
Under-Counted Tensor Completion with Neural Incorporation of Attributes

Shahana Ibrahim¹ Xiao Fu¹ Rebecca Hutchinson¹ Eugene Seo²

Abstract

Systematic under-counting effects are observed in data collected across many disciplines, e.g., epidemiology and ecology. *Under-counted tensor completion* (UC-TC) is well-motivated for many data analytics tasks, e.g., inferring the case numbers of infectious diseases at unobserved locations from under-counted case numbers in neighboring regions. However, existing methods for similar problems often lack supports in theory, making it hard to understand the underlying principles and conditions beyond empirical successes. In this work, a low-rank Poisson tensor model with an expressive unknown *nonlinear* side information extractor is proposed for under-counted multi-aspect data. A joint low-rank tensor completion and neural network learning algorithm is designed to recover the model. Moreover, the UC-TC formulation is supported by theoretical analysis showing that the fully counted entries of the tensor and each entry’s under-counting probability can be provably recovered from partial observations—under reasonable conditions. To our best knowledge, the result is the first to offer theoretical supports for under-counted multi-aspect data completion. Simulations and real-data experiments corroborate the theoretical claims.

1. Introduction

Tensor completion (TC) has been a workhorse for a large variety of data analytics tasks, e.g., image/video inpainting (Liu et al., 2013), hyperspectral denoising (Wang et al., 2018), sampling and recovery in magnetic resonance imaging (MRI) (Kanatsoulis et al., 2020), and multimodal data mining (Papalexakis et al., 2017)—just to name a few. In the past decade, theory and methods of TC progressed

significantly—aspects such as recoverability, sample complexity, optimal algorithm design were extensively studied under various tensor models, e.g., *Tucker decomposition* (Mu et al., 2014), *canonical polyadic decomposition* (CPD) (Kanatsoulis et al., 2020; Sørensen & De Lathauwer, 2019), and *block-term decomposition* (BTD) (Zhang et al., 2020a; Ding et al., 2020). The vast majority of classic TC methods were proposed for handling real-valued data, but TC algorithms designed for integer data, e.g., binary data (Ghadermarzy et al., 2018) and count data (Chi & Kolda, 2012), also exist. This is because such data often arise in real-world problems, e.g., movie recommendation (Karatzoglou et al., 2010), knowledge graph completion (Balazevic et al., 2019), and social network analysis (Kashima et al., 2009).

However, less attention has been paid to integer data TC problems where the observed entries suffer from systematic *under-counting*—but (strong) under-counting effects are encountered in many disciplines, e.g., epidemiology and ecology. In epidemiology, the cases of an infectious disease (e.g., COVID-19) may be under-counted due to the existence of symptom-free patients and the lack of testing (Schneider, 2020). In ecology, when an observer sees a species at a certain location, the observer likely just observes a small portion of this species’ population (MacKenzie et al., 2006; Tylianakis et al., 2010). Hence, taking under-counting effects into consideration for TC is meaningful in these domains. For example, recovering the “true” (fully counted) number of COVID-19 cases at a certain location and time helps estimate/understand the infection situation. The task of *under-counted tensor completion* (UC-TC) is naturally more challenging than the conventional (integer data) TC, as it implicitly involves an extra task of compensating the *unknown* under-counting effect.

1.1. Prior Work and Challenges

In the literature, there exist some limited efforts towards dealing with under-counted data completion/prediction. For example, the line of work for user exposure matrix completion considered link prediction under undetected links and unobserved links simultaneously (Liang et al., 2016). The work on “zero-inflated” matrix/tensor factorization and related models (Billio et al., 2022; Simchowit, 2013; Durif et al., 2019) considered a similar scenario, where the datasets’ observed entries are zeros with a certain non-zero

¹School of Electrical Engineering and Computer Science, Oregon State University, Corvallis, OR 97331, USA ²Department of Applied Mathematics, Brown University, Providence, RI 02912, USA. Correspondence to: Xiao Fu <xiao.fu@oregonstate.edu>.

probability. The recent work by (Fu et al., 2021) proposed an approach to tackle under-counted matrix completion (UC-MC) using a Poisson-Binomial matrix factorization model—which was inspired by the species’ abundance modeling in ecology (Joseph et al., 2009; Royle, 2004; Hutchinson et al., 2011).

Challenges. The existing works by (Liang et al., 2016; Billio et al., 2022; Simchowit, 2013; Durif et al., 2019; Fu et al., 2021) were shown effective on many UC-MC tasks. With proper modifications, these computational frameworks may be generalized to handle tensor data. However, some notable challenges remain. First, how to effectively model the under-counting effect has been an open question. For example, the works by (Liang et al., 2016; Billio et al., 2022; Simchowit, 2013; Durif et al., 2019) all considered zero-inflated MC models, which dealt with a special case of under-counting, i.e., the observed matrix entries are either the true counts or zeros. For more general under-counted data (i.e., the observed counts are smaller than or equal to the true counts), the recent work by (Fu et al., 2021) proposed to model the probability of missing a count (e.g., an interaction of a pollinator and a plant) as a *linear function* of side attributes, e.g., the location, temperature, and humidity—which is reminiscent of the classic species distribution models (SDMs) (Hutchinson et al., 2011). However, using simple linear functions is a compromise between model expressiveness and computational convenience. Second, the existing works demonstrated their effectiveness only *empirically*. It has been unclear under what conditions the underlying true counts and some key model parameters (e.g., the detection probabilities of the counts) could be *provably* recovered.

1.2. Contributions

In this work, we propose a UC-TC framework based on the Poisson-Binomial model (Joseph et al., 2009; Royle, 2004; Hutchinson et al., 2011) as in the MC work by (Fu et al., 2021). Unlike (Fu et al., 2021) that used a simple side information model and was largely empirical, we address a number of key challenges in modeling, computation, and theory. Our contribution is twofold:

UC-TC with Nonlinear Incorporation of Side Information. To model and learn the under-counting effect in tensor data, we propose a UC-TC framework with neural network-based incorporation of attributes. Specifically, we extend the UC-MC idea in (Fu et al., 2021) to the higher-order tensor regime and model the underlying fully counted data as a Poisson tensor that admits a low-rank *canonical polyadic decomposition* (CPD) (Harshman, 1970); the Poisson tensor is then under-counted via a Binomial detection procedure. Different from the existing work that uses a linear function to model the relationship between the attributes and the

detection probability of a count, we represent the relationship as an *unknown nonlinear function*, which substantially generalizes the modeling capacity. We recast the UC-TC problem as a *maximum likelihood estimation* (MLE) criterion, where the unknown nonlinear function is represented by a neural network (NN) for side information incorporation. We also design a *block coordinate descent* (BCD) algorithm (Razaviyayn et al., 2013) to tackle the formulated MLE effectively.

Recoverability Analysis of UC-TC. We present theoretical characterizations of the proposed UC-TC criterion under the Poisson-Binomial framework. We show that, under reasonable conditions, the fully counted tensor is recoverable, and the associated detection probabilities of all the entries are identifiable—both are up to a global scaling/counter-scaling ambiguity. Our analysis reveals interesting performance-deciding factors such as the similarity of the under-counting effects across some entries. The result also shows that the nonlinear incorporation of side information imposes an intuitive trade-off between model complexity and recovery accuracy, given a fixed amount of samples—which also justifies our proposal of using neural nonlinear models. To our best knowledge, this is the first provable learning criterion for under-counted factor analysis models.

We conduct extensive evaluation of the proposed approach on synthetic data to validate our theoretical findings. We also test our approach on real-world prediction tasks in epidemiology (Zhang et al., 2020b) and ecology (Seo et al., 2022; Daly et al., 2019).

2. Background

2.1. UC-TC: Problem Statement

A K th-order tensor $\underline{\mathbf{Y}} \in \mathbb{R}^{I_1 \times \dots \times I_K}$ is a multi-way array whose entries are indexed by K coordinates. An entry of such a tensor can be expressed as follows:

$$y_{\mathbf{i}} = [\underline{\mathbf{Y}}]_{\mathbf{i}}, \quad \mathbf{i} = (i_1, \dots, i_K) \in [I_1] \times \dots \times [I_K]. \quad (1)$$

Tensors are widely used to represent multi-aspect data. For example, in a four-aspect epidemic data, $y_{\mathbf{i}}$ can represent the number of recorded infected cases in city i_1 , during week i_2 , among age group i_3 and ethnicity group i_4 .

TC and MC tasks naturally arise in many applications, where multi-aspect data are only observed partially. A classic use case is recommender systems (Hu et al., 2008; Karatzoglou et al., 2010), where the ratings that users provide to products are used to predict the unseen ratings—which can be done via completing the rating matrix/tensor. TC problems were primarily considered for continuous data (Gandy et al., 2011; Montanari & Sun, 2016; Yuan & Zhang, 2016; Zhang & Aeron, 2017; Sørensen & De Lathauwer, 2019; Zhang et al., 2020a), but TC for other data formats (e.g., binary

and integer data) were also studied in the literature (Ghademarzy et al., 2018; Chi & Kolda, 2012; Lee & Wang, 2020)—as the latter found many applications in real data analytics problems, e.g., 5-star rating-based recommender systems (Karatzoglou et al., 2010), adjacency network analysis (Acar et al., 2009) and traffic flow prediction (Tan et al., 2016).

In this work, our interest lies in integer TC problems where the partially observed tensor data suffer from systematical under-counting. To be specific, we consider integer tensor data as in (Chi & Kolda, 2012; Lee & Wang, 2020). Like in conventional integer TC problems, the observed tensor $\underline{\mathbf{Y}}$ is highly incomplete; i.e., $\Omega \subseteq [I_1] \times \dots \times [I_K]$ are the indices of the observed entries, where $|\Omega| \ll \prod_{k=1}^K I_k$. The key difference from conventional integer TC is that in our setting, the observed entries are smaller than or equal to their actual values, i.e., $y_i \leq n_i$, $i \in \Omega$, where $n_i = [\underline{\mathbf{N}}]_i$, $\underline{\mathbf{N}} \in \mathbb{Z}_+^{I_1 \times \dots \times I_K}$ represents the (unobserved) true-count tensor, and $y_i \in \mathbb{Z}_+$ is the observed and under-counted version of n_i .

We assume that for each entry of $\underline{\mathbf{Y}}$, there is an associated attribute/feature vector $\mathbf{z}_i \in \mathbb{R}^D$, capturing side information regarding the observation. The attributes reflect the *detection probability* of each count in n_i . We note that such entry attributes are often available in real-world applications. Two examples are as follows:

- **Epidemiology Data:** Consider tensor data in epidemiology that records infected case counts. The count tensor could be indexed by ‘city’ \times ‘age group’ \times ‘profession’. Then, \mathbf{z}_i may contain observation-affecting features such as the testing capacity of city i_1 , population of age group i_2 , and the level of exposure to the virus for profession i_3 .
- **Ecology Data:** In ecology, the counts of interactions between pollinators and plants at different times could be represented by a tensor (i.e., a ‘pollinator’ \times ‘plant’ \times ‘time’ tensor). There, \mathbf{z}_i could contain the properties of the pollinator i_1 and plant i_2 , the temperature, humidity and other observation-affecting factors of site i_3 .

In this work, we assume that the entry features \mathbf{z}_i are available for a subset of entries indexed by $\Xi \subseteq [I_1] \times \dots \times [I_K]$. Note that Ξ and Ω need not be identical.

2.2. Existing Work: The Poisson-Binomial MC Model

The work in (Fu et al., 2021) proposed an MC framework with the following generative model for the under-counted entries of a pollinator-plant interaction count matrix:

$$n_{i,j} \sim \text{Poisson}(\lambda_{i,j}), \lambda_{i,j} = \mathbf{u}_i^\top \mathbf{v}_j, \mathbf{u}_i, \mathbf{v}_j \geq \mathbf{0}, \quad (2a)$$

$$y_{i,j} \sim \text{Binomial}(n_{i,j}, p_{i,j}), p_{i,j} = \mathbf{z}_{i,j}^\top \boldsymbol{\theta}, \quad (2b)$$

where $\mathbf{u}_i \in \mathbb{R}^F$, $\mathbf{v}_j \in \mathbb{R}^F$ and $\boldsymbol{\theta} \in \mathbb{R}^D$ are the model parameters to estimate. To be more specific, \mathbf{u}_i and \mathbf{v}_j are the F -dimensional latent representations of pollinator i and plant j . The parameter $\lambda_{i,j}$ stands for the average number of interactions between pollinator i and plant j , and the ground-truth number of interactions $n_{i,j}$ is the realization of a Poisson sampling process with parameter $\lambda_{i,j}$. The observed data $y_{i,j}$ is under-counted through a Binomial detection process, with the detection probability $p_{i,j}$. The detection probability $p_{i,j}$ was modeled as a linear function of the entry features $\mathbf{z}_{i,j}$. Under this model, $\Lambda = \mathbf{U}\mathbf{V}^\top$ where $[\Lambda]_{i,j} = \lambda_{i,j}$ is a low-rank matrix model if F is relatively small. Hence, estimating \mathbf{u}_i ’s, \mathbf{v}_j ’s, and $\boldsymbol{\theta}$ can be regarded as a variant of Poisson matrix completion problem.

2.3. Challenges

To extend the Poisson-Binomial model to tensor cases and more general settings, there are a couple of key challenges to be addressed. First, the linear model in (2b), i.e., $p_{i,j} = \mathbf{z}_{i,j}^\top \boldsymbol{\theta}$, may not be expressive enough, as the relation between the features and the detection probability is highly likely to be *nonlinear*. The linear model assumption was used for computational convenience, but is a performance-limiting factor in the model of (Fu et al., 2021) (as will be seen in the experiments). Second, perhaps more importantly, there were no theoretical characterizations of the Poisson-Binomial model, despite of its popularity in ecological data analysis (Fu et al., 2021; Hutchinson et al., 2011; Royle, 2004; Dennis et al., 2015). In the context of UC-MC, it is unclear if the underlying Poisson parameters $\lambda_{i,j}$ for all (i, j) could be recovered from a subset of observations $y_{i,j}$, $(i, j) \in \Omega$. The recoverability analysis can not be covered by existing TC theory, as the ‘‘extra’’ task of estimating $p_{i,j}$ makes UC-TC a much harder learning problem compared to the traditional TC problems.

3. Problem Formulation

In this work, we generalize the Poisson-Binomial model to cover tensor data, with *nonlinear* attribute incorporation—and more importantly—with recoverability support.

3.1. Generative Model

We generalize the model (2) by considering under-counted tensors with a size of $I_1 \times \dots \times I_K$ generated as follows:

$$\lambda_i = \sum_{f=1}^F \prod_{k=1}^K U_k(i_k, f), \mathbf{U}_k \geq \mathbf{0}, \forall k \in [K], \quad (3a)$$

$$n_i \sim \text{Poisson}(\lambda_i), \quad (3b)$$

$$p_i = g(\mathbf{z}_i), 0 \leq p_i \leq 1, \quad (3c)$$

$$y_i \sim \text{Binomial}(n_i, p_i), \quad (3d)$$

where $\mathbf{i} = (i_1, \dots, i_K)$ is a shorthand notation as defined before, $\mathbf{U}_k(i_k, \cdot) \in \mathbb{R}^F$ denotes the F -dimensional latent representation of entity i_k of aspect k , y_i denotes the *observed* count indexed by \mathbf{i} , n_i is the corresponding true count, λ_i stands for the Poisson parameter (which is the expectation of the true count), and p_i represents the detection probability as in the model of (2). The vector $\mathbf{z}_i \in \mathbb{R}^D$ collects the features of entry \mathbf{i} . The *unknown* nonlinear function $g(\cdot) : \mathbb{R}^D \rightarrow [0, 1]$ is used to model the dependence between the detection probability p_i and \mathbf{z}_i . In (3), we have imposed nonnegativity constraints on \mathbf{U}_k 's in order to make sure that the Poisson parameters are all nonnegative.

Given y_i for $\mathbf{i} \in \Omega \subset [I_1] \times \dots \times [I_K]$ and \mathbf{z}_i for $\mathbf{i} \in \Xi \subset [I_1] \times \dots \times [I_K]$, our goal is to estimate λ_i and p_i for all \mathbf{i} .

3.2. Maximum Likelihood Estimation-Based UC-TC

Assuming that y_i 's are independently sampled from the generative model in (3), the log-likelihood of the observations can be expressed as follows:

$$\begin{aligned} & \log \prod_{\mathbf{i} \in \Omega} \Pr(y_i; \lambda_i, p_i) \\ &= \log \left(\prod_{\mathbf{i} \in \Omega} \sum_{n=y_i}^{\infty} \Pr(N_i = n; \lambda_i) \Pr(y_i | N_i = n; p_i) \right) \\ &= \sum_{\mathbf{i} \in \Omega} \log \left(\sum_{n=y_i}^{\infty} \left(\frac{\lambda_i^n e^{-\lambda_i} n! p_i^{y_i} (1-p_i)^{n-y_i}}{n! y_i! (n-y_i)!} \right) \right) \\ &= \sum_{\mathbf{i} \in \Omega} [y_i \log \lambda_i + y_i \log p_i - \lambda_i p_i - \log y_i!]; \quad (4) \end{aligned}$$

see the detailed derivation in the supplementary material in Sec. B. Under the generative model in (3), we have $p_i = g(\mathbf{z}_i)$. Since $g(\cdot)$ is unknown and nonlinear, we represent it using a neural network denoted as $g_\theta : \mathbb{R}^D \rightarrow [0, 1]$:

$$g_\theta(\mathbf{z}) = \sigma(\mathbf{w}_L^\top \sigma(\mathbf{W}_{L-1} \sigma(\dots \sigma(\mathbf{W}_1 \mathbf{z}))))), \quad (5)$$

where $\mathbf{W}_\ell \in \mathbb{R}^{d_\ell \times d_{\ell-1}}$ denotes the weight matrix at ℓ th layer ($d_0 = D$), $\mathbf{w}_L \in \mathbb{R}^{d_L}$ denotes the last layer's weights, θ denotes the parameters of the NN ($\{\mathbf{W}_\ell\}_{\ell=1}^{L-1}, \mathbf{w}_L$), and $\sigma(\cdot) = [\sigma(\cdot), \dots, \sigma(\cdot)]^\top$ denotes the activation function with a proper dimension. Hence, combining (4) with (3a) and (3c), we formulate the MLE criterion as follows:

$$\begin{aligned} \underset{\{\mathbf{U}_k\}, \theta, \{p_i\}}{\text{minimize}} \quad & L(\mathbf{U}, \theta, \mathbf{P}) := \sum_{\mathbf{i} \in \Omega} \left[\left(\sum_{f=1}^F \prod_{k=1}^K U_k(i_k, f) \right) p_i \right. \\ & \left. - y_i \log \left(\sum_{f=1}^F \prod_{k=1}^K U_k(i_k, f) \right) - y_i \log p_i \right], \quad (6a) \end{aligned}$$

$$\text{subject to } \mathbf{U}_k \in \mathcal{U}_k, \forall k \in [K], \quad (6b)$$

$$p_i = g_\theta(\mathbf{z}_i), \forall \mathbf{i} \in \Xi, g_\theta \in \mathcal{G}, \quad (6c)$$

where \mathcal{U}_k is a constraint set that encodes the prior knowledge of \mathbf{U}_k (e.g., nonnegativity), and \mathcal{G} represents the function class where g_θ is taken from. Given the formulated MLE, there are a couple of key aspects that are of interest. First, if one solves the formulated learning criterion in (6), can the *optimal* solution recover the latent parameters λ_i and p_i over all \mathbf{i} ? This question will be answered in Sec. 4. Second, the criterion (6) presents a challenging optimization problem that involves tensor decomposition and neural network learning—both are nontrivial problems. A BCD algorithm to effectively tackle this problem will be proposed in Sec. 5.

4. Recoverability Analysis

Our goal is to show that the optimal solution given by (6) provably recovers λ_i 's and p_i 's under the generative model (3) (up to certain reasonable ambiguities).

