

Noisy Matrix Completion under Sparse Factor Models

Akshay Soni, Swayambhoo Jain, Jarvis Haupt, and Stefano Gonella

Abstract

This paper examines a general class of noisy matrix completion tasks where the goal is to estimate a matrix from observations obtained at a subset of its entries, each of which is subject to random noise or corruption. Our specific focus is on settings where the matrix to be estimated is well-approximated by a product of two (a priori unknown) matrices, one of which is sparse. Such structural models – referred to here as “sparse factor models” – have been widely used, for example, in subspace clustering applications, as well as in contemporary sparse modeling and dictionary learning tasks. Our main theoretical contributions are estimation error bounds for sparsity-regularized maximum likelihood estimators for problems of this form, which are applicable to a number of different observation noise or corruption models. Several specific implications are examined, including scenarios where observations are corrupted by additive Gaussian noise or additive heavier-tailed (Laplace) noise, Poisson-distributed observations, and highly-quantized (e.g., one-bit) observations. We also propose a simple algorithmic approach based on the alternating direction method of multipliers for these tasks, and provide experimental evidence to support our error analyses.

Index Terms

Penalized maximum likelihood estimation, dictionary learning, matrix completion, subspace clustering.

I. INTRODUCTION

In recent years, there has been significant research activity aimed at the analysis and development of efficient *matrix completion* methods, which seek to “impute” missing elements of a matrix given possibly noisy or corrupted observations collected at a subset of its locations. Let $\mathbf{X}^* \in \mathbb{R}^{n_1 \times n_2}$ denote a matrix whose elements we wish to estimate, and suppose that we observe \mathbf{X}^* at only a subset $\mathcal{S} \subset [n_1] \times [n_2]$ of its locations, where $[n_1] = \{1, 2, \dots, n_1\}$ is the set of all positive integers less or equal to n_1 (and similarly for n_2), obtaining at each $(i, j) \in \mathcal{S}$ a noisy, corrupted, or inexact measurement denoted by $Y_{i,j}$. The overall aim is to estimate \mathbf{X}^* given \mathcal{S} and the observations $\{Y_{i,j}\}_{(i,j) \in \mathcal{S}}$. Of course, such estimation problems may be ill-posed without further assumptions, since the values of \mathbf{X}^* at the unobserved locations could in general be arbitrary. A common approach is to augment the inference method with an assumption that the underlying matrix to be estimated exhibits some form of intrinsic *low-dimensional* structure.

Manuscript submitted October 31, 2014. AS, SJ, and JH are with the Department of Electrical and Computer Engineering, and SG is with the Department of Civil, Environmental, and Geo- Engineering, at the University of Minnesota – Twin Cities. Author emails: {sonix022, jainx174, jdhaupt, sgonella}@umn.edu. The first two (student) authors contributed equally to this work. A preliminary version of this work will appear at the 2014 Global Conference on Signal and Information Processing (GlobalSIP). JH graciously acknowledges support from NSF Award AST-1247885, and the DARPA Young Faculty Award, Grant N66001-14-1-4047.

One application where such techniques have been successfully utilized is collaborative filtering (e.g., as in the well-known *Netflix Prize* competition [1]). There, the matrix to be estimated corresponds to an array of users' preferences or ratings for a collection of items (which could be quantized, e.g., to one of a number of levels); accurately inferring missing entries of the underlying matrix is a useful initial step in *recommending* items (here, movies or shows) to users deemed likely to rate them favorably. A popular approach to this problem utilizes a low-rank modeling assumption, which implicitly assumes that individual ratings depend on some unknown but nominally small number (say r) of features, so that each element of \mathbf{X}^* may be described as an inner product between two length- r vectors – one quantifying how well each of the features are embodied or represented by a given item, and the other describing a user's affinity for each of the features. Recent works examining the efficacy of low-rank models for matrix completion include [2]–[8].

Several other applications where analogous ideas have been employed, but which leverage different structural modeling assumptions, include:

- **Sparse Coding for Image Inpainting and Demosaicing:** Suppose that the underlying data to be estimated takes the form of an $n_1 \times n_2$ color image, which may be interpreted as an $n_1 \times n_2 \times 3$ array (the three levels correspond to values in three color planes). The image inpainting task amounts to estimating the image from a collection of (possibly noisy) observations obtained at individual pixel locations (so that at each pixel, either all or none of the color planes are observed), and the demosaicing task entails estimating the image from noisy measurements corresponding to only one of the 3 possible color planes at each pixel. The recent work [9] proposed estimation approaches for these tasks that leverage *local shared structure* manifesting at the patch level. Specifically, in that work, the overall image to be estimated is viewed equivalently as a matrix comprised of vectorized versions of its small (e.g., $5 \times 5 \times 3$ or $8 \times 8 \times 3$) blocks, and the missing values are imputed using a structural assumption that this patch-based matrix be well-approximated by a product of two matrices, one of which is sparse.
- **Sparse Models for Learning and Content Analytics:** A recent work [10] investigated a matrix completion approach to machine-based learning analytics. There, the elements of the $n_1 \times n_2$ matrix to be estimated, say \mathbf{X}^* , are related to the probability with which one of n_1 questions will be answered correctly by one of n_2 “learners” through a *link function* $\Phi : \mathbb{R} \rightarrow [0, 1]$, so that the value $\Phi(X_{i,j}^*)$ denotes the *probability* with which question i will be correctly answered by learner j . The observed data are a collection of some $m < n_1 n_2$ binary values, which may be interpreted as (random) Bernoulli($\Phi(X_{i,j}^*)$) variables. The approach proposed in [10] entails maximum-likelihood estimation of the unknown latent factors of \mathbf{X}^* , under an assumption that \mathbf{X}^* be well-approximated by a sum of two matrices, the first being product of a sparse non-negative matrix (relating questions to some latent “concepts”) and a matrix relating a learner's knowledge to the concepts, and the second quantifying the intrinsic difficulty of each question.
- **Subspace Clustering from Missing Data:** The general subspace clustering problem entails separating a collection of data points, using an assumption that similar points are described as points lying in the same subspace,

so that the overall collection of data are represented as points belonging generally to a union of (ostensibly, low-dimensional) subspaces. This general task finds application in image processing, computer vision, and disease detection, to name a few (see, e.g., [11]–[16], and the references therein). One direct way to perform clustering in such applications entails approximating the underlying matrix \mathbf{X}^* whose columns comprise the (uncorrupted) data points by a product of two matrices, the second of which is sparse, so that the support (the set of locations of the nonzero elements) of each column of the sparse matrix factor identifies the subspace to which the corresponding column of \mathbf{X}^* belongs.

While these examples all seem qualitatively similar in scope, their algorithmic and analytical tractability can vary significantly depending on the type of structural model adopted. In the collaborative filtering application, for example, a desirable aspect of adopting low-rank models is that the associated inference (imputation) procedures can be relaxed to efficient convex methods that are amenable to precise performance analyses. Indeed, the statistical performance of convex methods for low-rank matrix completion are now well-understood in noise-free settings (see, e.g., [2]–[5]), in settings where observations are corrupted by some form of additive uncertainty [17]–[21], and even in settings where the observations may be interpreted as nonlinear (e.g., highly-quantized) functions of the underlying matrix entries [22]–[24]. In contrast, the aforementioned inference methods based on general bilinear (and sparse) factor models are difficult to solve to global optimality, and are instead replaced by tractable alternating minimization methods. More fundamentally, the statistical performance of inference methods based on these more general bilinear models, in scenarios where the observations could arise from general (perhaps nonlinear) corruption models or could even be multi-modal in nature, has not (to our knowledge) been fully characterized.

This work provides some initial results in this direction. We establish a general-purpose estimation error guarantee for matrix completion problems characterized by any of a number of structural data models and observation noise/corruption models. For concreteness, we instantiate our main result here for the special case where the matrix to be estimated adheres to a *sparse factor model*, meaning that it is well-approximated by the product of two matrices, one of which is sparse (or approximately so). Sparse factor models are inherent in the modeling assumptions adopted in the aforementioned works on image denoising/demosaicing, content analytics, and subspace clustering, and are also at the heart of recent related efforts in dictionary learning [25]–[27]. Sparse factor models may also serve as a well-motivated extension to the low-rank models often utilized in collaborative filtering tasks. There, while it is reasonable to assume that users’ preferences will depend on a small number of abstract features, it may be that any particular user’s preference relies heavily on only a subset of the features, and that the features that are most influential in forming a rating may vary from user to user. Low rank models alone are insufficient for capturing this “higher order” structure on the latent factors, while this behavior may be well-described using the sparse factor models we consider here.

A. Our Contributions

We address general problems of matrix completion under sparse factor modeling assumptions using the machinery of *complexity-regularized maximum likelihood* estimation. Our main contributions come in the form of estimation

error bounds that are applicable in settings where the available data correspond to an incomplete collection of noisy observations of elements of the underlying matrix (obtained at random locations), and under general (random) noise/corruption models. We examine several specific implications of our main result, including for scenarios characterized by additive Gaussian noise or additive heavier-tailed (Laplace) noise, Poisson-distributed observations, and highly-quantized (e.g., one-bit) observations. Where possible, we draw direct comparisons with existing results in the low-rank matrix completion literature, to illustrate the potential benefit of leveraging additional structure in the latent factors. We also propose an efficient unified algorithmic approach based on the alternating direction method of multipliers [28] for obtaining a local solution to the (non-convex) optimizations prescribed by our analysis, and provide experimental evidence to support our error results.

B. Connections with Existing Works

As alluded above, our theoretical analyses here are based on the framework of complexity regularized maximum likelihood estimation [29], [30], which has been utilized in a number of works to establish error bounds for Poisson estimation problems using multi scale models [31], [32], transform domain sparsity models [33], and dictionary-based matrix factorization models [34]. Here, our analysis extends that framework to the “missing data” scenarios inherent in matrix completion tasks (and also provides a missing-data extension of our own prior work on dictionary learning from 1-bit data [35]).

Our proposed algorithmic approach is based on the alternating direction method of multipliers (ADMM) [28]. ADMM-based methods for related tasks in *dictionary learning* (DL) were described recently in [36], and while our algorithmic approach here is qualitatively similar to that work, we consider missing data scenarios as well as more general loss functions that arise as negative log-likelihoods for our various probabilistic corruption models (thus generalizing these techniques beyond common squared error losses). In addition, our algorithmic framework also allows for direct incorporation of constraints not only on estimates of the matrix factors, but also on the estimate of \mathbf{X}^* itself to account for entry-wise structural constraints that could arise naturally in many matrix completion scenarios. Several other recent efforts in the DL literature have proposed algorithmic procedures for coping with missing data [37], [38], and a survey of algorithmic approaches to generalized low-rank modeling tasks is given in the recent work [39].

Our inference tasks here essentially entail learning two factors in a bilinear model. With a few notable exceptions (e.g., low-rank matrices, and certain non-negative matrices [40]–[43]), the joint non-convexity of these problems can complicate their analysis. Recently, several efforts in the dictionary learning literature have established theoretical guarantees on identifiability, as well as local correctness of a number of factorization methods [44]–[49], including in noisy settings [50]. Our efforts here may be seen as a complement to those works, providing additional insight into the achievable *statistical* performance of similar methods under somewhat general noise models.

The factor models we employ here essentially enforce that each column of \mathbf{X}^* lie in a union of linear subspaces. In this sense our efforts here are also closely related to problems in sparse principal component analysis [51], which seek to decompose the (sample) covariance matrix of a collection of data points as a sum of rank-one factors

expressible as outer products of sparse vectors. Several efforts have examined algorithmic approaches to the sparse PCA problem based on greedy methods [52] or convex relaxations [53]–[55], and very recently several efforts have examined the statistical performance of cardinality- (or ℓ_0 -) constrained methods for identifying the first sparse principal component [56], [57]. These latter approaches are related to our effort here, as our analysis below pertains to the performance of matrix completion methods utilizing an ℓ_0 penalty on one of the matrix factors.

Finally, we note that problems of subspace clustering from missing or noisy data have received considerable attention in recent years. Algorithmic approaches to subspace clustering with missing data were proposed in [16], [58], [59], and several recent works have identified sufficient conditions under which tractable algorithms will provably recover the unknown subspaces in missing data (but noise-free) scenarios [60]. Robustness of subspace clustering methods to missing data, additive noise, and potentially large-valued outliers were examined recently in [15], [16], [61].

C. Outline

The remainder of this paper is organized as follows. Following a brief discussion of several preliminaries (below), we formalize our problem in Section II and present our main result establishing estimation error guarantees for a general class of estimation problems characterized by incomplete and noisy observations. In Section III we discuss implications of this result for several specific noise models. In Section IV we discuss a unified algorithmic approach to problems of this form, based on the alternating direction method of multipliers, and provide a brief experimental investigation that partially validates our theoretical analyses. We conclude with a brief discussion in Section V. Auxiliary material and detailed proofs are relegated to the appendix.

D. Preliminaries

To set the stage for the statement of our main result, we remind the reader of a few key concepts. First, recall that for $p \leq 1$ a vector $\mathbf{x} \in \mathbb{R}^n$ is said to belong to a weak- ℓ_p ball of radius $R > 0$, denoted $\mathbf{x} \in w\ell_p(R)$, if its ordered elements $|x_{(1)}| \geq |x_{(2)}| \geq \dots \geq |x_{(n)}|$ satisfy

$$|x_{(i)}| \leq Ri^{-1/p} \quad \text{for all } i \in \{1, 2, \dots, n\}, \quad (1)$$

see e.g., [62]. Vectors in weak- ℓ_p balls may be viewed as approximately sparse; indeed, it is well-known (and easy to show, using standard results for bounding sums by integrals) that for a vector $\mathbf{x} \in w\ell_p(R)$, the ℓ_q error associated with approximating \mathbf{x} by its best k -term approximation obtained by retaining its k largest entries in amplitude (denoted here by $\mathbf{x}^{(k)}$) satisfies

$$\|\mathbf{x} - \mathbf{x}^{(k)}\|_q \triangleq \left(\sum_{i=1}^n |x_i - x_i^{(k)}|^q \right)^{1/q} \leq R C_{p,q} k^{1/q-1/p}, \quad (2)$$

for any $q > p$, where $C_{p,q}$ is given by

$$C_{p,q} = \left(\frac{p}{q-p} \right)^{1/q}. \quad (3)$$

For the special case $q \geq 2p$, we have $C_{p,q} \leq 1$, and so

$$\|\mathbf{x} - \mathbf{x}^{(k)}\|_q \leq R k^{1/q-1/p}. \quad (4)$$

We also recall several information-theoretic preliminaries. When $p(Y)$ and $q(Y)$ denote the pdf (or pmf) of a real-valued random variable Y , the Kullback-Leibler divergence (or KL divergence) of q from p is denoted $D(p\|q)$ and given by

$$D(p\|q) = \mathbb{E}_p \left[\log \frac{p(Y)}{q(Y)} \right]$$

where the logarithm is taken to be the natural log. By definition, $D(p\|q)$ is finite only if the support of p is contained in the support of q . Further, the KL divergence satisfies $D(p\|q) \geq 0$ and $D(p\|q) = 0$ when $p(Y) = q(Y)$. We also use the Hellinger affinity denoted by $A(p, q)$ and given by

$$A(p, q) = \mathbb{E}_p \left[\sqrt{\frac{q(Y)}{p(Y)}} \right] = \mathbb{E}_q \left[\sqrt{\frac{p(Y)}{q(Y)}} \right]$$

Note that $A(p, q) \geq 0$ essentially by definition, and a simple application of the Cauchy-Schwarz inequality gives that $A(p, q) \leq 1$, implying overall that $0 \leq A(p, q) \leq 1$. When p and q are parameterized by elements $X_{i,j}$ and $\tilde{X}_{i,j}$ of matrices \mathbf{X} and $\tilde{\mathbf{X}}$, respectively, so that $p(Y_{i,j}) = p_{X_{i,j}}(Y_{i,j})$ and $q(Y_{i,j}) = q_{\tilde{X}_{i,j}}(Y_{i,j})$, we use the shorthand notation $D(p_{\mathbf{X}}\|q_{\tilde{\mathbf{X}}}) \triangleq \sum_{i,j} D(p_{X_{i,j}}\|q_{\tilde{X}_{i,j}})$ and $A(p_{\mathbf{X}}, q_{\tilde{\mathbf{X}}}) \triangleq \prod_{i,j} A(p_{X_{i,j}}, q_{\tilde{X}_{i,j}})$.

Finally, for a matrix \mathbf{M} we denote by $\|\mathbf{M}\|_0$ its number of nonzero elements, $\|\mathbf{M}\|_1$ the sum of absolute values of its elements, $\|\mathbf{M}\|_{\max}$ the magnitude of its largest element (in absolute value), and $\|\mathbf{M}\|_*$ its nuclear norm (sum of singular values).

II. PROBLEM STATEMENT, APPROACH, AND A GENERAL RECOVERY RESULT

As above, we let $\mathbf{X}^* \in \mathbb{R}^{n_1 \times n_2}$ denote the unknown matrix whose entries we seek to estimate. Our focus is on cases where the unknown matrix \mathbf{X}^* admits a factorization of the form

$$\mathbf{X}^* = \mathbf{D}^* \mathbf{A}^*, \quad (5)$$

where for some integer $r \leq n_2$, $\mathbf{D}^* \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{A}^* \in \mathbb{R}^{r \times n_2}$ are *a priori unknown* factors. For pragmatic reasons, we assume that the elements of \mathbf{D}^* , \mathbf{A}^* , and \mathbf{X}^* are bounded, in the sense that

$$\|\mathbf{D}^*\|_{\max} \leq 1, \quad \|\mathbf{A}^*\|_{\max} \leq A_{\max}, \quad \text{and} \quad \|\mathbf{X}^*\|_{\max} \leq X_{\max}/2 \quad (6)$$

for some constants $0 < A_{\max} \leq (n_1 \vee n_2) = \max\{n_1, n_2\}$ and $X_{\max} \geq 1$. Bounds on the amplitudes of the elements of the matrix to be estimated often arise naturally in practice¹, while our assumption that the entries of the factor matrices be bounded is essentially to fix scaling ambiguities associated with the bilinear model. Our particular focus here will be on cases where (in addition to the entry-wise bounds) the matrix \mathbf{A}^* is *sparse* (having

¹Here, the factor of 1/2 in the bound on $\|\mathbf{X}^*\|_{\max}$ is somewhat arbitrary – any factor in $(0, 1)$ would suffice – and is chosen to facilitate our subsequent analysis.

no more than $k < rn_2$ nonzero elements), or *approximately sparse*, in the sense that for some $p \leq 1$, all of its columns lie in a weak- ℓ_p ball of radius A_{\max} .

Rather than acquire all of the elements of \mathbf{X}^* directly, we assume here that we only observe \mathbf{X}^* at a known *subset* of its locations, obtaining for each observation a noisy or corrupted version of the underlying matrix entry. Here, we will interpret the notion of “noise” somewhat generally in an effort to make our analysis amenable to any of a number of different corruption models; in what follows, we will model each entry-wise observation as a random quantity (either continuous or discrete-valued) whose probability density (or mass) function is parameterized by the true underlying matrix entry. We denote by $\mathcal{S} \subseteq [n_1] \times [n_2]$ the set of locations at which observations are collected, and assume that the sampling locations are random in the sense that for an integer m satisfying $4 \leq m \leq n_1 n_2$ and $\gamma = m(n_1 n_2)^{-1}$, \mathcal{S} is generated according to the independent Bernoulli(γ) model so that each $(i, j) \in [n_1] \times [n_2]$ is included in \mathcal{S} independently with probability γ . Then, given \mathcal{S} , we model the collection of $|\mathcal{S}|$ measurements of \mathbf{X}^* in terms of a collection $\{Y_{i,j}\}_{(i,j) \in \mathcal{S}} \triangleq \mathbf{Y}_{\mathcal{S}}$ of conditionally (on \mathcal{S}) independent random quantities. Formally, we write the joint pdf (or pmf) of the observations as

$$p_{\mathbf{X}_{\mathcal{S}}^*}(\mathbf{Y}_{\mathcal{S}}) \triangleq \prod_{(i,j) \in \mathcal{S}} p_{X_{i,j}^*}(Y_{i,j}), \quad (7)$$

where $p_{X_{i,j}^*}(Y_{i,j})$ denotes the corresponding scalar pdf (or pmf), and we use the shorthand $\mathbf{X}_{\mathcal{S}}^*$ to denote the collection of elements of \mathbf{X}^* indexed by $(i, j) \in \mathcal{S}$. In terms of this model, our task may be described concisely as follows: given \mathcal{S} and corresponding noisy observations $\mathbf{Y}_{\mathcal{S}}$ of \mathbf{X}^* distributed according to (7), our goal is to estimate \mathbf{X}^* under the assumption that it admits a sparse factor model decomposition.

