Banach Spaces: Isoperimetry and Processes*, vol. 23, Springer-Verlag, 1991.
- [27] L.-H. LIM AND P. COMON, *Nonnegative approximations of nonnegative tensors*, Journal of Chemometrics, 23 (2009), pp. 432–441.
- [28] H. MOGOUIE, G. A. RAISSI ARDALI, A. AMIRI, AND E. BAHRAMI SAMANI, *Monitoring attributed social networks based on count data and random effects*, Scientia Iranica, (2020), pp. –.
- [29] A. K. MUOKA, O. O. NGESE, AND A. G. WAITITU, *Statistical models for count data*, Science Journal of Applied Mathematics and Statistics, 4 (2016), pp. 256–262. PMID: 19205933.
- [30] I. J. MYUNG, *Tutorial on maximum likelihood estimation*, Journal of Mathematical Psychology, 47 (2003), pp. 90–100.
- [31] A. NOUY, *Low-Rank Tensor Methods for Model Order Reduction*, Springer International Publishing, Cham, 2017, pp. 857–882.
- [32] D. POLLARD, *A User’s Guide to Measure Theoretic Probability*, vol. 8, Cambridge University Press, 2002.
- [33] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, 2015.
- [34] K. SANTRA, E. A. SMITH, J. W. PETRICH, AND X. SONG, *Photon counting data analysis: Application of the maximum likelihood and related methods for the determination of lifetimes in mixtures of rose bengal and rhodamine b*, Journal of Physical Chemistry. A, Molecules, Spectroscopy, Kinetics, Environment, and General Theory, 121 (2016).
- [35] J. TACHELLA, Y. ALTMANN, AND N. MELLADO, *Real-time 3d reconstruction from single-photon lidar data using plug-and-play point cloud denoisers*, Nature Communications, 10 (2019).
- [36] M. YUAN AND C. ZHANG, *On tensor completion via nuclear norm minimization*, Found. Comput. Math, 16 (2016), pp. 1031–1068.
- [37] O. B. YUSUF, R. F. AFOLABI, AND A. S. AGBAJE, *Modelling excess zeros in count data with application to antenatal care utilisation*, International Journal of Statistics and Probability, 7 (2018).
- [38] N. S. N. ZAMRI AND Z. H. ZAMZURI, *A review on models for count data with extra zeros*, AIP Conference Proceedings, 1830 (2017), p. 080010.