4.1. Analysis Setup

To better present the analysis, we use “ $\hat{\cdot}$ ” to denote the ground-truth parameters in the model (3)—e.g., $\lambda_i^{\hat{}}$ and $p_i^{\hat{}}$ denote the ground-truth Poisson parameter and ground-truth detection probability of entry \mathbf{i} , respectively. We have the corresponding tensor notations $[\hat{\mathbf{A}}]_{\mathbf{i}} = \lambda_i^{\hat{}}$ and $[\hat{\mathbf{P}}]_{\mathbf{i}} = p_i^{\hat{}}$. The “hat” notation is used to denote the terms constructed from the optimal solution given by (6). For instance, $\{\hat{\mathbf{U}}_k\}_{k=1}^K, \hat{g}_\theta$ and $\hat{p}_i, \forall \mathbf{i}$ denote the estimates given by the solution obtained from solving (6). We consider the following assumptions:

Assumption 4.1 (Bounded Parameters). There exist scalars $p_{\min}, p_{\max}, \beta_u, \alpha_u > 0$ such that $p_i^{\hat{}} = g^{\hat{}}(\mathbf{z}_i)$, $0 < p_{\min} \leq p_i^{\hat{}} \leq p_{\max}, \forall \mathbf{i}$ and $0 < \beta_u \leq \mathbf{U}_k^{\hat{}}(i_k, f) \leq \alpha_u, \forall i_k \in [I_k], f \in [F], k \in [K]$. In addition, the constraints in (6), namely, \mathcal{U}_k and \mathcal{G} satisfy $\mathcal{U}_k = \{\mathbf{U} \in \mathbb{R}^{I_k \times F} \mid \beta_u \leq \mathbf{U}(i_k, f) \leq \alpha_u, \forall i_k, f\}$ and $\mathcal{G} = \{g_\theta : \mathbb{R}^D \rightarrow [p_{\min}, p_{\max}]\}$, respectively.

Assumption 4.2 (Approximation Error). Assume that there exists $\tilde{g}_\theta \in \mathcal{G}$ such that $|\tilde{g}_\theta(\mathbf{z}) - g^{\hat{}}(\mathbf{z})| \leq \nu$ for all $\mathbf{z} \in \mathbb{R}^D$, where $0 \leq \nu < \infty$. In addition, assume that the class \mathcal{G} has a complexity measure $\mathcal{R}_{\mathcal{G}}$.

Assumption 4.3 (Similar Attribute Subset). There exists an index set Θ whose elements are sampled uniformly at random (without replacement) from $[I_1] \times \dots \times [I_K]$ and a scalar $\zeta > 0$ such that $\max_{\mathbf{i}, \mathbf{j} \in \Theta} \|\mathbf{z}_i - \mathbf{z}_j\|_2 \leq \zeta$.

Assumption 4.4 (Lipschitz Continuity). Assume that the ground-truth function $g^{\hat{}}$ and any function $g_\theta \in \mathcal{G}$ are Lipschitz continuous, i.e., for certain $L_g, L_\theta > 0$ and for any pair of $\mathbf{z}_i, \mathbf{z}_j$, $|g^{\hat{}}(\mathbf{z}_i) - g^{\hat{}}(\mathbf{z}_j)| \leq L_g \|\mathbf{z}_i - \mathbf{z}_j\|_2$, and $|g_\theta(\mathbf{z}_i) - g_\theta(\mathbf{z}_j)| \leq L_\theta \|\mathbf{z}_i - \mathbf{z}_j\|_2$ hold.

Assumption 4.2 requires that the neural network class \mathcal{G} contains a function \tilde{g}_θ that is closer to the ground-truth func-

tion g^\natural . In practice, if deeper/wider neural networks are employed, the value of ν is smaller. For the complexity measure \mathcal{R}_G used in Assumption 4.2, we use the *sensitive complexity* parameter introduced in (Lin & Zhang, 2019). For deeper and wider neural network function classes, the parameter \mathcal{R}_G gets larger—also see supplementary material Sec. M for more details. Assumptions 4.3-4.4 together imply that there exists a set of p_i 's that are similar. For example, in an epidemic dataset (e.g., the COVID-19 dataset (Zhang et al., 2020b)) that records the number of detected infectious disease cases, when the testing capacity and the population of a number of locations are similar with each other, it is reasonable to assume that the corresponding detection probabilities p_i 's are similar. The Lipschitz continuity assumption on g_θ given by Assumption 4.4 requires that the learned g_θ is smooth enough, which can be achieved via bounding (or penalizing) the norm of the neural network parameter θ during the learning process.

4.2. Main Result

Our main result is as follows:

Theorem 4.5. *Suppose that the Assumptions 4.1-4.4 hold true. Assume that the entries of index sets Ω and Ξ are sampled from $[I_1] \times \dots \times [I_K]$ uniformly at random with replacement and satisfies $\Omega \subseteq \Xi$. In addition, let $\mathbf{z}_i, \dots, \mathbf{z}_{i_S}$ be the set of observed features. Assume that each \mathbf{z}_i is drawn from a distribution \mathcal{D} . Then, for $\delta > 0$, the following hold with probability of at least $1 - 5\delta - 3e^{-\alpha(e^2-3)}$:*

$$\left\| \frac{\hat{\Lambda}^\natural - \widehat{\xi} \widehat{\Lambda}}{\prod_k I_k} \right\|_F \leq \mathcal{O} \left(\frac{1}{p_{\min}^2} \varrho_1 \right), \quad (7a)$$

$$\mathbb{E}_{\mathbf{z}_i \sim \mathcal{D}} \left[(p_i^\natural - 1/\widehat{\xi} \widehat{g}_\theta(\mathbf{z}_i))^2 \right] \leq \mathcal{O} \left(\frac{p_{\max}}{\beta p_{\min}^2} \varrho_1 + \varrho_2 \right), \quad (7b)$$

where $\varrho_1 = F\alpha_u^K \left(\eta + \zeta^2(L_g + L_\theta)^2 + \sqrt{\frac{\log(1/\delta)}{|\Theta|}} \right)$,

$$\varrho_2 = \sqrt{\frac{\log(1/\delta)}{S}} + \frac{(\|\mathbf{Z}\|_F \mathcal{R}_G)^{1/4}}{\sqrt{S}},$$

$$\eta^2 = \frac{K \sum_k I_k + \|\mathbf{Z}\|_F \mathcal{R}_G}{\sqrt{T}} + c \sqrt{\frac{\log(1/\delta)}{T}} + \alpha \nu,$$

$c = \alpha \max\{|\log \beta|, \log \alpha\}$, $\alpha = F\alpha_u^K p_{\max}$, $\beta = F\beta_u^K p_{\min}$, $\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}] \in \mathbb{R}^{D \times S}$, $T = |\Omega|$, and $\widehat{\xi}$ denote the global scaling ambiguity.

The proof of Theorem 4.5 is relegated to the supplementary material in Sec. D. The result reveals several interesting insights. First, the results in (7) show that the proposed MLE criterion can correctly recover the average true counts λ_i^\natural 's and the detection probabilities p_i^\natural 's, up to a certain global scaling ambiguity between the entries. Second, the bounds

indicate that the estimation accuracy is better when there are more observations (i.e., larger $|\Omega|$) and more similar features (i.e., smaller ζ and larger $|\Theta|$). Third, the result suggests that an appropriate choice of the neural network class plays a role in ensuring good estimation accuracy. This is because if one uses deeper/wider neural networks, the approximation error ν becomes smaller, but the complexity parameter \mathcal{R}_G is larger. Hence, there is a trade-off to strike while choosing the network architecture of g_θ .

5. Algorithm Design

The proposed MLE in (6) is a combination of Poisson tensor decomposition and neural network learning—both are hard optimization problems. To tackle this problem, we design a BCD-based (Razaviyayn et al., 2013; Bertsekas, 1999) algorithm. We first re-formulate (6) using a regularized form by considering nonnegativity constraints for \mathbf{U}_k 's:

$$\underset{\{\mathbf{U}_k\}, \theta, \{p_i\}}{\text{minimize}} \frac{1}{|\Omega|} L(\mathbf{U}, \theta, \mathbf{P}) + \frac{\mu}{|\Xi|} \sum_{i \in \Xi} \ell(g_\theta(\mathbf{z}_i), p_i), \quad (8a)$$

$$\text{subject to } \mathbf{U}_k \geq \mathbf{0}, \forall k \in [K], \quad (8b)$$

$$0 \leq p_i \leq 1, \forall i \in \Omega \cup \Xi, g_\theta \in \mathcal{G}, \quad (8c)$$

where $\mu > 0$ is a regularization parameter and $\ell(\cdot, \cdot)$ denotes a certain distance/divergence measure, e.g., the least squares function $\ell(x, y) = (x - y)^2$.

Note that in the re-formulated problem (8), we did not explicitly use the bounds on \mathcal{U}_k and \mathcal{G} in Assumption 4.1. The reason is twofold: First, the bounds serve for analytical purposes, yet not easy to know exactly in practice. Second, ignoring the bounds is inconsequential in terms of performance (as will be seen in experiments), but substantially simplifies the algorithm design. Nonetheless, as one will see, the update of \mathbf{U}_k naturally results in positive and bounded solutions, under reasonable conditions (cf. the comments after Eq. 10). In addition, as g_θ can be constructed to have a sigmoid output layer, the value of $g_\theta(\mathbf{z}_i)$ is naturally bounded away from 0 and 1.

Also note that we adopt this regularized optimization design since it gives flexibility to handle different cases such as $\Omega \subseteq \Xi$ and $\Omega \supset \Xi$ —as we will see in the following section. Our BCD updates are designed as follows:

The \mathbf{U}_k -Subproblem. When $\theta, p_i, \mathbf{U}_j, j \neq k$ are fixed, the subproblem w.r.t. \mathbf{U}_k is given by

$$\underset{\mathbf{U}_k \geq \mathbf{0}}{\text{minimize}} \sum_{i \in \Omega} \left[\left(\sum_{f=1}^F \mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f) \right) p_i - y_i \log \left(\sum_{f=1}^F \mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f) \right) \right]. \quad (9)$$

From (9), the update for U_k is as follows:

$$U_k(i_k, f) \leftarrow \frac{\sum_{i \in \Omega} y_i \alpha_i^{(f)}}{\sum_{i \in \Omega} \prod_{j \neq k} U_j(i_j, f) p_i}, \forall i, f, \quad (10)$$

where $\alpha_i^{(f)} = \frac{\bar{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f)}{\sum_{f=1}^F \bar{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f)}$ and \bar{U}_k denotes U_k from the previous iteration. Note that the above solution is always strictly positive if $U_k > \mathbf{0}$ for all k and $p_i > 0$ for all i . Our update rule design for U_k follows the ideas in (Chi & Kolda, 2012; Fu et al., 2021), which is essentially a *majorization minimization* step w.r.t. U_k ; see Sec. C in the supplementary material.

The p -Subproblem. To update p_i , we fix U_k 's and θ and consider different cases.

Case 1. When both y_i and z_i are available, i.e., $i \in \Omega \cap \Xi$, p_i is updated via optimizing the following subproblem:

$$\min_{p_i \in [0,1]} \frac{1}{|\Omega|} (\lambda_i p_i - y_i \log(p_i)) + \frac{\mu}{|\Xi|} \ell(g_\theta(z_i), p_i), \quad (11)$$

where $\lambda_i = \sum_{f=1}^F \prod_{k=1}^K U_k(i_k, f)$ is obtained using the current estimates of U_k 's.

Case 2. Suppose that the observation y_i is not available, but z_i is available, i.e., $i \in \Xi - \Omega \cap \Xi$. Then, we consider the following subproblem to update such p_i 's:

$$\min_{p_i \in [0,1]} \ell(g_\theta(z_i), p_i). \quad (12)$$

If we choose a convex $\ell(\cdot)$ in (11)-(12), the problems can be optimally solved using any standard techniques such as projected gradient decent. Moreover, if we set $\ell(\cdot)$ to be one of the popular choices such Euclidean distance and KL divergence, closed-form solutions for the problems in (11)-(12) exist—see Table 6 in the supplementary material.

*Case 3.*¹ When y_i is observed, but z_i is not accessible, i.e., $i \in \Omega - \Omega \cap \Xi$, we update p_i via solving the following subproblem:

$$\min_{p_i \in (0,1]} \lambda_i p_i - y_i \log(p_i), \quad (13)$$

which gives the update rule for p_i as $p_i \leftarrow [y_i/\lambda_i]_{[0,1]}$.

The θ -Subproblem. By fixing the U_k 's and all the p_i 's, the subproblem for updating θ is as follows:

$$\min_{\theta} \sum_{i \in \Xi} \ell(g_\theta(z_i), p_i).$$

¹Our recoverability analysis is built upon the assumption $\Omega \subseteq \Xi$; see Theorem 4.5. However, in practice, there are cases where y_i is observed, but z_i is not available (i.e., $\Omega - \Omega \cap \Xi \neq \emptyset$). Our algorithm can still operate under such cases.

This is an unconstrained neural network learning problem. Many off-the-shelf neural network learning algorithms that use gradient, e.g., Adam (Kingma & Ba, 2015) and Adagrad (Duchi et al., 2011), can be used to handle this subproblem.

More design details and the description of the algorithm are provided in supplementary material in Sec. C. The proposed algorithm is referred to as the *Under-Counted Data Prediction Via Learner-Aided Tensor Completion* (UncleTC) algorithm.

A remark is that our algorithm uses a batch (instead of stochastic) update rule for the U_k 's. This can also be replaced by stochastic tensor decomposition algorithms (see, e.g., (Fu et al., 2020a; Pu et al., 2022) and the references in (Fu et al., 2020b)). Using stochastic algorithms for the U_k -subproblems may further improve efficiency.

6. Experiments

In this section, we evaluate our proposed method through a series of synthetic and real data experiments. The source code is available at <https://github.com/shahanaibrahimosu/undercounted-tensor-completion>.

Baselines. We employ a number of low-rank tensor completion-based baselines, namely NTF-CPD-KL (Chi & Kolda, 2012), HaLRTC (Liu et al., 2013), NTF-CPD-LS (Shashua & Hazan, 2005), BPTF-CPD (Schein et al., 2015), and NTF-Tucker-LS (Kim & Choi, 2007). Among them, both NTF-CPD-KL and BPTF consider Poisson modeling, but they do not have a Binomial detection stage in their models to accommodate the under-counting effect. We also consider two recent NN-based tensor completion methods, CostCo (Liu et al., 2019) and POND (Tillinghast et al., 2020), as baselines. In order to incorporate the side features, the CostCo method is slightly modified from its original implementation by concatenating side features with the indices of the observed entries as the input to their networks. In addition, we also include a tensor version of (Fu et al., 2021), where a linear function $g_\theta(z_i) = \theta_1^\top z_i + \theta_2$ is used to learn g function. This baseline is referred to as UncleTC (Linear). More details on the implementation are provided in the supplementary material in Sec. N.

6.1. Synthetic-Data Simulations

Data Generation. The data generating process follows (3). We first generate $U_k \in \mathbb{R}^{I_k \times F}$, $\forall k$ by randomly sampling its entries from a uniform distribution between 0 and κ . We fix $K = 3$, $F = 3$ and $I_k = 20$, $\forall k$, unless specified otherwise. We also fix $\kappa = 15$ so that λ_i 's are not unreasonably small. The feature vectors $z_i \in \mathbb{R}^D$

Table 1. Average MAE_p and MAE_λ over 20 random trials under various values of γ_Ω ; $\gamma_\Xi = 0.3, \gamma_\Theta = 0.2, \text{SNR} = 40\text{dB}$ with $\Theta \cup \Omega \subseteq \Xi$.

Algorithm	Metric	$\gamma_\Omega = 0.05$	$\gamma_\Omega = 0.1$	$\gamma_\Omega = 0.3$
UncleTC	MAE _p	0.156 ± 0.055	0.111 ± 0.038	0.063 ± 0.047
	MAE _λ	0.104 ± 0.063	0.063 ± 0.041	0.037 ± 0.031
UncleTC(Linear)	MAE _p	0.231 ± 0.072	0.189 ± 0.046	0.193 ± 0.180
	MAE _λ	0.135 ± 0.069	0.092 ± 0.039	0.090 ± 0.089
NTF-CPD-KL	MAE _λ	0.306 ± 0.119	0.258 ± 0.128	0.190 ± 0.124
BPTF	MAE _λ	0.547 ± 0.143	0.493 ± 0.161	0.345 ± 0.104
HaLRTC	MAE _λ	0.503 ± 0.382	0.495 ± 0.391	0.436 ± 0.282
NTF-LS	MAE _λ	0.742 ± 0.462	0.684 ± 0.421	0.403 ± 0.099
NTF-Tucker-LS	MAE _λ	0.639 ± 0.494	0.526 ± 0.299	0.341 ± 0.072
CostCo	MAE _λ	0.964 ± 0.235	0.853 ± 0.213	0.809 ± 0.147
POND	MAE _λ	0.841 ± 0.146	0.877 ± 0.098	0.885 ± 0.087

Table 2. Average MAE_p and MAE_λ over 20 random trials under various values of ζ ; $\gamma_\Omega = 0.2, \gamma_\Xi = 0.3, \gamma_\Theta = 0.2$ with $\Theta \cup \Omega \subseteq \Xi$.

Algorithm	Metric	SNR = 0dB	SNR = 10dB	SNR = 40dB
UncleTC	MAE _p	0.154 ± 0.007	0.047 ± 0.018	0.035 ± 0.001
	MAE _λ	0.089 ± 0.010	0.026 ± 0.007	0.023 ± 0.001
UncleTC(Linear)	MAE _p	0.304 ± 0.012	0.109 ± 0.041	0.129 ± 0.049
	MAE _λ	0.193 ± 0.054	0.052 ± 0.016	0.065 ± 0.020
NTF-CPD-KL	MAE _λ	0.299 ± 0.055	0.111 ± 0.001	0.107 ± 0.001
BPTF	MAE _λ	0.420 ± 0.051	0.414 ± 0.047	0.427 ± 0.070
HaLRTC	MAE _λ	0.423 ± 0.068	0.441 ± 0.109	0.440 ± 0.126
NTF-LS	MAE _λ	0.479 ± 0.062	0.474 ± 0.057	0.472 ± 0.046
NTF-Tucker-LS	MAE _λ	0.366 ± 0.047	0.375 ± 0.050	0.378 ± 0.045
CostCo	MAE _λ	0.987 ± 0.256	0.876 ± 0.432	0.890 ± 0.357
POND	MAE _λ	0.824 ± 0.096	0.785 ± 0.131	0.837 ± 0.087

with $D = 10$ are generated by randomly sampling its entries from the standard normal distribution. We use the ground-truth nonlinear function (unknown to our algorithm) $g(\mathbf{z}) = \text{sigmoid}(\boldsymbol{\nu}^\top (0.5\mathbf{z}^3 + 0.2\mathbf{z}))$, where the vector $\boldsymbol{\nu} \in \mathbb{R}^D$ is generated by randomly sampling its entries from the uniform distribution between 0 and 1. We observe only a subset of y_i 's such that each i belongs to Ω with probability $\gamma_\Omega \in (0, 1]$. Similarly, only a portion of \mathbf{z}_i 's are observed such that each i belongs to Ξ with probability $\gamma_\Xi \in (0, 1]$. We also make a subset of \mathbf{z}_i 's similar such that $\mathbf{z}_i = \boldsymbol{\varphi} + \mathbf{n}_i, i \in \Theta, \Theta \subseteq \Xi$, with each i belongs to Θ with probability $\gamma_\Theta \in (0, 1]$. Here, $\boldsymbol{\varphi} \in \mathbb{R}^D$ is a random vector whose entries are drawn from the standard normal distribution and \mathbf{n}_i denotes additive noise whose entries are zero mean Gaussian with variance σ^2 . Under this setting, we define the *signal-to-noise* (SNR) ratio $\text{SNR} = 10 \log_{10} (\|\boldsymbol{\varphi}\|_2^2 / D\sigma^2)$ dB. As the value of SNR increases, the feature vectors $\mathbf{z}_i \in \Theta$ are more similar.