Our approach will be to estimate \mathbf{X}^* via sparsity-penalized maximum likelihood methods; we consider estimates of the form

$$\hat{\mathbf{X}} = \arg \min_{\mathbf{X} = \mathbf{D}\mathbf{A} \in \mathcal{X}} \{-\log p_{\mathbf{X}_{\mathcal{S}}^*}(\mathbf{Y}_{\mathcal{S}}) + \lambda \cdot \|\mathbf{A}\|_0\}, \quad (8)$$

where $\lambda > 0$ is a user-specified regularization parameter, $\mathbf{X}_{\mathcal{S}}$ is shorthand for the collection $\{X_{i,j}\}_{(i,j) \in \mathcal{S}}$ of entries of \mathbf{X} indexed by \mathcal{S} , and \mathcal{X} is an appropriately constructed class of candidate estimates. To facilitate our analysis here, we take \mathcal{X} to be a countable class of estimates constructed as follows: first, for a specified $\beta \geq 1$, we set $L_{\text{lev}} = 2^{\lceil \log_2(n_1 \vee n_2)^\beta \rceil}$ and construct \mathcal{D} to be the set of all matrices $\mathbf{D} \in \mathbb{R}^{n_1 \times r}$ whose elements are discretized to one of L_{lev} uniformly-spaced levels in the range $[-1, 1]$ and \mathcal{A} to be the set of all matrices $\mathbf{A} \in \mathbb{R}^{r \times n_2}$ whose elements either take the value zero, or are discretized to one of L_{lev} uniformly-spaced levels in the range $[-A_{\max}, A_{\max}]$. Then, we let

$$\mathcal{X}' \triangleq \{\mathbf{X} = \mathbf{D}\mathbf{A} : \mathbf{D} \in \mathcal{D}, \mathbf{A} \in \mathcal{A}, \|\mathbf{X}\|_{\max} \leq X_{\max}\}, \quad (9)$$

and take \mathcal{X} to be any subset of \mathcal{X}' . This general formulation will allow us to easily and directly handle additional constraints (e.g., non-negativity constraints on the elements of \mathbf{X} , as arise in our treatment of the Poisson-distributed observation model), within the same unified analytical framework.

Our first main result establishes error bounds for sparse factor model matrix completion problems under general noise or corruption models, where the corruption is described by any generic likelihood model. We state the result

here as a theorem; its proof appears in Appendix A and utilizes a key lemma that extends a main result of [29] to “missing data” scenarios inherent in completion tasks.

Theorem II.1. *Let the sample set \mathcal{S} be drawn from the independent Bernoulli model with $\gamma = m(n_1 n_2)^{-1}$ as described above, and let $\mathbf{Y}_{\mathcal{S}}$ be described by (7). If $C_{\mathcal{D}}$ is any constant satisfying*

$$C_{\mathcal{D}} \geq \max_{\mathbf{X} \in \mathcal{X}} \max_{i,j} D(p_{X_{i,j}^*} \| p_{X_{i,j}}), \quad (10)$$

where \mathcal{X} is as above for some $\beta \geq 1$, then for any

$$\lambda \geq 2 \cdot (\beta + 2) \cdot \left(1 + \frac{2C_{\mathcal{D}}}{3}\right) \cdot \log(n_1 \vee n_2), \quad (11)$$

the complexity penalized maximum likelihood estimator (8) satisfies the (normalized, per-element) error bound

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} [-2 \log A(p_{\hat{\mathbf{X}}}, p_{\mathbf{X}^*})]}{n_1 n_2} &\leq \frac{8C_{\mathcal{D}} \log m}{m} \\ &+ 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{D(p_{\mathbf{X}^*} \| p_{\mathbf{X}})}{n_1 n_2} + \left(\lambda + \frac{4C_{\mathcal{D}}(\beta + 2) \log(n_1 \vee n_2)}{3} \right) \left(\frac{n_1 p + \|\mathbf{A}\|_0}{m} \right) \right\}. \end{aligned} \quad (12)$$

In the next section we consider several specific instances of this result, but we first note a few salient points about this result in its general form. First, as alluded above, our result is not specific to any one observation model; thus, our general result will allow us to analyze the error performance of sparse factor matrix completion methods under a variety of different noise or corruption models. Specialization to a given noise model requires us to only compute (or appropriately bound) the KL divergences and negative log Hellinger affinities of the corresponding probability densities or probability mass functions. Second, our error bound is a kind of *oracle* bound, in that it is specified in terms of a minimum over $\mathbf{X} \in \mathcal{X}$. In practice, we may evaluate this oracle term for *any* $\mathbf{X} \in \mathcal{X}$ and still obtain a valid upper bound (since our guarantee is in terms of the minimum). In our analyses that follows we will impose assumptions on β and \mathbf{X}^* that ensure \mathbf{X}^* be sufficiently “close” to some element \mathbf{X} of \mathcal{X} . This will enable us to obtain non-trivial bounds on the first term in the oracle expression, and to subsequently quantify the corresponding normalized, per-element error (as described in terms of the corresponding negative log Hellinger affinity) by judiciously “balancing” the terms in the oracle expression. This approach will be illustrated in the following section.

Finally, it is worth noting that the estimation strategies prescribed by our analysis are not computationally tractable. Indeed, as written, formation of our estimators would require solving a combinatorial optimization, because of the ℓ_0 penalty, as well as the optimization over the discrete set \mathcal{X} . However, it is worth noting that inference in the bilinear models we consider here is fundamentally challenging on account of the fact that these inference problems cannot directly be cast as (jointly) convex optimizations in the matrix factors. In that sense, our results here may be interpreted as quantifying the performance of one (benchmark) estimation approach for sparse factor matrix completion under various corruption models. (We discuss several extensions, including potential avenues for convexification, in Section V.)

III. IMPLICATIONS FOR SPECIFIC NOISE MODELS

In this section we consider the implications of Theorem II.1 in four unique scenarios, characterized by additive Gaussian noise, additive heavier-tailed (Laplace) noise, Poisson-distributed observations, and quantized (one-bit) observations. In each case, our aim is to identify the scaling behavior of the estimation error as a function of the key problem parameters. To that end, we consider for each case the fixed choice

$$\beta = \max \left\{ 1, 1 + \frac{\log(8rA_{\max}/X_{\max})}{\log(n_1 \vee n_2)} \right\} \quad (13)$$

for describing the number of discretization levels in the elements of each of the matrix factors. Then, for each scenario (characterized by its own unique likelihood model) we consider a specific choice of \mathcal{X} , and an estimate obtained according to (8) with the specific choice

$$\lambda = 2 \left(1 + \frac{2C_D}{3} \right) (\beta + 2) \cdot \log(n_1 \vee n_2), \quad (14)$$

(where C_D depends on the particular likelihood model), and simplify the resulting oracle bounds for both sparse and approximately sparse factors. In what follows, we will make use of the fact that our assumption $X_{\max} \geq 1$ implies $\beta = \mathcal{O}(\log(r \vee A_{\max})/\log(n_1 \vee n_2))$, and so $(\beta + 2) \log(n_1 \vee n_2) = \mathcal{O}(\log(n_1 \vee n_2))$, on account of the fact that $r < n_2$ and $A_{\max} < (n_1 \vee n_2)$ by assumption.

A. Additive Gaussian Noise

We first examine the implications of Theorem II.1 in a setting where observations are corrupted by independent additive zero-mean Gaussian noise with known variance. In this case, the observations \mathbf{Y}_S are distributed according to a multivariate Gaussian density of dimension $|\mathcal{S}|$ whose mean corresponds to the collection of matrix parameters at the sample locations, and with covariance matrix $\sigma^2 \mathbf{I}_{|\mathcal{S}|}$, where $\mathbf{I}_{|\mathcal{S}|}$ is the identity matrix of dimension $|\mathcal{S}|$, so

$$p_{\mathbf{X}_S^*}(\mathbf{Y}_S) = \frac{1}{(2\pi\sigma^2)^{|\mathcal{S}|/2}} \exp \left(-\frac{1}{2\sigma^2} \|\mathbf{Y}_S - \mathbf{X}_S^*\|_F^2 \right), \quad (15)$$

where we have used the representative shorthand notation $\|\mathbf{Y}_S - \mathbf{X}_S^*\|_F^2 \triangleq \sum_{(i,j) \in \mathcal{S}} (Y_{i,j} - X_{i,j}^*)^2$. In this setting we have the following result; its proof appears in Appendix B.

Corollary III.1 (Sparse Factor Matrix Completion with Gaussian Noise). *Let β be as in (13), let λ be as in (14) with $C_D = 2X_{\max}^2/\sigma^2$, and let $\mathcal{X} = \mathcal{X}'$. The estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies*

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left((\sigma^2 + X_{\max}^2) \left(\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \right) \log(n_1 \vee n_2) \right) \quad (16)$$

when \mathbf{A}^* is exactly sparse, having $\|\mathbf{A}^*\|_0$ nonzero elements. If, instead, the columns of \mathbf{A}^* are approximately sparse in the sense that for some $p \leq 1$ each belongs to a weak- ℓ_p ball of radius A_{\max} , then the estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(A_{\max}^2 \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} + (\sigma^2 + X_{\max}^2) \left(\frac{n_1 r}{m} + \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} \right) \log(n_1 \vee n_2) \right), \quad (17)$$

where $\alpha = 1/p - 1/2$.

Remark III.1. We utilize Big-Oh notation to suppress leading constants for clarity of exposition, and to illustrate the dependence of the bounds on the key problem parameters. Our proofs for each of the specific results provides the explicit constants.

A few comments are in order regarding these error guarantees. First, we note that our analysis provides some useful (and intuitive) understanding of how the estimation error decreases as a function of the number of measurements obtained, as well as the dimension and sparsity parameters associated with the matrix to be estimated. Consider, for instance, the case when \mathbf{A}^* is sparse and where $\log m < n_1 r + \|\mathbf{A}^*\|_0$ (which should often be the case, since $\log(m) \leq \log(n_1 n_2)$). In this setting, our error bound shows that the dependence of the estimation error on the dimension (n_1, n_2, r) and sparsity $(\|\mathbf{A}^*\|_0)$ parameters, as well as the (nominal) number of measurements m is

$$\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \log(n_1 \vee n_2). \quad (18)$$

We may interpret the quantity $n_1 r + \|\mathbf{A}^*\|_0$ as the number of *degrees of freedom* in the matrix \mathbf{X}^* to be estimated, and in this sense we see that the error rate of the penalized maximum likelihood estimator exhibits characteristics of the well-known parametric rate (modulo the logarithmic factor). Along related lines, note that in the case where columns of \mathbf{A}^* are approximately sparse, the $(n_2/m)^{\frac{2\alpha}{2\alpha+1}}$ term that arises in the error rate is reminiscent of error rates that arise when estimating approximately sparse vectors in noisy compressive sensing (e.g., see [63], [64]). Indeed, since $(n_2/m)^{\frac{2\alpha}{2\alpha+1}} \leq n_2 m^{-\frac{2\alpha}{2\alpha+1}}$, we see that the overall matrix estimation error may be interpreted as being comprised of errors associated with approximating the n_2 nearly-sparse columns of \mathbf{A}^* in this noisy setting, each of which would contribute a (normalized) error on the order of $m^{-\frac{2\alpha}{2\alpha+1}}$.

Next, our error bounds provide some guidelines for identifying in which scenarios accurate estimation may be possible. Consider a *full sampling* scenario where the matrix $\mathbf{X}^* = \mathbf{D}^* \mathbf{A}^*$ has a coefficient matrix with no more than k nonzero elements per column (thus, $\|\mathbf{A}^*\|_0 \leq n_2 k$). Now, to ensure that

$$\frac{n_1 r + \|\mathbf{A}^*\|_0}{n_1 n_2} \log(n_1 \vee n_2) \preceq 1, \quad (19)$$

(where the notation \preceq suppresses leading constants) it is sufficient to have $n_1 n_2 \succeq n_1 r \log(n_1 \vee n_2)$ and $n_1 n_2 \succeq \|\mathbf{A}^*\|_0 \log(n_1 \vee n_2)$. Simplifying a bit, we see that the first sufficient condition is satisfied when $n_2 \succeq 2r \log(n_1 \vee n_2)$, or when the number of columns of the matrix \mathbf{X}^* exceeds (by a multiplicative constant and logarithmic factor) the number of columns of its dictionary factor \mathbf{D}^* . Further, the second sufficient condition holds when $n_1 \succeq k \log(n_1 \vee n_2)$, or when the number of measurements of each column exceeds (again, by a multiplicative constant and logarithmic factor) the number of nonzeros in the sparse representation of each column. This latter condition is reminiscent of the sufficient conditions arising in sparse inference problems inherent in noisy compressive sensing (see, e.g., [63], [65]). Analogous insights may be derived from our results for the subsampled regimes that comprise our main focus here (i.e., when $m < n_1 n_2$).

Further, we comment on the presence of the X_{\max}^2 term present in the error bounds for both the sparse and nearly-sparse settings. Readers familiar with the literature on matrix completion under low rank assumptions will recall that various forms of ‘‘incoherence’’ assumptions have been utilized to date as a means to ensure identifiability

under various sampling models, and that the form of the resulting error bounds depend on the particular type of assumption employed. For example, the authors of [19] consider an additive noise model similar to here but employ incoherence assumptions that essentially enforce that the row and column spaces of the matrix to be estimated not be overly aligned with the canonical bases (reminiscent of initial works on noise-free matrix completion [2]) and obtain estimation error bounds that do not depend on max-norm bounds of the matrix to be estimated (though the necessary conditions on the number of samples obtained do depend on the incoherence parameters). The work [21] also examines matrix completion problems with additive noises but utilizes a different form of incoherence assumption formulated in terms of the “spikiness” of the matrix to be estimated (and quantified in terms of the ratio between the max norm and Frobenius norm). There, the estimation approach entails optimization over a set of candidates that each satisfy a “spikiness” constraint, and the bounds so obtained scale in proportion to the max-norm of the matrix to be estimated (similar to here). Incoherence assumptions manifesting as an assumed bound on the largest matrix element also arise in [23], [24].

One direct point of comparison to our result here is [20], which considers matrix completion problems characterized by entry-wise observations obtained at locations chosen uniformly at random (with replacement), each of which may be modeled as corrupted by independent additive noise, and estimates obtained by nuclear norm penalized estimators. Casting the results of that work (specifically, [20, Corollary 2]) to the setting we consider here, we observe that those results imply rank- r matrices may be accurately estimated in the sense that

$$\frac{\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2}{n_1 n_2} \leq c (\sigma \vee X_{\max})^2 \left(\frac{(n_1 \vee n_2) r}{m} \right) \log(n_1 + n_2) \quad (20)$$

$$\leq c' (\sigma^2 + X_{\max}^2) \left(\frac{(n_1 + n_2)r}{m} \right) \log(n_1 \vee n_2) \quad (21)$$

with high probability, where c, c' are positive constants. Comparing this last result with our result (16), we see that our guarantees exhibit the same effective scaling with the max-norm bound X_{\max} , but can have an (perhaps significantly) improved error performance in the case where $\|\mathbf{A}^*\|_0 \ll n_2 r$ – precisely what we sought to identify by considering sparse factor models in our analyses. The two bounds roughly coincide in the case where \mathbf{A}^* is not sparse, in which case we may take $\|\mathbf{A}^*\|_0 = n_2 r$ in our error bounds.

B. Additive Laplace Noise

As another example, suppose that the observations \mathbf{Y}_S are corrupted by independent additive heavier-tailed noises, each of which we model using a Laplace distribution with parameter $\tau > 0$. In this scenario, we have that

$$p_{\mathbf{X}_S^*}(\mathbf{Y}_S) = \left(\frac{\tau}{2} \right)^{|S|} \exp(-\tau \|\mathbf{Y}_S - \mathbf{X}_S^*\|_1), \quad (22)$$

where we use $\|\mathbf{Y}_S - \mathbf{X}_S^*\|_1 \triangleq \sum_{(i,j) \in S} |Y_{i,j} - X_{i,j}^*|$ for shorthand. The following result holds; its proof appears in Appendix C.

Corollary III.2 (Sparse Factor Matrix Completion with Laplace Noise). *Let β be as in (13), let λ be as in (14)*

with $C_D = 2\tau X_{\max}$, and let $\mathcal{X} = \mathcal{X}'$. The estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(\left(\frac{1}{\tau} + X_{\max} \right)^2 \tau X_{\max} \left(\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \right) \log(n_1 \vee n_2) \right), \quad (23)$$

when \mathbf{A}^* is exactly sparse, having $\|\mathbf{A}^*\|_0$ nonzero elements. If, instead, for some $p \leq 1/2$ the columns of \mathbf{A}^* belong to a weak- ℓ_p ball of radius A_{\max} , then the estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\begin{aligned} & \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \\ & \mathcal{O} \left(\left(\frac{1}{\tau} + X_{\max} \right)^2 \tau A_{\max} \left(\frac{n_2}{m} \right)^{\frac{\alpha'}{\alpha'+1}} + \left(\frac{1}{\tau} + X_{\max} \right)^2 \tau X_{\max} \left(\frac{n_1 r}{m} + \left(\frac{n_2}{m} \right)^{\frac{\alpha'}{\alpha'+1}} \right) \log(n_1 \vee n_2) \right), \end{aligned} \quad (24)$$

where $\alpha' = 1/p - 1$.

A few comments are in order regarding these results. First, recall that our main theorem naturally provides error guarantees in terms of KL divergences and negative log Hellinger affinities. However, here we state our bounds in terms of the average per element squared error, and draw comparisons with the previous case (and, perhaps, to make the results more amenable to interpretation). To achieve this we employed a series of bounds – quadratic (in the parameter difference) lower bounds on the negative log Hellinger affinities, and upper bounds on the KL divergences that are proportional to the absolute deviations between the parameters (see the proof for details). It is interesting to note that this bounding approach, the error performance that we obtain for the case where \mathbf{A}^* is sparse again exhibits characteristics of the parametric rate, while we do obtain different error behavior as compared to the Gaussian noise case for the case where \mathbf{A}^* is nearly sparse. As one specific example, consider the case where the coefficients of \mathbf{A}^* exhibit the ordered decay with $p = 1/3$ (a parameter that is valid for both Corollaries III.1 and III.2). The error rate for the Gaussian noise setting in this case contains a term that decays on the order of $(m/n_2)^{-5/6}$, while here when the noise is heavy-tailed, the analogous term decays at a slower rate, like $(m/n_2)^{-2/3}$. Overall, casting the error bounds all in terms of the same loss metric (here, ℓ_2) makes our results directly amenable to such comparisons.

Along related lines, it is interesting to note that the estimation error bound here is slightly “inflated” relative to the Gaussian-noise counterparts (albeit with constants suppressed in each case). Recall that the variance of a Laplace(τ) random variable is $2/\tau^2$; thus, the leading term $(1/\tau + X_{\max})^2 = \mathcal{O}(2/\tau^2 + X_{\max}^2)$ here is somewhat analogous to the $(\sigma^2 + X_{\max}^2)$ factor arising in the Gaussian-noise error bounds. In this sense, we see that the factor of τX_{\max} in the Laplace-noise case appears to be “extra.” Here, this factor is effectively introduced by our attempt to cast the “natural” error guarantees arising from our analysis (which manifest in terms of negative log Hellinger affinities) into more interpretable squared-error bounds.

C. Poisson-distributed Observations

We now consider an example motivated by applications where the observed data may correspond to discrete “counts” (e.g., in imaging applications). Suppose that the entries of the matrix \mathbf{X}^* are all non-negative and that

our observation at each location $(i, j) \in \mathcal{S}$ is a Poisson random variable with rate $X_{i,j}^*$. In this setting, our matrix completion problem amounts to a kind of Poisson denoising task; we have that $\mathbf{Y}_{\mathcal{S}} \in \mathbb{N}^{|\mathcal{S}|}$ and

$$p_{\mathbf{X}_{\mathcal{S}}^*}(\mathbf{Y}_{\mathcal{S}}) = \prod_{(i,j) \in \mathcal{S}} \frac{(X_{i,j}^*)^{Y_{i,j}} e^{-X_{i,j}^*}}{(Y_{i,j})!}. \quad (25)$$

In this case, we employ Theorem II.1 to obtain the following result; a sketch of the proof is provided in Appendix D.

