A Channel Coding Perspective of Collaborative Filtering

S. T. Aditya

Department of Electrical Engineering

Indian Institute of Technology Bombay

Mumbai, India

Email: staditya@ee.iitb.ac.in

Onkar Dabeer

School of Technology and Computer Science

Tata Institute of Fundamental Research

Mumbai, India

Email: onkar@tcs.tifr.res.in

Bikash Kumar Dey

Department of Electrical Engineering

Indian Institute of Technology Bombay

Mumbai, India

Email: bikash@ee.iitb.ac.in

Abstract

We consider the problem of collaborative filtering from a channel coding perspective. We model the underlying rating matrix as a finite alphabet matrix with block constant structure. The observations are obtained from this underlying matrix through a discrete memoryless channel with a noisy part representing noisy user behavior and an erasure part representing missing data. Moreover, the clusters over which the underlying matrix is constant are *unknown*. We establish a sharp threshold result for this model: if

the largest cluster size is smaller than $C_1 \log(mn)$ (where the rating matrix is of size $m \times n$), then the underlying matrix cannot be recovered with any estimator, but if the smallest cluster size is larger than $C_2 \log(mn)$, then we show a polynomial time estimator with diminishing probability of error. In the case of uniform cluster size, not only the order of the threshold, but also the constant is identified.

I. INTRODUCTION

As new content mushrooms at a brisk pace, finding relevant information is increasingly a challenge. Consequently, recommendation systems are commonly being used to assist users: Amazon recommends books, Netflix recommends movies, LinkedIn recommends professional contacts, Google recommends webpages for a given query, etc. Such recommendation systems exploit various aspects to make suggestions: popularity amongst peers, similarity of content, available user-item ratings, etc. This paper is about collaborative filtering using the rating matrix: we are interested in making recommendations using only available ratings given by users to the items they have experienced. In a practical system, such a rating based collaborative filter is typically complemented by content-based analysis specific to the data.

There is vast literature on recommendation systems and collaborative filtering; see for example the special issue [9] and the survey paper [3]. Given the massive datasets and the lack of good statistical model of user behavior, the dominant stream of work has been to propose methods and demonstrate their scalability on real data sets. However, recently the Netflix Prize [1] has popularized the problem to other research communities and several researchers have started exploring *provably* good methods. This paper falls in the latter category: we deal with fundamental limits of collaborative filters. In the remainder of this section, we first discuss related models and results, and then outline our model and results.

A. Related Work

The Netflix data consists of rating matrix where the rows correspond to movies and the columns correspond to users. Only a small fraction of the entries are known and the goal is to estimate the missing entries, that is, this is a matrix completion problem. Several algorithms have been proposed and tested on this data set; see for example [13]. Mathematically, without any further restriction, this is an ill-posed problem. Motivated by this, some authors have recently considered the matrix completion problem under the restriction of low-rank matrices. (This problem also arises in other contexts such as location estimation in sensor networks.) This problem has attracted much attention, and in the past year a number of results have been reported. In [5], using nuclear norm minimization proposed in [16], an upper bound on the number of samples needed for recovery asymptotically is derived in terms of the

size and rank of the matrix. In [6], a lower bound is established on the number of samples needed by any algorithm. The order of this lower bound is shown to be achievable in [12]. In [14], the problem of matrix recovery from linear measurements (of which sampling is a special case) is considered and a new algorithm is proposed. In [4], the problem of matrix completion under bounded noise is considered. A semi-definite programming based algorithm is proposed and shown to have recovery error proportional to the noise magnitude.

In this paper, we take an alternative channel coding viewpoint of the problem. Our results differ from the above works in several aspects outlined below.

- We consider finite alphabet for the ratings and a different model for the rating matrix based on row and column clusters.
- We consider noisy user behavior, and our goal is not to complete the missing entries, but to estimate an underlying "block constant" matrix (in the limit as the matrix size grows).
- Since we consider a finite alphabet, even in the presence of noise, error free recovery is asymptotically feasible. Hence, unlike [4], which considers real-valued matrices, we do not allow any distortion.

We next outline our model and results.

B. Summary of Our Model and Results

We consider a finite alphabet for the ratings. In this section, we briefly outline our model and results without any mathematical details; the details can be found in subsequent sections.

To motivate our model, consider an ideal situation where every user rates every item without any noise. In this ideal scenario, it is reasonable to expect that similar users rate similar items by the same value. We therefore assume that the users (items) are clustered into groups of similar users (items, respectively). The rating matrix in this ideal situation (say $\mathbf{X}$ with size $m \times n$) is then a block constant matrix (where the blocks correspond to cartesian product of row and column clusters). The observations are obtained from $\mathbf{X}$ by passing its entries through a discrete memoryless channel (DMC) consisting of an erasure channel modeling missing data and a noisy DMC representing noisy user behavior. Moreover, the row and column clusters are *unknown*. The goal is to make recommendations by estimating $\mathbf{X}$ based on the observations. The performance metric we use is the probability of block error: we make an error if any of the entries in the estimate is erroneous. Our goal is to identify conditions under which error free recovery is possible in the limit as the matrix size grows large. Thus we view the recommendation system problem as a channel coding problem.

The cluster sizes in our model represent the resolution: the larger the cluster, the smaller are the degrees of freedom (or rate of the channel code). If the channel is more noisy and the erasures are high, then we can only support a small number of codewords. The challenge is to find the exact order. For our model, we show that if the largest cluster size (defined precisely in Section III) is smaller than $C_1 \log(mn)$, where C_1 is a constant dependent on the channel parameters, then for any estimator the probability of error approaches one. On the other hand, if the smallest cluster size (defined precisely in Section III) is larger than $C_2 \log(mn)$, where C_2 is a constant dependent on the channel parameters, then we give a polynomial time algorithm that has diminishing probability of error. Thus we identify the order of the threshold exactly. In the case of uniform cluster size, the constants C_1 and C_2 are identical and thus in this special case, even the constant is identified precisely. Moreover, for the special case of binary ratings and uniform cluster size, the algorithm used to show the achievability part does nor depend on the cluster size, erasure parameter, and needs knowledge of a worst case parameter for the noisy part of the channel. These results are obtained by averaging over $\mathbf{X}$ (as per the probability law specified in Section II).

The achievability part of our result is shown by first clustering the rows and columns, and then estimating the matrix entries assuming that the clustering is correct. The clustering is done by computing a normalized Hamming metric for every pair of rows and comparing with a threshold to determine if the rows are in same cluster or not. The converse is proved by considering the case when the clusters are known exactly. Our results for the average case show that the threshold is determined by the problem of estimating entries, and relatively, clustering is an easier task (see Figure 1 for an illustration).

C. Organization of the Paper

The precise model for $\mathbf{X}$ and the observations is stated in Section II. The case of uniform cluster size and binary ratings leads to sharper bounds and results. Hence results for this case are given in Section III. The case of general alphabets and non-uniform cluster sizes is considered in Section IV. The conclusion is given in Section V, while all the proofs are collected together in Section VI.

D. Notation

All the logarithms are to the natural base unless specified otherwise. $D(\mu\|\nu)$ denotes the KL divergence ([8]) between probability mass functions μ and ν. By $T = \Omega(f(n))$ we mean that for n large enough, $T \geq \text{constant} \cdot f(n)$. By $1(A)$ we denote the indicator variable, which is 1 if A is true and 0 otherwise.
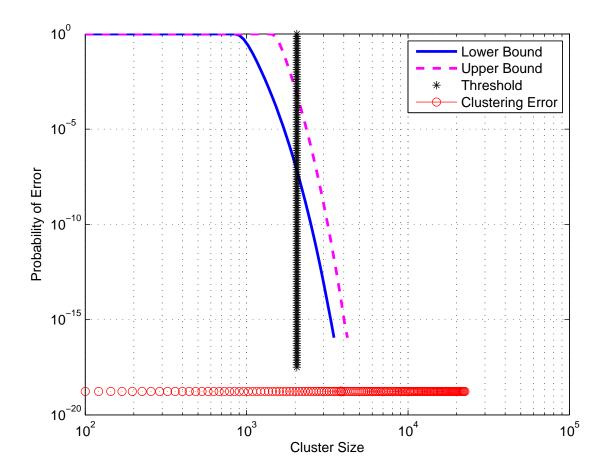
Fig. 1. The figure shows lower and upper bounds for the probability of error under known clustering (Theorem 2), the asymptotic cluster size threshold from Theorem 1, and an upper bound on the clustering error (Theorem 3) for the case $m = n = 10^6$, erasure probability $\epsilon = 0.9$, and binary symmetric channel with error $p = 0.25$. The threshold in the clustering algorithm is chosen to be $d_0 = (2p(1-p) + 1)/3 = 0.4583$.