Metric. To characterize the performance of recovering p_i 's and λ_i 's, we employ *mean absolute error* (MAE) which is defined as $\text{MAE}_\lambda = 1/\prod_k I_k \sum_i |\hat{\lambda}_i / \lambda_i^{(\text{avg})} - \lambda_i^\dagger / \lambda_i^{(\text{avg})}|$, where $\hat{\lambda}_i$ is the estimated value corresponding to λ_i^\dagger , and the subscript (avg) denotes the average across all values, e.g., $\lambda_i^{(\text{avg})} = 1/\prod_k I_k \sum_i \lambda_i$. Note that the metric MAE_λ involves entry-wise normalization of the estimates and the ground truth by the average in order to remove the scaling ambiguity. The metric MAE_p is defined in the same way,

Table 3. Average MAE_p and MAE_λ over 20 random trials under various values of γ_Θ ; $\gamma_\Omega = 0.2, \gamma_\Xi = 0.7, \text{SNR} = 40\text{dB}$ with $\Theta \cup \Omega \subseteq \Xi$.

Algorithm	Metric	$\gamma_\Theta = 0$	$\gamma_\Theta = 0.3$	$\gamma_\Theta = 0.5$
UncleTC	MAE _p	0.163 ± 0.007	0.126 ± 0.044	0.104 ± 0.076
	MAE _λ	0.062 ± 0.003	0.055 ± 0.028	0.054 ± 0.148
UncleTC(Linear)	MAE _p	0.231 ± 0.005	0.252 ± 0.029	0.201 ± 0.081
	MAE _λ	0.078 ± 0.005	0.104 ± 0.041	0.214 ± 0.174
NTF-CPD-KL	MAE _λ	0.185 ± 0.006	0.107 ± 0.002	0.156 ± 0.078
BPTF	MAE _λ	0.381 ± 0.024	0.445 ± 0.092	0.473 ± 0.174
HaLRTC	MAE _λ	0.436 ± 0.030	0.533 ± 0.212	0.561 ± 0.352
NTF-LS	MAE _λ	0.434 ± 0.030	0.472 ± 0.068	0.644 ± 0.412
NTF-Tucker-LS	MAE _λ	0.345 ± 0.039	0.411 ± 0.080	0.564 ± 0.337
CostCo	MAE _λ	0.876 ± 0.178	0.814 ± 0.301	0.987 ± 0.239
POND	MAE _λ	0.926 ± 0.158	0.851 ± 0.165	0.859 ± 0.112

with all the λ -terms replaced by their p counterparts..

Implementation Settings. To learn the nonlinear function, we use a neural network $g_\theta(\cdot)$ with 3 hidden layers and 20 ReLU activation functions in each hidden layer. The regularization parameter μ is fixed to be 2000. More implementation details and experiments using different μ 's are provided in the supplementary material in Sec. N.

Results. Table. 1 shows the performance under various γ_Ω 's of all the methods under test. Note that for the baselines that do not take detection probability into their models, only the MAE_λ results are presented. The results show that all the non-Poisson baselines including the NN-based methods struggle to estimate the true counts λ_i . This shows the importance of tailored modeling for integer data. Among the algorithms using Poisson models, the proposed UncleTC method estimates λ_i 's with much higher accuracy relative to NTF-CPD-KL, as our method explicitly models the undercounting effect. One can also note that UncleTC significantly outperforms UncleTC(Linear), in terms of both MAE_p and MAE_λ. This is due to the ability of UncleTC to better handle the nonlinear relationships between the detection probabilities p_i and the side information \mathbf{z}_i , by using its NN-based learning. In addition, we note that as γ_Ω is higher, i.e., more observations are available, the performance becomes better, as expected and indicated in our theoretical claim.

In Table 2 and Table 3, we show the performance of the algorithms under various SNR and γ_Θ , respectively. In Table 2, as SNR increases from 0 to 40dB, i.e., ζ in Assumption 4.3 becomes smaller and the feature vectors \mathbf{z}_i 's become more similar, the MAEs decrease noticeably. The impact on MAE_p is obviously spelled out, as a reduction of around 7 times is observed. Table 3 shows that as the size of Θ gets larger, the estimation accuracy of the proposed method is positively impacted. Both tables echo the results presented by Theorem 4.5.

Table 4 validates a key result from Theorem 4.5. That is, there exists a global scaling factor and a counter-scaling

Table 4. Statistics of $\hat{\lambda}_i/\lambda_i^{\dagger}$ and \hat{p}_i/p_i^{\dagger} for different trials of the proposed UncleTC with settings $\gamma_{\Omega} = 0.3, \gamma_{\varepsilon} = 0.3, \gamma_{\Theta} = 0.2, \zeta = 40\text{dB}$, and $g(\mathbf{z}) = \text{sigmoid}(\nu^{\top} (0.1 \log(\mathbf{z}^2) + 0.1 \mathbf{z}^2))$

#Trial	mean \pm std of $\lambda_i/\lambda_i^{\dagger}$	mean \pm std of \hat{p}_i/p_i^{\dagger}	mean \pm std of $\lambda_i \hat{p}_i / \lambda_i^{\dagger} p_i^{\dagger}$
1	1.16 \pm 0.04	0.86 \pm 0.08	0.99 \pm 0.08
2	1.29 \pm 0.04	0.78 \pm 0.08	1.00 \pm 0.07
3	0.98 \pm 0.03	1.02 \pm 0.06	0.99 \pm 0.07
4	1.33 \pm 0.05	0.76 \pm 0.10	1.00 \pm 0.09
5	1.09 \pm 0.03	0.93 \pm 0.06	1.00 \pm 0.06
6	1.16 \pm 0.05	0.86 \pm 0.11	0.99 \pm 0.12
7	1.07 \pm 0.03	0.94 \pm 0.07	1.00 \pm 0.06
8	1.24 \pm 0.03	0.81 \pm 0.05	0.99 \pm 0.05
9	0.83 \pm 0.02	1.21 \pm 0.05	1.00 \pm 0.04
10	0.80 \pm 0.02	1.23 \pm 0.06	0.99 \pm 0.07

factor for λ_i and p_i , respectively. In other words, assuming perfect estimation for $\hat{\lambda}_i$ and \hat{p}_i by the proposed UncleTC, we have $\hat{\lambda}_i = \xi \lambda_i^{\dagger}$, $\hat{p}_i = 1/\xi p_i^{\dagger}$, $\forall i$ for a certain scalar $\xi > 0$. To see if the UncleTC algorithm attains good estimation of λ_i and p_i up to a global scaling/counter-scaling ambiguity, we report the mean and standard deviation of $\hat{\lambda}_i/\lambda_i^{\dagger}$ and \hat{p}_i/p_i^{\dagger} taken over all the i 's in each random trial. One can see that only small standard deviations exist in all the trials, suggesting that the scaling/counter-scaling factors are approximately identical over all the indices. Similar observations are made on \hat{p}_i/p_i^{\dagger} . One can also note that the scaling factors of $\hat{\lambda}_i$ and \hat{p}_i are approximately reciprocal to each other (cf. the last column of Table 4). This observations corroborate the findings in Theorem 4.5.

6.2. Real-Data Experiments

COVID-19 Data. The dataset is obtained from the COVID-19 impact analysis platform hosted by the University of Maryland, College Park, United States (Zhang et al., 2020b). The dataset includes the number of reported COVID-19 cases of about 2270 US counties during 2020-2021, with 25 different attributes such as social distancing index, population density, testing capacity and so on. In order to extract a count-type tensor from the COVID-19 dataset, we represent each county using two coordinates, i.e., the rounded value of latitude and longitude of its geographical center. This way, we obtain an integer tensor of size $43 \times 80 \times 60$ (with 18.42% of entries observed), recording the COVID-19 case numbers of 642 counties over 60 days from January 1, 2021 to March 1, 2021.

Pollinator-Plant Interaction (PPI) Data. The Plant-Pollinator Interaction (PPI) dataset is a publicly available dataset, collected by the researchers at the H.J. Andrews (HJA) Long-term Ecological Research site in Oregon, USA (Seo et al., 2022; Daly et al., 2019). We consider a subset of the dataset by selecting 50 plants and 50 pollinators that have the highest numbers of interactions. This leads to a $50 \times 50 \times 37$ integer tensor with 22.06% observed entries.

Table 5. The rRMSE of various methods on the COVID-19 and PPI datasets.

Method	COVID-19	PPI
UncleTC	3.52 \pm 0.38	9.89 \pm 1.67
UncleTC (Linear)	3.90 \pm 0.33	10.21 \pm 1.45
HaLRTC	6.14 \pm 0.45	10.77 \pm 1.25
BPTF	6.72 \pm 0.42	11.02 \pm 1.39
NTF-CPD-KL	3.61 \pm 0.52	10.17 \pm 1.12
NTF-CPD-LS	7.00 \pm 0.38	11.25 \pm 1.27
NTF-Tucker-LS	6.82 \pm 0.48	11.22 \pm 1.27
CostCo	7.12 \pm 0.34	11.01 \pm 1.14
POND	6.92 \pm 0.51	10.95 \pm 1.31




Figure 1. The top-10 counties that have the highest UncleTC-estimated average detection probabilities (left) and lowest average detection probabilities (right) for COVID-19. Inside the bar, we show the average numbers of COVID-19 tests done per 1000 people during the selected period.

The dataset also includes about 37 entry attributes corresponding to each observation, such as flower color of the plant, pollinator tongue length, weather conditions, and so on.

For both datasets, since the features are in different numerical units and range, we normalize each feature value using min-max normalization (Jain et al., 2005), across all observations. More details on the datasets and tensor construction are given in the supplementary material in Sec. N.

Quantitative Evaluation. Since the ground-truth λ_i^{\dagger} 's and p_i^{\dagger} 's are unknown for real data, we employ a count data prediction metric, i.e., *relative root mean square* (rRMSE) (Fu et al., 2021) for evaluation. Let $\{y_i\}_{i \in \Delta}$ be the set of actual holdout observations and $\{\hat{y}_i\}_{i \in \Delta}$ be the corresponding predicted values. Then, the rRMSE is defined as $\text{rRMSE} = 1/\bar{y} \sqrt{\sum_{i \in \Delta} |y_i - \hat{y}_i|^2 / |\Delta|}$, where \bar{y} is the mean value of the actual observations $\{y_i\}_{i \in \Delta}$. Note that, for the proposed method, the observations are predicted via $\hat{y}_i = \hat{\lambda}_i \hat{p}_i$.

Table 5 presents the performance of various approaches on the COVID-19 and PPI datasets. One can see that the pro-

posed `UncleTC` exhibits promising performance on both datasets. One can also see that the performance of the Poisson model-based method `NTF-CPD-KL` has a comparable `rRMSE` performance relative to our method. This is expected, as the Poisson-Binomial model with parameters λ and p can be considered as a Poisson model with the parameter λp (see Lemma B.1). Nonetheless, our method still outperforms `NTF-CPD-KL`, as the incorporation of the entry attributes in our framework may be beneficial. Compared to the non-Poisson methods, our approach stands out with obviously lower `rRMSEs`, which is similar to what was observed in the simulations. It can also be noted that `UncleTC` outperforms its linear function-based counterpart `UncleTC (Linear)` in both cases, which is also consistent with our simulation results.

Qualitative Evaluation. In Fig. 1, we analyze the detection probabilities output by the proposed `UncleTC` for COVID-19 dataset. Here, we present the top 10 counties having the highest and lowest detection probabilities for COVID-19 cases found by the proposed `UncleTC` (averaged over all days). Note that the estimated detection probabilities have a global scaling ambiguity—see Theorem 4.5—but their relative ranking order is not affected by this ambiguity. We also report the average number of COVID-19 tests done per 1000 people during the selected period (we removed this feature while training the model to avoid correlated results). One can see that the counties with more tests done are mostly aligned with higher detection probabilities estimated by our algorithm. Also, most counties reported with high detection probabilities are populous, urban counties, which enjoy better infrastructure. These observations suggest that the proposed algorithm outputs plausible results for the detection probabilities—also see the detection probability results for the PPI dataset in the supplementary material in Sec. N.

7. Conclusion

We proposed an under-counted tensor completion framework, which is motivated by the commonly seen under-counting effects in data acquisition across many domains, e.g., ecology and epidemiology. Our model uses a Poisson-Binomial data-generating perspective as often advocated in computational ecology, where the true integer data are under-counted by a Binomial detection process. Unlike existing approaches that treat the under-counting effects using relatively simple models, we used a neural network-based design to account for a wide range of unknown nonlinear relations between the side information (attributes) and the under-counting effects. More importantly, we showed that the proposed under-counted tensor completion criterion can *provably* recover the underlying Poisson tensor and its entries’ under-counting characteristics, under reasonable

conditions—which is the first theoretical result of its kind. We validated our theorem and algorithm using simulated and real data and observed intuitively plausible results from our real data results.

Limitations. A key limitation lies in the model assumption that the under-counted probabilities are functions of observed side attributes. However, such attributes may not always be available. Provable UC-TC criteria that do not rely on such assumption are highly desirable, yet the current formulation and analysis cannot cover such settings without substantial changes and re-design.

Acknowledgement. This work is supported in part by the National Science Foundation (NSF) under Project NSF IIS-1910118. The PPI dataset is provided by the HJ Andrews Experimental Forest research program, funded by the National Science Foundation’s Long-Term Ecological Research Program (DEB 2025755), US Forest Service Pacific Northwest Research Station, and Oregon State University. The authors thank Dr. Julia A. Jones for helpful guidance with the PPI data and associated features.

References

- Acar, E., Dunlavy, D. M., and Kolda, T. G. Link prediction on evolving data using matrix and tensor factorizations. In *IEEE International Conference on Data Mining Workshop*, pp. 262–269, 2009.
- Balazevic, I., Allen, C., and Hospedales, T. TuckER: Tensor factorization for knowledge graph completion. In *Proceedings of Empirical Methods in Natural Language Processing and International Joint Conference on Natural Language Processing*, pp. 5185–5194, 2019.
- Bartlett, P. L., Foster, D. J., and Telgarsky, M. J. Spectrally-normalized margin bounds for neural networks. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Bertsekas, D. P. *Nonlinear programming*. Athena Scientific, 1999.
- Billio, M., Casarin, R., and Iacopini, M. Bayesian markov-switching tensor regression for time-varying networks. *Journal of the American Statistical Association*, pp. 1–13, 2022.
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.
- Cao, Y. and Xie, Y. Poisson matrix recovery and completion. *IEEE Transactions on Signal Processing*, 64(6):1609–1620, 2015.
- Chacoff, N. P., Vázquez, D. P., Lomáscolo, S. B., Stevani, E. L., Dorado, J., and Padrón, B. Evaluating sampling

- completeness in a desert plant–pollinator network. *Journal of Animal Ecology*, 81(1):190–200, 2012.
- Chi, E. C. and Kolda, T. G. On tensors, sparsity, and non-negative factorizations. *SIAM Journal on Matrix Analysis and Applications*, 33(4):1272–1299, 2012.
- Daly, C., Schulze, M. D., and McKee, W. A. Meteorological data from benchmark stations at the HJ Andrews experimental forest, 1957 to present, 2019. URL <http://andlter.forestry.oregonstate.edu/data/abstract.aspx?dbcode=MS001>. <https://doi.org/10.6073/pasta/c021a2ebf1f91adf0ba3b5e53189c84f>.
- Dennis, E. B., Morgan, B. J., and Ridout, M. S. Computational aspects of N-mixture models. *Biometrics*, 71(1):237–246, 2015.
- Ding, M., Fu, X., Huang, T.-Z., Wang, J., and Zhao, X.-L. Hyperspectral super-resolution via interpretable block-term tensor modeling. *IEEE Journal of Selected Topics in Signal Processing*, 15(3):641–656, 2020.
- Dorado, J., Vázquez, D. P., Stevani, E. L., and Chacoff, N. P. Rareness and specialization in plant–pollinator networks. *Ecology*, 92(1):19–25, 2011.
- Duchi, J., Hazan, E., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12(Jul):2121–2159, 2011.
- Durif, G., Modolo, L., Mold, J. E., Lambert-Lacroix, S., and Picard, F. Probabilistic count matrix factorization for single cell expression data analysis. *Bioinformatics*, 35(20):4011–4019, 2019.
- Fan, J., Ding, L., Yang, C., and Udell, M. Low-rank tensor recovery with euclidean-norm-induced Schatten-p quasi-norm regularization. *arXiv preprint arXiv:2012.03436*, 2020.
- Fu, X., Ibrahim, S., Wai, H.-T., Gao, C., and Huang, K. Block-randomized stochastic proximal gradient for low-rank tensor factorization. *IEEE Transactions on Signal Processing*, 68:2170–2185, 2020a.
- Fu, X., Vervliet, N., De Lathauwer, L., Huang, K., and Gillis, N. Computing large-scale matrix and tensor decomposition with structured factors: A unified nonconvex optimization perspective. *IEEE Signal Processing Magazine*, 37(5):78–94, 2020b.
- Fu, X., Seo, E., Clarke, J., and Hutchinson, R. A. Link prediction under imperfect detection: Collaborative filtering for ecological networks. *IEEE Transactions on Knowledge and Data Engineering*, 33(8):3117–3128, 2021.
- Gandy, S., Recht, B., and Yamada, I. Tensor completion and low-n-rank tensor recovery via convex optimization. *Inverse Problems*, 27:025010, 2011.
- Ghadermarzy, N., Plan, Y., and Yilmaz, O. Learning tensors from partial binary measurements. *IEEE Transactions on Signal Processing*, 67(1):29–40, 2018.
- Harshman, R. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis. *UCLA Working Papers in Phonetics*, 16, 1970.
- Hu, Y., Koren, Y., and Volinsky, C. Collaborative Filtering for Implicit Feedback Datasets. *IEEE International Conference on Data Mining*, 2008.
- Hutchinson, R., Liu, L.-P., and Dietterich, T. Incorporating boosted regression trees into ecological latent variable models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 25, pp. 1343–1348, 2011.
- Jain, A., Nandakumar, K., and Ross, A. Score normalization in multimodal biometric systems. *Pattern Recognition*, 38(12):2270–2285, 2005.
- Joseph, L. N., Elkin, C., Martin, T. G., and Possingham, H. P. Modeling abundance using n-mixture models: the importance of considering ecological mechanisms. *Ecological Applications*, 19(3):631–642, 2009.
- Kanatsoulis, C. I., Fu, X., Sidiropoulos, N. D., and Akcakaya, M. Tensor completion from regular sub-nyquist samples. *IEEE Transactions on Signal Processing*, 68:1–16, 2020.
- Karatzoglou, A., Amatriain, X., Baltrunas, L., and Oliver, N. Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering. In *Proceedings of the ACM Conference on Recommender Systems*, pp. 79–86, 2010.
- Kashima, H., Kato, T., Yamanishi, Y., Sugiyama, M., and Tsuda, K. Link propagation: A fast semi-supervised learning algorithm for link prediction. In *Proceedings of the SIAM international conference on data mining*, pp. 1100–1111, 2009.
- Kim, Y.-D. and Choi, S. Nonnegative tucker decomposition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1–8, 2007.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2015.
- Lee, C. and Wang, M. Tensor denoising and completion based on ordinal observations. In *International Conference on Machine Learning*, pp. 5778–5788. PMLR, 2020.