Corollary III.3 (Sparse Factor Matrix Completion with Poisson Noise). *Suppose that the elements of the matrix \mathbf{X}^* to be estimated satisfy $\min_{i,j} |X_{i,j}^*| \geq X_{\min}$ for some constant $X_{\min} > 0$. Let β be as in (13), let λ be as in (14) with $C_D = 4X_{\max}^2/X_{\min}$, and let \mathcal{X}' be the subset of \mathcal{X}' comprised of all candidate estimates having non-negative entries. The estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies*

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(\left(X_{\max} + \frac{X_{\max}}{X_{\min}} \cdot X_{\max}^2 \right) \left(\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \right) \log(n_1 \vee n_2) \right), \quad (26)$$

when \mathbf{A}^* is exactly sparse, having $\|\mathbf{A}^*\|_0$ nonzero elements. If, instead, for some $p \leq 1$ the columns of \mathbf{A}^* belong to a weak- ℓ_p ball of radius A_{\max} , then the estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(A_{\max}^2 \left(\frac{X_{\max}}{X_{\min}} \right) \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} + \left(X_{\max} + \frac{X_{\max}}{X_{\min}} \cdot X_{\max}^2 \right) \left(\frac{n_1 r}{m} + \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} \right) \log(n_1 \vee n_2) \right), \quad (27)$$

where $\alpha = 1/p - 1/2$.

As in the previous case, our analysis approach here entails bounding (appropriately) the KL divergence and negative log Hellinger affinities each in terms of squared Frobenius norms; similar bounding methods were employed in [33], which analyzed a compressive sensing sparse vector reconstruction task under a Poisson observation model. Overall, we observe an interesting behavior relative to the preceding two cases. Recall that the bounds for the setting where \mathbf{A}^* is exactly sparse, for each of the previous two cases, exhibited a leading factor that was essentially the sum of the variance and X_{\max}^2 . In each of those cases, the per-observation noise variances were independent of the underlying matrix entry; in contrast, Poisson-distributed observations exhibit a variance equal to the underlying rate parameter. So, in this sense, we might interpret the $(X_{\max} + (X_{\max}/X_{\min})X_{\max}^2)$ term as roughly corresponding to a “worst-case” variance plus X_{\max}^2 . Indeed, when X_{\max}/X_{\min} is upper-bounded by a (small) constant; then, this leading factor is $\mathcal{O}(X_{\max} + X_{\max}^2)$, somewhat analogously to the leading factor arising in the Laplace-noise and Gaussian-noise bounds. More generally, that the error behavior in Poisson denoising tasks be similar to the Gaussian case is perhaps not surprising. Indeed, a widely used approach in Poisson inference tasks is to employ a *variance stabilizing transformations*, such as the Anscombe transform [66], so that the transformed data distribution be “approximately” Gaussian.

It is worth commenting a bit further on our *minimum* rate assumption on the elements of \mathbf{X}^* , that each be no smaller than some constant $X_{\min} > 0$. Similar assumptions were employed in [33], as well as other works that examine Poisson denoising tasks using the penalized ML analysis framework (e.g., [31]). Here, this X_{\min}

parameter shows up in the denominator of a leading factor in our bound, suggesting that the bounds become more loose as the estimation task transitions closer to scenarios characterized by “low-rate” Poisson sources. Indeed, closer inspection of our error bounds as stated above shows that they diverge (tend to $+\infty$) as X_{\min} tends to zero, suggesting that the estimation task becomes more difficult in “low rate” settings. Contrast this with classical analyses of scalar Poisson rate estimation problems show that the Cramer-Rao lower bound associated with estimating the rate parameter θ of a Poisson random variable using n iid observations is θ/n , and this error is achievable with the sample average estimator. This suggests that the estimation problem actually becomes *easier* as the rate decreases, at least for the scalar estimation problem. On this note, we briefly mention several recent works that rectify this apparent discrepancy, for matrix estimation tasks as here [34], and for sparse vector estimation from Poisson-distributed compressive observations [67]. The ideas underlying those works might have applicability for the completion problems we consider here, but this extension may require imposing different (or even much stronger) forms of incoherence assumptions on the matrix to be estimated as compared to the bounded-entry condition we adopt here. We do not pursue those extensions here, opting instead to state our result as a direct instantiation of our main result in Theorem II.1.

D. Quantized (One-bit) Observation Models

We may also utilize our main result to assess the estimation performance in scenarios where entry-wise observations of the matrix are quantized to few bits, or even a *single bit*, each. Such quantized observations are natural in collaborative filtering applications such as the aforementioned *Netflix* problem, where users’ ratings are quantized to fixed levels. One may also envision applications in distributed estimation tasks where one seeks to estimate some underlying matrix from highly-quantized observations of a subset of its entries; here, the quantization could serve as a mechanism for enforcing global communication rate constraints (e.g., when the data is transmitted to a centralized location for inference). Our general framework would facilitate analysis of observations quantized to any of a number of levels; here, for concreteness, we consider a one-bit observation model.

Formally, given a sampling set \mathcal{S} we suppose that our observations are conditionally (on \mathcal{S}) independent random variables described by

$$Y_{i,j} = \mathbf{1}_{\{Z_{i,j} \geq 0\}}, \quad (i,j) \in \mathcal{S}, \quad (28)$$

where

$$Z_{i,j} = X_{i,j}^* - W_{i,j}, \quad (29)$$

the $\{W_{i,j}\}_{i \in [m], j \in [n]}$ are some iid continuous zero-mean real scalar “noises” having probability density function and cumulative distribution function $f(w)$ and $F(w)$, respectively, for $w \in \mathbb{R}$, and $\mathbf{1}_{\{\mathcal{E}\}}$ denotes the indicator of the event \mathcal{E} that takes the value 1 when \mathcal{E} occurs and zero otherwise. Note that in this model, we assume that the individual noise realizations $\{W_{i,j}\}_{(i,j) \in \mathcal{S}}$ are unknown (but we assume that the noise *distribution* is known). Stated another way, we may interpret the observations modeled as above essentially as quantized noisy versions of the true matrix parameters (the minus sign on the $W_{i,j}$ ’s is merely a modeling convenience here, and is intended

to simplify the exposition). Under this model, it is easy to see that each $Y_{i,j}$ is a Bernoulli random variable whose parameter is related to the true parameter through the cumulative distribution function. Specifically, note that for any fixed $(i, j) \in \mathcal{S}$, we have that $\Pr(Y_{i,j} = 1) = \Pr(W_{i,j} \leq X_{i,j}^*) = F(X_{i,j}^*)$. Thus, in this scenario, we have that $\mathbf{Y}_S \in \{0, 1\}^{|\mathcal{S}|}$ and

$$p_{\mathbf{X}_S^*}(\mathbf{Y}_S) = \prod_{(i,j) \in \mathcal{S}} [F(X_{i,j}^*)]^{Y_{i,j}} [1 - F(X_{i,j}^*)]^{1-Y_{i,j}} \quad (30)$$

We will also assume here that X_{\max} and $F(\cdot)$ are such that $F(X_{\max}) < 1$ and $F(-X_{\max}) > 0$; it follows that the true Bernoulli parameters (as well as the Bernoulli parameters associated with candidate estimates $\mathbf{X} \in \mathcal{X}$) are bounded away from 0 and 1; these assumptions will allow us to avoid some pathological scenarios in our analysis.

Given the above model and assumptions, we may establish the following result; the proof is provided in Appendix E.

Corollary III.4 (Sparse Factor Matrix Completion from One-bit Observations). *Let β be as in (13), let $\mathcal{X} = \mathcal{X}'$, and let $p_{\mathbf{X}_S^*}$ be of the form in (30) with $F(X_{\max}) < 1$ and $F(-X_{\max}) > 0$. Define*

$$c_{F, X_{\max}} \triangleq \left(\sup_{|t| \leq X_{\max}} \frac{1}{F(t)(1-F(t))} \right) \cdot \left(\sup_{|t| \leq X_{\max}} f^2(t) \right), \quad (31)$$

and

$$c'_{F, X_{\max}} \triangleq \inf_{|t| \leq X_{\max}} \frac{f^2(t)}{F(t)(1-F(t))}, \quad (32)$$

and let λ be as in (14) with $C_D = 2c_{F, X_{\max}} X_{\max}^2$. The estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(\left(\frac{c_{F, X_{\max}}}{c'_{F, X_{\max}}} \right) \left(\frac{1}{c_{F, X_{\max}}} + X_{\max}^2 \right) \left(\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \right) \log(n_1 \vee n_2) \right), \quad (33)$$

when \mathbf{A}^* is exactly sparse, having $\|\mathbf{A}^*\|_0$ nonzero elements. If, instead, for any $p \leq 1$ the columns of \mathbf{A}^* belong to a weak- ℓ_p ball of radius A_{\max} , then the estimate $\widehat{\mathbf{X}}$ obtained via (8) satisfies

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} = \mathcal{O} \left(\left(\frac{c_{F, X_{\max}}}{c'_{F, X_{\max}}} \right) A_{\max}^2 \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} + \left(\frac{c_{F, X_{\max}}}{c'_{F, X_{\max}}} \right) \left(\frac{1}{c_{F, X_{\max}}} + X_{\max}^2 \right) \left(\frac{n_1 r}{m} + \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}} \right) \log(n_1 \vee n_2) \right). \quad (34)$$

where $\alpha = 1/p - 1/2$.

It is interesting to compare the results of (33) and (34) with the analogous results (16) and (17). Specifically, we see that our estimation error guarantees for each case exhibit the same fundamental dependence on the dimension, sparsity, and (nominal) number of measurements, with the primary difference overall arising in the form of the leading factors (that in the one-bit case depend on the specific distribution of the $W_{i,j}$ terms). That the estimation errors for rate-constrained tasks approximately mimic that of their Gaussian-corrupted counterparts was observed in earlier works on rate-constrained parameter estimation (see, e.g., [68], [69]), and more recently in [23], which considered low-rank matrix completion from one-bit measurements, using a generative model analogous to the model we consider here.

It is also worth noting that the cdf $F(\cdot)$ that we specify here could be replaced by any of a number of commonly-used *link functions*. For example, choosing $F(x) = \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-t^2/2} dt$ to be the cdf of a standard Gaussian random variable gives rise to the well-known probit model, while taking $F(x)$ to be the logistic function, $F(x) = \frac{1}{1+e^{-x}}$, leads to the logit regression model. In this sense, our results are related to classical methods on inference in generalized linear models (see, e.g., [70]); a key distinction here is that we assume both of the factors in the bilinear form to be unknown.

Finally, we briefly compare our results with the results of [23] for low-rank matrix completion from one-bit observations. In that work, the authors consider maximum-likelihood optimizations over a (convex) set of max-norm and nuclear-norm constrained matrices, and show that the estimates so-obtained satisfy

$$\frac{\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2}{n_1 n_2} = \mathcal{O} \left(C_{F, X_{\max}} X_{\max} \sqrt{\frac{(n_1 + n_2)r}{m}} \right) \quad (35)$$

with high probability, where $C_{F, X_{\max}}$ is a parameter that depends on the max-norm constraint cdf $F(\cdot)$ and pdf $f(\cdot)$, somewhat analogously to the leading factor of $(c_{F, X_{\max}}/c'_{F, X_{\max}})$ in our bounds. It is interesting to note a main qualitative difference between that result and ours. For concreteness, let us consider the case where \mathbf{A}^* is not sparse, so that we may set $\|\mathbf{A}^*\|_0 = n_2 r$ in (33). In that case, it is easy to see that the overall estimation error behavior predicted by our bound (33) scales in proportion to ratio between the number of degrees of freedom $((n_1 + n_2)r)$ and the nominal number of measurements m , while the bound in [23] scales according to the square root of that ratio. The authors of [23] proceed to show that the estimation error rate they obtain is minimax optimal over their set of candidate estimates; on the other hand, our bound appears (at least up to leading factors) to be tighter for this case where \mathbf{X}^* is exactly low-rank (e.g., setting $\|\mathbf{A}^*\|_0 = n_2 r$ in our bound) and $m \geq c(n_1 + n_2)r$ for a constant $c > 1$. That said, our approach also enjoys the benefit of having the rank or an upper bound for it be known (and being combinatorial in nature!), while the procedure in [23] assumes only a bound on the nuclear norm of the unknown matrix. Whether our bounds here exhibit minimax-optimal estimation error rates for matrix completion under sparse factor models and for the several various likelihood models we consider here is still an open question since (to our knowledge) lower bounds for these problems have not yet been established (but are a topic of our ongoing efforts).

IV. EXPERIMENTAL EVALUATION

In this section we provide experimental evidence to validate the error rates established by our theoretical results. Recall (as noted above) that the original problem (8) we aim to solve has multiple sources of non-convexity, including the bilinear matrix factor model (i.e., $\mathbf{X} = \mathbf{D}\mathbf{A}$), the presence of the ℓ_0 penalty, and the discretized sets \mathcal{D} and \mathcal{A} . In what follows, we undertake a slight relaxation of (8) replacing the sets \mathcal{D} and \mathcal{A} by their convex hulls (and with slight overloading of notation in what follows, we refer to these new sets also as \mathcal{D} and \mathcal{A}). With this relaxation, the set \mathcal{X} becomes a set of all matrices $\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}$ with bounded entries. Note that with these simplifying relaxations, the sets \mathcal{X} , \mathcal{D} and \mathcal{A} are convex.

Now, for each likelihood model, our aim is to solve a constrained maximum likelihood problem of the form

$$\begin{aligned} \min_{\mathbf{D} \in \mathbb{R}^{n_1 \times r}, \mathbf{A} \in \mathbb{R}^{r \times n_2}} \quad & \sum_{i,j} s_{i,j} \ell(Y_{i,j}, X_{i,j}) + I_{\mathcal{X}}(\mathbf{X}) + I_{\mathcal{D}}(\mathbf{D}) + I_{\mathcal{A}}(\mathbf{A}) + \lambda \|\mathbf{A}\|_0 \\ \text{s.t.} \quad & \mathbf{X} = \mathbf{D}\mathbf{A}. \end{aligned} \quad (36)$$

where $\ell(Y_{i,j}, X_{i,j}) = -\log(p_{X_{i,j}}(Y_{i,j}))$ is the negative log-likelihood for the corresponding noise model, $s_{i,j}$ is a selector taking the value 1 when $(i, j) \in \mathcal{S}$ and 0 otherwise, $\lambda \geq 0$ is a regularization parameter, and each of $I_{\mathcal{X}}(\cdot)$, $I_{\mathcal{D}}(\cdot)$, and $I_{\mathcal{A}}(\cdot)$ are the indicator functions of the sets \mathcal{X} , \mathcal{D} and \mathcal{A} respectively². Here, we have that each of the indicator functions is separable in the individual entries of its argument, e.g. $I_{\mathcal{X}}(\mathbf{X}) = \sum_{i,j} I_{\mathcal{X}_{i,j}}(X_{i,j})$, and similarly for the indicator functions of \mathcal{D} and \mathcal{A} .

We propose a solution approach based on the Alternating Direction Method of Multipliers (ADMM) [28]. First we write the augmented Lagrangian of (36) as

$$\mathcal{L}(\mathbf{D}, \mathbf{A}, \mathbf{X}, \mathbf{\Lambda}) = \sum_{i,j} s_{i,j} \ell(Y_{i,j}, X_{i,j}) + I_{\mathcal{X}}(\mathbf{X}) + I_{\mathcal{D}}(\mathbf{D}) + I_{\mathcal{A}}(\mathbf{A}) + \lambda \|\mathbf{A}\|_0 + \text{tr}(\mathbf{\Lambda}(\mathbf{X} - \mathbf{D}\mathbf{A})) + \frac{\rho}{2} \|\mathbf{X} - \mathbf{D}\mathbf{A}\|_{\mathbb{F}}^2, \quad (37)$$

where $\mathbf{\Lambda}$ is a matrix of Lagrange multiplier parameters and $\rho > 0$ is a parameter. Then, starting with some feasible $\mathbf{A}^{(0)}, \mathbf{D}^{(0)}, \mathbf{\Lambda}^{(0)}$ we iteratively update \mathbf{X} , \mathbf{A} , \mathbf{D} , and $\mathbf{\Lambda}$ according to

$$(\mathbf{S1} :) \mathbf{X}^{(k+1)} := \arg \min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} \mathcal{L}(\mathbf{D}^{(k)}, \mathbf{A}^{(k)}, \mathbf{X}, \mathbf{\Lambda}^{(k)}) \quad (38)$$

$$(\mathbf{S2} :) \mathbf{A}^{(k+1)} := \arg \min_{\mathbf{A} \in \mathbb{R}^{r \times n_2}} \mathcal{L}(\mathbf{D}^{(k)}, \mathbf{A}, \mathbf{X}^{(k+1)}, \mathbf{\Lambda}^{(k)}) \quad (39)$$

$$(\mathbf{S3} :) \mathbf{D}^{(k+1)} := \arg \min_{\mathbf{D} \in \mathbb{R}^{n_1 \times r}} \mathcal{L}(\mathbf{D}, \mathbf{A}^{(k+1)}, \mathbf{X}^{(k+1)}, \mathbf{\Lambda}^{(k)}) \quad (40)$$

$$(\mathbf{S4} :) \mathbf{\Lambda}^{(k+1)} = \mathbf{\Lambda}^{(k)} + \rho(\mathbf{X}^{(k+1)} - \mathbf{D}^{(k+1)}\mathbf{A}^{(k+1)}), \quad (41)$$

until convergence, which here is quantified in terms of when norms of primal and dual residuals become sufficiently small (along the lines of the criteria described in [28]). Next we describe how to solve each of these steps.

Solving **S1** involves the following optimization problem

$$\min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} \sum_{i,j} s_{i,j} \ell(Y_{i,j}, X_{i,j}) + I_{\mathcal{X}}(\mathbf{X}) + \text{tr}(\mathbf{\Lambda}^{(k)}(\mathbf{X} - \mathbf{D}^{(k)}\mathbf{A}^{(k)})) + \frac{\rho}{2} \|\mathbf{X} - \mathbf{D}^{(k)}\mathbf{A}^{(k)}\|_{\mathbb{F}}^2, \quad (42)$$

which after completing the square and ignoring constant terms is equivalent to

$$\min_{\mathbf{X} \in \mathbb{R}^{n_1 \times n_2}} \sum_{i,j} s_{i,j} \ell(Y_{i,j}, X_{i,j}) + I_{\mathcal{X}}(\mathbf{X}) + \frac{\rho}{2} \left\| \mathbf{X} - \mathbf{D}^{(k)}\mathbf{A}^{(k)} + \frac{\mathbf{\Lambda}^{(k)}}{\rho} \right\|_{\mathbb{F}}^2. \quad (43)$$

Due to the assumed separability of the indicator function, the above problem is separable in each entry $X_{i,j}$ and the entries can be updated in parallel by solving the following scalar convex optimization problem for each entry.

²Recall that the indicator function is defined as a function that takes values 0 or ∞ depending on whether its argument is an element of the set described as the subscript.

	$\ell(y, x)$	$\text{prox}_\ell(z; \rho, y)$
Gaussian	$\frac{(y-x)^2}{2\sigma^2}$	$\frac{y+\sigma^2\rho z}{1+\sigma^2\mu}$
Poisson	$x - y \log(x)$	$\frac{\rho z - 1 + \sqrt{(\rho z - 1)^2 + 4\rho y}}{2\rho}$
Laplace	$\lambda y - x $	$y - \text{Soft}(z - y, \lambda/\rho)$
One-bit	$-y \log(F(x)) - (1 - y) \log(1 - F(x))$	<i>(Newton's method)</i>

TABLE I: Expressions for $\text{prox}_\ell(z; \rho, y) = \arg \min_{x \in \mathbb{R}} \ell(y, x) + \frac{\rho}{2} (x - z)^2$ for different $\ell(y, x)$, corresponding to negative log-likelihoods for the models we examine. Here, for $\lambda > 0$, $\text{soft}(x, \lambda) = \text{sgn}(x) \max\{|x| - \lambda, 0\}$

When $\ell(y, x)$ is a convex function of x , the solution is given by

$$\begin{aligned}
 X_{i,j}^{(k+1)} &= \text{Proj}_{\mathcal{X}_{i,j}} \left[\text{prox}_{s_{i,j}, \ell} \left((\mathbf{D}^{(k)} \mathbf{A}^{(k)})_{i,j} - \frac{(\mathbf{\Lambda}^{(k)})_{i,j}}{\rho}; \rho, Y_{i,j} \right) \right], \\
 &= \begin{cases} \text{Proj}_{\mathcal{X}_{i,j}} \left[\text{prox}_\ell \left((\mathbf{D}^{(k)} \mathbf{A}^{(k)})_{i,j} - \frac{(\mathbf{\Lambda}^{(k)})_{i,j}}{\rho}; \rho, Y_{i,j} \right) \right], & \text{if } s_{i,j} = 1 \\ \text{Proj}_{\mathcal{X}_{i,j}} \left[(\mathbf{D}^{(k)} \mathbf{A}^{(k)})_{i,j} - \frac{(\mathbf{\Lambda}^{(k)})_{i,j}}{\rho} \right], & \text{otherwise} \end{cases} \quad (44)
 \end{aligned}$$

where $\text{prox}_\ell(z; \rho, y) = \arg \min_{x \in \mathbb{R}} \ell(y, x) + \frac{\rho}{2} (x - z)^2$ is the *proximal operator* of the loss function $\ell(y, \cdot)$ and $\text{Proj}_{\mathcal{X}_{i,j}}(x)$ is the projection³ of the scalar x onto the set $\mathcal{X}_{i,j}$. For several of the loss functions we consider, the proximal operator can be computed in closed form; for the one-bit settings we can use Newton's second order method (or gradient descent) to numerically evaluate it as described later. A table of proximal operators for the various losses we consider here is provided in Table I.