II. MODEL AND ASSUMPTIONS

The main elements of our model are a block constant ensemble of rating matrices (whose blocks of constancy are not known) and an observation matrix obtained from the underlying rating matrix via a noisy channel and erasures. The noise in the observations represents the inherent noise in user-item ratings as well as the error in our model. The erasures denote missing entries. To be more precise, suppose $\mathbf{X}$ is the unknown $m \times n$ rating matrix with entries from a finite alphabet, where n is the number of buyers and m is the number of items. Let $\mathcal{A} = \{A_i\}_{i=1}^r$ and $\mathcal{B} = \{B_j\}_{j=1}^t$ be partitions of $[1 : m]$ and $[1 : n]$ respectively. We call the sets $A_i \times B_j$ clusters and we call A_i's (B_j's) the row (column) clusters. We

denote the corresponding row and column cluster sizes by m_i and n_j, and the number of row clusters and the number of column clusters by r and t respectively. Thus $\sum_{i=1}^{r} m_i = m$, $\sum_{j=1}^{t} n_j = n$.

We state our results under two sets of conditions - the set of conditions A1)-A4) and B1)-B3) below. Conditions A1)-A4) are a special case of conditions B1)-B3). The results under A1)-A4) are sharper and illustrate the important concepts more easily. Hence they are stated separately. We begin by stating and discussing A1)-A4) first and then we state B1)-B3). (A few additional conditions needed in the results are stated at appropriate places.)

Conditions A1)-A4): The conditions A1)-A4) below correspond to binary rating matrix with equal size clusters and uniform probability of sampling entries.

A1) The entries of $\mathbf{X}$ are from $\{0, 1\}$.

A2) The row (column) clusters are of equal size: $m_i = m_0$, $n_i = n_0$ for all i.

A3) $\mathbf{X}$ is constant over the cluster $A_i \times B_j$ and the entries are i.i.d. Bernoulli(1/2) across the clusters.

A4) The observed data $\mathbf{Y} \in \{0, 1, e\}$ (e denotes erasure) is obtained by passing the entries of $\mathbf{X}$ through the cascade of a binary symmetric channel (BSC) with probability of error p and an erasure channel with erasure probability ϵ.

The cluster sizes are representative of the *resolution* of $\mathbf{X}$ - large cluster sizes correspond to a coarse structure with fewer degrees of freedom in choosing $\mathbf{X}$, while small cluster size corresponds to a fine structure. Condition A2) suggests that we can think of the cluster size $m_0 n_0$ as representative of the resolution of $\mathbf{X}$ and it plays a central role in our results. If we think of all permissible $\mathbf{X}$ as a channel code, then a higher $m_0 n_0$ corresponds to a smaller rate code. However, in order to interpret $m_0 n_0$ precisely, we also need to take into account condition A3). When the entries of the cluster are filled with i.i.d. Bernoulli(1/2) random variables as per A3), it is likely that rows in two clusters turn out to be the same, and hence these two row clusters can be merged to form a single bigger cluster. The following lemma shows that if the number of clusters is $\Omega(\log(n))$, then this happens with small probability and hence we should think of $m_0 n_0$ as *the* representative cluster size.

Lemma 1: If $t \geq (2 + \delta) \log_2(n)$, $\delta > 0$, then

$$P\left(\text{Rows in two different clusters are same}\right) \leq \frac{1}{m^{\delta}}$$

and a similar result holds for the column clusters.

Proof: Each row is uniformly distributed over 2^t possibilities and rows in different clusters are independent. Hence the probability that any given pair of rows is same is $1/2^t$. Since there are $\binom{r}{2}$ pairs,

we then have

$$P\left(\text{Rows in two different clusters are same}\right) \leq \frac{\binom{r}{2}}{2^t},$$

Since $r \leq m$, we have

$$P\left(\text{Rows in two different clusters are same}\right) \leq \frac{m^2}{2^t}.$$

Hence if $t > (2 + \delta) \log_2 m$ for some $\delta > 0$, then

$$P\left(\text{Rows in two different clusters are same}\right) \leq \frac{1}{m^\delta}.$$

■

Condition A3) also implies that in any row or column, for large matrices, roughly the number of 0s and 1s is same. This essentially implies that the opinions are diverse for any user or item. While this may seem unrealistic (and can indeed be fixed), we prefer the Bernoulli(1/2) model for the following reason: under this assumption no recommendations can be extracted from any row or column alone and thus collaborative filtering is necessary. Such a model is desirable for evaluation of collaborative filtering schemes. Moreover, one can pre-process data so that rows and columns with fraction of 1s far from 1/2 are removed (because they are relatively easy to recommend) and then assumption A3) is reasonable. We note that in condition A3), we only specify the probability law of $\mathbf{X}$ given the clusters; the clusters are deterministic, even though they are unknown.

The BSC in A4) models the inherent noise in user-item ratings as well as modeling error, while the erasure channel models the missing data.

Conditions B1)-B3): These conditions are more general allowing any finite alphabet and non-uniform cluster sizes.

B1) The entries of $\mathbf{X}$ are from a finite alphabet A.

B2) $\mathbf{X}$ is constant over the cluster $A_i \times B_j$ and the entries across the clusters are i.i.d. with a uniform distribution over A.

B3) The observed data $\mathbf{Y} \in A \cup \{e\}$ (e denotes erasure) is obtained from $\mathbf{X}$ as follows

    a) The entries of $\mathbf{X}$ are passed through a DMC with probability law $q(.|.)$ and output alphabet A, resulting in $\tilde{\mathbf{X}}$.

    b) The entries $\tilde{\mathbf{X}}_{ij}$ are then passed through an erasure channel with erasure probability ϵ.

III. BINARY RATING MATRIX

In this section, we state our results under conditions A1)-A4). The main result of this section appears in Section III-A. It is obtained by studying two quantities: probability of error when the clustering is *known* (Section III-B) and probability error in clustering for a specific algorithm (Section III-C).

A. Main Result

Our main result stated below identifies a threshold on the cluster size above which error free recovery is asymptotically feasible but below which error free recovery is not possible.

Theorem 1: Suppose conditions A1)-A4) are true and the clusters are *unknown*. Let $p_1 = \epsilon + 2(1 - \epsilon)\sqrt{p(1-p)}$. Suppose that $\epsilon < 1$ and $p \in [0, p_0]$, $p_0 < 1/2$.

1) **Converse:** If

$$m_0 n_0 < (1 - \delta)\frac{\ln(mn)}{\ln(1/p_1)}, \quad \delta > 0,$$

then $\bar{P}_e \to 1$ for *any* estimator.

2) **Achievability:** If $t = \Omega(\log(n))$, $r = \Omega(\log(m))$, $\limsup m/n < \infty$, $\limsup n/m < \infty$ and

$$m_0 n_0 > \frac{\ln(mn)}{\ln(1/p_1)},$$

then $\bar{P}_e \to 0$ for the following polynomial time estimator:

- Cluster rows and columns using the algorithm of Section III-C using the threshold $d_0 \in (2p_0(1 - p_0), 1/2)$ (which does not depend on ϵ, m_0, n_0).

- Employ majority decoding in a cluster (as in Section III-B) assuming the clustering to be correct.

Proof: The proof is given in Section VI-A. ∎

The result identifies $\ln(mn)/\ln(1/p_1)$ as the cluster size threshold. The first part states that if the cluster size is too small, then any estimator makes an error with high probability. The second part states that if the cluster size is large enough, then diminishing probability of error can be achieved with a polynomial time estimator, which does not need knowledge of ϵ, m_0, n_0 and needs only knowledge of a worst case bound on p. The result is reminiscent of the channel coding theorem in the context of our model.

The proof of Part 1) of Theorem 1 relies on lower bounding $\bar{P}_e$ by considering the case of known clustering (see Theorem 2 in Section III-B). The proof of Part 2) of Theorem 1 relies on showing that for the average case, the probability of error in clustering is much smaller than the probability of error in

filling values when the clusters are known (see Theorem 3 in Section III-C). We illustrate this in Figure 1 by plotting various bounds: for $m = n = 10^6$, $m_0 = n_0$ ranging from 10 to 150, $p = 0.25$ and $\epsilon = 0.9$, we plot

- upper and lower bounds for probability of error when clustering is known (from Theorem 2),
- upper bound on probability of clustering error (from Theorem 3),
- and the asymptotic threshold $\ln(mn)/\ln(1/p_1)$ (from Theorem 1).

It is seen that around the asymptotic threshold, the probability of clustering error is dominated by the probability of error in filling values under known clustering.

B. Known Clustering

In this section, we consider the case when the clusters are known. Under this assumption, the decoder only has to estimate the value in a cluster, and the minimum probability of error estimator under A3) is just a majority decoder. The analysis of this decoder is elementary and we state a stronger result for a fixed $\mathbf{X}$ with possibly unequal cluster sizes. Let

$$s_*(\mathbf{X}) := \min_{i,j} m_i(\mathbf{X}) n_j(\mathbf{X}), \quad s^*(\mathbf{X}) := \max_{i,j} m_i(\mathbf{X}) n_j(\mathbf{X}),$$

where $\{m_i(\mathbf{X})\}$ and $\{n_j(\mathbf{X})\}$ are the row and column cluster sizes in $\mathbf{X}$.