- Liang, D., Charlin, L., McInerney, J., and Blei, D. M. Modeling user exposure in recommendation. In *Proceedings of the International Conference on World Wide Web*, pp. 951–961, 2016.
- Lin, S. and Zhang, J. Generalization bounds for convolutional neural networks. *arXiv preprint arXiv:1910.01487*, 2019.
- Liu, H., Li, Y., Tsang, M., and Liu, Y. Costco: A neural tensor completion model for sparse tensors. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 324–334, 2019.
- Liu, J., Musialski, P., Wonka, P., and Ye, J. Tensor completion for estimating missing values in visual data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):208–220, 2013.
- MacKenzie, D. I., Nichols, J. D., Royle, J. A., Pollock, K. H., Bailey, L. L., and Hines, J. E. *Occupancy estimation and modeling: Inferring patterns and dynamics of species occurrence*. Elsevier, San Diego, USA, 2006.
- Montanari, A. and Sun, N. Spectral algorithms for tensor completion. *Communications on Pure and Applied Mathematics*, 71, 2016.
- Mu, C., Huang, B., Wright, J., and Goldfarb, D. Square deal: Lower bounds and improved relaxations for tensor recovery. In *Proceedings of the International Conference on Machine Learning*, volume 32, pp. 73–81, 2014.
- Papalexakis, E. E., Faloutsos, C., and Sidiropoulos, N. D. Tensors for data mining and data fusion: Models, applications, and scalable algorithms. *ACM Transactions on Intelligent Systems and Technology*, 8(2):16, 2017.
- Pu, W., Ibrahim, S., Fu, X., and Hong, M. Stochastic mirror descent for low-rank tensor decomposition under non-euclidean losses. *IEEE Transactions on Signal Processing*, 70:1803–1818, 2022.
- Razaviyayn, M., Hong, M., and Luo, Z.-Q. A unified convergence analysis of block successive minimization methods for nonsmooth optimization. *SIAM Journal on Optimization*, 23(2):1126–1153, 2013.
- Royle, J. A. N-mixture models for estimating population size from spatially replicated counts. *Biometrics*, 60(1): 108–115, 2004.
- Schein, A., Paisley, J., Blei, D. M., and Wallach, H. Bayesian poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1045–1054, 2015.
- Schneider, E. C. Failing the test — the tragic data gap undermining the U.S. pandemic response. *New England Journal of Medicine*, 383(4):299–302, 2020.
- Seo, E. and Hutchinson, R. Predicting links in plant-pollinator interaction networks using latent factor models with implicit feedback. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018.
- Seo, E., Jones, J. A., Hutchinson, R. A., and Pfeiffer, V. W. Plant pollinator data at HJ Andrews experimental forest, 2011 to 2021, 2022. URL <http://andlter.forestry.oregonstate.edu/data/abstract.aspx?dbcode=SA026>. <https://doi.org/10.6073/pasta/eec55cbc3dbfc56428629773737ab3e5>.
- Serfling, R. J. Probability Inequalities for the Sum in Sampling without Replacement. *The Annals of Statistics*, 2(1):39 – 48, 1974.
- Shalev-Shwartz, S. and Ben-David, S. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Shashua, A. and Hazan, T. Non-negative tensor factorization with applications to statistics and computer vision. In *Proceedings of the International Conference on Machine Learning*, pp. 792–799, 2005.
- Simchowitz, M. Zero-inflated poisson factorization for recommendation systems. *Junior Independent Work (advised by D. Blei)*, Princeton University, Department of Mathematics, 2013.
- Sørensen, M. and De Lathauwer, L. Fiber sampling approach to canonical polyadic decomposition and application to tensor completion. *SIAM Journal on Matrix Analysis and Applications*, 40(3):888–917, 2019.
- Tan, H., Wu, Y., Shen, B., Jin, P. J., and Ran, B. Short-term traffic prediction based on dynamic tensor completion. *IEEE Transactions on Intelligent Transportation Systems*, 17(8):2123–2133, 2016.
- Tillinghast, C., Fang, S., Zhang, K., and Zhe, S. Probabilistic neural-kernel tensor decomposition. In *IEEE International Conference on Data Mining*, pp. 531–540, 2020.
- Tylianakis, J. M., Laliberté, E., Nielsen, A., and Bascompte, J. Conservation of species interaction networks. *Biological Conservation*, 143(10):2270–2279, 2010.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing: Theory and Applications*, pp. 210–268. Cambridge University Press, 2012.

- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Wang, Y., Peng, J., Zhao, Q., Leung, Y., Zhao, X.-L., and Meng, D. Hyperspectral image restoration via total variation regularized low-rank tensor decomposition. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 11(4):1227–1243, 2018.
- Yuan, M. and Zhang, C.-H. On tensor completion via nuclear norm minimization. *Foundations of Computational Mathematics*, 16(4):1031–1068, 2016.
- Zhang, G., Fu, X., Wang, J., Zhao, X.-L., and Hong, M. Spectrum cartography via coupled block-term tensor decomposition. *IEEE Transactions on Signal Processing*, 68:3660–3675, 2020a.
- Zhang, L., Ghader, S., Pack, M. L., Xiong, C., Darzi, A., Yang, M., Sun, Q., Kabiri, A., and Hu, S. An interactive COVID-19 mobility impact and social distancing analysis platform. *medRxiv*, 2020b.
- Zhang, Z. and Aeron, S. Exact tensor completion using t -SVD. *IEEE Transactions on Signal Processing*, 65(6): 1511–1526, 2017.

Supplementary Material of “Under-Counted Tensor Completion with Neural Incorporation of Attributes”

A. Notation

The notations x , \mathbf{x} , \mathbf{X} and $\underline{\mathbf{X}}$ represent a scalar, a vector, a matrix, and a tensor, respectively. x_i denotes the i th element of the vector \mathbf{x} . $[\mathbf{X}]_{i,j}$ and $\mathbf{X}(i,j)$ both mean the (i,j) th entry of \mathbf{X} . Similarly, $[\underline{\mathbf{X}}]_i$ belongs to the i th entry of the tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I_1 \times \dots \times I_K}$ where $\mathbf{i} = (i_1, \dots, i_K)$. \mathbb{Z}_+ denote the set of all nonnegative integers. $\mathbf{x} \geq \mathbf{0}$, $\mathbf{X} \geq \mathbf{0}$, and $\underline{\mathbf{X}} \geq \mathbf{0}$ imply that all the associated entries are greater than 0. $\|\mathbf{x}\|_2$, $\|\mathbf{X}\|_F$ and $\|\underline{\mathbf{X}}\|_F$ all mean the Euclidean (Frobenius) norm of the augment. $\|[\underline{\mathbf{X}}]_\Omega\|_F$ indicates the Euclidean norm of the vector formed by concatenating the entries indexed by all unique \mathbf{i} 's belonging to Ω , i.e., $\|[\underline{\mathbf{X}}]_\Omega\|_F^2 = \sum_{\mathbf{i} \in \text{Supp}(\Omega)} x_{\mathbf{i}}^2$, where $\text{Supp}(\Omega)$ denotes the set of all unique entries in Ω . $\|\underline{\mathbf{X}}\|_\infty$ denotes the max norm of $\underline{\mathbf{X}}$ such that $\|\underline{\mathbf{X}}\|_\infty = \max_{\mathbf{i}} [\underline{\mathbf{X}}]_{\mathbf{i}}$. $[I]$ means an integer set $\{1, 2, \dots, I\}$. \top and \dagger denote transpose and pseudo-inverse, respectively. $\mathbf{x} \circ \mathbf{y}$ denotes the outer product of two vectors \mathbf{x} and \mathbf{y} , i.e., $\mathbf{x} \circ \mathbf{y} = \mathbf{x}\mathbf{y}^\top$. \otimes denotes the element-wise product (also known as Hadamard product). The function $\text{sigmoid}(x)$ represents $\frac{1}{1+\exp(-x)}$. The constant e denotes the Euler's number. $[x]_{[0,1]}$ denotes $\min\{\max\{x, 0\}, 1\}$.

B. MLE Formulation

Assuming that y_i 's are independently sampled from the generative model in (3), the log-likelihood of the observations can be expressed as follows:

$$\begin{aligned} \log \prod_{i \in \Omega} \Pr(Y_i = y_i; \lambda_i, p_i) &= \log \left(\prod_{i \in \Omega} \sum_{n=y_i}^{\infty} \Pr(N_i = n; \lambda_i) \Pr(y_i | N_i = n; p_i) \right) \\ &= \sum_{i \in \Omega} \log \left(\sum_{n=y_i}^{\infty} \left(\frac{\lambda_i^n e^{-\lambda_i}}{n!} \frac{n! p_i^{y_i} (1-p_i)^{n-y_i}}{y_i! (n-y_i)!} \right) \right), \end{aligned} \quad (14)$$

where Y_i denotes the random variable associated with the observation y_i and N_i denotes the random variable associated with the true count n_i . To circumvent the infinite sum in the log-likelihood, we invoke the following lemma:

Lemma B.1. (Dennis et al., 2015) Assume that the observation y follows the below generative model:

$$n \sim \text{Poisson}(\lambda), \quad (15a)$$

$$y \sim \text{Binomial}(n, p). \quad (15b)$$

Then, the model in (15) is equivalent to $y \sim \text{Poisson}(\lambda p)$.

Lemma B.1 states that the Poisson-Binomial model admits an equivalent Poisson model whose Poisson parameter is the multiplication of the original Poisson parameter and the detection probability. The derivation is based on an interesting equality:

$$\sum_{n=y}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} \frac{n! p^y (1-p)^{n-y}}{y! (n-y)!} = \frac{(p\lambda)^y e^{-\lambda p}}{y!};$$

see the derivation in (Fu et al., 2021).

Applying the above equality in (14), we obtain:

$$\begin{aligned} \log \prod_{i \in \Omega} \Pr(Y_i = y_i; \lambda_i, p_i) &= \sum_{i \in \Omega} \log \left(\frac{(p_i \lambda_i)^{y_i} e^{-\lambda_i p_i}}{y_i!} \right) \\ &= \sum_{i \in \Omega} [y_i \log \lambda_i + y_i \log p_i - \lambda_i p_i - \log y_i!]. \end{aligned} \quad (16)$$

Here, the MLE is formulated and derived under the same spirit of the MC case in (Fu et al., 2021).

C. More Details on Algorithm Design

In this section, we provide more details on the implementation of the proposed algorithm `UncleTC`. Our formulation is given below:

$$\underset{\{\mathbf{U}_k\}_{k=1}^K, \theta, \{p_i\}}{\text{minimize}} \frac{1}{|\Omega|} \sum_{i \in \Omega} \left[\left(\sum_{f=1}^F \prod_{k=1}^K \mathbf{U}_k(i_k, f) \right) p_i - y_i \log \left(\sum_{f=1}^F \prod_{k=1}^K \mathbf{U}_k(i_k, f) \right) - y_i \log p_i \right] + \frac{\mu}{|\Xi|} \sum_{i \in \Xi} \ell(g_\theta(\mathbf{z}_i), p_i), \quad (17a)$$

$$\text{subject to } \mathbf{U}_k \geq \mathbf{0}, \forall k \in [K], \quad (17b)$$

$$0 \leq p_i \leq 1, \forall i \in \Omega \cup \Xi, \quad (17c)$$

where $\mu > 0$ is a regularization parameter and $\ell(\cdot, \cdot)$ denotes a certain distance/divergence measure, e.g., the least squares function $\ell(x, y) = (x - y)^2$.

C.1. The U_k -Subproblem

The subproblem for each U_k is given below:

$$\underset{\mathbf{U}_k \geq \mathbf{0}}{\text{minimize}} w(\mathbf{U}_k) := \sum_{i \in \Omega} \left[\left(\sum_{f=1}^F \mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f) \right) p_i - y_i \log \left(\sum_{f=1}^F \mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f) \right) \right]. \quad (18)$$

To obtain the update rule for U_k from (18), we have the following result:

Lemma C.1. *Let $\bar{\mathbf{U}}_k$ denote the current estimate of \mathbf{U}_k . Then, the objective function in (18) can be upper-bounded by the following function:*

$$s(\mathbf{U}_k; \bar{\mathbf{U}}_k) = \sum_{i \in \Omega} \left[\left(\sum_{f=1}^F \mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f) \right) p_i - y_i \sum_{f=1}^F \alpha_i^{(f)} \log \left(\frac{\mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f)}{\alpha_i^{(f)}} \right) \right], \quad (19)$$

$$\text{where } \alpha_i^{(f)} = \frac{\bar{\mathbf{U}}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f)}{\sum_{f=1}^F \bar{\mathbf{U}}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f)}.$$

The construction of the surrogate function $s(\mathbf{U}_k; \bar{\mathbf{U}}_k)$ follows the idea in (Chi & Kolda, 2012; Fu et al., 2021); The differences include that the model in (Chi & Kolda, 2012) did not have p_i terms and that the work in (Fu et al., 2021) did not consider the tensor case. Here, we use the Jensen’s inequality to construct a “tight” upper bound, i.e., $w(\mathbf{U}_k) \leq s(\mathbf{U}_k; \bar{\mathbf{U}}_k)$ and the equality can be attained at $\mathbf{U}_k = \bar{\mathbf{U}}_k$; see Sec. L for the proof of Lemma C.1. To be more precise, the update of U_k is through solving the following subproblem:

$$\mathbf{U}_k \leftarrow \arg \min_{\mathbf{U}_k \geq \mathbf{0}} s(\mathbf{U}_k; \bar{\mathbf{U}}_k). \quad (20)$$

By taking the derivative of $s(\cdot; \bar{\mathbf{U}}_k)$ w.r.t. \mathbf{U}_k and equating the derivative to zero, we obtain:

$$\mathbf{U}_k(i_k, f) \leftarrow \frac{\sum_{i \in \Omega} y_i \alpha_i^{(f)}}{\sum_{i \in \Omega} \prod_{j \neq k} \mathbf{U}_j(i_j, f) p_i}. \quad (21)$$

For efficient implementation of the above updates in the presence of missing data, we adopt a strategy proposed in (Chi & Kolda, 2012) for handling sparse tensors.

C.2. The p -Subproblem

Table 6 summarizes the update rules for the p_i -subproblem in Sec. 5 under different choices of the regularization function $\ell(\cdot)$. The complete description of the proposed `UncleTC` algorithm is provided in Algorithm 1.

Table 6. Updates for p_i subproblems for different loss functions

	$\ell(x, y) = (x - y)^2$	$\ell(x, y) = x \log \frac{x}{y} - x + y$
Case 1: $\forall i \in \Omega \cap \Xi$	$p_i \leftarrow \frac{(2\bar{\mu}g_\theta(\mathbf{z}_i) - \bar{\lambda}_i) + \sqrt{(2\bar{\mu}g_\theta(\mathbf{z}_i) - \bar{\lambda}_i)^2 + 8\bar{\mu}\bar{y}_i}}{4\bar{\mu}}$	$p_i \leftarrow \frac{\bar{y}_i + \bar{\mu}g_\theta(\mathbf{z}_i)}{\bar{\lambda}_i + \bar{\mu}}$
Case 2: $\forall i \in \Xi - \Omega \cap \Xi$	$p_i \leftarrow g_\theta(\mathbf{z}_i)$	$p_i \leftarrow g_\theta(\mathbf{z}_i)$
Case 3: $\forall i \in \Omega - \Omega \cap \Xi$	$p_i \leftarrow [y_i/\lambda_i]_{[0,1]}$	$p_i \leftarrow [y_i/\lambda_i]_{[0,1]}$

$\bar{\mu} = \mu/|\Xi|, \bar{\lambda}_i = \lambda_i/|\Omega|, \bar{y}_i = y_i/|\Omega|$

Algorithm 1 UncleTC

input : Observations $y_i, \forall i \in \Omega$, feature vectors $\mathbf{z}_i, \forall i \in \Xi$, Initializations $\theta^{(0)}, p_i^{(0)}, \forall i \in \Xi$, and $\mathbf{U}_k^{(0)}, \forall k$.

- 1 $t \leftarrow 0$;
- 2 **repeat**
- 3 $\lambda_i^{(t+1)} \leftarrow \sum_{f=1}^F \prod_{k=1}^K \mathbf{U}_k^{(t)}(i_k, f), \forall i$;
- 4 $\tilde{p}_i^{(t+1)} \leftarrow g_\theta(\mathbf{z}_i), i \in \Xi$;
- 5 $p_i^{(t+1)} \leftarrow \arg \min_{p_i \in [0,1]} \frac{1}{|\Omega|} (\lambda_i^{(t+1)} p_i^{(t)} - y_i \log(p_i^{(t)})) + \frac{\mu}{|\Xi|} \ell(\tilde{p}_i^{(t+1)}, p_i^{(t)}), \forall i \in \Omega \cap \Xi$;
- 6 $p_i^{(t+1)} \leftarrow \arg \min_{p_i \in [0,1]} \ell(\tilde{p}_i^{(t+1)}, p_i^{(t)}), \forall i \in \Xi - \Omega \cap \Xi$;
- 7 $p_i^{(t+1)} \leftarrow [y_i/\lambda_i^{(t+1)}]_{[0,1]}, \forall i \in \Omega - \Omega \cap \Xi$;
- 8 $\theta^{(t+1)} \leftarrow \arg \min_\theta \sum_{i \in \Xi} \ell(g_\theta(\mathbf{z}_i), p_i^{(t+1)})$;
- 9 **for** $k = 1$ **to** K **do**
- 10 $r \leftarrow 0$;
- 11 **repeat**
- 12 $\alpha_i^{(f)} \leftarrow \frac{\mathbf{U}_k^{(r)}(i_k, f) \prod_{j \neq k} \mathbf{U}_j^{(t)}(i_j, f)}{\sum_{f=1}^F \mathbf{U}_k^{(r)}(i_k, f) \prod_{j \neq k} \mathbf{U}_j^{(t)}(i_j, f)}, \forall i, f$;
- 13 $\mathbf{U}_k^{(r+1)}(i_k, f) \leftarrow \frac{\sum_{i \in \Omega} y_i \alpha_i^{(f)}}{\sum_{i \in \Omega} \prod_{j \neq k} \mathbf{U}_j^{(t)}(i_j, f) p_i^{(t+1)}}, \forall i_k \in [I_k], f$;
- 14 $r \leftarrow r + 1$;
- 15 **until** the stopping criterion is reached;
- 16 $\mathbf{U}_k^{(t+1)} \leftarrow \mathbf{U}_k^{(r)}$;
- 17 $\mathbf{U}_k^{(0)} \leftarrow \mathbf{U}_k^{(r)}$;
- 18 **end**
- 19 $t \leftarrow t + 1$;
- 20 **until** the stopping criterion is reached;

output : estimates $\{\hat{\mathbf{U}}_k\}_{k=1}^K, \hat{\theta}$, and $\hat{p}_i, \forall i \in \Omega \cup \Xi$

D. Proof of the Main Result (Theorem 4.5)

Let us first construct the following constraint sets:

$$\mathcal{L} = \{\underline{\mathbf{A}} \mid \underline{\mathbf{A}} = \sum_{f=1}^F \mathbf{U}_1(:, f) \circ \dots \circ \mathbf{U}_K(:, f), \|\mathbf{U}_k(:, f)\|_2 \leq u, \mathbf{U}_k \geq \mathbf{0}, \forall k, f\}, \quad (22a)$$

$$\mathcal{M} = \{\underline{\mathbf{M}} = \underline{\mathbf{A}} \circledast \underline{\mathbf{P}} \mid \beta \leq [\underline{\mathbf{M}}]_i \leq \alpha, \forall i, \underline{\mathbf{A}} \in \mathcal{L}, [\underline{\mathbf{P}}]_i = g_\theta(\mathbf{z}_i), g_\theta \in \mathcal{G}, \mathbf{z}_i \in \mathbb{R}^D, \forall i \in \Xi\}. \quad (22b)$$

Under Assumption 4.1, we can choose the parameters u, α , and β as follows:

$$u = \sqrt{F} \alpha_u, \alpha = F \alpha_u^K p_{\max}, \beta = F \beta_u^K p_{\min}. \quad (23)$$

Using these notations, the MLE in (6) can be re-expressed as follows:

$$\hat{\underline{\mathbf{M}}} = \arg \min_{\underline{\mathbf{M}} \in \mathcal{M}} \sum_{i \in \Omega} f(m_i; y_i, \mathbf{z}_i), \quad (24)$$

where $f(m_i; y_i, \mathbf{z}_i) = m_i - y_i \log m_i$, where $m_i = \lambda_i g_\theta(\mathbf{z}_i)$.