Completing the square and ignoring the constant terms the subproblem **S2** is equivalent to

$$\mathbf{A}^{(k+1)} = \arg \min_{\mathbf{A} \in \mathbb{R}^{r \times n_2}} I_{\mathcal{A}}(\mathbf{A}) + \lambda \|\mathbf{A}\|_0 + \frac{\rho}{2} \left\| \mathbf{X}^{(k+1)} - \mathbf{D}^{(k)} \mathbf{A} + \frac{\mathbf{\Lambda}^{(k)}}{\rho} \right\|_F^2. \quad (45)$$

In order to solve this problem we adopt the constrained iterative hard thresholding approach from [71], as outlined in Algorithm 1. Finally, after completing the square and ignoring the constant terms, we see that the subproblem **S3** is equivalent to

$$\mathbf{D}^{(k+1)} = \arg \min_{\mathbf{D} \in \mathbb{R}^{n_1 \times r}} I_{\mathcal{D}}(\mathbf{D}) + \frac{\rho}{2} \left\| \mathbf{X}^{(k+1)} - \mathbf{D} \mathbf{A}^{(k+1)} + \frac{\mathbf{\Lambda}^{(k)}}{\rho} \right\|_F^2, \quad (46)$$

which we solve here by projected Newton gradient descent algorithm, described in Algorithm 2. Our overall algorithmic approach is summarized in Algorithm 3.

A. Experiments

We perform experimental validation of our theoretical results on synthetic data for two different scenarios, corresponding to when the columns of the matrix \mathbf{A}^* are k -sparse, and when each belongs to a weak- l_p ball. For

³Here, this set is just an interval and the projection operator returns x if $x \in \mathcal{X}_{i,j}$ or the nearest endpoint of the interval $\mathcal{X}_{i,j}$ otherwise.

Algorithm 1 A_IHT($\mathbf{X}, \mathbf{D}, \mathbf{Z}, \epsilon$) – For solving $\min_{\mathbf{A} \in \mathbb{R}^{n_1 \times p}} I_{\mathcal{A}}(\mathbf{A}) + \lambda \|\mathbf{A}\|_0 + \frac{\rho}{2} \|\mathbf{Z} - \mathbf{D}\mathbf{A}\|_F^2$

Inputs: $\mathbf{X}, \mathbf{D}, \mathbf{Z}, \epsilon, \rho$

Initialize: $\mathbf{A}^{(0)} = \mathbf{0}$

repeat

$$\mathbf{Y}^{(k+1)} = \mathbf{A}^{(k)} - \mathbf{D}^T(\mathbf{D}\mathbf{A}^{(k)} - \mathbf{Z}) / \|\mathbf{D}\|_2^2$$

Update: $Y_{i,j}^{(k+1)} = 0$ if $|Y_{i,j}^{(k+1)}| \leq \sqrt{\frac{2\lambda}{\rho \|\mathbf{D}\|_2^2}}$.

if $Y_{i,j}^{(k+1)} \in \mathcal{A}_{i,j}$:

$$A_{i,j}^{(k+1)} = Y_{i,j}^{(k+1)};$$

else:

$$A_{i,j}^{(k+1)} = \arg \min_{x \in \mathcal{A}_{i,j}} \left\{ x^2 - 2x \left[A_{i,j}^{(k)} - \frac{((\mathbf{D}^T \mathbf{D} \mathbf{A}^{(k)})_{i,j} - (\mathbf{D}^T \mathbf{Z})_{i,j})}{\|\mathbf{D}\|_2^2} \right] \right\}$$

until $\frac{\|\mathbf{A}^{(k+1)} - \mathbf{A}^{(k)}\|_F}{\|\mathbf{A}^{(k)}\|_F} \leq \epsilon$

Output: $\mathbf{A} = \mathbf{A}^{(k+1)}$

Algorithm 2 D_Newton($\mathbf{X}, \mathbf{A}, \mathbf{Z}, \epsilon$) – For solving $\min_{\mathbf{D} \in \mathbb{R}^{n_1 \times p}} I_{\mathcal{D}}(\mathbf{D}) + \frac{\rho}{2} \|\mathbf{Z} - \mathbf{D}\mathbf{A}\|_F^2$

Inputs: $\mathbf{X}, \mathbf{A}, \mathbf{Z}, \epsilon, \rho$

Initialize: $\mathbf{D}^{(0)} = \mathbf{0}$

repeat

$$\mathbf{D}^{(k+1)} = \text{Proj}_{\mathcal{D}} \left[\mathbf{D}^{(k)} - \rho (\mathbf{D}^{(k)} \mathbf{A} - \mathbf{Z}) \mathbf{A}^T (\rho \mathbf{A} \mathbf{A}^T + \delta \mathbf{I})^{-1} \right]$$

until $\frac{\|\mathbf{D}^{(k+1)} - \mathbf{D}^{(k)}\|_F}{\|\mathbf{D}^{(k)}\|_F} \leq \epsilon$

Output: $\mathbf{D} = \mathbf{D}^{(k+1)}$

each scenario we construct the true data matrices $\mathbf{X}^* = \mathbf{D}^* \mathbf{A}^*$ by individually constructing the matrices \mathbf{D}^* and \mathbf{A}^* (as described below), where the entries of the true matrices \mathbf{X}^* , \mathbf{D}^* , and \mathbf{A}^* are bounded in $[X_{\min}^*, X_{\max}^*]$, $[D_{\min}^*, D_{\max}^*]$ and $[A_{\min}^*, A_{\max}^*]$ respectively.

We generate the \mathbf{D}^* matrix by first generating a Gaussian random matrix of size $n_1 \times r$ whose entries are distributed as $\mathcal{N}(0, 1)$, then multiplying each element by $(D_{\max}^* - D_{\min}^*)$ to avoid pathological scaling issues. Finally, we project the resulting scaled matrix onto the set \mathcal{D} , which here is done by truncating all the entries bigger than D_{\max}^* to D_{\max}^* and truncating all the entries smaller than D_{\min}^* to D_{\min}^* . We construct sparse \mathbf{A}^* by generating a Gaussian random matrix of size $r \times n_2$, multiplying it by $(A_{\max}^* - A_{\min}^*)/3$, and projecting it onto the set \mathcal{A} . Then we randomly select $r - k$ locations from each column of the resulting matrix and set the corresponding entries to 0. For the approximately sparse \mathbf{A}^* , we generate each column to be a randomly permuted version of $\{A_{\max}^* \cdot i^{-1/p}\}_{i=1}^r$ with random signs (except for the Poisson likelihood case, where each column of \mathbf{A}^* has nonnegative elements).

We define the set \mathcal{X} such that each entry of \mathbf{X} is bounded in the range $[X_{\min}, X_{\max}]$, \mathcal{D} and \mathcal{A} are the set of all matrices $\mathbf{D} \in \mathbb{R}^{n_1 \times r}$ and $\mathbf{A} \in \mathbb{R}^{r \times n_2}$ whose entries are bounded in the range $[D_{\min}, D_{\max}]$ and $[A_{\min}, A_{\max}]$

Algorithm 3 ADMM algorithm for solving problem (36)

Inputs: $\epsilon_1, \epsilon_2, \Delta_1, \Delta_2, \Delta_1^{\text{stop}}, \Delta_2^{\text{stop}}, \eta, \rho^{(0)} > 0$

Initialize: $\mathbf{D}^{(0)} \in \mathcal{D}, \mathbf{A}^{(0)} \in \mathcal{A}, \mathbf{\Lambda}^{(0)}$.

repeat

$$\mathbf{X}_{i,j}^{(k+1)} = \text{Proj}_{\mathcal{X}} \left[\text{prox}_{s_{i,j}\ell} \left((\mathbf{D}^{(k)} \mathbf{A}^{(k)})_{i,j} - \frac{(\mathbf{\Lambda}^{(k)})_{i,j}}{\rho^{(k)}}; \rho^{(k)}, Y_{i,j} \right) \right]$$

$$\mathbf{A}^{(k+1)} := \mathbf{A_IHT} \left(\mathbf{X}_{(k+1)}, \mathbf{D}^{(k)}, \mathbf{X}^{(k+1)} + \mathbf{\Lambda}^{(k)} / \rho^{(k)}, \epsilon_1 \right)$$

$$\mathbf{D}^{(k+1)} := \mathbf{D_Newton} \left(\mathbf{X}_{(k+1)}, \mathbf{A}^{(k+1)}, \mathbf{X}^{(k+1)} + \mathbf{\Lambda}^{(k)} / \rho^{(k)}, \epsilon_2 \right)$$

$$\mathbf{\Lambda}^{(k+1)} = \mathbf{\Lambda}^{(k)} + \rho^{(k)} (\mathbf{X}^{(k+1)} - \mathbf{D}^{(k+1)} \mathbf{A}^{(k+1)})$$

$$\text{Set } \Delta_1 = \|\mathbf{X}^{(k+1)} - \mathbf{D}^{(k+1)} \mathbf{A}^{(k+1)}\|_F \text{ and } \Delta_2 = \rho^{(k)} \cdot \|\mathbf{D}^{(k)} \mathbf{A}^{(k)} - \mathbf{D}^{(k+1)} \mathbf{A}^{(k+1)}\|_F$$

$$\rho^{(k+1)} = \begin{cases} \eta \cdot \rho^{(k)}, & \text{if } \Delta_1 \geq 10 \cdot \Delta_2 \\ \rho^{(k)} / \eta, & \text{if } \Delta_2 \geq 10 \cdot \Delta_1 \\ \rho^{(k)}, & \text{otherwise} \end{cases}$$

until $\Delta_1 \leq \Delta_1^{\text{stop}}$ and $\Delta_2 \leq \Delta_2^{\text{stop}}$

Output: $\mathbf{D} = \mathbf{D}^{(k+1)}$ and $\mathbf{A} = \mathbf{A}^{(k+1)}$

respectively. It is important to note that in general the actual bounds on the magnitude of the entries of true matrices (for e.g., $\mathbf{D}_{\min}^*, \mathbf{A}_{\max}^*$ etc.) are unknown, and therefore during optimization we might have to use their approximations (which here are denoted as $\mathbf{D}_{\min}, \mathbf{A}_{\max}$ etc.) to define the feasible sets \mathcal{X}, \mathcal{D} and \mathcal{A} . Our specific choices of parameters for the four different likelihoods considered in this paper are summarized in Table II.

Parameters \ Likelihood	Gaussian	Laplace	Poisson	One-bit
$n_1 \times n_2$	100×1000	100×1000	100×1000	1000×1000
r, k, p	20, 8, 1/3	20, 8, 1/3	20, 8, 1/3	5, 2, 1/3
$[\mathbf{D}_{\min}^*, \mathbf{D}_{\max}^*]$	$[-1, 1]$	$[-1, 1]$	$[0.1, 1]$	$[-1, 1]$
$[\mathbf{A}_{\min}^*, \mathbf{A}_{\max}^*]$	$[-20, 20]$	$[-20, 20]$	$[0, 40]$	$[-20, 20]$
$[\mathbf{D}_{\min}, \mathbf{D}_{\max}]$	$[-2, 2]$	$[-2, 2]$	$[-2, 2]$	$[-2, 2]$
$[\mathbf{A}_{\min}, \mathbf{A}_{\max}]$	$[-40, 40]$	$[-40, 40]$	$[-80, 80]$	$[-40, 40]$
$[\mathbf{X}_{\min}, \mathbf{X}_{\max}]$	$[-2 \cdot \mathbf{X}_{\min}^*, 2 \cdot \mathbf{X}_{\max}^*]$	$[-2 \cdot \mathbf{X}_{\min}^*, 2 \cdot \mathbf{X}_{\max}^*]$	$[0, 2 \cdot \mathbf{X}_{\max}^*]$	$[-2 \cdot \mathbf{X}_{\min}^*, 2 \cdot \mathbf{X}_{\max}^*]$

TABLE II: Experimental parameters for different likelihood models we examine. Here $\mathbf{X}_{\min}^* = \min_{i,j} X_{i,j}^*$ and $\mathbf{X}_{\max}^* = \|\mathbf{X}^*\|_{\max}$.

Now, our experimental approach is as follows. For sparse and nearly-sparse (with columns belonging to a weak ℓ_p ball with $p = 1/3$) coefficient matrices \mathbf{A}^* we generate a corresponding matrix \mathbf{X}^* as above. Then, for each of a number of regularization parameters $\lambda > 0$ and sampling rates $\gamma \in (0, 1]$ we perform 20 trials of the following experiment: we generate \mathcal{S} according to the independent Bernoulli(γ) model, obtain noisy observations

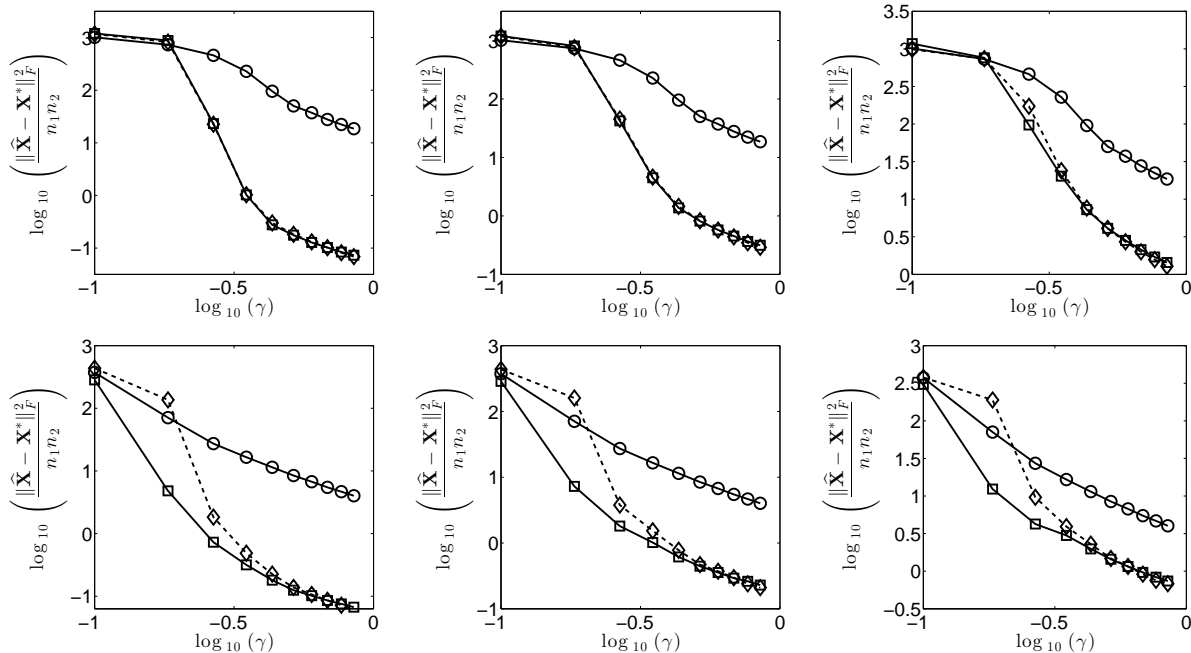


Fig. 1: Results of synthetic experiments for matrix completion with Gaussian, Laplace and Poisson likelihoods: — \square — is Gaussian, — \diamond — is Laplace, and — \circ — is Poisson. Top row corresponds to sparse-factor model with $k = 8$ while the bottom row corresponds to weak- l_p model with $p = 1/3$. Column 1 corresponds to $\sigma^2 = (0.5)^2$ (for Laplace $\tau = \sqrt{8}$), column 2 corresponds to $\sigma^2 = (1)^2$ (for Laplace $\tau = \sqrt{2}$) and column 3 corresponds to $\sigma^2 = (2)^2$ (for Laplace $\tau = 1/\sqrt{2}$). Here $n_1 = 100$, $n_2 = 1000$ and $r = 20$.

of \mathbf{X}^* according to the (7), use Algorithm 3 to obtain⁴ an estimate $\widehat{\mathbf{X}} = \widehat{\mathbf{D}}\widehat{\mathbf{A}}$, and compute its approximation error $\frac{\|\widehat{\mathbf{X}} - \mathbf{X}^*\|_F^2}{n_1 n_2}$. We then compute the empirical average of the errors over the 20 trials for each setting. Fig. 1 shows the results of this experiment for the Gaussian, Laplace and Poisson likelihood models. The plots depict the empirical average (over 20 trials) per-element error as a function of sampling rate on a log-log scale; the curves shown are corresponding to the best (lowest) errors achieved over all of the regularization parameters λ we examined. The first row corresponds to exactly sparse \mathbf{A}^* matrices and the plots in the second row corresponds to settings where \mathbf{A}^* is approximately sparse. The three columns correspond to three different regimes for the Gaussian and Laplace settings (we chose the parameters σ and τ for the Gaussian and Laplace settings, respectively, to yield identical variances; the first column corresponds to $\sigma = 0.5$ and $\tau = \sqrt{8}$, the second column corresponds to $\sigma = 1$ and $\tau = \sqrt{2}$, and the third column corresponds to $\sigma = 2$ and $\tau = 1/\sqrt{2}$).

A few interesting points are worth noting here. First, for the case where \mathbf{A}^* is exactly sparse, our theoretical results predict the error decay be inversely proportional to the nominal sampling rate γ ; viewed on a log-log scale, this would correspond to the error decay having slope -1. Our experimental results provide some evidence to validate

⁴For Algorithm 3 we set $\epsilon_1 = \epsilon_2 = 10^{-7}$, $\Delta_1^{\text{stop}} = \Delta_2^{\text{stop}} = 10$, $\eta = 1.05$ and $\rho^{(0)} = 0.001$.




Fig. 2: Results of synthetic experiments for one-bit matrix completion under sparse factor models, using the logistic link function. The left panel corresponds to the case where the \mathbf{A}^* matrix is exactly sparse; the right when columns of \mathbf{A}^* lie in a weak- ℓ_p ball with $p = 1/3$.

our analysis, at least in the settings where the sampling rate $\gamma > 0.4$ – there, the slopes of the error decays for each of the likelihood models is indeed approximately -1. For the settings where the columns of \mathbf{A}^* belong to a weak- ℓ_p (with $p = 1/3$) our theory predicts that the slope of the error decay (on a log-log scale) be at least $-5/6$ for the Gaussian-noise and Poisson-distributed cases, and at least $(-2/3)$ for the Laplace-noise case. For our experiments here, it appears that the error decay in these approximately-sparse settings is actually a bit faster than predicted by the theory, as the error appears to decay with a slope of approximately -1. That said, it is worth noting that our predicted rate in these cases was obtained essentially by a (squared) bias-variance tradeoff, so quantify a kind of worst-case behavior that may not always be observed in practice.

We also evaluated the performance in this setting for a one-bit observation model, using an analogous experimental setting as above. Here, we used the logistic cumulative distribution function as the link function, i.e., $F(x) = \frac{1}{1+e^{-x/s}}$ where $s = \frac{\sqrt{3}\cdot\sigma}{\pi}$ and σ is a parameter that could be viewed as additive noise standard deviation⁵, for the specific choice $\sigma = 0.1$. Fig. 2 shows the error results for this case, with the first plot corresponds to sparse \mathbf{A}^* and the second to when each column of \mathbf{A}^* lies in a weak- ℓ_p ball with $p = 1/3$. As in the previous experiments, it appears here that the slope of the error decay is approximately -1 in each case. Note that we adapted the experimental setting here to be more amenable to this more difficult estimation regime (specifically, we consider slightly larger matrices but having smaller rank and fewer nonzeros per column of the factor \mathbf{A}^* , as outlined in Table II, so that the number of observations per parameter to be estimated is larger than in the previous three experimental settings).

V. DISCUSSION AND CONCLUSIONS

We conclude with a brief discussion of our results and potentially interesting future directions.

⁵For this link function the proximal operator is $\text{prox}_\ell(z; \rho, y) = \arg \min_{x \in \mathbb{R}} -y \log(F(x)) - (1-y) \log(1-F(x)) + \frac{\rho}{2} (x-z)^2$, which in general is not solvable in closed form. Here, we resort to Newton’s gradient descent algorithm – rewriting the problem as $\text{prox}_\ell(z; \rho, y) = \arg \min_{x \in \mathbb{R}} G(x)$, where $G(x) = -y \log(F(x)) - (1-y) \log(1-F(x)) + \frac{\rho}{2} (x-z)^2$, it is easy to show that the gradient is $\nabla G(x) = -\frac{y}{s} + \frac{F(x)}{s} + \rho(x-z)$ and the Hessian is $\nabla^2 G(x) = \frac{F(x)(1-F(x))}{s^2} + \rho$. We can then iteratively solve for $\text{prox}_\ell(z; \rho, y)$ by Newton steps (starting from a random $x^{(0)}$) of the form $x^{(k+1)} = x^{(k)} - \frac{\nabla G(x^{(k)})}{\nabla^2 G(x^{(k)})}$ until convergence (here, until $\|x^{(k+1)} - x^{(k)}\| \leq 10^{-7}$).