Theorem 2: Suppose conditions A1), A3) are true and in addition assume that the clusters are known. Let

$$s_*(\mathbf{X}) \geq \frac{\ln(2)}{\ln(1/p_1)}.$$

Then the probability of error in filling in values satisfies

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \exp\left(-\frac{1}{4}\sqrt{\frac{p}{1-p}\frac{mnp_1^{s^*(\mathbf{X})}}{s^*(\mathbf{X})(s^*(\mathbf{X})+1)}}\right),$$

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \leq 1 - \exp\left(-\frac{2\ln(2)mnp_1^{s_*(\mathbf{X})}}{s_*(\mathbf{X})}\right). \tag{1}$$

Suppose we are given a sequence of rating matrices of increasing size, that is, $mn \to \infty$. Then the following are true.

1) If

$$s_*(\mathbf{X}) \geq \frac{\ln(mn)}{\ln(1/p_1)}$$

then $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 0$.

2) If

$$s^*(\mathbf{X}) \leq \frac{(1-\delta)\ln(mn)}{\ln(1/p_1)} \text{ for some } \delta > 0,$$

then $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 1$.

Proof: The proof is given in Section VI-B. ∎

We note that when all the clusters are of the same size (which happens with high probability as per Lemma 1), then the above result states that there is a sharp threshold: if the cluster size is smaller than $\ln(mn)/\ln(1/p_1)$, then exact recovery is not possible, but if it is larger, then we can make probability of error as small as we wish.

Example: For $m = n = 10^6$, $m_i = n_j = n_0$, $\epsilon = 0.9$, $p = 0.25$, this threshold corresponds to clusters of size about $45 \times 45 = 2025$. We plot the lower and upper bounds for $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})$ from Theorem 2 and the threshold in Figure 1.

Remark: A finer analysis reveals that we can refine Part 2) of Theorem 2 (and hence also Part 1) of Theorem 1) by letting δ approach zero as $m, n \to \infty$. The result holds as long as $\delta_{m,n} \ln(mn) - 2\ln\ln(mn) \to \infty$.

C. Probability of Clustering Error

To get an upper bound on the probability of error $\bar{P}_e$, in this section we analyze a specific collaborative filter: we first cluster the rows and columns using the algorithm described below and then we fill in values using the majority decoder assuming that the clustering is correct. The majority decoder has already been analyzed in Section III-B and for proving Part 2 of Theorem 1, we only need to analyze the probability of error in clustering.

Clustering Algorithm: We cluster rows and columns separately. For rows i, j, the normalized Hamming distance over commonly sampled entries is

$$d_{ij} = \frac{1}{N_{ij}} \sum_{k=1}^{n} \mathbb{1}\left(Y_{ik} \neq e, Y_{jk} \neq e, Y_{ik} \neq Y_{jk}\right),$$

where N_{ij} is the number of commonly sampled positions in rows i and j, given by

$$N_{ij} = \sum_{k=1}^{n} \mathbb{1}\left(Y_{ik} \neq e, Y_{jk} \neq e\right).$$

Let I_{ij} be equal to 1 if rows i, j belong to the same cluster and let it be 0 otherwise. The algorithm gives an estimate:

$$\hat{I}_{ij} = \begin{cases} 1, & d_{ij} < d_0, \\ 0, & d_{ij} \geq d_0, \end{cases}$$

where d_0 is a treshold whose choice will be discussed later. A similar algorithm is used to cluster columns. We are interested in the probability that we make an error in row clustering averaged over the probability law on the rating matrices defined as

$$\bar{P}_{e,rc} = \Pr\left(\hat{I}_{ij} \neq I_{ij} \text{ for some } i,j\right).$$

We note that this is a conservative definition of clustering error. As seen in Lemma 1, there is a small chance that rows in different clusters may be the same resulting in the merging of two clusters into a larger one. The above definition of error does not account for this and declares more errors. We use this conservative definition of clustering error to simplify analysis.

Theorem 3: Suppose conditions A1)-A4) are true. Let $r_1 > 1$, $r_2 \in (0,1)$ be constants and let h_* be the smaller root of the quadratic equation

$$2\mu\nu(1 - d_0)h^2 + (2d_0 - 2\mu\nu - 1)h + 1 - 2d_0 = 0, \tag{2}$$

where $\mu := 2p(1-p), \nu = 1 - \mu$. Suppose the threshold $d_0 \in (\mu, \mu + 1/2)$. Let

$$\alpha_1 = D\left(r_1(1-\epsilon)^2 \| (1-\epsilon)^2\right),$$

$$\alpha_2 = D\left(r_2(1-\epsilon)^2 \| (1-\epsilon)^2\right),$$

$$\lambda_1(n_0) = \frac{1}{2}\left(1 + \left(\frac{1 - \nu h^*}{1 - \mu h^*}\right)^{n_0 r_2(1-\epsilon)^2}\right), \text{ and,}$$

$$\lambda_2(n_0) = \frac{1}{2}\left(1 + 2^{-n_0\alpha_2}\right).$$

Then for the above clustering algorithm,

$$\bar{P}_{e,rc} \leq \frac{m(m-1)}{2} \max\left\{P(\hat{I}_{ij} = 0 | I_{ij} = 1), P(\hat{I}_{ij} = 1 | I_{ij} = 0)\right\},$$

where

$$P(\hat{I}_{ij} = 0 | I_{ij} = 1) \leq \exp\left(-n \min\left\{r_2(1-\epsilon)^2 D(d_0 \| \mu), \alpha_2\right\}\right) \tag{3}$$

and

$$P(\hat{I}_{ij} = 1 | I_{ij} = 0) \leq \min\{P_1, P_2\} \tag{4}$$

$$P_1 = \left(\frac{1 - \mu h_*}{(1 - h_*)^{d_0}}\right)^{n r_1 (1-\epsilon)^2} \lambda_1(n_0)^t + \exp(-\alpha_1 n) + \lambda_2(n_0)^t \tag{5}$$

$$P_2 = \exp(-\alpha_3 t) \tag{6}$$

for a positive constant α_3.

Proof: The proof is given in Section VI-C. ■

The proof uses the union bound and considers pairwise errors. The pairwise errors consists of two cases: error when the pair of rows is in the same cluster and error when they are in different clusters. The probability of the first kind of error is exponentially decaying in n. The probability of the second kind of error is upper bounded by the minimum of P_1 and P_2: while P_1 is tight for finite n and large p, ϵ, the bound P_2 is useful for establishing asymptotic results (like Theorem 1) for *all* p, ϵ. For example, in Figure 1, the upper bound on clustering error is dominated by P_1, while the proof of Part 2) of Theorem 1 uses P_2. We note that both P_1 and P_2 have terms that decay exponentially in n as well as t. The terms decaying exponentially in t are related to Lemma 1 and the conservative definition of clustering error as discussed before the statement of Theorem 3. These terms are the origin of the $t = \Omega(\log(n))$ condition in Part 2) of Theorem 1 and can perhaps be avoided with more sophisticated analysis; however, we prefer to work with this condition since as per Lemma 1, the condition $t = \Omega(\log(n))$ is anyway needed for interpreting $m_0 n_0$ as the representative cluster size.

IV. GENERAL FINITE ALPHABET AND NON-UNIFORM CLUSTERS

In this section, we consider a general finite alphabet A and non-uniform cluster sizes. We work with assumptions B1)-B3) described in the Section II and generalize the results in Section III. To state our results, we first introduce some notation. For $p, q \in$ A, define

$$\mu_{pq} := \sum_{y_i \neq y_j} q(y_i | p) q(y_j | q). \tag{7}$$

If A_1, A_2 are i.i.d. uniform on A and we pass them through the DMC $q(\cdot | \cdot)$ to get outputs $\tilde{A}_1$, $\tilde{A}_2$, then

$$\Pr(\tilde{A}_1 \neq \tilde{A}_2 | A_1 = A_2) = \frac{1}{|A|} \sum_{p \in A} \mu_{pp} =: d_{lb}$$

$$\Pr(\tilde{A}_1 \neq \tilde{A}_2) = \frac{1}{|A|^2} \sum_{p,q \in A} \mu_{pq} =: d_{ub}.$$

The following useful lemma sheds light on the relationship between d_{lb} and d_{ub}.

Lemma 2: For any DMC, $d_{ub} \geq d_{lb}$, with equality iff $q(y|p) = q(y|q) \ \forall \ p, q, y \in \mathsf{A}$.

Proof: The proof is given in Section VI-E. ■

We next state our main result for general finite alphabet and non-uniform cluster size.

Theorem 4: Suppose conditions B1)-B3) are true and the clusters are *unknown*. Then there exist constants $p_1, p_2 \in (0, 1)$, $p_1 > p_2$ such that

1) **Converse:** If

$$\max_{i,j} m_i n_j < (1 - \eta) \frac{\ln(mn)}{\ln(1/p_2)}, \quad \eta > 0,$$

then $\bar{P}_e \to 1$ for *any* estimator.