Estimating $\|\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}\|_{\mathbb{F}} / \prod_k I_k$: Our goal is to characterize the estimation accuracies of λ_i^{\natural} and p_i^{\natural} . To achieve this goal, we first characterize the term $\|\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}\|_{\mathbb{F}}$. We have the following result:

Theorem D.1. *Suppose that the assumptions of Theorem 4.5 hold true. Let $T = |\Omega|$. Assume that $\{y_i\}_{i \in \Omega}$ are i.i.d. samples under the generative model (3) and that $\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}$ are the set of observed features. Then, under (24), with probability at least $1 - 2\delta - 3e^{-\alpha(e^2-3)}$, we have the following relation, for any $0 < \delta < 1$:*

$$\frac{\|\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}\|_{\mathbb{F}}^2}{\prod_k I_k} \leq C(\alpha, \beta) \left(2\mathfrak{R}_T(\mathcal{F}) + (5f_{\max} - f_{\min}) \sqrt{\frac{2 \log(4/\delta)}{T}} + F\alpha_u^K \left(1 + \frac{p_{\max}(e^2 - 2)}{p_{\min}} \right) \nu \right), \quad (25)$$

where $C(\alpha, \beta) = \frac{4\alpha\gamma}{1-e^{-\gamma}}$, $\gamma = \frac{1}{8\beta}(\alpha - \beta)^2$, $f_{\max} \leq \alpha(1 + (e^2 - 2) \max\{\log \beta, \log \alpha\})$, $f_{\min} = \beta - \alpha(e^2 - 2) \log \alpha$,

$$\mathfrak{R}_T \leq \left(\frac{4}{\sqrt{T}} + \frac{12}{\sqrt{T}} \sqrt{F \sum_k I_k \log \left(6K L_f \sqrt{T} (F\alpha_u)^K \right) + \left(F\alpha_u^K L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{G}} \right)} \right),$$

$L_f = 1 + \frac{\alpha(e^2-2)}{\beta}$, $\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}] \in \mathbb{R}^{D \times S}$, and $\mathcal{R}_{\mathcal{G}}$ denotes the complexity measure of the class \mathcal{G} .

The proof is relegated to Appendix E.

Estimating $\|[\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}]_{\Theta}\|_{\mathbb{F}} / |\Theta|$: Next, we proceed to estimate $\|[\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}]_{\Theta}\|_{\mathbb{F}} / |\Theta|$ using the result in Theorem D.1.

Consider the following result:

Lemma D.2. *Assume that $i \in \Theta$ is drawn from $[I_1] \times \dots \times [I_K]$ uniformly at random (without replacement). Consider that $\underline{\mathbf{M}}_1, \underline{\mathbf{M}}_2 \in \{\underline{\mathbf{M}} \mid \underline{\mathbf{M}} \in \mathbb{R}_+^{I_1 \times \dots \times I_K}, \|\underline{\mathbf{M}}\|_{\infty} \leq \alpha\}$ holds. Let us define*

$$\tau(\Theta; \underline{\mathbf{M}}_1, \underline{\mathbf{M}}_2) = \left| \frac{1}{\sqrt{|\Theta|}} \|(\underline{\mathbf{M}}_1 - \underline{\mathbf{M}}_2)_{\Theta}\|_{\mathbb{F}} - \frac{1}{\sqrt{I^K}} \|\underline{\mathbf{M}}_1 - \underline{\mathbf{M}}_2\|_{\mathbb{F}} \right|.$$

Then, with probability greater than $1 - \delta$, we have

$$\tau(\Theta; \underline{\mathbf{M}}_1, \underline{\mathbf{M}}_2) \leq \alpha \left(\log \left(\frac{2}{\delta} \right) \left(1 - \frac{(|\Theta| - 1)}{\prod_k I_k} \right) \frac{1}{2|\Theta|} \right)^{1/4}.$$

The proof is provided in Appendix J.

Combining Lemma D.2 (here, $\alpha = F\alpha_u^K p_{\max}$ using Assumption 4.1) with Theorem D.1 and denoting $\eta = \frac{1}{\sqrt{\prod_k I_k}} \|\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}\|_{\mathbb{F}}$, we have the following with probability at least $1 - 3\delta - 3e^{-\alpha(e^2-3)}$:

$$\frac{\|[\underline{\mathbf{M}}^{\natural} - \widehat{\underline{\mathbf{M}}}]_{\Theta}\|_{\mathbb{F}}}{\sqrt{|\Theta|}} \leq \eta + \tau(\Theta; \underline{\mathbf{M}}^{\natural}, \widehat{\underline{\mathbf{M}}}), \quad (26)$$

where η^2 is given by (25) from Theorem D.1.

Recovering λ_i^{\natural} 's: Next, we proceed to characterize the estimation accuracies of λ_i^{\natural} 's using the result in (26). Under Assumptions 4.3 and 4.4, we have the following representation:

$$\begin{aligned} p_i^{\natural} &= g^{\natural}(\mathbf{z}_i) = \xi + s_i, \quad |s_i| \leq \zeta L_g, \quad \forall i \in \Theta, \\ \hat{p}_i &= \hat{g}_{\Theta}(\mathbf{z}_i) = \xi' + s'_i, \quad |s'_i| \leq \zeta L_{\Theta}, \quad \forall i \in \Theta, \end{aligned}$$

where $0 \leq \xi, \xi' \leq 1$ are certain scalars.

Hence, we get the following relation, with probability at least $1 - 3\delta - 3e^{-\alpha(e^2-3)}$:

$$\begin{aligned}
 (\eta + \tau(\Theta; \underline{\mathbf{M}}^\natural, \widehat{\underline{\mathbf{M}}})\sqrt{|\Theta|} &\geq \left\| \left[\underline{\mathbf{M}}^\natural - \widehat{\underline{\mathbf{M}}} \right]_{\Theta} \right\|_{\text{F}} = \sqrt{\sum_{i \in \Theta} (m_i^\natural - \widehat{m}_i)^2} \\
 &= \sqrt{\sum_{i \in \Theta} (\lambda_i^\natural \widehat{p}_i^\natural - \widehat{\lambda}_i \widehat{p}_i)^2} = \sqrt{\sum_{i \in \Theta} (\lambda_i^\natural (\xi + s_i) - \widehat{\lambda}_i (\xi' + s'_i))^2} \\
 &= \sqrt{\sum_{i \in \Theta} (\xi \lambda_i^\natural - \xi' \widehat{\lambda}_i + s_i \lambda_i^\natural - s'_i \widehat{\lambda}_i)^2} \\
 &\geq \left| \sqrt{\sum_{i \in \Theta} (\xi \lambda_i^\natural - \xi' \widehat{\lambda}_i)^2} - \sqrt{\sum_{i \in \Theta} (s_i \lambda_i^\natural - s'_i \widehat{\lambda}_i)^2} \right|.
 \end{aligned}$$

Since $|s_i| \leq \zeta L_g$, $|s'_i| \leq \zeta L_\theta$, and $\lambda_i \leq F\alpha_u^K$, $\forall \underline{\Lambda} \in \mathcal{L}$ by (23), the above relation can be further expressed as

$$\begin{aligned}
 (\eta + \tau(\Theta; \underline{\mathbf{M}}^\natural, \widehat{\underline{\mathbf{M}}})\sqrt{|\Theta|} + \zeta(L_g + L_\theta)F\alpha_u^K\sqrt{|\Theta|} &\geq \xi \sqrt{\sum_{i \in \Theta} (\lambda_i^\natural - \xi'/\xi \widehat{\lambda}_i)^2} \\
 &= \xi \left\| \left[\underline{\Lambda}^\natural - \xi'/\xi \widehat{\underline{\Lambda}} \right]_{\Theta} \right\|_{\text{F}} \\
 \Rightarrow \frac{\left\| \left[\underline{\Lambda}^\natural - \xi'/\xi \widehat{\underline{\Lambda}} \right]_{\Theta} \right\|_{\text{F}}}{\sqrt{|\Theta|}} &\leq \frac{(\eta + \tau(\Theta; \underline{\mathbf{M}}^\natural, \widehat{\underline{\mathbf{M}}}) + \zeta(L_g + L_\theta)F\alpha_u^K)}{\xi} \\
 &\leq \frac{(\eta + \tau(\Theta; \underline{\mathbf{M}}^\natural, \widehat{\underline{\mathbf{M}}}) + \zeta(L_g + L_\theta)F\alpha_u^K)}{p_{\min}}. \tag{27}
 \end{aligned}$$

Let $\xi'/\xi = \widehat{\xi}$. Our next goal is to characterize $\left\| \underline{\Lambda}^\natural - \widehat{\xi} \widehat{\underline{\Lambda}} \right\|_{\text{F}}$ using the result in (27). To proceed, consider the term

$$\tau(\Theta; \underline{\Lambda}^\natural, \widehat{\xi} \widehat{\underline{\Lambda}}) = \left| \frac{\left\| \left[\underline{\Lambda}^\natural - \widehat{\xi} \widehat{\underline{\Lambda}} \right]_{\Theta} \right\|_{\text{F}}}{\sqrt{|\Theta|}} - \frac{\left\| \underline{\Lambda}^\natural - \widehat{\xi} \widehat{\underline{\Lambda}} \right\|_{\text{F}}}{\sqrt{\prod_k I_k}} \right|. \tag{28}$$

We again invoke Lemma D.2 for characterizing the term $\tau(\Theta; \underline{\Lambda}^\natural, \widehat{\xi} \widehat{\underline{\Lambda}})$, and obtain that with probability of at least $1 - \delta$,

$$\tau(\Theta; \underline{\Lambda}^\natural, \widehat{\xi} \widehat{\underline{\Lambda}}) \leq F\alpha_u^K \left(\log \left(\frac{2}{\delta} \right) \left(1 - \frac{(|\Theta| - 1)}{\prod_k I_k} \right) \frac{1}{2|\Theta|} \right)^{1/4}. \tag{29}$$

Combining the result in (28) and (29) with (27), we obtain the following result with probability at least $1 - 4\delta - 3e^{-\alpha(e^2-3)}$:

$$\frac{\left\| \underline{\Lambda}^\natural - \widehat{\xi} \widehat{\underline{\Lambda}} \right\|_{\text{F}}}{\sqrt{\prod_k I_k}} \leq \underbrace{\tau(\Theta; \underline{\Lambda}^\natural, \widehat{\xi} \widehat{\underline{\Lambda}}) + \frac{(\eta + \tau(\Theta; \underline{\mathbf{M}}^\natural, \widehat{\underline{\mathbf{M}}}) + \zeta(L_g + L_\theta)F\alpha_u^K)}{p_{\min}}}_{\eta'}.$$

Recovering p_i 's: Next, we proceed to bound the estimation accuracies for p_i 's. Consider the following chain of inequalities:

$$\begin{aligned}
 \|\underline{\mathbf{M}}^\natural - \widehat{\underline{\mathbf{M}}}\|_{\text{F}} &= \left\| \underline{\mathbf{\Lambda}}^\natural \otimes \underline{\mathbf{P}}^\natural - \widehat{\underline{\mathbf{\Lambda}}} \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} \\
 &= \left\| \underline{\mathbf{\Lambda}}^\natural \otimes \underline{\mathbf{P}}^\natural - \underline{\mathbf{\Lambda}}^\natural \otimes 1/\widehat{\underline{\mathbf{P}}} + \underline{\mathbf{\Lambda}}^\natural \otimes 1/\widehat{\underline{\mathbf{P}}} - \widehat{\underline{\mathbf{\Lambda}}} \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} \\
 &\geq \left\| \underline{\mathbf{\Lambda}}^\natural \otimes (\underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}}) \right\|_{\text{F}} - \left\| (\underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}}) \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} \\
 &\geq \lambda_{\min} \left\| \underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} - \left\| (\underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}}) \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} \\
 &\geq \frac{\beta}{p_{\max}} \left\| \underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} - \left\| (\underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}}) \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}},
 \end{aligned}$$

where $\lambda_{\min} = \min_i \lambda_i = \frac{\min_i m_i}{\max_i p_i} = \frac{\beta}{p_{\max}}$. The above inequalities imply that with probability at least $1 - 4\delta - 3e^{-\alpha(e^2-3)}$

$$\begin{aligned}
 \frac{\left\| \underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}}}{\sqrt{\prod_k I_k}} &\leq \frac{p_{\max}}{\beta} \left(\frac{\left\| \underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}} \right\|_{\text{F}}}{p_{\min} \sqrt{\prod_k I_k}} + \frac{\|\underline{\mathbf{M}}^\natural - \widehat{\underline{\mathbf{M}}}\|_{\text{F}}}{\sqrt{\prod_k I_k}} \right), \\
 &= \frac{p_{\max}}{\beta} \left(\frac{\eta'}{p_{\min}} + \eta \right), \tag{30}
 \end{aligned}$$

where we have used the result that $\left\| (\underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}}) \otimes 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}} \leq \frac{1}{p_{\min}} \left\| \underline{\mathbf{\Lambda}}^\natural - \widehat{\underline{\mathbf{\Lambda}}} \right\|_{\text{F}}$, since $[1/\widehat{\underline{\mathbf{P}}}]_i = 1/\widehat{p}_i = \xi/\xi' \widehat{p}_i \leq 1/p_{\min}$.

Next, we proceed to characterize the generalization error of the learned nonlinear function \widehat{g}_θ . Towards this, we first bound the term $\frac{\left\| [\underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}}]_{\Xi} \right\|_{\text{F}}}{\sqrt{S}}$ using the result in (30). Let us consider

$$\tau(\Xi; \underline{\mathbf{P}}^\natural, 1/\widehat{\underline{\mathbf{P}}}) = \left| \frac{\left\| [\underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}}]_{\Xi} \right\|_{\text{F}}}{\sqrt{S}} - \frac{\left\| \underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}} \right\|_{\text{F}}}{\sqrt{\prod_k I_k}} \right|.$$

By applying Lemma D.2 combined with (30), we get that with a probability greater than $1 - 5\delta - 3e^{-\alpha(e^2-3)}$

$$\frac{\left\| [\underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}}]_{\Xi} \right\|_{\text{F}}}{\sqrt{S}} \leq \tau(\Xi; \underline{\mathbf{P}}^\natural, 1/\widehat{\underline{\mathbf{P}}}) + \frac{p_{\max}}{\beta} \left(\frac{\eta'}{p_{\min}} + \eta \right), \tag{31}$$

where

$$\tau(\Xi; \underline{\mathbf{P}}^\natural, 1/\widehat{\underline{\mathbf{P}}}) \leq p_{\max} \left(\log \left(\frac{2}{\delta} \right) \left(1 - \frac{(S-1)}{\prod_k I_k} \right) \frac{1}{2S} \right)^{1/4}.$$

Using the result in (31), we proceed to bound the generalization performance of the learned nonlinear function \widehat{g}_θ . Towards this goal, consider the following result:

Lemma D.3. *Under the Assumptions given by Theorem 4.5, the following holds with probability greater than $1 - \delta$:*

$$\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[(g^\natural(\mathbf{z}) - 1/\widehat{\xi} \widehat{g}_\theta(\mathbf{z}))^2 \right] \leq \frac{\left\| [\underline{\mathbf{P}}^\natural - 1/\widehat{\underline{\mathbf{P}}}]_{\Xi} \right\|_{\text{F}}^2}{S} + \frac{32(2\|\mathbf{Z}\|_{\text{F}} \mathcal{R}_{\mathcal{G}})^{\frac{1}{4}}}{S^{5/8}} + 4\sqrt{\frac{2\log(4/\delta)}{S}}, \tag{32}$$

where $\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}]$ and $\mathcal{R}_{\mathcal{G}}$ denotes the complexity measure as given by Assumption 4.2.

The proof is given in Appendix K.

Combining Lemma D.3 with (31), we get the final bound for $\mathbb{E}_{\mathbf{z} \sim \mathcal{D}} \left[(g^\natural(\mathbf{z}) - 1/\widehat{\xi} \widehat{g}_\theta(\mathbf{z}))^2 \right]$ with probability greater than $1 - 6\delta - 3e^{-\alpha(e^2-3)}$.

E. Proof of Theorem D.1

For a set $\Omega = \{i_1, \dots, i_T\}$, we assume that $\Omega \sim \Pi$ where Π is uniform, i.e., each y_i is observed independently at random with probability $1/(\prod_k I_k)^2$. Hence, we define the expected risk as follows:

$$D_{\Pi}(\underline{\mathbf{M}}; \underline{\mathbf{Y}}) := \mathbb{E}_{\Omega \sim \Pi} [f(m_i; y_i, \mathbf{z}_i)] = \sum_i \frac{1}{\prod_k I_k} f(m_i; y_i, \mathbf{z}_i). \quad (33)$$

Accordingly, we define the empirical risk as follows:

$$D_{\Omega}(\underline{\mathbf{M}}; \underline{\mathbf{Y}}) = \frac{1}{T} \sum_{i \in \Omega} f(m_i; y_i, \mathbf{z}_i). \quad (34)$$

Eq.(24) implies that

$$D_{\Omega}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) \leq D_{\Omega}(\widetilde{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) \quad (35)$$

where $[\widehat{\underline{\mathbf{M}}}]_i = \lambda_i^{\natural} \widetilde{g}_{\theta}(\mathbf{z}_i)$ and \widetilde{g}_{θ} is given by Assumption 4.2.

To proceed, consider the following chain of inequalities:

$$\begin{aligned} \mathbb{E}_y [D_{\Pi}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Pi}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})] &= \mathbb{E}_y [D_{\Pi}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}})] - D_{\Omega}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) + D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}}) - \mathbb{E}_y [D_{\Pi}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})] \\ &\quad + D_{\Omega}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Omega}(\widetilde{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) + D_{\Omega}(\widetilde{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}}) \\ &\leq \mathbb{E}_y [D_{\Pi}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}})] - D_{\Omega}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) + D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}}) - \mathbb{E}_y [D_{\Pi}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})] \\ &\quad + |D_{\Omega}(\widetilde{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})| \\ &\leq \sup_{\underline{\mathbf{M}} \in \mathcal{M}} |D_{\Omega}(\underline{\mathbf{M}}; \underline{\mathbf{Y}}) - \mathbb{E}_y [D_{\Pi}(\underline{\mathbf{M}}; \underline{\mathbf{Y}})]| + |D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}}) - \mathbb{E}_y [D_{\Pi}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})]| \\ &\quad + |D_{\Omega}(\widetilde{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Omega}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})|, \end{aligned} \quad (36)$$

where expectation is taken w.r.t. y_i 's and the first inequality is by (35).

Let us consider the L.H.S. of (36):

$$\begin{aligned} \mathbb{E}_y [D_{\Pi}(\widehat{\underline{\mathbf{M}}}; \underline{\mathbf{Y}}) - D_{\Pi}(\underline{\mathbf{M}}^{\natural}; \underline{\mathbf{Y}})] &= \frac{1}{\prod_k I_k} \sum_i \mathbb{E}_y \left[(\widehat{m}_i - y_i \log \widehat{m}_i) - (m_i^{\natural} - y_i \log m_i^{\natural}) \right] \\ &= \frac{1}{\prod_k I_k} \sum_i \left[(\widehat{m}_i - m_i^{\natural} \log \widehat{m}_i) - (m_i^{\natural} - m_i^{\natural} \log m_i^{\natural}) \right] \\ &= \frac{1}{\prod_k I_k} \sum_i \left[m_i^{\natural} \log \frac{m_i^{\natural}}{\widehat{m}_i} - (m_i^{\natural} - \widehat{m}_i) \right] \\ &:= \text{KL}(\underline{\mathbf{M}}^{\natural} \parallel \widehat{\underline{\mathbf{M}}}), \end{aligned} \quad (37)$$

where the second equality utilizes Lemma B.1 to get $\mathbb{E}[y_i] = \lambda_i^{\natural} p_i^{\natural} = m_i^{\natural}$.

Upper-bounding the first term on the R.H.S of (36). Next, we characterize the first term on the R.H.S. of (36). Towards this, we invoke the following theorem (Theorem 26.5 in (Shalev-Shwartz & Ben-David, 2014)) :

Theorem 1. (Shalev-Shwartz & Ben-David, 2014, Theorem 26.5) Assume that for all y and for all x , we have $|f(x; y)| \leq f_{\max}$. Then for any $\underline{\mathbf{M}} \in \mathcal{M}$, the following holds with probability greater than $1 - \delta$:

$$|D_{\Omega}(\underline{\mathbf{M}}; \underline{\mathbf{Y}}) - \mathbb{E}_y [D_{\Pi}(\underline{\mathbf{M}}; \underline{\mathbf{Y}})]| \leq 2\mathfrak{R}_T(\mathcal{F}) + 4f_{\max} \sqrt{\frac{2 \log(4/\delta)}{T}}, \quad (38)$$

²Here, the set Ω can be a multiset, meaning it can contain multiple copies of an index since the assumed sampling scheme is with replacement.

where \mathcal{F} denotes the set

$$\mathcal{F} \triangleq \{[f(m_{i_1}; y_{i_1}, \mathbf{z}_{i_1}), \dots, f(m_{i_T}; y_{i_T}, \mathbf{z}_{i_T})]^\top \mid \underline{\mathbf{M}} \in \mathcal{M}, \mathbf{i}_t \in \Omega\}$$

and $\mathfrak{R}_T(\mathcal{F})$ denotes the empirical Rademacher complexity of the set \mathcal{F} .