A. Extensions to Other Data Models

Each of our theoretical results above follow essentially from the specialization of a more general result (appearing below as Lemma A.1) to the case of sparse factor models. It is interesting to note that this lemma may also be specialized (in a straightforward manner) to any of a number of other interesting factor models (e.g., non-negative matrix factorizations, factorizations where each factor may be sparse, etc.) under the same general observation models we consider here. Further, while we provide Lemma A.1 specifically for the case of matrix completion, the essential analysis extends (simply) to higher-order structures (i.e., tensors) as well.

B. Convexification?

As discussed in several points in the preceding sections, the optimization associated with the estimators we consider here is non-convex on account of several factors, including the presence of the ℓ_0 term in the objective, our optimization over a discretized set, and more fundamentally, the fact that we perform inference in a general bilinear model, where both factors are unknown. Resolving ourselves, then, to seek only local optima of the corresponding optimizations allows us to bring to bear alternating direction method of multipliers techniques, in which the ℓ_0 -based optimization *subproblems* may be solved efficiently. Interestingly, within this framework we may also directly incorporate the constraints that the matrix factor elements each come from a discretized set (indeed, this would correspond to choosing set indicator functions that take the value ∞ outside of the discretized sets over which we seek to optimize). We did not pursue this latter condition in our simulations, assuming instead that the discretization of each of the elements be “sufficiently fine” so that we may solve the optimization numerically at machine precision (and replace the discretized sets for the candidate matrix factors by their convex hulls).

The fact that we can (locally) handle the ℓ_0 constraints within the ADMM framework notwithstanding, it is interesting to consider whether there is any benefit to relaxing this constraint to a convex surrogate (e.g., replacing the ℓ_0 penalty with an ℓ_1 penalty). The resulting procedure would still be jointly non-convex in the matrix factors, but could be addressed within a similar algorithmic framework to the one we propose above. Analytically, methods that prescribe optimization over a convex set comprised of the Cartesian product of a set \mathcal{D} of matrices whose elements satisfy a max-norm constraint and a set \mathcal{A} of matrices whose columns satisfy an ℓ_1 -constraint may be amenable to analysis using entropy-based methods that can be employed to analyze estimation error performance by bounding suprema of empirical processes indexed by elements of the feasible set of candidate estimates – see, e.g., [72]–[74]. It would be interesting to see whether analyses along these yield substantially different results than our analysis here; analyses along these lines are a subject of our ongoing work and will be reported in a subsequent effort.

In the meantime, it is interesting to examine (albeit, empirically) whether our algorithmic approach yields significantly different performance if we replace the ℓ_0 regularization term by an ℓ_1 term. To provide some insight into this, we consider a problem of completing a 50×500 matrix $\mathbf{X}^* = \mathbf{D}^* \mathbf{A}^*$, where \mathbf{D}^* is 50×10 and \mathbf{A}^* is 10×500 and sparse, having 4 nonzero elements per column. We consider Gaussian noise-corrupted observations obtained at a subset of locations of \mathbf{X}^* (generated according to the independent Bernoulli model), and three different




Fig. 3: Comparison between sparse-factor and nuclear-norm-regularized matrix completion methods. The curves are: our proposed procedure with ℓ_0 regularizer (\square), the ℓ_1 regularized variant of our approach (\triangleright), and nuclear norm regularized low-rank matrix completion (\circ). The sparse factor completion methods perform similarly, and both achieve a lower error than the best nuclear-norm regularized estimate for sampling rates $\gamma \geq 10^{-0.5} \approx 30\%$.

reconstruction approaches: the first is the algorithmic approach described in the previous section, the second is a slight variation of our proposed approach where we replace the ℓ_0 penalty by an ℓ_1 penalty (and replace the corresponding inference step with an accelerated first-order method as in [75]), and the third method is a more standard low-rank recovery obtained via *nuclear-norm* regularization, as $\hat{\mathbf{X}} = \arg \min_{\mathbf{X}} \|\mathbf{Y}_S - \mathbf{X}_S\|_F^2 + \lambda \|\mathbf{X}\|_*$. For each method, we examined a range of possible values for the regularization parameter, and selected the reconstruction corresponding (clairvoyantly) to the best choice for each method. The results, provided in Figure 3, show that the best-performing ℓ_0 and ℓ_1 regularized sparse factor completion methods perform comparably, while both achieve (slightly) lower error than the best nuclear norm regularized completion estimate. Of course, as noted above, our algorithmic approach identifies (at best) a local minimum of the overall non-convex problem we aim to solve, but even at that, it is encouraging to see that the ADMM-based optimization(s) identify good-quality estimates.

It is also interesting to consider an alternative, more essential, convexification of our problem of interest here, using the machinery of *atomic norms* as introduced in [76]. Specifically, one may view matrices adhering to the sparse factor models we investigate here as sums of rank-one matrices formed as outer products between a (non-sparse) $n_1 \times 1$ vector and a (sparse) $n_2 \times 1$ vector. Following [76], one can consider the convex hull of the set of all such rank-one atoms having unit (Frobenius) norm as the unit-ball for a norm that serves as a regularizer for matrices representable by weighted sums of only a few atoms. A very recent work [77] has begun to identify properties of atomic norms so-formed, and extensions to the cases where both of the vectors may be sparse, and have established some estimation guarantees for recovering simple matrices (comprised of a single rank-one outer product of sparse vectors) from a collection of Gaussian measurements. Interestingly, the authors of [77] note that resulting inference procedures using their so-called (k, q) -norm (formed from atoms that are rank-one outer products between k -sparse and q -sparse vectors), while convex, may still be computationally intractable (even NP-hard)! At any rate, it would be quite interesting to extend this approach to the entry-wise sampling models and various likelihood models we

consider here, and we defer investigations along these lines to a future work.

C. Lower Bounds

Our error bounds here provide some insight into the performance of sparsity-penalized maximum likelihood estimation approaches to sparse factor matrix completion tasks. To the best of our knowledge, lower bounds on the achievable mean-square estimation error for these tasks have not been established, but would be a valuable complement to place our results here into a broader context. Efforts along these lines are ongoing, and will be reported in a future work.

APPENDIX

A. Proof of Theorem II.1

Our proof of Theorem II.1 is based on an application of the following general lemma, which we prove in Appendix F.

Lemma A.1. *Let \mathbf{X}^* be an $n_1 \times n_2$ matrix whose elements we aim to estimate, and let \mathcal{X} be a countable collection of candidate reconstructions \mathbf{X} of \mathbf{X}^* , each with corresponding penalty $\text{pen}(\mathbf{X}) \geq 1$, so that the collection of penalties satisfies the summability condition $\sum_{\mathbf{X} \in \mathcal{X}} 2^{-\text{pen}(\mathbf{X})} \leq 1$.*

*Fix an integer m with $4 \leq m \leq n_1 n_2$, let $\gamma = m(n_1 n_2)^{-1}$, generate a sampling set \mathcal{S} according to the independent Bernoulli(γ) model so that each $(i, j) \in [n_1] \times [n_2]$ is included in \mathcal{S} independently with probability γ , and obtain corresponding observations $\mathbf{Y}_{\mathcal{S}} \sim p_{\mathbf{X}^*_{\mathcal{S}}} = \prod_{(i,j) \in \mathcal{S}} p_{X^*_{i,j}}$, which are assumed to be conditionally independent given \mathcal{S} . Then, if C_{D} is any constant satisfying*

$$C_{\text{D}} \geq \max_{\mathbf{X} \in \mathcal{X}} \max_{(i,j) \in [n_1] \times [n_2]} \text{D}(p_{X^*_{i,j}} \| p_{X_{i,j}}), \quad (47)$$

we have that for any

$$\xi \geq \left(1 + \frac{2C_{\text{D}}}{3}\right) \cdot 2 \log 2, \quad (48)$$

the complexity penalized maximum likelihood estimator

$$\widehat{\mathbf{X}}^{\xi} = \widehat{\mathbf{X}}^{\xi}(\mathcal{S}, \mathbf{Y}_{\mathcal{S}}) = \arg \min_{\mathbf{X} \in \mathcal{X}} \{-\log p_{\mathbf{X}_{\mathcal{S}}}(\mathbf{Y}_{\mathcal{S}}) + \xi \cdot \text{pen}(\mathbf{X})\}, \quad (49)$$

satisfies the (normalized, per-element) error bound

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} [-2 \log A(p_{\widehat{\mathbf{X}}^{\xi}}, p_{\mathbf{X}^*})]}{n_1 n_2} \leq 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{\text{D}(p_{\mathbf{X}^*} \| p_{\mathbf{X}})}{n_1 n_2} + \left(\xi + \frac{4C_{\text{D}} \log 2}{3} \right) \frac{\text{pen}(\mathbf{X})}{m} \right\} + \frac{8C_{\text{D}} \log m}{m}, \quad (50)$$

where, as denoted, the expectation is with respect to the joint distribution of \mathcal{S} and $\mathbf{Y}_{\mathcal{S}}$.

In order to use this result here, we need to define penalties $\text{pen}(\mathbf{X}) \geq 1$ on candidate reconstructions \mathbf{X} of \mathbf{X}^* , so that for every subset \mathcal{X} of the set \mathcal{X}' specified in the conditions of Theorem II.1 the summability condition $\sum_{\mathbf{X} \in \mathcal{X}} 2^{-\text{pen}(\mathbf{X})} \leq 1$ holds. To this end, we will use the fact that for any $\mathcal{X} \subseteq \mathcal{X}'$ we always have

$\sum_{\mathbf{X} \in \mathcal{X}} 2^{-\text{pen}(\mathbf{X})} \leq \sum_{\mathbf{X} \in \mathcal{X}'} 2^{-\text{pen}(\mathbf{X})}$; thus, it suffices for us to show that for the specific set \mathcal{X}' described in Section II,

$$\sum_{\mathbf{X} \in \mathcal{X}'} 2^{-\text{pen}(\mathbf{X})} \leq 1. \quad (51)$$

Note that the condition (51) is the well-known Kraft-McMillan Inequality for coding elements of \mathcal{X}' with an alphabet of size 2, which is satisfied automatically if we choose the penalties to be *code lengths* for some uniquely decodable binary code for the elements $\mathbf{X} \in \mathcal{X}'$; see [78]. This interpretation will provide us with a *constructive* approach to designing penalties, as we will see below.

Now, consider any discretized matrix factors $\mathbf{D} \in \mathcal{D}$ and $\mathbf{A} \in \mathcal{A}$, as described in Section II. Let us fix an ordering of the indices of elements of \mathbf{D} and encode the amplitude of each element using $\log_2 L_{\text{lev}}$ bits, and for $L_{\text{loc}} \triangleq 2^{\lceil \log_2 r n_2 \rceil}$ we encode each *nonzero* element of \mathbf{A} using $\log_2 L_{\text{loc}}$ bits to denote its location and $\log_2 L_{\text{lev}}$ bits for its amplitude. With this strategy, a total of $n_1 r \log_2 L_{\text{lev}}$ bits are used to encode \mathbf{D} and matrices \mathbf{A} having $\|\mathbf{A}\|_0$ nonzero entries are encoded using $\|\mathbf{A}\|_0 (\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}})$ bits. Now, we let \mathcal{X}'' be the set of all such $\mathbf{X} = \mathbf{D}\mathbf{A}$, and let the code for each \mathbf{X} be the concatenation of the (fixed-length) code for \mathbf{D} followed by the (variable-length) code for \mathbf{A} . It follows that we may assign penalties $\text{pen}(\mathbf{X})$ to all $\mathbf{X} \in \mathcal{X}''$ whose lengths satisfy

$$\text{pen}(\mathbf{X}) = n_1 r \log_2 L_{\text{lev}} + \|\mathbf{A}\|_0 (\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}}). \quad (52)$$

It is easy to see that such codes are (by construction) uniquely decodable, so we have that $\sum_{\mathbf{X} \in \mathcal{X}''} 2^{-\text{pen}(\mathbf{X})} \leq 1$. Now, the set \mathcal{X}' specified in the theorem is a subset of \mathcal{X}'' (or perhaps \mathcal{X}'' itself, if all elements satisfy the max norm bound condition $\|\mathbf{X}\|_{\text{max}} \leq X_{\text{max}}$), so (51) holds for \mathcal{X}' as specified in the theorem.

Now let \mathcal{X} be any subset of \mathcal{X}' . By the above argument the summability condition holds for \mathcal{X} , so we may apply the results of Lemma A.1. For randomly subsampled and noisy observations \mathbf{Y}_S our estimates take the form

$$\begin{aligned} \widehat{\mathbf{X}}^\xi &= \arg \min_{\mathbf{X}=\mathbf{D}\mathbf{A} \in \mathcal{X}} \{-\log p_{\mathbf{X}_S}(\mathbf{Y}_S) + \xi \cdot \text{pen}(\mathbf{X})\} \\ &= \arg \min_{\mathbf{X}=\mathbf{D}\mathbf{A} \in \mathcal{X}} \{-\log p_{\mathbf{X}_S}(\mathbf{Y}_S) + \xi \cdot (\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}}) \cdot \|\mathbf{A}\|_0\}. \end{aligned} \quad (53)$$

where the last line follows by disregarding additive constants in the optimization arising from terms that do not depend on \mathbf{X} (or more specifically, on \mathbf{D} or \mathbf{A}) in the penalty. Further, when ξ satisfies (48), we have

$$\begin{aligned} \frac{\mathbb{E}_{S, \mathbf{Y}_S} [-2 \log A(p_{\widehat{\mathbf{X}}^\xi}, p_{\mathbf{X}^*})]}{n_1 n_2} &\leq \frac{8C_D \log m}{m} + \\ 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} &\left\{ \frac{D(p_{\mathbf{X}^*}, p_{\mathbf{X}})}{n_1 n_2} + \left(\xi + \frac{4C_D \log 2}{3} \right) (\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}}) \left(\frac{n_1 r + \|\mathbf{A}\|_0}{m} \right) \right\}, \end{aligned} \quad (54)$$

Finally, letting

$$\lambda = \xi \cdot (\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}}) \quad (55)$$

and using the fact that

$$\log_2 L_{\text{loc}} + \log_2 L_{\text{lev}} \leq (\beta + 2) \cdot \log(n_1 \vee n_2) \cdot 2 \log 2 \quad (56)$$

which follows by our selection of L_{lev} and L_{loc} and the fact that $r < n_2$, it follows (after some straightforward simplification) that for

$$\lambda \geq 2(\beta + 2) \left(1 + \frac{2C_D}{3}\right) \log(n_1 \vee n_2) \quad (57)$$

the estimate

$$\widehat{\mathbf{X}}^\lambda = \arg \min_{\mathbf{X} \in \mathcal{X}} \{-\log p_{\mathbf{X}_S}(\mathbf{Y}_S) + \lambda \cdot \|\mathbf{A}\|_0\}$$

satisfies

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [-2 \log A(p_{\widehat{\mathbf{X}}^\lambda}, p_{\mathbf{X}^*})]}{n_1 n_2} &\leq \frac{8C_D \log m}{m} + \\ &3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{D(p_{\mathbf{X}^*} \| p_{\mathbf{X}})}{n_1 n_2} + \left(\lambda + \frac{4C_D(\beta + 2) \log(n_1 \vee n_2)}{3} \right) \left(\frac{n_1 r + \|\mathbf{A}\|_0}{m} \right) \right\}, \end{aligned} \quad (58)$$

as claimed.

B. Proof of Corollary III.1

We first establish a general error bound, which we then specialize to the case stated in the corollary. Note that for \mathbf{X}^* as specified and any $\mathbf{X} \in \mathcal{X}$, using the model (15) we have

$$D(p_{X_{i,j}^*} \| p_{X_{i,j}}) = \frac{(X_{i,j}^* - X_{i,j})^2}{2\sigma^2} \quad (59)$$

for any fixed $(i, j) \in S$. It follows that $D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) = \|\mathbf{X}^* - \mathbf{X}\|_F^2 / 2\sigma^2$, and using the fact that the amplitudes of entries of \mathbf{X}^* and all $\mathbf{X} \in \mathcal{X}$ are no larger than X_{\max} , it is clear that we may choose $C_D = 2X_{\max}^2 / \sigma^2$. Further, for any $\mathbf{X} \in \mathcal{X}$ and any fixed $(i, j) \in S$ it is easy to show that in this case

$$-2 \log A(p_{X_{i,j}}, p_{X_{i,j}^*}) = \frac{(X_{i,j}^* - X_{i,j})^2}{4\sigma^2}, \quad (60)$$

so that $-2 \log A(p_{\mathbf{X}}, p_{\mathbf{X}^*}) = \|\mathbf{X}^* - \mathbf{X}\|_F^2 / 4\sigma^2$. It follows that

$$\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [-2 \log A(p_{\widehat{\mathbf{X}}}, p_{\mathbf{X}^*})] = \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2]}{4\sigma^2}. \quad (61)$$

Incorporating this into Theorem II.1, we obtain that for any

$$\lambda \geq \left(1 + \frac{4X_{\max}^2}{3\sigma^2}\right) \cdot 2(\beta + 2) \cdot \log(n_1 \vee n_2), \quad (62)$$

the sparsity penalized ML estimate satisfies the per-element mean-square error bound

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2]}{n_1 n_2} &\leq \frac{64X_{\max}^2 \log m}{m} \\ &+ 6 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{\|\mathbf{X}^* - \mathbf{X}\|_F^2}{n_1 n_2} + \left(2\sigma^2 \lambda + \frac{16X_{\max}^2 (\log 2)^2 (\beta + 1) \log(n_1 \vee n_2)}{3} \right) \left(\frac{n_1 p + \|\mathbf{A}\|_0}{m} \right) \right\}. \end{aligned} \quad (63)$$

We now establish the error bound for the case where the coefficient matrix \mathbf{A}^* is exactly sparse and λ is fixed to the value specified in (14). Consider a candidate reconstruction of the form $\mathbf{X}_Q^* = \mathbf{D}_Q^* \mathbf{A}_Q^*$, where the elements of \mathbf{D}_Q^* are the closest discretized surrogates of the entries of \mathbf{D}^* , and the entries of \mathbf{A}_Q^* are the closest discretized

surrogates of the *nonzero* entries of \mathbf{A}^* (and zero otherwise). Denote $\mathbf{D}_Q^* = \mathbf{D}^* + \Delta_{\mathbf{D}^*}$ and $\mathbf{A}_Q^* = \mathbf{A}^* + \Delta_{\mathbf{A}^*}$. Then it is easy to see that

$$\mathbf{D}_Q^* \mathbf{A}_Q^* - \mathbf{D}^* \mathbf{A}^* = \mathbf{D}^* \Delta_{\mathbf{A}^*} + \Delta_{\mathbf{D}^*} \mathbf{A}^* + \Delta_{\mathbf{D}^*} \Delta_{\mathbf{A}^*}. \quad (64)$$

Given the range limits on allowable \mathbf{D} and \mathbf{A} and that each range is quantized to L_{lev} levels, we have that $\|\Delta_{\mathbf{D}^*}\|_{\max} \leq 1/(L_{\text{lev}} - 1)$ and $\|\Delta_{\mathbf{A}^*}\|_{\max} \leq A_{\max}/(L_{\text{lev}} - 1)$. Now, we can obtain a bound on the magnitudes of the elements of $\mathbf{D}_Q^* \mathbf{A}_Q^* - \mathbf{D}^* \mathbf{A}^*$ that hold uniformly over all i, j , as follows

$$\begin{aligned} \|\mathbf{D}_Q^* \mathbf{A}_Q^* - \mathbf{D}^* \mathbf{A}^*\|_{\max} &= \max_{i,j} |(\mathbf{D}^* \Delta_{\mathbf{A}^*} + \Delta_{\mathbf{D}^*} \mathbf{A}^* + \Delta_{\mathbf{D}^*} \Delta_{\mathbf{A}^*})_{i,j}| \\ &\leq \max_{i,j} |(\mathbf{D}^* \Delta_{\mathbf{A}^*})_{i,j}| + |(\Delta_{\mathbf{D}^*} \mathbf{A}^*)_{i,j}| + |(\Delta_{\mathbf{D}^*} \Delta_{\mathbf{A}^*})_{i,j}| \\ &\leq \frac{rA_{\max}}{L_{\text{lev}} - 1} + \frac{rA_{\max}}{L_{\text{lev}} - 1} + \frac{2rA_{\max}}{(L_{\text{lev}} - 1)^2} \\ &\leq \frac{8rA_{\max}}{L_{\text{lev}}}, \end{aligned} \quad (65)$$

where the first inequality follows from the triangle inequality, the second from the bounds on $\|\Delta_{\mathbf{D}^*}\|_{\max}$ and $\|\Delta_{\mathbf{A}^*}\|_{\max}$ and the entry-wise bounds on elements of allowable \mathbf{D} and \mathbf{A} , and the last because $L_{\text{lev}} \geq 2$. Now, it is straight-forward to show that our choice of β in (13) implies $L_{\text{lev}} \geq 16rA_{\max}/X_{\max}$, so each entry of $\mathbf{D}_Q^* \mathbf{A}_Q^* - \mathbf{D}^* \mathbf{A}^*$ is bounded in magnitude by $X_{\max}/2$. It follows that each element of the candidate \mathbf{X}_Q^* constructed above is bounded in magnitude by X_{\max} , so \mathbf{X}_Q^* is indeed a valid element of the set \mathcal{X} .