2) **Achievability:** Suppose that there exist some $y, p, q \in \mathsf{A}$ such that $p \neq q$ and $q(y|p) \neq q(y|q)$. (By Lemma 2, this ensures that $d_{lb} < d_{ub}$.) If $n^2/(n_1^2 + n_2^2 + \ldots + n_t^2) = \Omega(\log(m))$, $m^2/(m_1^2 + m_2^2 + \ldots + m_r^2) = \Omega(\log(n))$, $\limsup m/n < \infty$, $\limsup n/m < \infty$ and

$$\min_{i,j} m_i n_j > \frac{\ln(mn)}{\ln(1/p_1)},$$

then $\bar{P}_e \to 0$ for the following polynomial time estimator:

- Cluster rows and columns using the algorithm of Section III-C using the threshold $d_0 \in (d_{lb}, d_{ub})$ (which does not depend on ϵ, m_i, n_j).
- Employ maximum likelihood decoding in a cluster assuming the clustering is correct.

Proof: The proof is similar to Theorem 1; we now use Theorems 5 and 6 in place of Theorems 2 and 3 respectively. ■

The above result again identifies $\ln(mn)$ as the exact order of the cluster size threshold for asymptotic recovery. Similar to the binary alphabet and uniform cluster size case in Section III, the constants p_1, p_2 arise from the case when the clusters are known (see Theorem 5 below). The gap between the constants p_1, p_2 can be made arbitrarily small: the proof of Theorem 5 identifies a constant C_1 (see equation (29)) such that for any $\delta > 0$,

$$p_1 = \epsilon + (1 - \epsilon) \exp(-C_1 + \delta), \quad p_2 = \epsilon + (1 - \epsilon) \exp(-C_1 - \delta)$$

is a valid choice in Theorem 4.

We next consider the case when the clusters are known and extend Theorem 2.

Theorem 5: Suppose conditions B1)-B3) are true and in addition assume that the clusters are *known*. Also let

$$s_*(\mathbf{X}) \geq \frac{\ln(1/2|\mathsf{A}|)}{\ln(\epsilon/p_2)}, \tag{8}$$

where p_1, p_2 are as defined above. Then for a sequence of rating matrices of increasing size $mn \to \infty$, the following are true.

1) If
$$s_*(\mathbf{X}) \geq \frac{\ln(mn)}{\ln(1/p_1)},$$
then $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 0$.

2) If
$$s^*(\mathbf{X}) \leq \frac{(1 - \zeta)\ln(mn)}{\ln(1/p_2)}, \quad \text{for some } \zeta > 0,$$
then $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 1$.

Proof: The proof is given in Section VI-D. ∎

Finally, we study the performance of the clustering algorithm and extend Theorem 3.

Theorem 6: Suppose conditions B1)-B3) are true and in addition suppose that there exist some $y, p, q \in$ A such that $p \neq q$ and $q(y|p) \neq q(y|q)$. (By Lemma 2, this ensures that $d_{lb} < d_{ub}$.) If we choose the threshold $d_0 \in (d_{lb}, d_{ub})$, then

$$\bar{P}_{e,rc} \leq c' \frac{m(m-1)}{2} \exp\left(-cn^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right), \tag{9}$$

for some positive constants c, c'. Consequently, if $n^2/(n_1^2 + n_2^2 + \ldots + n_t^2) = \Omega(\log(m))$, then $\bar{P}_{e,rc} \to 0$ as $m, n \to \infty$.

Proof: The proof is given in Section VI-F. ∎

V. CONCLUSION

We take a channel coding perspective of collaborative filtering and identify the threshold on cluster size for perfect reconstruction of the underlying rating matrix. The result is similar in flavor to some recent results in completion of real-valued matrices. The advantage of our model is that the proofs are relatively simple relying on Chernoff bounds and noisy user behavior can be easily handled.

In the typical applications of recommendation systems, there is a lack of good models. We believe that our model has two characteristics that make it suitable for analytical comparison of various methods: a) in our model the user opinions are diverse and no single user/item reveals much information about itself, that is, collaborative filtering is necessary; b) as we have shown, the model is analytically tractable. There are several directions where this model may turn out to be useful: analysis of bit error probability instead of block error probability, analysis of local popularity based mechanisms, etc.

VI. PROOFS OF RESULTS

A. Proof of Theorem 1

The proof is based on Theorems 2 and 3.

When $\mathcal{A}, \mathcal{B}$ are known, under our model all feasible rating matrices are equally likely. Hence the ML decoder gives the minimum probability of error and so we have $\bar{P}_e \geq E[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})]$. To prove Part 1), we lower bound $E[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})]$. Let T be the event that $s^*(\mathbf{X}) > m_0 n_0$. Proceeding as in Lemma 1, we have for $t \geq (2+\delta)\log_2(n)$, $r \geq (2+\delta)\log_2(m)$, $\delta > 0$,

$$\Pr(T) \leq \frac{1}{m^\delta} + \frac{1}{n^\delta}.$$

Hence $\Pr(T) \to 0$. Now,

$$E[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})] \geq E[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}); T^c].$$

But on the event T^c, $s^*(\mathbf{X}) = m_0 n_0$ and hence we get

$$\bar{P}_e \geq E[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})] \geq (1 - \Pr(T))P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}). \tag{10}$$

But from Part 1) of Theorem 2, $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 1$ for

$$m_0 n_0 < (1-\delta)\frac{\ln(mn)}{\ln(1/p_1)}, \quad \delta > 0.$$

This proves Part 1).

Next we prove Part 2). Let D denote the event that the clustering is identified correctly. We note that the probability of error in estimating $\mathbf{X}$ averaged over the probability law on the block constant matrices satisfies

$$\bar{P}_e \leq E\left[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})\Pr(D) + \Pr(D^c)\right]$$
$$\leq E\left[P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})\right] + \left(\bar{P}_{e,rc} + \bar{P}_{e,cc}\right)$$

where $\bar{P}_{e,cc}$ is the probability of error in column clustering. The desired result follows from Part 2) of Theorem 2 and Theorem 3. ■

B. Proof of Theorem 2

Suppose in cluster $A_i \times B_j$ we have s non erased samples. Then the probability of correct decision in this cluster is given by

$$\Pr(E_{i,j,s}^c) = \sum_{q=0}^{\lfloor \frac{s}{2} \rfloor} \binom{s}{q} p^q (1-p)^{s-q} \text{ if } s \text{ is odd}$$

$$= \sum_{q=0}^{\frac{s}{2}-1} \binom{s}{q} p^q (1-p)^{s-q} \tag{11}$$

$$+ \frac{1}{2} \binom{s}{\frac{s}{2}} p^{\frac{s}{2}} (1-p)^{\frac{s}{2}} \text{ if } s \text{ is even.}$$

Averaging over the number of non erased samples, the probability of correct decision in cluster $A_i \times B_j$ is given by

$$\Pr(E_{i,j}^c) = \sum_{s=0}^{m_i n_j} \binom{m_i n_j}{s} \epsilon^{m_i n_j - s} (1-\epsilon)^s \Pr(E_{i,j,s}^c). \tag{12}$$

Since the erasure and BSC are memoryless

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) = \Pr\left(\cup_{i=1,j=1}^{r,t} E_{i,j}\right)$$

$$= 1 - \prod_{i=1,j=1}^{r,t} \Pr\left(E_{i,j}^c\right). \tag{13}$$

Equations (11), (12), and (13) specify the probability of error.

Upper Bound: The desired upper bound is obtained by deriving a lower bound on $\Pr(E_{i,j,s}^c)$. First we note that from (11),

$$1 - \Pr(E_{i,j,s}^c) \le \sum_{\lceil \frac{s}{2} \rceil}^{s} \binom{s}{q} p^q (1-p)^{s-q}.$$

But for $0 \le p \le \frac{1}{2}$ and $q \ge \frac{s}{2}$, $p^q(1-p)^{s-q} \le p^{\frac{s}{2}}(1-p)^{\frac{s}{2}}$. Substituting this in the previous equation, we have

$$\Pr(E_{i,j,s}^c) \ge 1 - (2\sqrt{p(1-p)})^s. \tag{14}$$

From Equations (12) and (14), we have $\Pr(E_{i,j}^c) \ge 1 - p_1^{m_i n_j}$ and so from (13),

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \le 1 - \prod_{i=1,j=1}^{r,t} \left(1 - p_1^{m_i n_j}\right).$$

We note that for $x \in [0, 1/2]$, $1 - x \ge \exp(-2\ln(2)x)$. Hence

$$\exp\left(-2\ln(2) \sum_{i=1,j=1}^{r,t} p_1^{m_i n_j}\right) \le \prod_{i=1,j=1}^{r,t} \left(1 - p_1^{m_i n_j}\right).$$

Where the first inequality holds for $p_1^{m_i n_j} \leq 1/2$. This is true since $s_*(\mathbf{X}) \geq \ln(2)/\ln(1/p_1)$. The upper bound follows by noting that

$$\sum_{i=1,j=1}^{r,t} p_1^{m_i n_j} \leq rt p_1^{s_*(\mathbf{X})} \leq \frac{mn}{s_*(\mathbf{X})} p_1^{s_*(\mathbf{X})}.$$

Lower Bound: The lower bound on $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X})$ is obtained from an upper bound on $\Pr(E_{i,j,s}^c)$. From (11),

$$1 - \Pr(E_{i,j,s}^c) \geq \frac{1}{2} \binom{s}{\lceil s/2 \rceil} p^{\lceil s/2 \rceil} (1-p)^{s-\lceil s/2 \rceil},$$

$$\geq \frac{1}{2(s+1)} 2^{sh(\lceil s/2 \rceil/s)} p^{\lceil s/2 \rceil} (1-p)^{s-\lceil s/2 \rceil}.$$