Applying Theorem 1, with probability greater than $1 - \delta$, we have

$$\sup_{\underline{\mathbf{M}} \in \mathcal{M}} |D_\Omega(\underline{\mathbf{M}}; \underline{\mathbf{Y}}) - \mathbb{E}_y[D_\Pi(\underline{\mathbf{M}}; \underline{\mathbf{Y}})]| \leq 2\mathfrak{R}_T(\mathcal{F}) + 4f_{\max} \sqrt{\frac{2 \log(4/\delta)}{T}}. \quad (39)$$

Let us characterize f_{\max} in (39).

Lemma E.1. *Let $f(m_i; y_i, \mathbf{z}_i) = m_i - y_i \log m_i$ where $m_i = \lambda_i p_i = \lambda_i g_\theta(\mathbf{z}_i)$. Then, under Assumption 4.1, the maximum and minimum value taken by the function $f(m_i; y_i, \mathbf{z}_i)$, denoted as f_{\max} and f_{\min} are given by*

$$\begin{aligned} f_{\max} &= \alpha + c \max\{|\log \beta|, \log \alpha\} \\ f_{\min} &= \beta - c \log \alpha, \end{aligned}$$

with probability greater than $1 - e^{-\alpha(e^2-3)}$ where $c = \alpha(e^2 - 2)$.

The proof of the lemma is given in Appendix G.

Next, we aim to characterize the Rademacher complexity $\mathfrak{R}_T(\mathcal{F})$. Towards this, we have the following result:

Lemma E.2. *Suppose that the assumptions of Theorem 4.5 hold true. Let $\Omega = \{\mathbf{i}_1, \dots, \mathbf{i}_T\}$ and $T = |\Omega|$. Assume that the observations y_{i_t} 's are i.i.d. samples under the generative model (3) and that $\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}$ are the set of observed features. Consider the set*

$$\mathcal{F} = \{\mathbf{f} \in \mathbb{R}^T \mid \mathbf{f} = [f(m_{i_1}; y_{i_1}, \mathbf{z}_{i_1}), \dots, f(m_{i_T}; y_{i_T}, \mathbf{z}_{i_T})]^\top \mid m_{i_t} = [\underline{\mathbf{M}}]_{i_t}, \mathbf{i}_t \in \Omega, \underline{\mathbf{M}} \in \mathcal{M}\}.$$

With probability greater than $1 - e^{-\alpha(e^2-3)}$, the empirical Rademacher complexity of \mathcal{F} is bounded by

$$\mathfrak{R}_T(\mathcal{F}) \leq \left(\frac{4}{\sqrt{T}} + \frac{12}{\sqrt{T}} \sqrt{F \sum_k I_k \log(6KL_f \sqrt{T} (F\alpha_u)^K) + (F\alpha_u^K L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{G}})} \right),$$

where $L_f = 1 + \frac{\alpha(e^2-2)}{\beta}$, $\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}] \in \mathbb{R}^{D \times S}$ and $\mathcal{R}_{\mathcal{G}}$ denotes the complexity measure of the class \mathcal{G} .

The proof is provided in Appendix F. By combining Lemma E.1 and Lemma E.2, we completely characterize (39).

Upper-bounding the second term on the R.H.S of (36). Next, we upper bound the second term on the R.H.S of (36).

Let us consider the Hoeffding's inequality

Lemma E.3. *Let F_1, \dots, F_T be independent bounded random variables with $F_t \in [f_{\min}, f_{\max}]$ for all t where $-\infty < f_{\min} \leq f_{\max} < \infty$. Then for all $t \geq 0$,*

$$\Pr \left(\frac{1}{T} \sum_{t=1}^T (F_t - \mathbb{E}[F_t]) \geq q \right) \leq \exp \left(-\frac{2Tq^2}{(f_{\max} - f_{\min})^2} \right).$$

To use Lemma E.3, let us define the random variable F_t as follows:

$$F_t \triangleq f(m_{i_t}; y_{i_t}, \mathbf{z}_{i_t}).$$

Then, invoking Lemma E.3, one can obtain

$$\Pr \left(\left| D_\Omega(\underline{\mathbf{M}}^\dagger; \underline{\mathbf{Y}}) - \mathbb{E}_y[D_\Pi(\underline{\mathbf{M}}^\dagger; \underline{\mathbf{Y}})] \right| \geq q \right) \leq \exp \left(-\frac{2Tq^2}{(f_{\max} - f_{\min})^2} \right). \quad (40)$$

where f_{\max} and f_{\min} are given by Lemma E.1. Hence, by substituting $q = (f_{\max} - f_{\min}) \sqrt{\frac{\log(\frac{1}{\delta})}{2T}}$, where $\delta \in (0, 1)$ in (40), we get that with probability greater than $1 - \delta$

$$\left| D_\Omega(\underline{\mathbf{M}}^\dagger; \underline{\mathbf{Y}}) - \mathbb{E}_y[D_\Pi(\underline{\mathbf{M}}^\dagger; \underline{\mathbf{Y}})] \right| \leq (f_{\max} - f_{\min}) \sqrt{\frac{\log(\frac{1}{\delta})}{2T}}. \quad (41)$$

Upper-bounding the third term on the R.H.S of (36). Consider the following chain of equations:

$$\begin{aligned}
 \left| D_{\Omega}(\widetilde{\mathbf{M}}; \mathbf{Y}) - D_{\Omega}(\mathbf{M}^{\natural}; \mathbf{Y}) \right| &= \left| \frac{1}{T} \sum_{i \in \Omega} f(\widetilde{m}_i; y_i, \mathbf{z}_i) - \frac{1}{T} \sum_{i \in \Omega} f(m_i^{\natural}; y_i, \mathbf{z}_i) \right| \\
 &= \frac{1}{T} \left| \sum_{i \in \Omega} (\widetilde{m}_i - y_i \log \widetilde{m}_i) - (m_i^{\natural} - y_i \log m_i^{\natural}) \right| \\
 &= \frac{1}{T} \left| \sum_{i \in \Omega} (\lambda_i^{\natural}(\widetilde{g}_{\theta}(\mathbf{z}_i) - g_{\theta}^{\natural}(\mathbf{z}_i)) - y_i (\log \lambda_i^{\natural} \widetilde{g}_{\theta}(\mathbf{z}_i) - \log \lambda_i^{\natural} g_{\theta}^{\natural}(\mathbf{z}_i))) \right| \\
 &= \frac{1}{T} \left| \sum_{i \in \Omega} (\lambda_i^{\natural}(\widetilde{g}_{\theta}(\mathbf{z}_i) - g_{\theta}^{\natural}(\mathbf{z}_i)) - y_i (\log \widetilde{g}_{\theta}(\mathbf{z}_i) - \log g_{\theta}^{\natural}(\mathbf{z}_i))) \right| \\
 &\leq \max_i \lambda_i^{\natural} |\widetilde{g}_{\theta}(\mathbf{z}_i) - g_{\theta}^{\natural}(\mathbf{z}_i)| + \max_i y_i \frac{|\widetilde{g}_{\theta}(\mathbf{z}_i) - g_{\theta}^{\natural}(\mathbf{z}_i)|}{p_{\min}}, \\
 &\leq (c_{\lambda} + c/p_{\min})\nu, \\
 &= F\alpha_u^K (1 + p_{\max}(e^2 - 2)/p_{\min})\nu, \tag{42}
 \end{aligned}$$

where the first inequality employs the triangle inequality and the Lipschitz continuity of the log function. The first inequality also employs the Assumptions 4.1 that the lowerbound of both $\widetilde{g}_{\theta}(\mathbf{z}_i)$ and $g_{\theta}^{\natural}(\mathbf{z}_i)$ are p_{\min} . The second inequality is obtained by applying $c_{\lambda} = F\alpha_u^K$ from Assumption 4.1, $c = \alpha + \alpha(e^2 - 3)$ from Lemma H.1 with probability greater than $1 - e^{-\alpha(e^2-3)}$ and $\alpha = F\alpha_u^K p_{\max}$.

Putting Together. Combining (36), (37) with the upperbounds (39), (41), and (42), we get with probability greater than $1 - 2\delta - 3e^{-\alpha(e^2-3)}$.

$$\text{KL}(\mathbf{M}^{\natural} \|\widehat{\mathbf{M}}) \leq 2\mathfrak{R}_T(\mathcal{F}) + (5f_{\max} - f_{\min})\sqrt{\frac{2\log(4/\delta)}{T}} + F\alpha_u^K \left(1 + \frac{p_{\max}(e^2 - 2)}{p_{\min}}\right)\nu, \tag{43}$$

where $f_{\max} = \alpha + c \max\{|\log \beta|, \log \alpha\}$ and $f_{\min} = \beta - c \log \alpha$ given by Lemma E.1 .

By (Cao & Xie, 2015, Lemma 8), we have

$$\text{KL}(\mathbf{M}^{\natural} \|\widehat{\mathbf{M}}) \geq \frac{\|\mathbf{M}^{\natural} - \widehat{\mathbf{M}}\|_{\text{F}}^2}{C(\alpha, \beta) \prod_k I_k}, \tag{44}$$

where $C(\alpha, \beta) = \frac{4\alpha\gamma}{1-e^{-\gamma}}$ and $\gamma = \frac{1}{8\beta}(\alpha - \beta)^2$. Hence, combining (43) and (44), we have the following with probability greater than $1 - 2\delta - 3e^{-\alpha(e^2-3)}$:

$$\frac{\|\mathbf{M}^{\natural} - \widehat{\mathbf{M}}\|_{\text{F}}^2}{\prod_k I_k} \leq C(\alpha, \beta) \left(2\mathfrak{R}_T(\mathcal{F}) + (5f_{\max} - f_{\min})\sqrt{\frac{2\log(4/\delta)}{T}} + F\alpha_u^K \left(1 + \frac{p_{\max}(e^2 - 2)}{p_{\min}}\right)\nu \right). \tag{45}$$

F. Proof of Lemma E.2

Let us consider the following vector \mathbf{f} :

$$\mathbf{f} = [f(m_{i_1}; y_{i_1}, \mathbf{z}_{i_1}), \dots, f(m_{i_T}; y_{i_T}, \mathbf{z}_{i_T})]^{\top} \in \mathcal{F},$$

where $i_t \in \Omega$, $\forall t$ and

$$\mathcal{F} = \{\mathbf{f} \in \mathbb{R}^T \mid \mathbf{f} = [f(m_{i_1}; y_{i_1}, \mathbf{z}_{i_1}), \dots, f(m_{i_T}; y_{i_T}, \mathbf{z}_{i_T})]^{\top} \mid m_{i_t} = [\mathbf{M}]_{i_t}, \mathbf{M} \in \mathcal{M}\}. \tag{46}$$

The Rademacher complexity of set \mathcal{F} , denoted as $\mathfrak{R}_T(\mathcal{F})$, is defined as follows:

Definition F.1. (Shalev-Shwartz & Ben-David, 2014) The empirical Rademacher complexity of a set of vectors $\mathcal{F} \subset \mathbb{R}^T$ is defined as follows:

$$\mathfrak{R}_T(\mathcal{F}) = \frac{1}{T} \mathbb{E} \left[\sup_{\mathbf{f} \in \mathcal{F}} \sum_{i=1}^T \epsilon_i f_i \right], \quad (47)$$

where expectation is w.r.t. ϵ_i 's which are i.i.d. Rademacher random variables taking values from $\{-1, 1\}$.

By Definition F.1, the empirical Rademacher complexity of the set \mathcal{F} defined in (46) is given by

$$\mathfrak{R}_T(\mathcal{F}) = \frac{1}{T} \mathbb{E}_\sigma \left[\sup_{\mathbf{f} \in \mathcal{F}} \sum_{i=1}^T \sigma_i f(m_i; y_i, \mathbf{z}_i) \right]. \quad (48)$$

To characterize the Rademacher complexity $\mathfrak{R}_T(\mathcal{F})$, we proceed to characterize the covering number of \mathcal{F} which is defined as follows:

Definition F.2. (Vershynin, 2012) The ε -net of a set \mathcal{F} represented by $\overline{\mathcal{F}}_\varepsilon$ is a finite set such that for any $\mathbf{f} \in \mathcal{F}$, there is a $\overline{\mathbf{f}} \in \overline{\mathcal{F}}_\varepsilon$ satisfying

$$\|\mathbf{f} - \overline{\mathbf{f}}\|_2 \leq \varepsilon.$$

The covering number of \mathcal{F} is $N(\mathcal{F}, \varepsilon) = |\overline{\mathcal{F}}_\varepsilon|$.

Consider the following:

$$\begin{aligned} \|\mathbf{f} - \overline{\mathbf{f}}\|_2^2 &= \sum_{i \in \Omega} (f(m_i, y_i, \mathbf{z}_i) - f(\overline{m}_i; y_i, \mathbf{z}_i))^2 \\ &\leq \sum_{i \in \Omega} L_f^2 (m_i - \overline{m}_i)^2 \\ &\leq L_f^2 T \|\underline{\mathbf{M}} - \overline{\mathbf{M}}\|_{\Omega}^2 \\ &\leq L_f^2 T \|\underline{\mathbf{M}} - \overline{\mathbf{M}}\|_{\Xi}^2. \end{aligned} \quad (49)$$

where $\overline{\mathbf{f}} \in \overline{\mathcal{F}}_\varepsilon$ belongs to an ε -net for \mathcal{F} , $\overline{\mathbf{M}} \in \overline{\mathcal{M}}_\varepsilon$ belongs to an ε -net of \mathcal{M} , $\overline{m}_i = [\overline{\mathbf{M}}]_i$, and $T = |\Omega|$. The last inequality is by applying the assumption that $\Omega \subseteq \Xi$. To characterize the term L_f in (49), we have the following result:

Fact F.3. Assume that there exist two real numbers $\beta > 0$ and $\alpha < \infty$ such that $\beta \leq m_i \leq \alpha$ for all $i \in \Omega$. Then, the following hold for all $i \in \Omega$ with probability greater than $1 - e^{-\alpha(e^2-3)}$:

$$|f'_i(m_i; y_i, \mathbf{z}_i)| \leq 1 + \frac{\alpha(e^2 - 2)}{\beta}.$$

The proof is provided in Appendix H.

Hence, applying Fact F.3, in (49), we have the following relation with probability greater than $1 - e^{-\alpha(e^2-3)}$,

$$L_f = 1 + (\alpha(e^2 - 2))/\beta. \quad (50)$$

Next, we consider the term $\|\underline{\mathbf{M}} - \overline{\mathbf{M}}\|_{\Xi}^2$ in (49):

$$\begin{aligned} \|\underline{\mathbf{M}} - \overline{\mathbf{M}}\|_{\Xi}^2 &= \|\underline{\mathbf{A}} \otimes \underline{\mathbf{P}} - \overline{\mathbf{A}} \otimes \overline{\mathbf{P}}\|_{\Xi}^2 \\ &= \|\underline{\mathbf{A}} \otimes \underline{\mathbf{P}} + \overline{\mathbf{A}} \otimes \underline{\mathbf{P}} - \overline{\mathbf{A}} \otimes \underline{\mathbf{P}} - \overline{\mathbf{A}} \otimes \overline{\mathbf{P}}\|_{\Xi}^2 \\ &\leq \|(\underline{\mathbf{A}} - \overline{\mathbf{A}}) \otimes \underline{\mathbf{P}}\|_{\Xi}^2 + \|\overline{\mathbf{A}} \otimes (\underline{\mathbf{P}} - \overline{\mathbf{P}})\|_{\Xi}^2 \\ &\leq \max_{i \in \Xi} p_i \|\underline{\mathbf{A}} - \overline{\mathbf{A}}\|_{\Xi}^2 + \max_{i \in \Xi} \overline{\lambda}_i \|\underline{\mathbf{P}} - \overline{\mathbf{P}}\|_{\Xi}^2 \\ &\leq \max_i p_i \|\underline{\mathbf{A}} - \overline{\mathbf{A}}\|_{\mathbb{F}}^2 + \max_i \overline{\lambda}_i \|g_{\theta}([\mathbf{Z}]_{\Xi}) - \overline{g}_{\theta}([\mathbf{Z}]_{\Xi})\|_{\mathbb{F}}^2 \\ &\leq \|\underline{\mathbf{A}} - \overline{\mathbf{A}}\|_{\mathbb{F}}^2 + c_{\lambda} \|g_{\theta}([\mathbf{Z}]_{\Xi}) - \overline{g}_{\theta}([\mathbf{Z}]_{\Xi})\|_{\mathbb{F}}^2, \end{aligned} \quad (51)$$

where $g_\theta([\mathbf{Z}]\Xi) = [g_\theta(\mathbf{z}_{i_1}), \dots, g_\theta(\mathbf{z}_{i_S})]^\top$, $g_\theta \in \mathcal{G}$, $\bar{g}_\theta([\mathbf{Z}]\Xi) = [\bar{g}_\theta(\mathbf{z}_{i_1}), \dots, \bar{g}_\theta(\mathbf{z}_{i_S})]^\top$, $\bar{g}_\theta \in \bar{\mathcal{G}}_\varepsilon$ belongs to the ε -net of \mathcal{G} , $\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}$ are the set of observed features, and $\underline{\mathbf{A}} \in \bar{\mathcal{L}}_\varepsilon$ belongs to an ε -net of \mathcal{L} . The first inequality uses the triangle inequality and the last inequality uses the facts that $\max_i p_i \leq 1$ and $\max_i \bar{\lambda}_i \leq c_\lambda$ where we have $c_\lambda = F\alpha_u^K$ under Assumption 4.1.

Eq. (49) and (51) imply that that if there exists an $\varepsilon/2c_\lambda L_f \sqrt{T}$ -net covering for $\mathcal{G} \circ \Xi$, where

$$\mathcal{G} \circ \Xi = \{g_\theta([\mathbf{Z}]\Xi) = [g_\theta(\mathbf{z}_{i_1}), \dots, g_\theta(\mathbf{z}_{i_S})]^\top, g_\theta \in \mathcal{G}\}, \quad (52)$$

and an $\varepsilon/2L_f \sqrt{T}$ -net covering for \mathcal{L} , then we can construct an ε -net to cover \mathcal{F} . Since $\underline{\mathbf{A}}$ which is a low-rank tensor, we invoke the following result to get the covering number of for \mathcal{L} :

Lemma F.4. (Fan et al., 2020) Let $\mathcal{L} = \{\underline{\mathbf{X}} \in \mathbb{R}^{I_1 \times \dots \times I_K} \mid \underline{\mathbf{X}} = \sum_{f=1}^F \mathbf{U}_1(:, f) \circ \dots \circ \mathbf{U}_K(:, f), \|\mathbf{U}_k(:, f)\|_2 \leq u\}$. Then, the covering number of \mathcal{L} with respect to the Frobenius norm satisfies

$$\mathsf{N}(\mathcal{L}, \varepsilon) \leq \left(\frac{3K}{\varepsilon} (Fu^2)^{K/2} \right)^{F \sum_k I_k}. \quad (53)$$

The proof of Lemma F.4 is also detailed in Sec. I.