Further, the approximation error analysis above also implies directly that

$$\begin{aligned} \frac{\|\mathbf{X}^* - \mathbf{X}_Q^*\|_F^2}{n_1 n_2} &= \frac{1}{n_1 n_2} \sum_{i \in [n_1], j \in [n_2]} (\mathbf{D}_Q^* \mathbf{A}_Q^* - \mathbf{D}^* \mathbf{A}^*)_{i,j}^2 \\ &\leq \frac{64p^2 A_{\max}^2}{L_{\text{lev}}^2} \\ &\leq \frac{X_{\max}^2}{m}, \end{aligned} \quad (66)$$

where the last line follows from the fact that our specific choice of β in (13) also implies $L_{\text{lev}} \geq 8r\sqrt{m}A_{\max}/X_{\max}$. Now, evaluating the oracle term at the candidate $\mathbf{X}_Q^* = \mathbf{D}_Q^* \mathbf{A}_Q^*$, and using the fact that $\|\mathbf{A}_Q^*\|_0 = \|\mathbf{A}^*\|_0$, we have

$$\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} \leq \frac{70X_{\max}^2 \log m}{m} + 8(3\sigma^2 + 8X_{\max}^2)(\beta + 2) \log(n_1 \vee n_2) \left(\frac{n_1 r + \|\mathbf{A}^*\|_0}{m} \right). \quad (67)$$

Finally, we establish the error bound for the case where columns of \mathbf{A}^* are in a weak ℓ_p ball of radius A_{\max} , for $p \leq 1$. To that end, let us denote the columns of \mathbf{A}^* by \mathbf{a}_j^* for $j \in [n_2]$, and for any $k \in [r]$, we let $\mathbf{a}_j^{*,(k)}$ denote the best k -term approximation of \mathbf{a}_j^* , formed by retaining the largest (in magnitude) elements and setting the rest to zero. For shorthand, we denote by $\mathbf{A}^{*,(k)}$ the matrix with columns $\mathbf{a}_j^{*,(k)}$ for $j \in [n_2]$. Now, the approximation error incurred may be bounded as

$$\begin{aligned} \|\mathbf{X}^* - \mathbf{X}^{*,(k)}\|_F^2 &= \sum_{i,j} (\mathbf{D}^* (\mathbf{A}^* - \mathbf{A}^{*,(k)}))_{i,j}^2 \\ &\leq \sum_{i,j} \|\mathbf{a}_j^* - \mathbf{a}_j^{*,(k)}\|_2^2, \end{aligned} \quad (68)$$

where the inequality follows from the fact that each $(\mathbf{D}^*(\mathbf{A}^* - \mathbf{A}^{*,(k)}))_{i,j}$ may be expressed as an inner product between the i -th row of \mathbf{D}^* (whose elements are no larger than 1 in magnitude) and the j -th column of $\mathbf{A}^* - \mathbf{A}^{*,(k)}$. To simplify further, we use the fact that $p \leq 1$ (and $q \geq 2p$), and the approximation behavior of vectors in weak ℓ_p balls (discussed in the preliminaries) to obtain that $\|\mathbf{a}_j^* - \mathbf{a}_j^{*,(k)}\|_2^2 \leq A_{\max}^2 k^{-2(1/p-1/2)}$. Letting $\alpha = 1/p - 1/2$, we have that the approximation error associated with approximating \mathbf{A}^* by its best k -term approximation satisfies $\|\mathbf{X}^* - \mathbf{X}^{*,(k)}\|_F^2 \leq n_1 n_2 A_{\max}^2 k^{-2\alpha}$.

Now, we consider a candidate reconstruction of the form $\mathbf{X}_Q^{*,(k)} = \mathbf{D}_Q^* \mathbf{A}_Q^{*,(k)}$ where \mathbf{D}_Q^* is as above and where the nonzero elements of $\mathbf{A}_Q^{*,(k)}$ are taken to be the closest quantized surrogates of the corresponding nonzero elements of $\mathbf{A}^{*,(k)}$. Using the fact that

$$\begin{aligned} \frac{\|\mathbf{X}^* - \mathbf{X}_Q^{*,(k)}\|_F^2}{n_1 n_2} &\leq \frac{4 \left(\|\mathbf{X}^* - \mathbf{X}^{*,(k)}\|_F^2 + \|\mathbf{X}^{*,(k)} - \mathbf{X}_Q^{*,(k)}\|_F^2 \right)}{n_1 n_2} \\ &\leq 4A_{\max}^2 k^{-2\alpha} + \frac{4X_{\max}^2}{m}, \end{aligned} \quad (69)$$

where the first term on the bottom results from the approximation error analysis above and the second from our analysis of the first result of the corollary, we evaluate the oracle bound at the candidate $\mathbf{X}_Q^{(k)}$ to obtain

$$\begin{aligned} &\frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} \\ &\leq \frac{88X_{\max}^2 \log m}{m} + \min_{k \geq 1} \left\{ 24A_{\max}^2 k^{-2\alpha} + 8(3\sigma^2 + 8X_{\max}^2)(\beta + 2) \log(n_1 \vee n_2) \left(\frac{n_1 r + k n_2}{m} \right) \right\}. \end{aligned} \quad (70)$$

Finally, we choose $k = (m/n_2)^{1/(1+2\alpha)}$ to balance the decay rates on the $k^{-2\alpha}$ and $k n_2/m$ terms, and thus obtain

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} &\leq \frac{88X_{\max}^2 \log m}{m} + 8(3\sigma^2 + 8X_{\max}^2)(\beta + 2) \log(n_1 \vee n_2) \frac{n_1 r}{m} \\ &\quad + [24A_{\max}^2 + 8(3\sigma^2 + 8X_{\max}^2)(\beta + 2) \log(n_1 \vee n_2)] \left(\frac{n_2}{m} \right)^{\frac{2\alpha}{2\alpha+1}}. \end{aligned} \quad (71)$$

The stated bounds in each case follow from some straight-forward bounding, as well as the fact mentioned in Section III, that under our assumptions, $(\beta + 2) \log(n_1 \vee n_2) = \mathcal{O}(\log(n_1 \vee n_2))$.

C. Proof of Corollary III.2

We follow a similar approach as in the proof of Corollary III.1, and first establish the general error bound. For \mathbf{X}^* as specified and any fixed $\mathbf{X} \in \mathcal{X}$. We have by (relatively) straight-forward calculation that for any fixed $(i, j) \in \mathcal{S}$,

$$\begin{aligned} D(p_{X_{i,j}^*} \| p_{X_{i,j}}) &= \tau |X_{i,j}^* - X_{i,j}| - (1 - e^{-\tau |X_{i,j}^* - X_{i,j}|}) \\ &\leq \tau |X_{i,j}^* - X_{i,j}| \end{aligned} \quad (72)$$

where the inequality follows from the fact that $(1 - e^{-\tau |X_{i,j}^* - X_{i,j}|}) \geq 0$, and

$$\begin{aligned} -2 \log A(p_{X_{i,j}}, p_{X_{i,j}^*}) &= \tau |X_{i,j}^* - X_{i,j}| - 2 \log \left(1 + \tau \frac{|X_{i,j}^* - X_{i,j}|}{2} \right) \\ &\geq \frac{\tau^2}{4(\tau X_{\max} + 1)^2} (X_{i,j}^* - X_{i,j})^2, \end{aligned} \quad (73)$$

where the inequality follows from the convexity of the negative log Hellinger affinity along with an application of Taylor's theorem⁶. It follows from this that $D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) \leq \tau \|\mathbf{X}^* - \mathbf{X}\|_1$, and

$$\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [-2 \log A(p_{\widehat{\mathbf{X}}}, p_{\mathbf{X}^*})] \geq \frac{\tau^2}{4(\tau X_{\max} + 1)^2} \mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2]. \quad (74)$$

Further, we may choose $C_D = 2\tau X_{\max}$. Incorporating this into Theorem II.1, we have that for any

$$\lambda \geq 2(\beta + 2) \left(1 + \frac{4\tau X_{\max}}{3}\right) \log(n_1 \vee n_2), \quad (75)$$

the sparsity-penalized ML estimate satisfies

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2]}{n_1 n_2} &\leq \frac{1}{\tau} \cdot \frac{64(\tau X_{\max} + 1)^2 X_{\max} \log m}{m} + \\ &\frac{12(\tau X_{\max} + 1)^2}{\tau} \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{\|\mathbf{X}^* - \mathbf{X}\|_1}{n_1 n_2} + \left(\frac{\lambda}{\tau} + \frac{8X_{\max}(\beta + 2) \log(n_1 \vee n_2)}{3} \right) \left(\frac{n_1 p + \|\mathbf{A}\|_0}{m} \right) \right\}. \end{aligned} \quad (76)$$

We now establish the error bound for the case where the coefficient matrix \mathbf{A}^* is sparse and λ is fixed to the value (14). We again consider a candidate reconstruction of the form $\mathbf{X}_Q^* = \mathbf{D}_Q^* \mathbf{A}_Q^*$, where the elements of \mathbf{D}_Q^* are the closest discretized surrogates of the entries of \mathbf{D}^* , and the entries of \mathbf{A}_Q^* are the closest discretized surrogates of the nonzero entries of \mathbf{A}^* (and zero otherwise). Now, since β is the same as in the proof of Corollary III.1, we can directly apply the bound of (65) (and use the fact that $L_{\text{lev}} \geq 16rA_{\max}/X_{\max}$) to conclude that

$$\frac{\|\mathbf{X}^* - \mathbf{X}_Q^*\|_1}{n_1 n_2} \leq \frac{X_{\max}}{2n_1 n_2} \leq \frac{X_{\max}}{m}. \quad (77)$$

Now, evaluating the oracle term at the candidate $\mathbf{X}_Q^* = \mathbf{D}_Q^* \mathbf{A}_Q^*$, and using the fact that $\|\mathbf{A}_Q^*\|_0 = \|\mathbf{A}^*\|_0$, we have

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2]}{n_1 n_2} &\leq \frac{76(\tau X_{\max} + 1)^2}{\tau^2} \cdot \frac{\tau X_{\max} \log m}{m} \\ &+ \frac{12(\tau X_{\max} + 1)^2}{\tau^2} \left(2 + \frac{16\tau X_{\max}}{3}\right) (\beta + 2) \log(n_1 \vee n_2) \left(\frac{n_1 p + \|\mathbf{A}^*\|_0}{m}\right). \end{aligned} \quad (78)$$

Finally, we establish the error bound for the case where the columns of \mathbf{A}^* are vectors in a weak ℓ_p ball for $p \leq 1/2$. By a similar analysis as above, we conclude that $\|\mathbf{X}^* - \mathbf{X}^{*,(k)}\|_1 \leq n_1 n_2 A_{\max} k^{-\alpha'}$, where $\alpha = 1/p - 1$. Now, we consider a candidate reconstruction of the form $\mathbf{X}_Q^{*,(k)} = \mathbf{D}_Q^* \mathbf{A}_Q^{*,(k)}$ where \mathbf{D}_Q^* is as above and where the nonzero elements of $\mathbf{A}_Q^{*,(k)}$ are taken to be the closest quantized surrogates of the corresponding nonzero elements of $\mathbf{A}^{*,(k)}$. Using the fact that

$$\begin{aligned} \frac{\|\mathbf{X}^* - \mathbf{X}_Q^{*,(k)}\|_1}{n_1 n_2} &\leq \frac{\|\mathbf{X}^* - \mathbf{X}^{*,(k)}\|_1 + \|\mathbf{X}^{*,(k)} - \mathbf{X}_Q^{*,(k)}\|_1}{n_1 n_2} \\ &\leq A_{\max} k^{-\alpha'} + \frac{X_{\max}}{m}, \end{aligned} \quad (79)$$

⁶Formally, letting $x \triangleq X_{i,j}^* - X_{i,j}$ and $f(x) = \tau|x| - 2 \log(1 + \tau|x|/2)$ we have

$$f'(x) = \frac{\tau^2}{2} \left(\frac{x}{1 + \tau|x|/2} \right) \quad \text{and} \quad f''(x) = \frac{\tau^2}{2(1 + \tau|x|/2)^2}.$$

Thus, $f(x)$ is twice differentiable (everywhere). The result follows from the fact that $f(0) = f'(0) = 0$ and

$$f''(x) \geq \frac{\tau^2}{2(1 + \tau X_{\max})^2}$$

for all x of the specified form, given the assumptions on \mathbf{X}^* and \mathbf{X} .

where the first term on the bottom results from the approximation error analysis above and the second from our analysis of the first result of the corollary, we evaluate the oracle bound at the candidate $\mathbf{X}_Q^{(k)}$ to obtain

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} &\leq \frac{76(\tau X_{\max} + 1)^2}{\tau^2} \cdot \frac{\tau X_{\max} \log m}{m} \\ &+ \frac{12(\tau X_{\max} + 1)^2}{\tau^2} \min_{k \geq 1} \left\{ \tau A_{\max} k^{-\alpha'} + \left(2 + \frac{16\tau X_{\max}}{3} \right) (\beta + 2) \log(n_1 \vee n_2) \left(\frac{n_1 p + n_2 k}{m} \right) \right\}. \end{aligned} \quad (80)$$

Finally, we choose $k = (m/n_2)^{1/(1+\alpha')}$ to balance the $k^{-\alpha'}$ and $n_2 k/m$ terms, and thus obtain

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} &\leq \frac{76(\tau X_{\max} + 1)^2}{\tau^2} \cdot \frac{\tau X_{\max} \log m}{m} \\ &+ \frac{12(\tau X_{\max} + 1)^2}{\tau^2} \left(2 + \frac{16\tau X_{\max}}{3} \right) (\beta + 2) \log(n_1 \vee n_2) \left(\frac{n_1 p}{m} \right) \\ &+ \frac{12(\tau X_{\max} + 1)^2}{\tau^2} \left(\tau A_{\max} + \left(2 + \frac{16\tau X_{\max}}{3} \right) (\beta + 2) \log(n_1 \vee n_2) \right) \left(\frac{n_2}{m} \right)^{\frac{\alpha'}{\alpha'+1}}. \end{aligned} \quad (81)$$

D. Proof of Corollary III.3 (Sketch)

We follow a similar approach as for the previous proofs, by first establishing a general error bound. We make use of intermediate results from [33] to bound the KL divergences and negative log Hellinger affinities for the Poisson pmf in terms of quadratic differences. Applying those techniques to our setting, we obtain that

$$D(p_{X_{i,j}^*} \| p_{X_{i,j}}) \leq \frac{(X_{i,j}^* - X_{i,j})^2}{X_{\min}} \quad (82)$$

and

$$-2 \log A(p_{X_{i,j}^*}, p_{X_{i,j}}) \geq \frac{(X_{i,j}^* - X_{i,j})^2}{4X_{\max}}. \quad (83)$$

It follows that $D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) \leq \|\mathbf{X}^* - \mathbf{X}\|_F^2 / X_{\min}$, $\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [-2 \log A(p_{\widehat{\mathbf{X}}^\lambda}, p_{\mathbf{X}^*})] \geq \mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} [\|\mathbf{X}^* - \widehat{\mathbf{X}}^\lambda\|_F^2] / 4X_{\max}$, and we may choose $C_D = 4X_{\max}^2 / X_{\min}$. Incorporating this into Theorem II.1, we obtain that for any

$$\lambda \geq \left(1 + \frac{8X_{\max}^2}{3X_{\min}} \right) 2(\beta + 2) \cdot \log(n_1 \vee n_2), \quad (84)$$

the sparsity penalized ML estimate satisfies

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_S} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} &\leq \frac{1}{X_{\min}} \cdot \frac{128X_{\max}^3 \log m}{m} + \\ &\frac{12X_{\max}}{X_{\min}} \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{\|\mathbf{X}^* - \mathbf{X}\|_F^2}{n_1 n_2} + \left(\lambda + \frac{16X_{\max}^2 (\beta + 2) \log(n_1 \vee n_2)}{3} \right) \left(\frac{n_1 r + \|\mathbf{A}\|_0}{m} \right) \right\}. \end{aligned} \quad (85)$$

Now, the approximation error term in the oracle bound is in terms of a squared Frobenius norm, so the analysis for the case where λ is fixed to the specified value proceeds in an analogous manner to that in Appendix B for both the sparse and approximately sparse settings. We omit the details.

E. Proof of Corollary III.4

For \mathbf{X}^* as above and any $\mathbf{X} \in \mathcal{X}$, and using the model (30), it is easy to show that

$$D(p_{X_{i,j}^*} \| p_{X_{i,j}}) = F(X_{i,j}^*) \cdot \log \left(\frac{F(X_{i,j}^*)}{F(X_{i,j})} \right) + (1 - F(X_{i,j}^*)) \cdot \log \left(\frac{1 - F(X_{i,j}^*)}{1 - F(X_{i,j})} \right) \quad (86)$$

for any fixed $(i, j) \in S$. Now, we make use of two results that follow directly from lemmata established in [79]. The first lemma provides quadratic bounds on the KL divergence in terms of the Bernoulli parameters; its proof relies on a straightforward application of Taylor's theorem.

Lemma A.2 (from [79]). *Let p_π and $p_{\pi'}$ be Bernoulli pmf's with parameters $\pi, \pi' \in (0, 1)$. The KL divergences satisfy*

$$D(p_{\pi'} \| p_\pi), D(p_\pi \| p_{\pi'}) \leq \frac{1}{2} \left(\sup_{|t| \leq X_{\max}} \frac{1}{F(t)(1 - F(t))} \right) (\pi - \pi')^2. \quad (87)$$

The second lemma we utilize establishes a bound on the squared difference between Bernoulli parameters in terms of the squared difference of the underlying matrix elements; its proof is straightforward, and essentially entails establishing the Lipschitz continuity of F .

Lemma A.3 (from [79]). *Let $\pi = \pi(X)$ and $\pi' = \pi'(X')$ be Bernoulli parameters that are related to some underlying real-valued parameters X and X' via $\pi(X) = F(X)$ and $\pi'(X') = F(X')$, where $F(\cdot)$ is the cdf of a continuous random variable with density $f(\cdot)$. If $|X|, |X'| \leq X_{\max}$, then*

$$(\pi(X) - \pi'(X'))^2 \leq \left(\sup_{|t| \leq X_{\max}} f(t) \right)^2 (X - X')^2, \quad (88)$$

$$= \left(\sup_{|t| \leq X_{\max}} f^2(t) \right) (X - X')^2. \quad (89)$$

Together, these results allow us to claim here that for

$$c_{F, X_{\max}} \triangleq \left(\sup_{|t| \leq X_{\max}} \frac{1}{F(t)(1 - F(t))} \right) \cdot \left(\sup_{|t| \leq X_{\max}} f^2(t) \right). \quad (90)$$

we have

$$D(p_{X_{i,j}^*} \| p_{X_{i,j}}) \leq \frac{1}{2} \cdot c_{F, X_{\max}} (X_{i,j}^* - X_{i,j})^2. \quad (91)$$

It follows that we may take $C_D = 2c_{F, X_{\max}} X_{\max}^2$, and we have $D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) \leq (c_{F, X_{\max}}/2) \|\mathbf{X}^* - \mathbf{X}\|_F^2$.