If s is even, we have

$$1 - \Pr(E_{i,j,s}^c) \geq \frac{1}{2(s+1)} \left(2\sqrt{p(1-p)} \right)^s. \tag{15}$$

For s odd,

$$h\left(\lceil s/2 \rceil / s \right) = h(1/2 + 1/2s) \geq 1 - 1/s^2,$$

and so

$$1 - \Pr(E_{i,j,s}^c) \geq \frac{2^{-1/s}}{2(s+1)} \sqrt{\frac{p}{1-p}} \left(2\sqrt{p(1-p)} \right)^s. \tag{16}$$

From (15) and (16), we have for all s,

$$1 - \Pr(E_{i,j,s}^c) \geq \frac{1}{4(s+1)} \sqrt{\frac{p}{1-p}} \left(2\sqrt{p(1-p)} \right)^s. \tag{17}$$

Now from (12),

$$\Pr(E_{i,j}^c) \leq 1 - \frac{1}{4}\sqrt{\frac{p}{1-p}} \sum_{s=0}^{m_i n_j} \binom{m_i n_j}{s} \epsilon^{m_i n_j - s} \left((1-\epsilon) 2\sqrt{p(1-p)} \right)^s \frac{1}{s+1}$$

$$\leq 1 - \frac{p_1^{m_i n_j}}{4(m_i n_j + 1)} \sqrt{\frac{p}{1-p}}.$$

Using this bound on $\Pr(E_{i,j}^c)$ in (13), we have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \prod_{i=1,j=1}^{r,t} \left(1 - \frac{p_1^{m_i n_j}}{4(m_i n_j + 1)} \sqrt{\frac{p}{1-p}} \right)$$

$$\geq 1 - \exp\left(-\frac{1}{4}\sqrt{\frac{p}{1-p}} \sum_{i=1,j=1}^{r,t} \frac{p_1^{m_i n_j}}{m_i n_j + 1} \right) \tag{18}$$

$$\geq 1 - \exp\left(-\frac{1}{4}\sqrt{\frac{p}{1-p}} rt \frac{p_1^{s^*(\mathbf{X})}}{s^*(\mathbf{X}) + 1} \right)$$

$$\geq 1 - \exp\left(-\frac{1}{4}\sqrt{\frac{p}{1-p}} mn \frac{p_1^{s^*(\mathbf{X})}}{s^*(\mathbf{X}) + 1} \right),$$

where in (18) we have used $1 - x \leq \exp(-x)$. This completes the proof of (1).

Asymptotics: Now consider a sequence of rating matrices of increasing size. The upper bound on error in (1) is a decreasing function of $s_*(\mathbf{X})$. Hence if

$$s_*(\mathbf{X}) \geq \frac{\ln(mn)}{\ln(1/p_1)},$$

then

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \leq \frac{2\ln(2)\ln(1/p_1)}{\ln(mn)} \to 0.$$

Now suppose

$$s^*(\mathbf{X}) \leq \frac{(1-\delta)\ln(mn)}{\ln(1/p_1)}, \quad \text{for some } \delta > 0.$$

The lower bound on error (1) is a decreasing function of $s^*(\mathbf{X})$, and hence substituting the above upper bound on $s^*(\mathbf{X})$, we have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \exp\left(-c_1 \frac{(mn)^\delta}{(c_2 + \ln(mn))\ln(mn)}\right),$$

where c_1, c_2 are some positive constants. Hence $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 1$ as $mn \to \infty$. ∎

C. Proof of Theorem 3

Recall that N_{ij} is the number of commonly sampled positions in rows i and j, given by

$$N_{ij} = \sum_{k=1}^{n} 1\left(Y_{ik} \neq e, Y_{jk} \neq e\right).$$

From the Chernoff bound [10, Theorem 1], we have

$$\Pr\left(N_{ij} > nr_1(1-\epsilon)^2\right) \leq \exp\left(-nD\left(r_1(1-\epsilon)^2 \| (1-\epsilon)^2\right)\right) = \exp(-n\alpha_1), \quad \text{and} \tag{19}$$

$$\Pr\left(N_{ij} < nr_2(1-\epsilon)^2\right) \leq \exp\left(-nD\left(r_2(1-\epsilon)^2 \| (1-\epsilon)^2\right)\right) = \exp(-n\alpha_2). \tag{20}$$

To get a handle on the probability of error, we first analyze it conditioned on the erasure sequence and $\mathbf{X}$. Let $\bar{E}$ denote the erasure matrix:

$$\bar{E} = [1(Y_{ij} = e)]_{m \times n} \in \{0, 1\}^{m \times n}.$$

Rows in Same Cluster: Consider rows i, j of $\mathbf{X}$ and suppose $I_{ij} = 1$, i.e. i, j are in the same cluster. We wish to evaluate the probability of error $\Pr(d_{ij} \geq d_0 | I_{ij} = 1, \bar{E}, \mathbf{X})$. In this case, the random variable $N_{ij}d_{ij}$ is given by

$$N_{ij}d_{ij} = \sum_{\substack{\mathbf{Y}_{ik}, \mathbf{Y}_{jk} \neq e \\ \mathbf{X}_{ik} = \mathbf{X}_{jk}}} 1\left(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk}\right).$$

For any column k such that $\mathbf{Y}_{ik} \neq e, \mathbf{Y}_{jk} \neq e$, the indicator $1(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk})$ has mean $\mu = 2p(1-p)$. Hence, the above summation has N_{ij} i.i.d Bernoulli random variables of mean μ. An application of Chernoff bound [10, Theorem 1] yields

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1, \bar{E}, \mathbf{X}\right) \leq \exp\left(-N_{ij}D(d_0\|\mu)\right). \tag{21}$$

The bound is independent of $\mathbf{X}$. We only need to take the average of (21) with respect to $\bar{E}$. Using (20), we have

$$\Pr\left(\hat{I}_{ij} = 0 \Big| I_{ij} = 1\right) = \Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1\right)$$

$$\leq \exp\left(-n(1-\epsilon)^2 r_2 D(d_0\|\mu)\right) + \exp(-n\alpha_2),$$

$$\leq 2\exp\left(-n\min\left\{r_2(1-\epsilon)^2 D(d_0\|\mu), \alpha_2\right\}\right). \tag{22}$$

Rows in Different Clusters: Next consider the case $I_{ij} = 0$, i.e. rows i and j are in different clusters. We wish to evaluate $\Pr(d_{ij} \leq d_0 | I_{ij} = 0, \bar{E}, \mathbf{X})$. For $I_{ij} = 0$ and fixed $\bar{E}$, $\mathbf{X}$, the random variable $N_{ij}d_{ij}$ is given by

$$N_{ij}d_{ij} = \sum_{\substack{\mathbf{Y}_{ik}, \mathbf{Y}_{jk} \neq e \\ \mathbf{X}_{ik} = \mathbf{X}_{jk}}} 1\left(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk}\right) + \sum_{\substack{\mathbf{Y}_{ik}, \mathbf{Y}_{jk} \neq e \\ \mathbf{X}_{ik} \neq \mathbf{X}_{jk}}} 1\left(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk}\right). \tag{23}$$

Note that for any column k such that $\mathbf{Y}_{ik} \neq e, \mathbf{Y}_{jk} \neq e$, the indicator $1(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk})$ has mean

- $2p(1-p) = \mu$ if $\mathbf{X}_{ik} = \mathbf{X}_{jk}$, and
- $p^2 + (1-p)^2 = \nu$ if $\mathbf{X}_{ik} \neq \mathbf{X}_{jk}$.

Define S_{ij} as the number of columns k such that $\mathbf{Y}_{ik} \neq e, \mathbf{Y}_{jk} \neq e$ and $\mathbf{X}_{ik} \neq \mathbf{X}_{jk}$. Then from (23),we observe that the first sum in (23) has $N_{ij} - S_{ij}$ i.i.d Bernoulli random variables of mean μ and the second sum has S_{ij} i.i.d Bernoulli random variables of mean ν, all the random variables being independent. Using the Chernoff bound, we may then write

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right)$$
$$\leq \frac{(1 - \nu + \nu e^\theta)^{S_{ij}}(1 - \mu + \mu e^\theta)^{N_{ij} - S_{ij}}}{e^{d_0 N_{ij}\theta}}, \text{ for } \theta \leq 0. \tag{24}$$

By substituting $h = 1 - \exp(\theta)$, we can rewrite the above bound as

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right) \leq \frac{(1 - \nu h)^{S_{ij}}(1 - \mu h)^{N_{ij} - S_{ij}}}{(1 - h)^{d_0 N_{ij}}}, \text{ for } 0 \leq h < 1. \tag{25}$$

We are free to choose $0 \leq h < 1$ in the above bound. We choose h such that the bound is optimized for the average case $S_{ij} = N_{ij}/2$. For this case, the bound in (25) reduces to

$$\left(\frac{(1 - \nu h)(1 - \mu h)}{(1 - h)^{2d_0}}\right)^{N_{ij}/2}.$$

The value of h that minimizes this bound can be checked to be the smaller root of the quadratic given by (2).