Invoking Lemma F.4, we have:

$$\mathsf{N}(\mathcal{L}, \varepsilon/2L_f \sqrt{T}) \leq \left(\frac{6KL_f \sqrt{T} (F\alpha_u)^K}{\varepsilon} \right)^{F \sum_k I_k}, \quad (54)$$

where we have applied $\|\mathbf{U}_k(:, f)\|_2 \leq \sqrt{F}\alpha_u$. The covering number of the set $\mathcal{G} \circ \Xi$ is given by Lemma 14 of (Lin & Zhang, 2019) (cf. Lemma M.1 in Sec. M) as below:

$$\mathsf{N}(\mathcal{G} \circ \Xi, \varepsilon) \leq \exp \left(\frac{\|\mathbf{Z}\|_{\mathbb{F}} \mathfrak{R}_{\mathcal{G}}}{\varepsilon} \right)^{1/2}, \quad (55)$$

where $\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}] \in \mathbb{R}^{D \times S}$ and $\mathfrak{R}_{\mathcal{G}}$ denotes the complexity measure of the class \mathcal{G} . Applying (55), we get

$$\mathsf{N}(\mathcal{G} \circ \Xi, \varepsilon/2c_\lambda L_f \sqrt{T}) \leq \exp \left(\frac{2c_\lambda L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}} \mathfrak{R}_{\mathcal{G}}}{\varepsilon} \right)^{1/2}. \quad (56)$$

From (54) and (56), combined with (49) and (51), one can obtain that

$$\begin{aligned} \mathsf{N}(\mathcal{F}, \varepsilon) &\leq \mathsf{N}(\mathcal{L}, \varepsilon/2L_f \sqrt{T}) \times \mathsf{N}(\mathcal{G}, \varepsilon/2c_\lambda L_f \sqrt{T}) \\ &\leq \left(\frac{6KL_f \sqrt{T} (F\alpha_u)^K}{\varepsilon} \right)^{F \sum_k I_k} \times \exp \left(\frac{2c_\lambda L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}} \mathfrak{R}_{\mathcal{G}}}{\varepsilon} \right)^{1/2}. \end{aligned} \quad (57)$$

To characterize the Rademacher complexity $\mathfrak{R}_T(\mathcal{F})$ using the covering number $\mathsf{N}(\mathcal{F}, \varepsilon)$, we invoke the following lemma:

Lemma F.5. (Bartlett et al., 2017, Lemma A.5) The empirical Rademacher complexity of the set $\mathcal{F} \subset \mathbb{R}^T$ is upper bounded as follows:

$$\mathfrak{R}_T(\mathcal{F}) \leq \inf_{a>0} \left(\frac{4a}{\sqrt{T}} + \frac{12}{T} \int_a^{\sqrt{T}} \sqrt{\log \mathsf{N}(\mathcal{F}, \mu)} d\mu \right). \quad (58)$$

Applying Lemma F.5, we have the following set of relations:

$$\begin{aligned}
 \mathfrak{R}_T(\mathcal{F}) &\leq \inf_{a>0} \left(\frac{4a}{\sqrt{T}} + \frac{12}{T} \int_a^{\sqrt{T}} \sqrt{\log \mathbf{N}(\mathcal{F}, \mu)} d\mu \right) \\
 &\leq \inf_{a>0} \left(\frac{4a}{\sqrt{T}} + \frac{12}{\sqrt{T}} \sqrt{\log \mathbf{N}(\mathcal{F}, a)} \right) \\
 &\leq \inf_{a>0} \left(\frac{4a}{\sqrt{T}} + \frac{12}{\sqrt{T}} \sqrt{F \sum_k I_k \log \left(\frac{6KL_f \sqrt{T} (F\alpha_u)^K}{a} \right) + \left(\frac{c_\lambda L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}\mathcal{R}_G}}{a} \right)} \right) \\
 &\leq \left(\frac{4}{\sqrt{T}} + \frac{12}{\sqrt{T}} \sqrt{F \sum_k I_k \log \left(6KL_f \sqrt{T} (F\alpha_u)^K \right) + \left(F\alpha_u^K L_f \sqrt{T} \|\mathbf{Z}\|_{\mathbb{F}\mathcal{R}_G} \right)} \right), \tag{59}
 \end{aligned}$$

where the second inequality is obtained since $\sqrt{\log \mathbf{N}(\mathcal{F}, \mu)}$ increases monotonically with the decrease of μ and hence

$$\int_a^{\sqrt{T}} \sqrt{\log \mathbf{N}(\mathcal{F}, \mu)} d\mu \leq \sqrt{T} \sqrt{\log \mathbf{N}(\mathcal{F}, a)},$$

the third inequality is by applying (57), the last inequality is by setting $a = 1$, applying $c_\lambda = F\alpha_u^K$ and L_f is given by (50) with probability greater than $1 - e^{-\alpha(e^2-3)}$.

G. Proof of Lemma E.1

We start by noting that $\beta \leq m_i \leq \alpha$ (see (22) and (23)). By Lemma H.1, we have

$$\Pr(y_i \leq c) \geq 1 - e^{-\alpha(e^2-3)}, \quad c = \alpha(e^2 - 2). \tag{60}$$

Also, if $\beta < 1$, we have $f_i = m_i - y_i \log m_i \leq \alpha + c|\log \beta|$. Suppose $\beta \geq 1$, we have $f_i = m_i - y \log m_i \leq m_i + y_i \log m_i \leq \alpha + c \log \alpha$. Similarly, we get that $f_i \geq \beta - c \log \alpha$. Hence, we get

$$\begin{aligned}
 f_{\max} &\leq \alpha + c \max\{|\log \beta|, \log \alpha\} \\
 f_{\min} &\leq \beta - c \log \alpha,
 \end{aligned}$$

which hold with probability greater than $1 - e^{-\alpha(e^2-3)}$.

H. Proof of Fact F.3

We have

$$f_i(m_i) = m_i - y_i \log m_i.$$

Hence, we get

$$f'_i(m_i) = 1 - \frac{y_i}{m_i}. \tag{61}$$

To proceed, we invoke the following lemma:

Lemma H.1. (Cao & Xie, 2015), Lemma 9 (Tail Bound of Poisson) For $y \sim \text{Poisson}(m)$ with $m \leq \alpha$, we have

$$\Pr(y - m \geq t) \leq e^{-t} \tag{62}$$

for all $t \geq t_0$ where $t_0 = \alpha(e^2 - 3)$, where e is the Euler's number.

Utilizing Lemma H.1, it can be seen that

$$\begin{aligned} \Pr(y_i - \alpha \geq \alpha(e^2 - 3)) &\leq e^{-\alpha(e^2-3)} \\ \implies \Pr(y_i \leq \alpha(e^2 - 3) + \alpha) &\geq 1 - e^{-\alpha(e^2-3)}. \end{aligned}$$

Combining this result with (61), we get that with probability greater than $1 - e^{-\alpha(e^2-3)}$,

$$|f'_i(m_i)| \leq 1 + \frac{y_i}{m_i} \leq 1 + \frac{\alpha(e^2 - 2)}{\beta},$$

where the last inequality uses the assumption that $m_i \geq \beta$.

I. Proof of Lemma F.4

Let $\bar{\mathcal{L}}_\varepsilon$ denote the ε -net of \mathcal{L} . Consider $\bar{\mathbf{X}} \in \bar{\mathcal{L}}_\varepsilon$. In addition, let us define the following set:

$$\mathcal{R}_k = \{\mathbf{U} \in \mathbb{R}^{I_k \times F} \mid \|\mathbf{U}(:, f)\|_2 \leq u\}.$$

Let $\bar{\mathbf{U}}_k$ belongs to ε -net of \mathcal{R}_k such that $\|\mathbf{U}_k - \bar{\mathbf{U}}_k\|_F \leq \varepsilon$. Then, the cardinality of the ε -net of \mathcal{R}_k is given by (Wainwright, 2019)

$$\mathbf{N}(\mathcal{R}_k, \varepsilon) \leq \left(\frac{3\sqrt{Fu^2}}{\varepsilon} \right)^{I_k F}. \quad (63)$$

In order to derive the covering number of the set \mathcal{L} , we consider the following chain of relations:

$$\|\underline{\mathbf{X}} - \bar{\mathbf{X}}\|_F = \left\| \sum_{f=1}^F \mathbf{U}_1(:, f) \circ \dots \circ \mathbf{U}_K(:, f) - \sum_{f=1}^F \bar{\mathbf{U}}_1(:, f) \circ \dots \circ \bar{\mathbf{U}}_K(:, f) \right\|_F \quad (64a)$$

$$\begin{aligned} &\leq \left\| \sum_{f=1}^F (\mathbf{U}_1(:, f) - \bar{\mathbf{U}}_1(:, f)) \circ \mathbf{U}_2(:, f) \dots \circ \mathbf{U}_K(:, f) \right\|_F \\ &\quad + \left\| \sum_{f=1}^F \bar{\mathbf{U}}_1(:, f) \circ \left(\sum_{f=1}^F \mathbf{U}_2(:, f) \dots \circ \mathbf{U}_K(:, f) - \bar{\mathbf{U}}_2(:, f) \dots \circ \bar{\mathbf{U}}_K(:, f) \right) \right\|_F \end{aligned} \quad (64b)$$

$$\begin{aligned} &\leq (\sqrt{Fu^2})^{K-1} \left\| \sum_{f=1}^F (\mathbf{U}_1(:, f) - \bar{\mathbf{U}}_1(:, f)) \right\|_F \\ &\quad + \sqrt{Fu^2} \underbrace{\left\| \sum_{f=1}^F \mathbf{U}_2(:, f) \circ \dots \circ \mathbf{U}_K(:, f) - \sum_{f=1}^F \bar{\mathbf{U}}_2(:, f) \circ \dots \circ \bar{\mathbf{U}}_K(:, f) \right\|_F}_{Q_{\mathbf{U}_1}} \end{aligned} \quad (64c)$$

$$\leq (\sqrt{Fu^2})^{K-1} \|\mathbf{U}_1 - \bar{\mathbf{U}}_1\|_F + \sqrt{Fu^2} Q_{\mathbf{U}_1}. \quad (64d)$$

In a similar way as followed in (64), we can derive upper-bounds for every $Q_{\mathbf{U}_k}, k \in [K]$. Then, we can finally establish the below relationship:

$$\|\underline{\mathbf{X}} - \bar{\mathbf{X}}\|_F \leq (\sqrt{Fu^2})^{K-1} \sum_{k=1}^K \|\mathbf{U}_k - \bar{\mathbf{U}}_k\|_F. \quad (65)$$

The result in (65) implies that if there exists an $\varepsilon/K(\sqrt{Fu^2})^{K-1}$ -net to cover each \mathcal{R}_k , then we can construct an ε -net to cover \mathcal{L} . Then, the cardinality of the ε -net of \mathcal{L} is given by

$$\begin{aligned} N(\mathcal{L}, \varepsilon) &\leq \prod_{k=1}^K N(\mathcal{R}_k, \varepsilon/K(\sqrt{Fu^2})^{K-1}) \\ &\leq \left(\frac{3K}{\varepsilon} (Fu^2)^{K/2} \right)^{F \sum_k I_k}, \end{aligned}$$

where the inequality is by applying (63). Hence the proof.

J. Proof of Lemma D.2

Let us define

$$\hat{u} = \frac{1}{|\Theta|} \|\underline{\mathbf{M}}_1 - \underline{\mathbf{M}}_2\|_{\Theta}^2, \quad (66a)$$

$$u = \frac{1}{\prod_k I_k} \|\underline{\mathbf{M}}_1 - \underline{\mathbf{M}}_2\|_{\mathbb{F}}^2. \quad (66b)$$

Next, we consider the following lemma which is the Serfling's sampling-without-replacement extension of the Hoeffding's inequality (Serfling, 1974):

Lemma J.1. *Let X_1, X_2, \dots, X_M be a set of samples taken without replacement from $\{x_1, x_2, \dots, x_N\}$ of mean μ . Denote $a = \min_i x_i$ and $b = \max_i x_i$. Then*

$$\begin{aligned} \Pr \left[\left| \frac{1}{M} \sum_{i=1}^M X_i - \mu \right| \geq t \right] \\ \leq 2 \exp \left(- \frac{2Mt^2}{(1 - (M-1)/N)(b-a)^2} \right). \end{aligned}$$

One can see that u in (66b) denotes the mean of $\prod_k I_k$ terms of $(\hat{m}_i - m_i^{\natural})^2$, \hat{u} denotes the mean of $|\Theta|$ samples randomly drawn from $\{(\hat{m}_i - m_i^{\natural})^2\}$ without replacement. Hence, using the assumption that $i \in \Theta$ is drawn uniformly at random from $[I_1] \times \dots \times [I_K]$ and invoking Lemma J.1, we have

$$\Pr [|\hat{u} - u| \geq t] \leq 2 \exp \left(- \frac{2|\Theta|t^2}{(1 - (|\Theta|-1)/\prod_k I_k)\alpha^4} \right).$$

where we also applied $(\hat{m}_i - m_i^{\natural})^2 \leq \alpha^2$.

Let $\delta = 2 \exp \left(- \frac{2|\Theta|t^2}{(1 - (|\Theta|-1)/\prod_k I_k)\alpha^4} \right)$, we have the following result with probability $1 - \delta$:

$$|\hat{u} - u| \leq \alpha^2 \sqrt{\log \left(\frac{2}{\delta} \right) \left(1 - \frac{(|\Theta|-1)}{\prod_k I_k} \right) \frac{1}{2|\Theta|}}. \quad (67)$$

Using $|\sqrt{a} - \sqrt{b}| \leq \sqrt{|a-b|}$ for nonnegative a and b , we have

$$|\sqrt{\hat{u}} - \sqrt{u}| \leq \sqrt{|\hat{u} - u|},$$

which implies that

$$|\sqrt{\hat{u}} - \sqrt{u}| \leq \alpha \left(\log \left(\frac{2}{\delta} \right) \left(1 - \frac{(|\Theta|-1)}{\prod_k I_k} \right) \frac{1}{2|\Theta|} \right)^{1/4}, \quad (68)$$

holds with probability greater than $1 - \delta$.

K. Proof of Lemma D.3

First, let us define the following notations w.r.t the loss function $\ell(p_1, p_2) = (p_1 - p_2)^2$, $p_1, p_2 \in [0, 1]$ and the observed features $\mathbf{Z} = [z_{i_1}, \dots, z_{i_S}]$ where each z_{i_s} is sampled i.i.d. from the distribution \mathcal{D} :

$$L_{\mathbf{Z}}(\widehat{g}_{\theta}) \triangleq \frac{1}{S} \sum_{s=1}^S \ell(g^{\natural}(z_{i_s}), 1/\widehat{\xi}\widehat{g}_{\theta}(z_{i_s})) = \frac{1}{S} \sum_{s=1}^S (g^{\natural}(z_{i_s}) - 1/\widehat{\xi}\widehat{g}_{\theta}(z_{i_s}))^2 = \frac{\left\| \left[\mathbf{P}^{\natural} - 1/\widehat{\xi}\widehat{\mathbf{P}} \right]_{\Xi} \right\|_{\mathbb{F}}^2}{S} \quad (69a)$$

$$L_{\mathcal{D}}(\widehat{g}_{\theta}) \triangleq \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [\ell(g^{\natural}(\mathbf{z}), 1/\widehat{\xi}\widehat{g}_{\theta}(\mathbf{z}))] = \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [(g^{\natural}(\mathbf{z}) - 1/\widehat{\xi}\widehat{g}_{\theta}(\mathbf{z}))^2]. \quad (69b)$$

We invoke Theorem 26.5 in (Shalev-Shwartz & Ben-David, 2014) and get that with probability greater than $1 - \delta$:

$$\begin{aligned} L_{\mathcal{D}}(\widehat{g}_{\theta}) &\leq L_{\mathbf{Z}}(\widehat{g}_{\theta}) + 2\mathfrak{R}_S(\ell \circ \mathcal{G} \circ \Xi) + 4\bar{c} \sqrt{\frac{2 \log(4/\delta)}{S}} \\ \implies \mathbb{E}_{\mathbf{z} \sim \mathcal{D}} [(g^{\natural}(\mathbf{z}) - 1/\widehat{\xi}\widehat{g}_{\theta}(\mathbf{z}))^2] &\leq \frac{\left\| \left[\mathbf{P}^{\natural} - 1/\widehat{\xi}\widehat{\mathbf{P}} \right]_{\Xi} \right\|_{\mathbb{F}}^2}{S} + 2\mathfrak{R}_S(\mathcal{G} \circ \Xi) + 4\sqrt{\frac{2 \log(4/\delta)}{S}} \end{aligned} \quad (70)$$

where the last inequality utilizes the definitions in (69), the contraction lemma (Lemma 26.9 from (Shalev-Shwartz & Ben-David, 2014)) and also applied $\bar{c} = 1$ since $|\ell(\mathbf{f}, \mathbf{y})| \leq 1$ in our case. The term $\mathfrak{R}_S(\mathcal{G} \circ \Xi)$ denotes the empirical Rademacher complexity of the function class $\mathcal{G} \circ \Xi$ (see its definition in (52)) which is upperbounded via the sensitive complexity parameter $\mathcal{R}_{\mathcal{G}}$ as follows (Lin & Zhang, 2019):

$$\mathfrak{R}_S(\mathcal{G} \circ \Xi) \leq 16S^{-5/8} (2\|\mathbf{Z}\|_{\mathbb{F}} \mathcal{R}_{\mathcal{G}})^{\frac{1}{4}}.$$

This completes the proof.

L. Proof of Lemma C.1

Consider the following fact due to the concavity of the log function and Jensen's inequality (Boyd & Vandenberghe, 2004):

Fact L.1. Assume that α_f 's are certain scalars such that $\alpha_f > 0$, $\sum_f \alpha_f = 1$. Then, for any $u_f > 0$,

$$\log \sum_f \alpha_f u_f \geq \sum_f \alpha_f \log u_f.$$

To proceed, we define the following scalars for our case:

$$\alpha_i^{(f)} := \frac{\overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f)}{\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f)}, \quad \forall f$$

One can see that $\sum_{f=1}^F \alpha_i^{(f)} = 1$, $\alpha_i^{(f)} > 0$, $\forall f$. Hence, we can use Fact L.1 to have the following inequality:

$$\log \left(\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f) \right) \geq \sum_{f=1}^F \alpha_i^{(f)} \log \left(\frac{\overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f)}{\alpha_i^{(f)}} \right).$$

Applying the above relation, we immediately have

$$w(\mathbf{U}_k) \leq s(\mathbf{U}_k; \overline{\mathbf{U}}_k), \quad \forall \mathbf{U}_k. \quad (71)$$

By substituting $\mathbf{U}_k = \overline{\mathbf{U}}_k$ in $s(\mathbf{U}_k; \overline{\mathbf{U}}_k)$, we have

$$\begin{aligned} s(\overline{\mathbf{U}}_k; \overline{\mathbf{U}}_k) &= \sum_{i \in \Omega} \left[\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f) p_i - y_i \sum_{f=1}^F \alpha_i^{(f)} \log \left(\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f) \right) \right] \\ &= \sum_{i \in \Omega} \left[\left(\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f) \right) p_i - y_i \log \left(\sum_{f=1}^F \overline{U}_k(i_k, f) \prod_{j \neq k} U_j(i_j, f) \right) \right] \\ &= w(\overline{\mathbf{U}}_k), \end{aligned} \quad (72)$$

where the second equality is due to $\sum_f \alpha_i^{(f)} = 1$.