We next obtain a (quadratic) lower bound on the negative log Hellinger affinity. To that end, we introduce the squared Hellinger distance between $p_{X_{i,j}^*}$ and $p_{X_{i,j}}$, denoted here by $H^2(p_{X_{i,j}^*}, p_{X_{i,j}})$ and given by

$$H^2(p_{X_{i,j}^*}, p_{X_{i,j}}) = \sum_{y \in \{0,1\}} \left(\sqrt{p_{X_{i,j}^*}(y)} - \sqrt{p_{X_{i,j}}(y)} \right)^2. \quad (92)$$

It is straightforward to see that $H^2(p_{X_{i,j}^*}, p_{X_{i,j}}) = 2(1 - A(p_{X_{i,j}^*}, p_{X_{i,j}}))$. Now, recall that the Hellinger affinity is always between 0 and 1, so using the fact that $\log(x) \leq x - 1$ for $x > 0$, we see directly that

$$H^2(p_{X_{i,j}^*}, p_{X_{i,j}}) \leq -2 \log A(p_{X_{i,j}^*}, p_{X_{i,j}}). \quad (93)$$

Now, a direct application of the result of [23, Lemma 2] derived for a similar subproblem to our problem here yields that for

$$c'_{F, X_{\max}} \triangleq \inf_{|t| \leq X_{\max}} \frac{f^2(t)}{F(t)(1-F(t))}, \quad (94)$$

we have that

$$H^2(p_{X_{i,j}^*}, p_{X_{i,j}}) \geq \frac{1}{8} c'_{F, X_{\max}} (X_{i,j}^* - X_{i,j})^2. \quad (95)$$

It follows that for any fixed $\mathbf{X} \in \mathcal{X}$, we have $-2 \log A(p_{\mathbf{X}^*}, p_{\mathbf{X}}) \geq (c'_{F, X_{\max}}/8) \|\mathbf{X}^* - \mathbf{X}\|_F^2$.

Incorporating all of the above into Theorem II.1 with

$$\lambda \geq 2(\beta + 2) \left(1 + \frac{4c_{F, X_{\max}} X_{\max}^2}{3} \right) \log(n_1 \vee n_2), \quad (96)$$

the sparsity penalized ML estimate satisfies the per-element mean-square error bound

$$\begin{aligned} \frac{\mathbb{E}_{\mathcal{S}, \mathbf{Y}_{\mathcal{S}}} \left[\|\mathbf{X}^* - \widehat{\mathbf{X}}\|_F^2 \right]}{n_1 n_2} &\leq \left(\frac{c_{F, X_{\max}}}{c'_{F, X_{\max}}} \right) \cdot \frac{128 X_{\max}^2 \log m}{m} + \\ &24 \left(\frac{c_{F, X_{\max}}}{c'_{F, X_{\max}}} \right) \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{\|\mathbf{X}^* - \mathbf{X}\|_F^2}{n_1 n_2} + \left(\frac{\lambda}{c_{F, X_{\max}}} + \frac{8 X_{\max}^2 (\beta + 2) \log(n_1 \vee n_2)}{3} \right) \frac{n_1 r + \|\mathbf{A}\|_0}{m} \right\}. \end{aligned} \quad (97)$$

Now, the approximation error term in the oracle bound is again in terms of a squared Frobenius norm, so the analysis for the case where λ is fixed to the specified value proceeds in an analogous manner to that in Appendix B for both the sparse and nearly sparse settings. We again omit the details.

F. Proof of Lemma A.1

Our estimation approach here is, at its essence, a constrained maximum likelihood method and our proof approach follows the general framework proposed in [29] (see also [30], [80], [81]) and utilized in [31]–[33]. Compared with these existing efforts, the main challenge in our analysis here arises because of the “missing data” paradigm, since we aim to establish consistency results that hold *globally* (at all locations of the unknown matrix) using observations obtained at only a subset of the locations. Our approach will be to identify conditions under which, for the purposes of our analysis, a set of sample locations is deemed “good,” in a manner to be made explicit below. The primary characteristic of good sets S of sample locations that we will leverage in our analysis is that they be such that KL divergences and (negative logarithms of) Hellinger affinities evaluated only at the locations in S be representative surrogates for the corresponding quantities were we to evaluate them at all $(i, j) \in [n_1] \times [n_2]$ (i.e., even at the unmeasured locations). Clearly, such conditions will inherently rely on certain properties of the matrices that we seek to estimate, somewhat analogously to how notions of incoherence facilitate matrix completion analyses under low rank matrix models. Here, we will see these conditions manifest not as properties of the singular vectors of the unknown matrix to be estimated as in existing matrix completion works, but instead, as conditions on the magnitude of the largest matrix entry.

Our approach will be as follows. First, we describe formally the notion of “good” sets of sample locations, and we show that sets of sample locations generated randomly according to an independent Bernoulli model are “good”

with high probability. Then, we establish error guarantees that hold conditionally on the event that the set of sample locations is “good.” Finally, we obtain our overall result using some simple conditioning arguments.

1) *“Good” Sample Set Characteristics:* We begin by characterizing, formally, the properties of certain sets of sample locations that will be useful for our analysis here. As above \mathbf{X}^* denotes the true (unknown) matrix that we aim to estimate, and \mathcal{X} is a countable set of candidate estimates \mathbf{X} , each with corresponding penalty $\text{pen}(\mathbf{X}) \geq 1$ chosen so the inequality (51) is satisfied. Also, recall that $X_{\max} > 0$ is a finite constant for which $\max_{i,j} |X_{i,j}^*| \leq X_{\max}/2$ and $\max_{\mathbf{X} \in \mathcal{X}} \max_{i,j} |X_{i,j}| \leq X_{\max}$. Finally, we let C_A and C_D be any upper bounds, respectively, on (twice) the negative log Hellinger affinities between $p_{X_{i,j}^*}$ and $p_{X_{i,j}}$, and the KL divergences of $p_{X_{i,j}}$ from $p_{X_{i,j}^*}$ that hold over all indices, and for all elements $\mathbf{X} \in \mathcal{X}$, so that

$$C_A \geq \max_{\mathbf{X} \in \mathcal{X}} \max_{i,j} -2 \log A(p_{X_{i,j}^*}, p_{X_{i,j}}) \quad (98)$$

and

$$C_D \geq \max_{\mathbf{X} \in \mathcal{X}} \max_{i,j} D(p_{X_{i,j}^*} \| p_{X_{i,j}}). \quad (99)$$

Note that the statement of Theorem II.1 only prescribed a condition on C_D ; our introduction of an additional constant C_A here is only to simplify the subsequent analysis. In the concluding steps of the proof we will claim that upon selecting a suitable C_D , one may always obtain a valid choice of C_A by taking $C_A = C_D$. This will enable us to eliminate the C_A terms that arise in our bound by bounding them in terms of the constant C_D .

Let $m \in [n_1 n_2]$ denote a nominal number of measurements, and let $\gamma = m/n_1 n_2 \in (0, 1]$ denote the corresponding nominal fraction of observed matrix elements. For this γ and any fixed $\delta \in (0, 1)$, we define the “good” set $\mathcal{G}_{\gamma, \delta} = \mathcal{G}_{\gamma, \delta}(\mathbf{X}^*, \mathcal{X})$ of possible sample location sets as

$$\mathcal{G}_{\gamma, \delta} \triangleq \left\{ S \subseteq [n_1] \times [n_2] : \bigcap_{\mathbf{X} \in \mathcal{X}} D(p_{\mathbf{X}_S^*} \| p_{\mathbf{X}_S}) \leq \frac{3\gamma}{2} D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + 2 \left(\frac{2C_D}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \log 2] \right. \\ \left. \cap \bigcap_{\mathbf{X} \in \mathcal{X}} (-2 \log A(p_{\mathbf{X}_S^*}, p_{\mathbf{X}_S})) \geq \frac{\gamma}{2} (-2 \log A(p_{\mathbf{X}^*}, p_{\mathbf{X}})) - 2 \left(\frac{2C_A}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \log 2] \right\}. \quad (100)$$

Directly certifying whether any fixed set S is an element of $\mathcal{G}_{\gamma, \delta}$ may be difficult in general. However, our observation model here assumes that the sample location set is generated randomly, according to an independent Bernoulli(γ) model, where each location is included in the set independently with probability $\gamma \in (0, 1]$. In this case, we have that random sample location sets S so generated satisfy $S \in \mathcal{G}_{\gamma, \delta}$ with high probability, as shown in the following lemma.

Lemma A.4. *Let \mathcal{X} be any countable collection of candidate estimates \mathbf{X} for \mathbf{X}^* , with corresponding penalties $\text{pen}(\mathbf{X})$ satisfying (51). For any fixed $\gamma \in (0, 1)$, let $S \subseteq [n_1] \times [n_2]$ be a random sample set generated according to the independent Bernoulli(γ) model. Then, for any $\delta \in (0, 1)$ we have $\Pr(S \notin \mathcal{G}_{\gamma, \delta}) \leq 2\delta$.*

Proof: Write $\{S \in \mathcal{G}_{\gamma, \delta}\} = \mathcal{E}_u \cap \mathcal{E}_l$, where

$$\mathcal{E}_u \triangleq \left\{ \bigcap_{\mathbf{X} \in \mathcal{X}} D(p_{\mathbf{X}_S^*} \| p_{\mathbf{X}_S}) \leq \frac{3\gamma}{2} D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + 2 \left(\frac{2C_D}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \log 2] \right\}, \quad (101)$$

and

$$\mathcal{E}_l \triangleq \left\{ \bigcap_{\mathbf{X} \in \mathcal{X}} (-2 \log A(p_{\mathbf{X}_S^*}, p_{\mathbf{X}_S})) \geq \frac{\gamma}{2} (-2 \log A(p_{\mathbf{X}^*}, p_{\mathbf{X}})) - 2 \left(\frac{2C_A}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \log 2] \right\}, \quad (102)$$

Then, by straight-forward union bounding, $\Pr(\mathcal{S} \notin \mathcal{G}_{\gamma, \delta}(\mathcal{X})) \leq \Pr(\mathcal{E}_u^c) + \Pr(\mathcal{E}_l^c)$. The proof of the lemma entails bounding each term on the right-hand side, in turn.

We focus first on bounding the probability of the complement of \mathcal{E}_u . To proceed, we will find it convenient to consider an alternative (but equivalent) representation of the sampling operator described explicitly in terms of a collection $\{B_{i,j}\}_{(i,j) \in [n_1] \times [n_2]}$ of independent Bernoulli(γ) random variables, so that $\mathcal{S} = \{(i, j) : B_{i,j} = 1\}$. On account of our assumption that the observations be conditionally independent given \mathcal{S} , we have that for any fixed $\mathbf{X} \in \mathcal{X}$,

$$D(p_{\mathbf{X}_S^*} \| p_{\mathbf{X}_S}) = \sum_{(i,j) \in \mathcal{S}} D(p_{X_{i,j}^*}, p_{X_{i,j}}) = \sum_{i,j} B_{i,j} \cdot D(p_{X_{i,j}^*}, p_{X_{i,j}}). \quad (103)$$

Thus, our analysis reduces to quantifying the concentration behavior of random sums of these forms. For this, we employ a powerful version of Bernstein's Inequality established by Craig [82] that, for our purposes, may be stated as follows: let $\{U_{i,j}\}$ be a collection of independent random variables indexed by (i, j) , each satisfying the moment condition that for some $h > 0$,

$$\mathbb{E} [|U_{i,j} - \mathbb{E}[U_{i,j}]|^k] \leq \frac{\text{var}(U_{i,j})}{2} k! h^{k-2},$$

for all integers $k \geq 2$. Then, for any $\tau > 0$ and $0 \leq \epsilon h \leq \theta < 1$, the probability that

$$\sum_{i,j} (U_{i,j} - \mathbb{E}[U_{i,j}]) \geq \frac{\tau}{\epsilon} + \frac{\epsilon \sum_{i,j} \text{var}(U_{i,j})}{2(1-\theta)} \quad (104)$$

is no larger than $e^{-\tau}$. A useful (and easy to verify) fact is that whenever $|U_{i,j} - \mathbb{E}[U_{i,j}]| \leq \beta$, the moment condition is satisfied by the choice $h = \beta/3$.

Now, fix $\mathbf{X} \in \mathcal{X}$, and let $U_{i,j}(\mathbf{X}) = B_{i,j} \cdot D(p_{X_{i,j}^*}, p_{X_{i,j}})$ and $\mathbb{E}[U_{i,j}(\mathbf{X})] = \gamma \cdot D(p_{X_{i,j}^*}, p_{X_{i,j}})$. Applying Craig's version of Bernstein's inequality with $\theta = 1/4$, $h = C_D/3$, and $\epsilon = \theta/h = 3/(4C_D)$, and using the fact that

$$\text{var}(U_{i,j}(\mathbf{X})) = \gamma(1-\gamma) \left(D(p_{X_{i,j}^*}, p_{X_{i,j}}) \right)^2 \leq \gamma \left(D(p_{X_{i,j}^*}, p_{X_{i,j}}) \right)^2 \quad (105)$$

we obtain that for any $\tau > 0$,

$$\Pr \left(\sum_{i,j} (B_{i,j} - \gamma) D(p_{X_{i,j}^*}, p_{X_{i,j}}) \geq \frac{4C_D\tau}{3} + \frac{\sum_{i,j} \gamma \left(D(p_{X_{i,j}^*}, p_{X_{i,j}}) \right)^2}{2C_D} \right) \leq e^{-\tau}. \quad (106)$$

Now, since $D(p_{X_{i,j}^*}, p_{X_{i,j}}) \leq C_D$ by definition, the above result ensures that for any $\tau > 0$,

$$\Pr \left(\sum_{i,j} (B_{i,j} - \gamma) D(p_{X_{i,j}^*}, p_{X_{i,j}}) \geq \frac{4C_D\tau}{3} + \frac{\gamma}{2} D(p_{\mathbf{X}^*}, p_{\mathbf{X}}) \right) \leq e^{-\tau}. \quad (107)$$

Letting $\delta = e^{-\tau}$ and simplifying a bit, we obtain that for any $\delta \in (0, 1)$,

$$\Pr \left(D(p_{\mathbf{X}_S^*}, p_{\mathbf{X}_S}) \geq \frac{4C_D \log(1/\delta)}{3} + \frac{3\gamma}{2} D(p_{\mathbf{X}^*}, p_{\mathbf{X}}) \right) \leq \delta. \quad (108)$$

Now, if for each $\mathbf{X} \in \mathcal{X}$ we let $\delta_{\mathbf{X}} = \delta \cdot 2^{-\text{pen}(\mathbf{X})}$, we can apply the union bound to obtain that

$$\Pr \left(\bigcup_{\mathbf{X} \in \mathcal{X}} D(p_{\mathbf{X}_S^*}, p_{\mathbf{X}_S}) \geq \frac{3\gamma}{2} D(p_{\mathbf{X}^*}, p_{\mathbf{X}}) + 2 \left(\frac{2C_D}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \cdot \log 2] \right) \leq \delta. \quad (109)$$

Following a similar approach for the affinity terms (with $U_{i,j}(\mathbf{X}) = -B_{i,j} \cdot (-2 \log A(p_{X_{i,j}^*}, p_{X_{i,j}}))$ for all i, j), we obtain that for any $\delta \in (0, 1)$,

$$\Pr \left(\bigcup_{\mathbf{X} \in \mathcal{X}} (-2 \log A(p_{\mathbf{X}_S^*}, p_{\mathbf{X}_S})) \leq \frac{\gamma}{2} (-2 \log A(p_{\mathbf{X}^*}, p_{\mathbf{X}})) - 2 \left(\frac{2C_A}{3} \right) [\log(1/\delta) + \text{pen}(\mathbf{X}) \cdot \log 2] \right) \leq \delta. \quad (110)$$

The overall result now follows by combining equations (109) and (110) using a union bound. \blacksquare

Next, we show how the implications of a sample set being ‘‘good’’ can be incorporated into the analysis of [29] to provide (conditional) error guarantees for completion tasks.

2) *A Conditional Error Guarantee:* Next, we establish the consistency of complexity penalized maximum likelihood estimators, conditionally on the event that the sample set S is a fixed set S , such that for fixed $\gamma \in (0, 1)$ and $\delta \in (0, 1)$, $S \in \mathcal{G}_{\gamma, \delta}$ (i.e., S is ‘‘good’’ according to the criteria outlined above). Our analysis then proceeds along the lines of the approach of [29], but with several key differences that arise because of our subsampling model.

As above, \mathcal{X} is a countable set of candidate estimates \mathbf{X} for \mathbf{X}^* , with corresponding penalties $\text{pen}(\mathbf{X})$ satisfying (51). Now, for any choice of μ satisfying $\mu \geq 1 + 2C_A/3$, we form an estimate $\hat{\mathbf{X}}^\mu = \hat{\mathbf{X}}^\mu(\mathbf{Y}_S)$ according to

$$\begin{aligned} \hat{\mathbf{X}}^\mu &= \arg \min_{\mathbf{X} \in \mathcal{X}} \{-\log p_{\mathbf{X}_S}(\mathbf{Y}_S) + 2\mu \cdot \text{pen}(\mathbf{X}) \log 2\} \\ &= \arg \max_{\mathbf{X} \in \mathcal{X}} \left\{ \sqrt{p_{\mathbf{X}_S}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\mathbf{X})} \right\}. \end{aligned} \quad (111)$$

By this choice, we have that for any $\mathbf{X} \in \mathcal{X}$,

$$\sqrt{p_{\hat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\hat{\mathbf{X}}^\mu)} \geq \sqrt{p_{\mathbf{X}_S}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\mathbf{X})}. \quad (112)$$

This implies that for the particular (deterministic, and μ -dependent) candidate

$$\tilde{\mathbf{X}}^\mu = \arg \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \cdot \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\}, \quad (113)$$

(whose specification will become clear shortly) we have

$$\frac{\sqrt{p_{\hat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\hat{\mathbf{X}}^\mu)}}{\sqrt{p_{\tilde{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\tilde{\mathbf{X}}^\mu)}} \geq 1. \quad (114)$$

Using this, along with some straight-forward algebraic manipulations, we have

$$\begin{aligned}
 -2 \log A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*}) &= 2 \log \left(\frac{1}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \right) \\
 &\leq 2 \log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu)}}{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \cdot 2^{-\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu)}} \cdot \frac{1}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \right) \\
 &= 2 \log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)}}{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)}} \cdot \frac{\sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}{\sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}} \cdot \frac{2^{-\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu)}}{2^{-\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu)}} \cdot \frac{1}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \right) \\
 &= \log \left(\frac{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \right) + 2\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 + 2 \log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)/p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu)} \right). \quad (115)
 \end{aligned}$$

At this point, we make our first use of the implications of the ‘‘good’’ sample set condition. In particular, since $S \in \mathcal{G}_{\gamma, \delta}$ and $\widehat{\mathbf{X}}^\mu \in \mathcal{X}$, we have that

$$-2 \log A(p_{\mathbf{X}_S^*}, p_{\widehat{\mathbf{X}}_S^\mu}) \geq \frac{\gamma}{2} (-2 \log A(p_{\mathbf{X}^*}, p_{\widehat{\mathbf{X}}^\mu})) - 2 \left(\frac{2C_A}{3} \right) \left[\log(1/\delta) + \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 \right]. \quad (116)$$

Incorporating this into (115), we have

$$\begin{aligned}
 \frac{\gamma}{2} (-2 \log A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}^*})) &\leq \log \left(\frac{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} \right) + 2\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 + 2 \left(\frac{2C_A}{3} \right) \log(1/\delta) \\
 &\quad + 2 \log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)/p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\widehat{\mathbf{X}}^\mu)} \right). \quad (117)
 \end{aligned}$$

Now, we take expectations (formally, with respect to the conditional distribution of \mathbf{Y}_S given $\{S = S, S \in \mathcal{G}_{\gamma, \delta}\}$) on both sides to obtain that

$$\begin{aligned}
 \frac{\gamma}{2} \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}^*}) \mid S = S, S \in \mathcal{G}_{\gamma, \delta} \right] &\leq D(p_{\mathbf{X}_S^*} \| p_{\widehat{\mathbf{X}}_S^\mu}) + 2\mu \cdot \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 + 2 \left(\frac{2C_A}{3} \right) \log(1/\delta) \\
 &\quad + 2 \mathbb{E} \left[\log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)/p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\widehat{\mathbf{X}}^\mu)} \right) \mid S = S, S \in \mathcal{G}_{\gamma, \delta} \right]. \quad (118)
 \end{aligned}$$

Using again the implications of $S \in \mathcal{G}_{\gamma, \delta}$, that

$$D(p_{\mathbf{X}_S^*} \| p_{\widehat{\mathbf{X}}_S^\mu}) \leq \frac{3\gamma}{2} D(p_{\mathbf{X}^*} \| p_{\widehat{\mathbf{X}}^\mu}) + 2 \left(\frac{2C_D}{3} \right) \left[\log(1/\delta) + \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 \right] \quad (119)$$

since $\widehat{\mathbf{X}}^\mu \in \mathcal{X}$, we have that

$$\begin{aligned}
 \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}^*}) \mid S = S, S \in \mathcal{G}_{\gamma, \delta} \right] &\leq \\
 &3D(p_{\mathbf{X}^*} \| p_{\widehat{\mathbf{X}}^\mu}) + \frac{4}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 + \frac{4}{\gamma} \left(\frac{2(C_A + C_D)}{3} \right) \log(1/\delta) \\
 &+ \frac{4}{\gamma} \mathbb{E} \left[\log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)/p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\widehat{\mathbf{X}}^\mu)} \right) \mid S = S, S \in \mathcal{G}_{\gamma, \delta} \right]. \quad (120)
 \end{aligned}$$