Next, we take expectation in (25) with respect to the erasure sequence $\bar{E}$. Let s_{ij} denote the number of columns k such that $\mathbf{X}_{ik} \neq \mathbf{X}_{jk}$. Then we have, from the Chernoff bound, as in (20),

$$\Pr\left(S_{ij} > s_{ij} r_1 (1 - \epsilon)^2\right) \leq \exp(-s_{ij} \alpha_1). \tag{26}$$

Now, from (25), we have

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right) \leq \left(\frac{1 - \mu h}{(1 - h)^{d_0}}\right)^{N_{ij}} \left(\frac{1 - \nu h}{1 - \mu h}\right)^{S_{ij}}, \text{ for } 0 \leq h < 1.$$

Now, since $\mu < \nu \leq 1$, we have

$$\frac{1 - \nu h}{1 - \mu h} < 1, \text{ for } h \geq 0.$$

First note that the function $f(h) = (1 - \mu h)/(1 - h)^{d_0}$ for $h \in [0, 1)$ has derivative

$$f'(h) = \frac{d_0 - \mu + \mu h (1 - d_0)}{(1 - h)^{d_0}}.$$

Since $\mu < d_0 < 1$, $f'(h) > 0$ and so $f(h) > f(0) = 1$. Hence $(1 - \mu h)/(1 - h)^{d_0} > 1$. Now if $S_{ij} > s_{ij} r_2 (1 - \epsilon)^2$ and $N_{ij} < n r_1 (1 - \epsilon)^2$, then

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right) \leq \left(\frac{1 - \mu h}{(1 - h)^{d_0}}\right)^{n r_1 (1 - \epsilon)^2} \left(\frac{1 - \nu h}{1 - \mu h}\right)^{s_{ij} r_2 (1 - \epsilon)^2}, \text{ for } 0 \leq h < 1.$$

Combining this with (26) and (19), we have

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \mathbf{X}\right) \leq \left(\frac{1 - \mu h}{(1 - h)^{d_0}}\right)^{n r_1 (1 - \epsilon)^2} \left(\frac{1 - \nu h}{1 - \mu h}\right)^{s_{ij} r_2 (1 - \epsilon)^2}$$

$$+ \Pr\left(S_{ij} < s_{ij} r_2 (1 - \epsilon)^2\right) + \Pr\left(N_{ij} > n r_1 (1 - \epsilon)^2\right),$$

$$\leq \left(\frac{1 - \mu h}{(1 - h)^{d_0}}\right)^{n r_1 (1 - \epsilon)^2} \left(\frac{1 - \nu h}{1 - \mu h}\right)^{s_{ij} r_2 (1 - \epsilon)^2} + \exp(-s_{ij} \alpha_2) + \exp(-n \alpha_1). \tag{27}$$

Since $s_{ij} = n_0 X$, where X is Binomial$(t, 1/2)$, we have

$$\mathsf{E}\left[\exp(\lambda s_{ij})\right] = \mathsf{E}\left[\exp(\lambda n_0 X)\right] = \left(\frac{1 + \exp(\lambda n_0)}{2}\right)^t.$$

Now taking expectation with respect to $\mathbf{X}$ in (27), we have

$$P(\hat{I}_{ij} = 1 | I_{ij} = 0) = \Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0\right)$$

$$\leq \left(\frac{1 - \mu h}{(1 - h)^{d_0}}\right)^{n r_1 (1 - \epsilon)^2} \lambda_1(n_0)^t + \lambda_2(n_0)^t + \exp(-\alpha_1 n) = P_1. \tag{28}$$

It remains to show that

$$P(\hat{I}_{ij} = 1 | I_{ij} = 0) \le P_2.$$

This result follows from (9) of Theorem 6 for the general case, which is proved in Section VI-F. ∎

D. Proof of Theorem 5

For simplicity let $Y_1, ..., Y_s$ denote the s samples in block $A_i \times B_j$. Let $\mu_a := q(\cdot|a)$, $a \in A$ be the transition law of the channel for input a and let μ_Y denote the empirical probability mass function (PMF) of $Y_1, ..., Y_s$. Let E_{ijs} be the error event when the (i,j)th block has s samples. For simplicity let $\mathcal{P}_s$ denote the set of types with denominator s [8, pp. 348] and define the set of PMFs:

$$U_{a,b} := \{\nu : D(\nu\|\mu_b) \le D(\nu\|\mu_a)\} \cap \mathcal{P}_s, \quad V_{a,b} := \{\nu : D(\nu\|\mu_b) < D(\nu\|\mu_a)\} \cap \mathcal{P}_s.$$

Upper Bound: Then

$$\Pr(E_{ijs}) = \frac{1}{|A|} \sum_{a \in A} \Pr(E_{ijs}|a)$$

$$\le \frac{1}{|A|} \sum_{a \in A} \sum_{b \in A, b \neq a} \Pr\left(D(\mu_Y\|\mu_b) \le D(\mu_Y\|\mu_a)|a\right)$$

$$\le \frac{1}{|A|} \sum_{a,b \in A, a \neq b} \sum_{\nu \in U_{a,b}} \exp\left(-sD(\nu\|\mu_a)\right),$$

where in the second step we have used the union bound and in the last step we have used [8, Theorem 11.1.4, pp. 354]. Let

$$C_1 := \lim_{s \to \infty} -\frac{1}{s} \ln\left(\sum_{a,b \in A, a \neq b} \sum_{\nu \in U_{a,b}} \exp\left(-sD(\nu\|\mu_a)\right)\right) = \min_{a \neq b} \min_{\{\nu : D(\nu\|\mu_b) \le D(\nu\|\mu_a)\}} D(\nu\|\mu_a). \quad (29)$$

Then for $\delta > 0$ small, for $s > s_0(\delta)$, we have

$$\Pr(E_{ijs}) \le \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|},$$

while for $s \le s_0$ we can bound this probability by 1. Hence we have from (12),

$$\Pr(E_{ij}^c) = E\left[\Pr(E_{ijs}^c)\right] \ge E\left[1(s \ge s_0)\left(1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right)\right]$$

$$= E\left[1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right] - E\left[1(s < s_0)\left(1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right)\right]$$

$$\ge E\left[1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right] - E\left[1(s < s_0)\right]$$

$$= E\left[1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right] - \Pr(s < s_0). \quad (30)$$

But for large enough $m_i n_j$ using the Chernoff bound [10, Theorem 1],

$$\Pr(s < s_0) \leq \exp\left(-m_i n_j D(s_0/m_i n_j || 1 - \epsilon)\right). \tag{31}$$

As $m_i n_j \to \infty$, $D(s_0/m_i n_j || 1 - \epsilon) \to \ln(1/\epsilon)$. Hence given any $\eta > 0$, for large enough $m_i n_j$, we have

$$D(s_0/m_i n_j || 1 - \epsilon) \geq \ln(1/(\epsilon + \eta)).$$

Hence, from (30) and (31),

$$\Pr(E_{ij}^c) \geq E\left(1 - \frac{\exp\left(-(C_1 - \delta)s\right)}{|A|}\right) - (\epsilon + \eta)^{m_i n_j}$$

$$= 1 - \frac{p_1^{m_i n_j}}{|A|} - (\epsilon + \eta)^{m_i n_j}$$

where we have used the fact that s is Binomial$(m_i n_j, 1 - \epsilon)$ and so the binomial expansion. Note that $\epsilon < p_1$, and hence we can choose η so that $\epsilon + \eta < p_1$. Hence we have

$$\Pr(E_{ij}^c) \geq 1 - \frac{2p_1^{m_i n_j}}{|A|}$$

Using (13), we then have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \leq 1 - \prod_{i=1,j=1}^{r,t} \left(1 - 2p_1^{m_i n_j}/|A|\right)$$

$$\leq 1 - \exp\left(-\frac{4\ln(2)}{|A|} \sum_{i=1,j=1}^{r,t} p_1^{m_i n_j}\right), \tag{32}$$

where in the last step we have used $1 - x \geq \exp(-2\ln(2)x)$ for $x \in [0, 1/2]$. Note that for large enough $m_i n_j$, we have $p_1^{m_i n_j} < 1/2$. But using

$$\sum_{i=1,j=1}^{r,t} p_1^{m_i n_j} \leq rt p_1^{s_*(\mathbf{X})} \leq \frac{mn}{s_*(\mathbf{X})} p_1^{s_*(\mathbf{X})},$$

we have,

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \leq 1 - \exp\left(-\frac{4\ln(2)mn p_1^{s_*(\mathbf{X})}}{|A| s_*(\mathbf{X})}\right). \tag{33}$$

The RHS in (33) is a decreasing function of $s_*(\mathbf{X})$. Hence if

$$s_*(\mathbf{X}) \geq \frac{\ln(mn)}{\ln(1/p_1)},$$

then

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \leq 1 - \exp\left(-\frac{4\ln(2)\ln(1/p_1)}{|A|\ln(mn)}\right) \to 0.$$