By taking the derivative of $s(\mathbf{U}_k; \bar{\mathbf{U}}_k)$ w.r.t. $\mathbf{U}_k(i_k, f)$, we have

$$\begin{aligned}
 \nabla_{\mathbf{U}_k(i_k, f)} s(\mathbf{U}_k; \bar{\mathbf{U}}_k) &= \sum_{\mathbf{i}_{-k} \in \Omega_{-k}} \left[\prod_{j \neq k} \mathbf{U}_j(i_j, f) p_{\mathbf{i}} \right. \\
 &\quad \left. - y_{\mathbf{i}} \alpha_{\mathbf{i}}^{(f)} \left(\frac{\alpha_{\mathbf{i}}^{(f)}}{\mathbf{U}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f)} \right) \left(\frac{\prod_{j \neq k} \mathbf{U}_j(i_j, f)}{\alpha_{\mathbf{i}}^{(f)}} \right) \right] \\
 &= \sum_{\mathbf{i}_{-k} \in \Omega_{-k}} \left[\prod_{j \neq k} \mathbf{U}_j(i_j, f) p_{\mathbf{i}} - y_{\mathbf{i}} \left(\frac{\alpha_{\mathbf{i}}^{(f)}}{\mathbf{U}_k(i_k, f)} \right) \right] \\
 &= \sum_{\mathbf{i}_{-k} \in \Omega_{-k}} \left[\prod_{j \neq k} \mathbf{U}_j(i_j, f) p_{\mathbf{i}} \right. \\
 &\quad \left. - y_{\mathbf{i}} \left(\frac{\bar{\mathbf{U}}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f)}{\mathbf{U}_k(i_k, f) (\sum_{f=1}^F \bar{\mathbf{U}}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f))} \right) \right], \tag{73}
 \end{aligned}$$

where $\mathbf{i}_{-k} = (i_1, \dots, i_{k-1}, i_{k+1}, \dots, i_K)$, Ω_{-k} collects all \mathbf{i}_{-k} 's such that the corresponding $y_{\mathbf{i}}$'s are observed and the last equality (73) is obtained by substituting the definition of $\alpha_{\mathbf{i}}^{(f)}$. Further, by taking the derivative of $w(\mathbf{U}_k)$ w.r.t. $\mathbf{U}_k(i_k, f)$, we get

$$\nabla_{\mathbf{U}_k(i_k, f)} w(\mathbf{U}_k) = \sum_{\mathbf{i}_{-k} \in \Omega_{-k}} \left[\prod_{j \neq k} \mathbf{U}_j(i_j, f) p_{\mathbf{i}} - y_{\mathbf{i}} \left(\frac{\prod_{j \neq k} \mathbf{U}_j(i_j, f)}{(\sum_{f=1}^F \bar{\mathbf{U}}_k(i_k, f) \prod_{j \neq k} \mathbf{U}_j(i_j, f))} \right) \right]. \tag{74}$$

From (73) and (74), one can see that

$$\nabla_{\mathbf{U}_k(i_k, f)} w(\bar{\mathbf{U}}_k) = \nabla_{\mathbf{U}_k(i_k, f)} s(\bar{\mathbf{U}}_k; \bar{\mathbf{U}}_k). \tag{75}$$

M. Characterization of the Complexity Measure $\mathcal{R}_{\mathcal{G}}$

The work in (Lin & Zhang, 2019) introduces a complexity measure, called the sensitive complexity and denoted as $\mathcal{R}_{\mathcal{G}}$, to assess the expressive power of various families of neural network classes. For example, regarding the fully connected neural network function classes, the following result is presented:

Lemma M.1. (Lin & Zhang, 2019, Lemma 14) Consider the following function class:

$$\mathcal{G} = \{g_{\theta} : \mathbb{R}^D \rightarrow [0, 1] \mid g_{\theta}(\mathbf{z}) = \sigma(\mathbf{w}_L^{\top} \sigma(\mathbf{W}_{L-1} \sigma(\dots \sigma(\mathbf{W}_1 \mathbf{z} + \mathbf{b}_1) + \dots) + \mathbf{b}_L)), \forall \mathbf{z} \in \mathbb{R}^d\},$$

where $\sigma(\cdot) = [\sigma(\cdot), \dots, \sigma(\cdot)]^{\top}$ denote the activation function, L denotes the number of layers, $\mathbf{W}_{\ell} \in \mathbb{R}^{d_{\ell} \times d_{\ell-1}}$, $d_0 = D$, $\mathbf{b}_{\ell} \in \mathbb{R}^{d_{\ell}}$, and θ represents the parameters such that $\theta = (\{\mathbf{W}_{\ell}, \mathbf{b}_{\ell}\}_{\ell=1}^{L-1}, \mathbf{w}_L, \mathbf{b}_L)$. Then, the covering number of $\mathcal{G} \circ \Xi$ with respect to the Frobenius norm satisfies

$$\mathsf{N}(\mathcal{G} \circ \Xi, \varepsilon) \leq \exp\left(\frac{\|\mathbf{Z}\|_{\text{F}} \mathcal{R}_{\mathcal{G}}}{\varepsilon}\right)^{1/2}, \tag{76}$$

where

$$\mathcal{G} \circ \Xi = \{g_{\theta}([\mathbf{Z}]_{\Xi}) = [g(\mathbf{z}_{i_1}), \dots, g(\mathbf{z}_{i_S})]^{\top}, g_{\theta} \in \mathcal{G}\},$$

$\mathbf{Z} = [\mathbf{z}_{i_1}, \dots, \mathbf{z}_{i_S}] \in \mathbb{R}^{D \times S}$ represents the training data and

$$\mathcal{R}_{\mathcal{G}} = \left(2 \prod_{\ell=1}^L \rho_{\ell} s_{\ell}\right) \left(\sum_{\ell=1}^L \frac{d_{\ell}^2 d_{\ell-1}^2 a_{\ell}}{s_{\ell}}\right) L^2,$$

where $\|\mathbf{W}_{\ell}\|_{\text{F}} \leq a_{\ell}$, $\|\mathbf{W}_{\ell}\|_2 \leq s_{\ell}$, the activation function σ_{ℓ} is ρ_{ℓ} -Lipschitz and $\sigma_{\ell}(0) = 0$.




Figure 2. Average MAE_p (top) and MAE_λ (bottom) over 20 random trials plotted against iterations for different settings $\Omega \subset \Xi$, $\Omega = \Xi$, and $\Omega \supset \Xi$. $K = 3, I_k = 20, \forall k, F = 3, D = 10, \text{SNR} = 40\text{dB}, g(\mathbf{z}) = \text{sigmoid}(\boldsymbol{\nu}^\top (3\mathbf{z}^3 + 0.2\mathbf{z}))$, where the vector $\boldsymbol{\nu} \in \mathbb{R}^D$ is generated by randomly sampling its entries a uniform distribution between 0 and 1.

N. More Details on Experiments

N.1. Synthetic Data Experiments - Implementation Settings.

To learn $g(\cdot)$, we use a neural network $g_\theta(\cdot)$ with 3 hidden layers and 20 ReLU activation functions in each hidden layer. The optimizer for handling the subproblems of θ is Adam (Kingma & Ba, 2015) with an initial learning rate of 10^{-3} . The optimizer uses a batch size of 1024. The subproblems are stopped when the relative change in the corresponding objective functions is smaller than 10^{-6} . Also, the algorithm is stopped if the relative change in the overall objective function is less than 10^{-6} or if 100 BCD iterations are completed. We fix the regularization term in the loss function to be $\ell(x, y) = (x - y)^2$. The regularization parameter μ is chosen to be a high value, i.e., $\mu = 2000$ across all experiments.

N.2. Additional Synthetic Data Experiments

In Fig. 2, we compare the performance of the proposed approach and two baselines by plotting MAE_p and MAE_λ over iterations, under different conditions such as $\Omega \subset \Xi$, $\Omega \supset \Xi$, and $\Omega = \Xi$. The baseline NTF-CPD-KL does not consider the detection probability in their model. Hence, we only plot its MAE_λ . It can be observed that UncleTC outperforms the baselines after around 10 iterations under all the settings under test.

Table 7 shows the results on a different synthetic dataset. Here, we consider another nonlinear function for the data generation, i.e.,

$$g(\mathbf{z}) = \text{sigmoid}(\boldsymbol{\nu}^\top (0.1 \log(\mathbf{z}^2) + 0.1\mathbf{z}^2)),$$

where the vector $\boldsymbol{\nu} \in \mathbb{R}^D$ is generated by randomly sampling its entries from the uniform distribution between 0 and 1. We fix $\gamma_\Omega = 0.2, \gamma_\Xi = 0.3, \gamma_\theta = 0.2, \text{SNR} = 40\text{dB}$, and vary the rank F of the ground-truth tensor $\underline{\mathbf{A}}^{\text{H}}$. One can note that under these settings, the proposed algorithm still consistently gives better performance relative to the baselines.

Table 8 shows the results where wrong tensor ranks \hat{F} 's were used in our algorithms. In practice, the underlying data generation process is unknown and hence, the true rank F is hard to know or estimate. Hence, it is of interest to understand

Table 7. Average MAE_p and MAE_λ over 20 random trials under different values of F . $K = 3, I_k = 20, \forall k, D = 10, \gamma_\Omega = 0.2, \gamma_\Xi = 0.3, \gamma_\Theta = 0.2, \text{SNR} = 40\text{dB}, g(\mathbf{z}) = \text{sigmoid}(\boldsymbol{\nu}^\top (0.1 \log(\mathbf{z}^2) + 0.1 \mathbf{z}^2))$.

Algorithm	Metric	$F = 5$	$F = 7$	$F = 10$
UncleTC	MAE _p	0.059 ± 0.001	0.049 ± 0.029	0.030 ± 0.002
	MAE _λ	0.058 ± 0.007	0.052 ± 0.021	0.074 ± 0.001
UncleTC (Linear)	MAE _p	0.148 ± 0.018	0.187 ± 0.026	0.077 ± 0.005
	MAE _λ	0.102 ± 0.024	0.202 ± 0.024	0.131 ± 0.014
NTF-CPD-KL	MAE _λ	0.166 ± 0.050	0.214 ± 0.085	0.120 ± 0.001

Table 8. Average MAE_p and MAE_λ over 20 random trials for different \hat{F} with true rank $F = 7$. $K = 3, I_k = 20, \forall k, D = 10, \gamma_\Omega = 0.2, \gamma_\Xi = 0.3, \gamma_\Theta = 0.2, \text{SNR} = 40\text{dB}, g(\mathbf{z}) = \text{sigmoid}(\boldsymbol{\nu}^\top (0.1 \log(\mathbf{z}^2) + 0.1 \mathbf{z}^2))$.

Algorithm	Metric	$\hat{F} = 3$	$\hat{F} = 5$	$\hat{F} = 7$	$\hat{F} = 10$	$\hat{F} = 12$
UncleTC	MAE _p	0.052 ± 0.015	0.042 ± 0.014	0.046 ± 0.024	0.051 ± 0.018	0.046 ± 0.019
	MAE _λ	0.125 ± 0.008	0.084 ± 0.011	0.060 ± 0.018	0.075 ± 0.019	0.076 ± 0.019
UncleTC (Linear)	MAE _p	0.110 ± 0.041	0.108 ± 0.057	0.130 ± 0.045	0.112 ± 0.041	0.116 ± 0.046
	MAE _λ	0.141 ± 0.017	0.120 ± 0.039	0.140 ± 0.046	0.155 ± 0.064	0.197 ± 0.143
NTF-CPD-KL	MAE _p	0.151 ± 0.029	0.134 ± 0.046	0.159 ± 0.071	0.157 ± 0.072	0.151 ± 0.069

the algorithm’s robustness to using a wrong \hat{F} . One can see that the MAE values are worse when $\hat{F} < F$ but become better when $\hat{F} \geq F$. This makes sense since underestimating the rank means that less information of the true model is captured by the algorithms.

Table 9 shows the performance under various μ ’s, i.e., the regularization parameter of ℓ in UncleTC. The settings follow those of Table 7 with rank $F = 3$. The results indicate that using reasonably large μ ($\mu \geq 1000$) may be preferable, as we hope that the regularization term can enforce equality. In practice, one may also use a validation set to choose μ .

N.3. Real Data Experiments - Dataset Details.

COVID-19 Data. The dataset (Zhang et al., 2020b) includes the number of reported COVID-19 cases of about 2270 US counties during 462 days in 2020-2021. The dataset has 25 attributes, including social distancing index, percentage of people staying home, trips per person, percentage of out-of-county trips, percentage of out-of-state trips, transit mode share, miles per person, work trips per person, non-work trips per person, COVID exposure per 1000 people, percentage of people older than 60, median income, percentage of African Americans, percentage of Hispanic Americans, population density, number of hot spots per 1000 people, total population, percentage of people working from home, imported COVID cases, hospital beds per 1000 people, testing capacity gap, tests done per 1000 people, unemployment rate, cumulative inflation rate, and unemployment claims per 1000 people.

In order to extract a count-type tensor from the COVID-19 dataset, we adopt the following procedure: Each county is represented using two coordinates, i.e., the rounded value of latitude and longitude of its geographic center (we collected this information from a publicly available website³). The rounded value of latitude points ranges between 20 and 69, out of which 43 values correspond to the latitudes of at least one county. On the other hand, the discretized longitude points ranges between -165 and -69, out of which 80 discrete points correspond to the longitudes of at least one county. The third coordinate of the tensor is the dates on which the COVID-19 cases were reported. Hence, each entry of the tensor represents the number of reported COVID-19 cases associated with a latitude point and a longitude point on a particular day. We use the data corresponding to 60 days from January 1, 2021 to March 1, 2021. We form a $43 \times 80 \times 60$ tensor, which has 19.73% of observed entries, out of which 15.43% are nonzeros entries and 4.30% are zeros.

Pollinator-Plant Interaction (PPI) Data. The Plant-Pollinator Interaction (PPI) dataset is a publicly available dataset, collected by the researchers at the H.J. Andrews (HJA) Long-term Ecological Research site in Oregon, USA (Seo et al., 2022; Daly et al., 2019). The dataset comprises plant-pollinator interactions observed at about 12 meadows over a period of seven years. Each meadow was visited 5-7 times per year by student researchers. The dataset contains interactions between 69 plant species and 215 pollinator species observed during 37 visits. Only 1.07% of the entries are recorded. The counts in

³<https://public.opendatasoft.com/explore/dataset/us-county-boundaries>

Table 9. Average MAE_p and MAE_λ for $UncleTC$ over 20 random trials under different μ values. We fix $K = 3, F = 3, D = 10, \gamma_\Omega = 0.2, \gamma_\Xi = 0.3, \gamma_\Theta = 0.2, SNR = 40dB, g(z) = \text{sigmoid}(\nu^\top (0.1 \log(z^2) + 0.1z^2))$

Settings	Metric	$\mu = 100$	$\mu = 500$	$\mu = 1000$	$\mu = 2000$
$I_k = 20, \gamma_\Omega = 0.3, \gamma_\Xi = 0.3$	MAE_p	0.184 ± 0.025	0.074 ± 0.024	0.052 ± 0.011	0.042 ± 0.004
	MAE_λ	0.171 ± 0.031	0.067 ± 0.029	0.042 ± 0.014	0.028 ± 0.004
$I_k = 25, \gamma_\Omega = 0.3, \gamma_\Xi = 0.3$	MAE_p	0.198 ± 0.054	0.083 ± 0.030	0.056 ± 0.015	0.054 ± 0.013
	MAE_λ	0.184 ± 0.049	0.069 ± 0.029	0.039 ± 0.013	0.033 ± 0.010
$I_k = 20, \gamma_\Omega = 0.1, \gamma_\Xi = 0.3$	MAE_p	0.238 ± 0.069	0.151 ± 0.085	0.108 ± 0.074	0.080 ± 0.044
	MAE_λ	0.256 ± 0.087	0.166 ± 0.097	0.113 ± 0.084	0.078 ± 0.047
$I_k = 25, \gamma_\Omega = 0.1, \gamma_\Xi = 0.3$	MAE_p	0.260 ± 0.038	0.156 ± 0.073	0.141 ± 0.079	0.121 ± 0.079
	MAE_λ	0.305 ± 0.045	0.203 ± 0.111	0.194 ± 0.126	0.201 ± 0.134




Figure 3. The top 10 plant-pollinator interactions having highest average detection probabilities (left) and lowest average average detection probabilities (right) as identified by the proposed $UncleTC$.

this dataset are widely believed to be under-counted (Fu et al., 2021; Seo & Hutchinson, 2018). We consider a subset of the dataset by selecting 50 plants and 50 pollinators that exhibit the highest number of interactions. This leads to a $50 \times 50 \times 37$ integer tensor with 22.06% observed entries, out of which 4.01% are nonzero entries and 18.05% are zero entries. The zero entries indicate that the plant and pollinator were both observed by the researcher during the visit, but they were not observed to be interacting with one another.

The dataset has 37 attributes:

- 12 plant species static traits: the estimate of reward per flower, the basic ecological life form of a plant, floral structure, peduncle feebleness, the level of exclusion, exclusion by flower color, exclusion by pendant position, time of diel availability, size of pollen grain, the ability of the inflorescence to support the arrival of the flower visitor, the possible affinity of plant species to edaphic characteristics, and dispersal mechanism.
- 11 pollinator species static traits: the length of a species body size, the width of a species body size, the depth of a species body size, individual biomass, size, energy requirement, a species behavior characterized by active time, the level of exclusion, an indicator showing whether the platform is required to approach flowers, tongue length, and tube size.
- 1 plant temporal feature: flower abundance
- 4 temporal features: month, day, wind, and cloud
- 7 temperature features: mean, max, min, max time, min time, cumulative degree days with 5 degrees, and cumulative degree days with 10 degrees.
- 2 precipitation features: daily precipitation, and antecedent with $k = 0.9$




Figure 4. (left) The detection probabilities of the interaction pair *Senecio integerrimus*–*Apis mellifera*, which is reported by our algorithm to have the highest average detection probability. (right) The detection probabilities of the pair *Aquilegia formosa*–*Bombus flavifrons*, which is reported by our algorithm to have the lowest average detection probability.

N.4. Real Data Experiments - Implementation Settings.

For the proposed `UncleTC`, we fix `ReLU` as the neural network activation function and `Adam` as the the optimizer. Other parameters such as rank F , the regularization parameter μ , batch size, learning rates, and neural network size are selected via the grid search hyper-parameter tuning strategy on the validation set where the `rRMSE` is used to measure the validation score. For the proposed `UncleTC`, we choose the rank F of the tensor from $[5, 10, 15, 20]$, the regularization parameter μ from $[500, 1000, 2000, 3000]$, the number of hidden units from $[10, 20, 30]$, the number of hidden layers from $[1, 3, 5]$, the learning rate from $[0.001, 0.005, 0.01]$ and the batch size from $[256, 512, 1024]$. For all the low-rank tensor decomposition-based baselines, we select the rank from $[5, 10, 15, 20]$. In addition, the sparsity parameter of the `BPTF-CPD` (Schein et al., 2015) is selected from $[0.02, 0.05, 0.1, 0.2]$ and the regularization parameter of the `HaLRTC` (Liu et al., 2013) is selected from $[10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}]$, based on the recommended range of values from the respective papers. The number of iterations to stop training all the iterative algorithms is also chosen using the validation sets.

In the experiments, we perform 5-fold cross-validation to characterize the prediction performance with k chosen to be 5, where 3 splits serve for training, and the other two for validation and testing, respectively. The roles of the sets are switched in a cyclical way in different trials.

N.5. Real Data Experiments - Additional Results.

Fig. 3 displays the top 10 interactions with the highest and lowest average detection probabilities, as discovered by our proposed `UncleTC` method. Generally, interactions involving larger, more common species are more likely to have higher detection probabilities, while smaller, rarer species tend to have lower detection probabilities. The species with the highest average detection probabilities have a larger number of observed interactions compared to those with the lowest average detection probabilities. However, it is worth noting that not all species follow this trend. Other factors such as species rareness, environmental conditions, or seasonal variability may contribute to these differences (Dorado et al., 2011; Chacoff et al., 2012).

In Fig. 4, we present the periodic variation of detection probabilities for specific species interactions over the years. Both figures illustrate temporal variation across visits and years for two different interactions—one with the highest average detection probability and the other with the lowest. Despite the variations, the overall range remains relatively consistent, with probabilities ranging from approximately 0.4 to 0.8 in the left figure and from 0.2 to 0.6 in the right figure, ultimately leading to the highest and lowest average detection probabilities.

Considering the exceptions where detection probability does not align with species activity levels, we further examined interactions that show the most and least significant variations in detection probabilities over time. Fig. 5 plots the detection probabilities over visits for interactions with the highest (left) and lowest (right) standard deviation of detection probabilities




Figure 5. (left) The detection probability of the interaction *Gilia capitata*–*Apis mellifera* over different years, which is reported by our algorithm to have the highest standard deviation of the detection probability across 37 visits. (right) The detection probability pattern for the interaction *Eriophyllum lanatum*–*Bombus mixtus*, which is reported by our method to have the lowest standard deviation of the detection probability across 37 visits.

across visits and annually. The left figure confirms that the detection probability of the interaction between two well-known common species (*Gilia capitata* and *Apis mellifera*) is significantly influenced by the time of year, resulting in a relatively lower average detection probability. The right figure provides another example of an interaction between another plant and pollinator species (*Eriophyllum lanatum* and *Bombus mixtus*), with relatively lower detection probabilities over time. Such variation may be a product of the phenology of the species (i.e., when flowers are blooming and when insects are flying).