Turning our attention to the last term on the right-hand side, we have that

$$\begin{aligned}
 & \mathbb{E} \left[\log \left(\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} / \sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\widehat{\mathbf{X}}^\mu)} \right) \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \\
 & \stackrel{(a)}{\leq} \log \left(\mathbb{E} \left[\frac{\sqrt{p_{\widehat{\mathbf{X}}_S^\mu}(\mathbf{Y}_S)} / \sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}}{A(p_{\widehat{\mathbf{X}}_S^\mu}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\widehat{\mathbf{X}}^\mu)} \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \right) \\
 & \stackrel{(b)}{\leq} \log \left(\mathbb{E} \left[\sum_{\mathbf{X} \in \mathcal{X}} \frac{\sqrt{p_{\mathbf{X}_S}(\mathbf{Y}_S)} / \sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}}{A(p_{\mathbf{X}_S}, p_{\mathbf{X}_S^*})} \cdot 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\mathbf{X})} \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \right) \\
 & = \log \left(\sum_{\mathbf{X} \in \mathcal{X}} 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\mathbf{X})} \mathbb{E} \left[\frac{\sqrt{p_{\mathbf{X}_S}(\mathbf{Y}_S)} / \sqrt{p_{\mathbf{X}_S^*}(\mathbf{Y}_S)}}}{A(p_{\mathbf{X}_S}, p_{\mathbf{X}_S^*})} \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \right) \\
 & \stackrel{(c)}{=} \log \left(\sum_{\mathbf{X} \in \mathcal{X}} 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\mathbf{X})} \right). \tag{121}
 \end{aligned}$$

In the above, (a) follows from Jensen's Inequality, (b) from the facts that $\widehat{\mathbf{X}}^\mu \in \mathcal{X}$ and each term in the sum is non-negative, and (c) from the definition of the Hellinger affinity. Now, because $\text{pen}(\mathbf{X}) \geq 1$ and $\mu \geq 1 + 2C_A/3$ we have that

$$\sum_{\mathbf{X} \in \mathcal{X}} 2^{-(\mu - \frac{2C_A}{3}) \text{pen}(\mathbf{X})} \leq \sum_{\mathbf{X} \in \mathcal{X}} 2^{-\text{pen}(\mathbf{X})} \leq 1. \tag{122}$$

Thus, since the expectation term on the right-hand side of (120) is not positive, we can disregard it in the upper bound to obtain that

$$\begin{aligned}
 & \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \\
 & \leq 3D(p_{\mathbf{X}^*} \| p_{\widehat{\mathbf{X}}^\mu}) + \frac{6}{\gamma} \left(\lambda + \frac{2C_D}{3} \right) \text{pen}(\widehat{\mathbf{X}}^\mu) \log 2 + \frac{4}{\gamma} \left(\frac{2(C_A + C_D)}{3} \right) \log(1/\delta), \tag{123}
 \end{aligned}$$

where we have also inflated (slightly) the leading constant on the second term on the right-hand side to simplify subsequent analysis. Now, recalling the definition of $\widehat{\mathbf{X}}^\mu$, we can state the result equivalently as an oracle bound, as

$$\begin{aligned}
 & \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \\
 & \leq 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\} + \frac{4}{\gamma} \left(\frac{2(C_A + C_D)}{3} \right) \log(1/\delta). \tag{124}
 \end{aligned}$$

3) *Putting the Pieces Together:* The last steps of the analysis entail straightforward applications of conditioning arguments, along with the use of a well-known (and easy to verify) information inequality. First, note that

$$\begin{aligned}
 & \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \middle| \mathcal{S} \in \mathcal{G}_{\gamma, \delta} \right] \\
 & = \sum_{S \in [n_1] \times [n_2]} \mathbb{E} \left[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \middle| \mathcal{S} = S, S \in \mathcal{G}_{\gamma, \delta} \right] \cdot \Pr(S = S | \mathcal{S} \in \mathcal{G}_{\gamma, \delta}) \\
 & \leq 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\} + \frac{4}{\gamma} \left(\frac{2(C_A + C_D)}{3} \right) \log(1/\delta), \tag{125}
 \end{aligned}$$

where the last step follows from using the bound in (124) and bringing that term outside of the sum since it does not depend on S , and using the fact that the conditional probability mass function $\Pr(S = S | S \in \mathcal{G})$ sums to 1. Now, using the fact that

$$\begin{aligned} \mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*})] &= \\ \mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \mid S \in \mathcal{G}_{\gamma, \delta}] \cdot \Pr(S \in \mathcal{G}_{\gamma, \delta}) &+ \mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \mid S \notin \mathcal{G}_{\gamma, \delta}] \cdot \Pr(S \notin \mathcal{G}_{\gamma, \delta}), \end{aligned} \quad (126)$$

where the expectation on the left-hand side is with respect to the joint distribution of \mathbf{Y}_S and S , we obtain that

$$\begin{aligned} \mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*})] &\leq \\ 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\} &+ \frac{4}{\gamma} \left(\frac{2(C_A + C_D)}{3} \right) \log(1/\delta) + 2\delta \cdot n_1 n_2 C_A, \end{aligned} \quad (127)$$

where we use the trivial upper bound $\mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*}) \mid S \notin \mathcal{G}_{\gamma, \delta}] \leq n_1 n_2 C_A$. Now, since the result holds for any choice of $\delta \in (0, 1)$, we can choose δ judiciously to “balance” the last two terms. The particular choice $\delta = m^{-1} = (\gamma n_1 n_2)^{-1}$ yields

$$\begin{aligned} \mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*})] &\leq \\ 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\} &+ \frac{8(C_A + C_D) \log m}{3\gamma} + \frac{2C_A}{\gamma}, \end{aligned} \quad (128)$$

which implies the simpler (but slightly looser) bound

$$\mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*})] \leq 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ D(p_{\mathbf{X}^*} \| p_{\mathbf{X}}) + \frac{2}{\gamma} \left(\mu + \frac{2C_D}{3} \right) \text{pen}(\mathbf{X}) \log 2 \right\} + \frac{4(C_A + C_D) \log m}{\gamma}. \quad (129)$$

Finally, we make use of the fact that for each i, j , we have $-2 \log A(p_{X_{i,j}^*}, p_{X_{i,j}}) \leq D(p_{X_{i,j}^*}, p_{X_{i,j}})$, which is readily verified with one application of Jensen’s inequality. It follows that upon identifying a suitable C_D , we may always take $C_A = C_D$. Thus, it is sufficient to choose $\mu > 1 + 2C_D/3$ when forming our complexity regularized maximum likelihood estimator. We conclude that the error of any estimator formed using an appropriate regularization parameter μ satisfies

$$\frac{\mathbb{E}[-2 \log A(p_{\widehat{\mathbf{X}}^\mu}, p_{\mathbf{X}^*})]}{n_1 n_2} \leq 3 \cdot \min_{\mathbf{X} \in \mathcal{X}} \left\{ \frac{D(p_{\mathbf{X}^*} \| p_{\mathbf{X}})}{n_1 n_2} + \left(\mu + \frac{2C_D}{3} \right) \frac{\text{pen}(\mathbf{X}) 2 \log 2}{m} \right\} + \frac{8C_D \log(m)}{m}, \quad (130)$$

where we have divided both sides by $n_1 n_2$ and used the fact that $m = \gamma n_1 n_2$. Finally, making the substitution $\xi = 2\mu \log 2$ yields the stated version of the result.

REFERENCES

- [1] R. M. Bell and Y. Koren, “Lessons from the netflix prize challenge,” *ACM SIGKDD Explorations Newsletter*, vol. 9, no. 2, pp. 75–79, 2007.
- [2] E. J. Candès and B. Recht, “Exact matrix completion via convex optimization,” *Foundations of Computational Mathematics*, vol. 9, no. 6, pp. 717–772, 2009.
- [3] E. J. Candès and T. Tao, “The power of convex relaxation: Near-optimal matrix completion,” *IEEE Transactions on Information Theory*, vol. 56, no. 5, pp. 2053–2080, 2010.
- [4] R. H. Keshavan, A. Montanari, and S. Oh, “Matrix completion from a few entries,” *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2980–2998, 2010.

- [5] B. Recht, "A simpler approach to matrix completion," *The Journal of Machine Learning Research*, vol. 12, pp. 3413–3430, 2011.
- [6] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [7] W. Dai and O. Milenkovic, "SET: An algorithm for consistent matrix completion," in *Proc. IEEE International Conference on Acoustics Speech and Signal Processing*, 2010, pp. 3646–3649.
- [8] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," *Mathematical Programming*, vol. 128, no. 1-2, pp. 321–353, 2011.
- [9] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Transactions on Image Processing*, vol. 17, no. 1, pp. 53–69, 2008.
- [10] A. S. Lan, A. E. Waters, C. Studer, and R. G. Baraniuk, "Sparse factor analysis for learning and content analytics," *arXiv preprint arXiv:1303.5685*, 2013.
- [11] R. Agrawal, J. Gehrke, D. Gunopulos, and P. Raghavan, "Automatic subspace clustering of high dimensional data for data mining applications," in *ACM SIGMOD Int. Conf. Management of Data*, 1998, pp. 94–105.
- [12] P. Tseng, "Nearest q -flat to m points," *Journal of Optimization Theory and Applications*, vol. 105, no. 1, pp. 249–252, 2000.
- [13] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (GPCA)," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 12, pp. 1945–1959, 2005.
- [14] M. Soltanolkotabi and E. J. Candès, "A geometric analysis of subspace clustering with outliers," *The Annals of Statistics*, vol. 40, no. 4, pp. 2195–2238, 2012.
- [15] E. Elhamifar and R. Vidal, "Sparse subspace clustering: Algorithm, theory, and applications," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 11, pp. 2765–2781, 2013.
- [16] M. Soltanolkotabi, E. Elhamifar, and E. J. Candès, "Robust subspace clustering," *arXiv preprint arXiv:1301.2603*, 2013.
- [17] R. H. Keshavan, A. Montanari, and S. Oh, "Matrix completion from noisy entries," *Journal of Machine Learning Research*, vol. 11, pp. 2057–2078, 2010.
- [18] K. Lee and Y. Bresler, "ADMIRA: Atomic decomposition for minimum rank approximation," *IEEE Transactions on Information Theory*, vol. 56, no. 9, pp. 4402–4416, 2010.
- [19] E. J. Candès and Y. Plan, "Matrix completion with noise," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 925–936, 2010.
- [20] V. Koltchinskii, K. Lounici, and A. B. Tsybakov, "Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion," *The Annals of Statistics*, vol. 39, no. 5, pp. 2302–2329, 2011.
- [21] S. Negahban and M. J. Wainwright, "Restricted strong convexity and weighted matrix completion: Optimal bounds with noise," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 1665–1697, 2012.
- [22] N. Srebro, N. Alon, and T. S. Jaakkola, "Generalization error bounds for collaborative prediction with low-rank matrices," in *Advances In Neural Information Processing Systems*, 2004, pp. 1321–1328.
- [23] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wootters, "1-bit matrix completion," *arXiv preprint arXiv:1209.3672*, 2012.
- [24] Y. Plan, R. Vershynin, and E. Yudovina, "High-dimensional estimation with geometric constraints," *arXiv preprint arXiv:1404.3749*, 2014.
- [25] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?," *Vision Research*, vol. 37, pp. 3311–3325, 1997.
- [26] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Proc.*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [27] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *Proc. ICML*, 2009.
- [28] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [29] J. Q. Li, *Estimation of Mixture Models*, Ph.D. thesis, Dept. of Statistics, Yale University, New Haven, CT, 1999.
- [30] A. Barron, L. Birgé, and P. Massart, "Risk bounds for model selection via penalization," *Probability theory and related fields*, vol. 113, no. 3, pp. 301–413, 1999.
- [31] E. D. Kolaczyk and R. D. Nowak, "Multiscale likelihood analysis and complexity penalized estimation," *Annals of Statistics*, pp. 500–527, 2004.

- [32] R. M. Willett and R. D. Nowak, "Multiscale Poisson intensity and density estimation," *IEEE Transactions on Information Theory*, vol. 53, no. 9, pp. 3171–3187, 2007.
- [33] M. Raginsky, R. Willett, Z. T. Harmany, and R. F. Marcia, "Compressed sensing performance bounds under Poisson noise," *IEEE Transactions on Signal Processing*, vol. 58, no. 8, pp. 3990–4002, 2010.
- [34] A. Soni and J. Haupt, "Estimation error guarantees for Poisson denoising with sparse and structured dictionary models," in *Proc. International Symposium on Information Theory*, 2014, to appear.
- [35] J. D. Haupt, N. D. Sidiropoulos, and G. B. Giannakis, "Sparse dictionary learning from 1-bit data," in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2014.
- [36] A. Rakotomamonjy, "Applying alternating direction method of multipliers for constrained dictionary learning," *Neurocomputing*, vol. 106, pp. 126–136, 2013.
- [37] Z. Xing, M. Zhou, A. Castrodad, G. Sapiro, and L. Carin, "Dictionary learning for noisy and incomplete hyperspectral images," *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 33–56, 2012.
- [38] M. Zhou, H. Chen, J. Paisley, L. Ren, L. Li, Z. Xing, D. Dunson, G. Sapiro, and L. Carin, "Nonparametric bayesian dictionary learning for analysis of noisy and incomplete images," *IEEE Transactions on Image Processing*, vol. 21, no. 1, pp. 130–144, 2012.
- [39] M. Udell, C. Horn, R. Zadeh, and S. Boyd, "Generalized low rank models," *arXiv preprint arXiv:1410.0342*, 2014.
- [40] D. Donoho and V. Stodden, "When does non-negative matrix factorization give a correct decomposition into parts?," in *Advances in Neural Information Processing Systems*, 2003.
- [41] S. Arora, R. Ge, R. Kannan, and A. Moitra, "Computing a nonnegative matrix factorization—provably," in *Proc. ACM Symp. on Theory of Computing*, 2012, pp. 145–162.
- [42] E. Esser, M. Moller, S. Osher, G. Sapiro, and J. Xin, "A convex model for nonnegative matrix factorization and dimensionality reduction on physical space," *IEEE Transactions on Image Processing*, vol. 21, no. 7, pp. 3239–3252, 2012.
- [43] B. Recht, C. Re, J. Tropp, and V. Bittorf, "Factoring nonnegative matrices with linear programs," in *Advances in Neural Information Processing Systems*, 2012, pp. 1214–1222.
- [44] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear algebra and its applications*, vol. 416, no. 1, pp. 48–67, 2006.
- [45] R. Gribonval and K. Schnass, "Dictionary identification – Sparse matrix-factorization via l_1 minimization," *IEEE Transactions on Information Theory*, vol. 56, no. 7, pp. 3523–3539, 2010.
- [46] Q. Geng, H. Wang, and J. Wright, "On the local correctness of ℓ^1 minimization for dictionary learning," *Submitted*, 2011, online at: arxiv.org/abs/1101.5672.
- [47] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*, 2013, pp. 3087–3090.
- [48] K. Schnass, "On the identifiability of overcomplete dictionaries via the minimisation principle underlying K-SVD," *Submitted*, 2013, online at: arxiv.org/abs/1301.3375.
- [49] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning sparsely used overcomplete dictionaries via alternating minimization," *arXiv preprint arXiv:1310.7991*, 2013.
- [50] R. Jenatton, R. Gribonval, and F. Bach, "Local stability and robustness of sparse dictionary learning in the presence of noise," *Submitted*, 2012, online at: arxiv.org/abs/1210.0685.
- [51] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *Journal of computational and graphical statistics*, vol. 15, no. 2, pp. 265–286, 2006.
- [52] B. Moghaddam, Y. Weiss, and S. Avidan, "Spectral bounds for sparse PCA: Exact and greedy algorithms," in *Advances in Neural Information Processing Systems*, 2005, pp. 915–922.
- [53] A. d'Aspremont, L. El Ghaoui, M. I. Jordan, and G. R. G. Lanckriet, "A direct formulation for sparse PCA using semidefinite programming," *SIAM review*, vol. 49, no. 3, pp. 434–448, 2007.
- [54] H. Shen and J. Z. Huang, "Sparse principal component analysis via regularized low rank matrix approximation," *Journal of Multivariate Analysis*, vol. 99, no. 6, pp. 1015–1034, 2008.
- [55] Y. Zhang, A. d'Aspremont, and L. El Ghaoui, "Sparse PCA: Convex relaxations, algorithms and applications," in *Handbook on Semidefinite, Conic and Polynomial Optimization*, pp. 915–940. Springer, 2012.

- [56] V. Q. Vu and J. Lei, "Minimax rates of estimation for sparse PCA in high dimensions," in *International Conference on Artificial Intelligence and Statistics*, 2012, pp. 1278–1286.
- [57] K. Lounici, "Sparse principal component analysis with missing observations," in *High Dimensional Probability VI*, pp. 327–356. Springer, 2013.
- [58] A. Gruber and Y. Weiss, "Multibody factorization with uncertainty and missing data using the EM algorithm," in *Proc. Computer Vision and Pattern Recognition*, 2004.
- [59] R. Vidal, R. Tron, and R. Hartley, "Multiframe motion segmentation with missing data using PowerFactorization and GPCA," *International Journal of Computer Vision*, vol. 79, no. 1, pp. 85–105, 2008.
- [60] B. Eriksson, L. Balzano, and R. Nowak, "High-rank matrix completion and subspace clustering with missing data," *arXiv preprint arXiv:1112.5629*, 2011.
- [61] A. Singh, A. Krishnamurthy, S. Balakrishnan, and M. Xu, "Completion of high-rank ultrametric matrices using selective entries," in *Proc. IEEE International Conference on Signal Processing and Communications*, 2012, pp. 1–5.
- [62] R. G. Baraniuk, V. Cevher, and M. B. Wakin, "Low-dimensional models for dimensionality reduction and signal recovery: A geometric perspective," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 959–971, 2010.
- [63] E. J. Candès and T. Tao, "The Dantzig selector: Statistical estimation when p is much larger than n ," *The Annals of Statistics*, pp. 2313–2351, 2007.
- [64] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Transactions on Information Theory*, vol. 52, no. 9, pp. 4036–4048, 2006.
- [65] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [66] F. J. Anscombe, "The transformation of Poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3-4, pp. 246–254, 1948.
- [67] X. Jiang, G. Raskutti, and R. Willett, "Minimax optimal rates for Poisson inverse problems with physical constraints," *arXiv preprint arXiv:1403.6532*, 2014.
- [68] Z.-Q. Luo, "Universal decentralized estimation in a bandwidth constrained sensor network," *IEEE Transactions on Information Theory*, vol. 51, no. 6, pp. 2210–2219, 2005.
- [69] A. Ribeiro and G. B. Giannakis, "Bandwidth-constrained distributed estimation for wireless sensor networks-part i: Gaussian case," *IEEE Transactions on Signal Processing*, vol. 54, no. 3, pp. 1131–1143, 2006.
- [70] P. McCullagh and J. A. Nelder, *Generalized linear models*, Chapman & Hall, 1989.
- [71] Z. Lu, "Iterative hard thresholding methods for ℓ_0 regularized convex cone programming," *Mathematical Programming*, pp. 1–30, 2012.
- [72] S. van de Geer, *Empirical Processes in M-estimation*, vol. 105, Cambridge University Press, 2000.
- [73] V. Koltchinskii, *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, vol. 2033, Springer, 2011.
- [74] S. Boucheron, G. Lugosi, and P. Massart, *Concentration inequalities: A nonasymptotic theory of independence*, Oxford University Press, 2013.
- [75] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM Journal on Imaging Sciences*, vol. 2, no. 1, pp. 183–202, 2009.
- [76] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational Mathematics*, vol. 12, no. 6, pp. 805–849, 2012.
- [77] E. Richard, G. Obozinski, and J.-P. Vert, "Tight convex relaxations for sparse matrix factorization," *arXiv preprint arXiv:1407.5158*, 2014.
- [78] T. M. Cover and J. A. Thomas, *Elements of information theory*, John Wiley & Sons, 2006.
- [79] J. D. Haupt, N. D. Sidiropoulos, and G. B. Giannakis, "Sparse dictionary learning from 1-bit data," in *Proc. International Conference on Acoustics, Speech and Signal Processing*, 2014.
- [80] J. Q. Li and A. R. Barron, "Mixture density estimation," in *Advances in Neural Information Processing Systems*, 1999.
- [81] P. D. Grünwald, *The minimum description length principle*, MIT press, 2007.
- [82] C. C. Craig, "On the Tchebycheff inequality of Bernstein," *The Annals of Mathematical Statistics*, vol. 4, no. 2, pp. 94–102, 1933.