Lower Bound: Next we give a lower bound on $\Pr(E_{ijs})$. If for each a we consider some $b \neq a$, then we get

$$
\begin{aligned}
\Pr(E_{ijs}) &= \frac{1}{|\mathsf{A}|} \sum_{a \in \mathsf{A}} \Pr(E_{ijs}|a) \\
&\geq \frac{1}{|\mathsf{A}|} \sum_{a \in \mathsf{A}} \Pr\left(D(\mu_Y \| \mu_b) < D(\mu_Y \| \mu_a)|a\right), \\
&\geq \frac{1}{|\mathsf{A}|} \sum_{a \in \mathsf{A}} \sum_{\nu \in V_{a,b}} \frac{\exp\left(-sD(\nu \| \mu_a)\right)}{(s+1)^{|\mathsf{A}|}}, \\
&\geq \frac{1}{|\mathsf{A}|(m_i n_j + 1)^{|\mathsf{A}|}} \sum_{a \in \mathsf{A}} \sum_{\nu \in V_{a,b}} \exp\left(-sD(\nu \| \mu_a)\right),
\end{aligned}
$$

where again we have used [8, Theorem 11.1.4, pp. 354] in the third step. Since we are free to choose b, we choose it such that

$$
b = \arg\min_{b \neq a} \min_{\{\nu : D(\nu \| \mu_b) < D(\nu \| \mu_a)\}} D(\nu \| \mu_a).
$$

Then we see that

$$
\lim_{s \to \infty} -\frac{1}{s} \ln \left(\sum_{a \in \mathsf{A}} \sum_{\nu \in V_{a,b}} \exp\left(-sD(\nu \| \mu_a)\right) \right) = C_1.
$$

Hence for $\delta > 0$, for $s > s_1(\delta)$,

$$
\Pr(E_{ijs}) \geq \frac{\exp\left(-(C_1 + \delta)s\right)}{|\mathsf{A}|(m_i n_j + 1)^{|\mathsf{A}|}}
$$

and for smaller s we use the trivial bound that the probability is non-negative. Hence we have from (12),

$$
\Pr(E_{ij}^c) = E\left(\Pr(E_{ijs}^c)\right) \leq E\left(1 - \frac{\exp\left(-(C_1 + \delta)s\right)}{|\mathsf{A}|(m_i n_j + 1)^{|\mathsf{A}|}}\right) + \Pr(s < s_1) \tag{34}
$$

$$
\leq 1 - \frac{p_2^{m_i n_j}}{|\mathsf{A}|(m_i n_j + 1)^{|\mathsf{A}|}} + \exp\left(-m_i n_j D(s_1/m_i n_j \| 1 - \epsilon)\right) \tag{35}
$$

$$
\leq 1 - \frac{p_2^{m_i n_j}}{|\mathsf{A}|(m_i n_j + 1)^{|\mathsf{A}|}} + \epsilon^{m_i n_j} \tag{36}
$$

$$
\leq 1 - \frac{p_2^{m_i n_j}}{|\mathsf{A}|} + \epsilon^{m_i n_j},
$$

where in (35) we have used the Chernoff Bound [10, Theorem 1], in (36) we have used the fact that $D(s_1/m_i n_j \| 1 - \epsilon) \to \ln(1/\epsilon)$ monotonically and in the last step we have used $m_i n_j \geq 0$. Further, from (8), we have

$$
\epsilon^{m_i n_j} \leq \frac{p_2^{m_i n_j}}{2|\mathsf{A}|},
$$

and hence

$$
\Pr(E_{ij}^c) \leq 1 - \frac{p_2^{m_i n_j}}{2|\mathsf{A}|}.
$$

Using (13), we then have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \prod_{i=1,j=1}^{r,t} \left(1 - \frac{p_2^{m_i n_j}}{2|\mathsf{A}|} \right)$$

$$\geq 1 - \exp\left(-\frac{1}{2|\mathsf{A}|} \sum_{i=1,j=1}^{r,t} p_2^{m_i n_j} \right), \tag{37}$$

where to obtain (37) we have used $1 - x \leq \exp(-x)$. Now since

$$\sum_{i=1,j=1}^{r,t} p_2^{m_i n_j} \geq rt p_2^{s^*(\mathbf{X})} \geq \frac{mn}{s^*(\mathbf{X})} p_2^{s^*(\mathbf{X})},$$

we have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \exp\left(-\frac{mn}{2|\mathsf{A}|s^*(\mathbf{X})} p_2^{s^*(\mathbf{X})} \right).$$

The RHS above is a decreasing function of $s^*(\mathbf{X})$, and hence if

$$s^*(\mathbf{X}) \leq \frac{(1-\zeta)\ln(mn)}{\ln(1/p_2)}, \quad \text{for some } \delta > 0,$$

we have

$$P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \geq 1 - \exp\left(-\frac{(mn)^\zeta}{2|\mathsf{A}|(1-\zeta)\ln(mn)} \right),$$

and hence $P_{e|\mathcal{A},\mathcal{B}}(\mathbf{X}) \to 1$ as $mn \to \infty$. ∎

E. Proof of Lemma 2

We recall

$$d_{ub} = \sum_{p,q} \mu_{pq}/|\mathsf{A}|^2 = \sum_{p,q} \sum_{y \neq z} q(y|p)q(z|q)/|\mathsf{A}|^2.$$

Adding and subtracting the terms corresponding to $y = z$, we have,

$$d_{ub} = \sum_{p,q} \sum_{y,z} q(y|p)q(z|q)/|\mathsf{A}|^2$$

$$- \sum_{p,q} \sum_{y} q(y|p)q(y|q)/|\mathsf{A}|^2$$

$$= \left(\sum_{p,y} q(y|p) \right)^2 /|\mathsf{A}|^2 - \sum_{p,q} \sum_{y} q(y|p)q(y|q)/|\mathsf{A}|^2.$$

Now, $\sum_{p,y} q(y|p)$ is the sum of all entries of the transition probability matrix, and hence is equal to $|A|$. So we have

$$d_{ub} = 1 - \sum_{p,q} \sum_y q(y|p)q(y|q)/|A|^2$$

$$= 1 - \sum_y \left(\sum_p q(y|p) \right)^2 /|A|^2. \tag{38}$$

Similarly

$$d_{lb} = \sum_p \mu_{pp}/|A| = \sum_p \sum_{y \neq z} q(y|p)q(z|p)/|A|.$$

Adding and subtracting the terms coresponding to $y = z$, we have,

$$d_{lb} = \sum_p \sum_{y,z} q(y|p)q(z|p)/|A| - \sum_p \sum_y q^2(y|p)/|A|$$

$$= \sum_p \left(\sum_y q(y|p) \right)^2 /|A| - \sum_p \sum_y q^2(y|p)/|A|$$

$$= 1 - \sum_y \sum_p q^2(y|p)/|A|. \tag{39}$$

In the last step we have used $\sum_y q(y|p) = 1$ for the first term. From (38) and (39), we have

$$d_{ub} - d_{lb} = \frac{1}{|A|^2} \sum_y \left(|A| \sum_p q^2(y|p) - \left(\sum_p q(y|p) \right)^2 \right).$$

From the Cauchy-Schwarz inequality,

$$\left(\sum_p q(y|p) \right)^2 \leq |A| \sum_p q^2(y|p),$$

with equality iff $q(y|p) = q(y|q)$ for all p, q. The result then follows. ∎

F. Proof of Theorem 6

We begin with a lemma that provides some useful upper bounds.

Lemma 3: Let $Z_1, Z_2, Z_3, \ldots, Z_t$ be i.i.d with mean μ such that $0 \leq Z_i \leq 1$, $\forall\ i \in [1 : t]$. Let $m_1, m_2, \ldots, m_t$ and m be positive integers such that $\sum_i m_i = m$. Let

$$\beta = \frac{1}{m} \sum_{i=1}^t m_i Z_i.$$

Then the following hold for sufficiently large n.

1) For $d_0 > \mu$,

$$\Pr(\beta > d_0)$$
$$\leq \exp\left(-2(d_0 - \mu)^2 m^2/(m_1^2 + m_2^2 + \ldots + m_t^2)\right). \tag{40}$$

2) For $d_0 < \mu$,

$$\Pr(\beta < d_0)$$
$$\leq \exp\left(-2(d_0 - \mu)^2 m^2/(m_1^2 + m_2^2 + \ldots + m_t^2)\right). \tag{41}$$

3) For any positive constant c, there exists a positive constant a such that

$$E\left(\exp(-cm(\beta - d_0)^2)\right)$$
$$\leq \exp\left(-a(d_0 - \mu)^2 m^2/(m_1^2 + m_2^2 + \ldots + m_t^2)\right). \tag{42}$$

Proof of Lemma 3: (40) and (41) are direct applications of the Chernoff bound [10, Theorem 2]. (This particular form is also known as Hoeffding's inequality.) To prove (42), first assume that $d_0 > \mu$. Then

$$E\left(\exp(-cm(\beta - d_0)^2)\right)$$
$$\leq \Pr(|\beta - d_0| < (d_0 - \mu)/2) + \exp(-cm(d_0 - \mu)^2/4)$$
$$\leq \Pr(\beta > d_0 - (d_0 - \mu)/2) + \exp(-cm(d_0 - \mu)^2/4)$$
$$\leq \exp\left(-(d_0 - \mu)^2 m^2/2(m_1^2 + m_2^2 + \ldots + m_t^2)\right)$$
$$+ \exp(-cm(d_0 - \mu)^2/4) \quad \text{from (40)}.$$

Now, from the Cauchy-Schwarz inequality, we have

$$m^2/(m_1^2 + m_2^2 + \ldots + m_t^2) \leq t \leq m.$$

This gives with $a < \min\{1/2, c/4\}$,

$$E\left(\exp(-cm(\beta - d_0)^2)\right)$$
$$\leq \exp\left(-a(d_0 - \mu)^2 m^2/(m_1^2 + m_2^2 + \ldots + m_t^2)\right), \tag{43}$$

for sufficiently large n.

To prove (42) in the case $d_0 < \mu$, first note that (40) and (41) hold even when the random variables take values in $[-1, 0]$. Then apply the above result for the random variables $-Z_i$, $i \in [1 : t]$. ∎

As in the proof of Theorem 3, we first analyze the probability of error conditioned on the erasure sequence and $\mathbf{X}$. Let $\bar{E}$ denote the erasure matrix. That is,

$$\bar{E} = (1(Y_{ij} = e))_{m \times n} \in \{0, 1\}^{m \times n}.$$

Rows in Same Cluster: First consider case when $I_{ij} = 1$, i.e. i, j are in the same cluster. We wish to evaluate the probability of error $\Pr(d_{ij} \geq d_0 | I_{ij} = 1, \bar{E}, \mathbf{X})$. Define $s_{ij}(p, p)$ as the number of columns k such that $\mathbf{Y}_{ik} \neq e, \mathbf{Y}_{jk} \neq e$ and $\mathbf{X}_{ik} = p$. Clearly,

$$\sum_p s_{ij}(p, p) = N_{ij}.$$

Note that for such k, the indicator $1(Y_{ik} \neq Y_{jk})$ has mean μ_{pp}. Hence, for $I_{ij} = 1$ and a fixed $\bar{E}$, the random variable $N_{ij}d_{ij}$ is given by

$$N_{ij}d_{ij} = \sum_{p \in \mathsf{A}} \sum_{\substack{\mathbf{Y}_{ik}, \mathbf{Y}_{jk} \neq e \\ \mathbf{X}_{ik} = p = \mathbf{X}_{jk}}} 1\left(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk}\right).$$

The above summation has $s_{ij}(p, p)$ i.i.d Bernoulli random variables of mean μ_{pp}, for each $p \in \mathsf{A}$, all the random variables being independent. Hence the charcterstic function of $N_{ij}d_{ij}$ (for $I_{ij} = 1$, fixed $\bar{E}$ and $\mathbf{X}$) is given by

$$\prod_{p \in \mathsf{A}} (1 - \mu_{pp} + \mu_{pp}e^\theta)^{s_{ij}(p,p)}.$$

Using the Chernoff Bound, we have

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1, \bar{E}, \mathbf{X}\right)$$
$$\leq \frac{\prod_p (1 - \mu_{pp} + \mu_{pp}e^\theta)^{s_{ij}(p,p)}}{e^{d_0 N_{ij}\theta}}, \quad \text{for any } \theta \geq 0.$$

By using the inequality $1 + x \leq e^x$, we obtain

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1, \bar{E}, \mathbf{X}\right)$$
$$\leq \exp\left(N_{ij}\beta_{ij}(e^\theta - 1) - N_{ij}d_0\theta\right), \quad \theta \geq 0,$$

where

$$\beta_{ij} = \frac{\sum_{p \in \mathsf{A}} \mu_{pp}s_{ij}(p,p)}{N_{ij}}. \tag{44}$$

Using

$$\theta = \begin{cases} \max\{\ln(d_0/\beta_{ij}), 0\} & \text{if } \beta_{ij} \neq 0 \\ \infty & \text{if } \beta_{ij} = 0, \end{cases}$$

we obtain

$$\Pr\left(d_{ij} \geq d_0 \middle| I_{ij} = 1, \bar{E}, \mathbf{X}\right)$$

$$\leq \begin{cases} \exp\left(N_{ij}(d_0 - \beta_{ij}) + N_{ij}d_0 \ln\left(\frac{\beta_{ij}}{d_0}\right)\right) & \text{if } 0 < \beta_{ij} \leq d_0 \\ 1 \text{ if } \beta_{ij} > d_0 \\ 0 \text{ if } \beta_{ij} = 0. \end{cases} \tag{45}$$

For tractability, we further simplify this bound. To do so we note that for $-1 < x \leq 0$ and $0 < c < 1/2$, the function $f(x) = \ln(1 + x) - x + cx^2$ is increasing. This can be seen by noting that

$$f'(x) = \frac{x(2c - 1 + 2cx)}{1 + x}.$$

Since $x/(1 + x) \leq 0$ and $2c - 1 + 2cx < 0$, we have $f'(x) \geq 0$. Hence $\ln(1 + x) - x + cx^2 \leq 0$ in the interval $-1 < x \leq 0$. Now for $0 < \beta_{ij} \leq d_0$, $-1 < (\beta_{ij} - d_0)/d_0 \leq 0$, and so

$$\ln\left(\frac{\beta_{ij}}{d_0}\right) \leq \frac{\beta_{ij} - d_0}{d_0} - c\left(\frac{\beta_{ij} - d_0}{d_0}\right)^2.$$

Using this in (45), for $d_0 \geq \beta_{ij}$, we have

$$\Pr\left(d_{ij} \geq d_0 \middle| I_{ij} = 1, \bar{E}, \mathbf{X}\right) \leq \exp\left(-cN_{ij}\frac{(\beta_{ij} - d_0)^2}{d_0}\right). \tag{46}$$

Taking expectation over $\mathbf{X}$, we obtain

$$\Pr\left(d_{ij} \geq d_0 \middle| I_{ij} = 1, \bar{E}\right)$$

$$\leq \Pr\left(\beta_{ij} > d_0 \middle| I_{ij} = 1, \bar{E}\right)$$

$$+ E\left[\exp\left(-cN_{ij}\frac{(\beta_{ij} - d_0)^2}{d_0}\right)\middle| I_{ij} = 1, \bar{E}\right] \tag{47}$$

$$=: T_1 + T_2.$$

We next bound T_1 and T_2.

For $l \in [1 : t]$, let $n_l(\bar{E})$ denote the number of commonly sampled positions for rows i and j in the l^{th} column cluster, i.e.

$$n_l(\bar{E}) = \sum_{k \text{ in cluster } l} 1\left(Y_{ik} \neq e, Y_{jk} \neq e\right).$$

Note that

$$\sum_{l=1}^{t} n_l(\bar{E}) = N_{ij}$$

and

$$s_{ij}(p, p) = \sum_{l=1}^{t} n_l(\bar{E})1(\mathbf{X}_{i\{l\}} = p),$$

where $\mathbf{X}_{i\{l\}}$ is the rating vector of user i in the l^{th} column cluster. From (44) and the above equation,

$$N_{ij}\beta_{ij} = \sum_{p \in \mathsf{A}} \mu_{pp} \sum_{l=1}^{t} n_l(\bar{E})1(\mathbf{X}_{i\{l\}} = p)$$

$$= \sum_{l=1}^{t} n_l(\bar{E}) \sum_{p \in \mathsf{A}} \mu_{pp}1(\mathbf{X}_{i\{l\}} = p), \tag{48}$$

where the random variable

$$Z_l = \sum_{p \in \mathsf{A}} \mu_{pp}1(\mathbf{X}_{i\{l\}} = p)$$

takes the value μ_{pp} with probability $1/|\mathsf{A}|$, for each $p \in \mathsf{A}$. The mean of Z_l is $d_{lb} = \sum_p \mu_{pp}/|\mathsf{A}|$. Further, Z_l's are i.i.d. From (48), Lemma 3 can be applied to $\beta = \beta_{ij}$, Z_l, $m_i = n_i(\bar{E})$. Using (40) of Lemma 3, we have

$$T_1 \leq \exp\left(-a_1(d_0 - d_{lb})^2 N_{ij}^2/(n_1^2(\bar{E}) + n_2^2(\bar{E}) + \ldots + n_t^2(\bar{E}))\right)$$

for some positive constant a_1. Similarly using (42) of Lemma 3, we have

$$T_2 \leq \exp\left(-a_2(d_0 - d_{lb})^2 N_{ij}^2/(n_1^2(\bar{E}) + n_2^2(\bar{E}) + \ldots + n_t^2(\bar{E}))\right)$$

for some positive constant a_2. From (47), we then have

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1, \bar{E}\right) \leq T_1 + T_2$$

$$\leq 2\exp\left(-a(d_0 - d_{lb})^2 N_{ij}^2/(n_1^2(\bar{E}) + n_2^2(\bar{E}) + \ldots + n_t^2(\bar{E}))\right)$$

for some positive constant a and for sufficiently large n. Using $n_l(\bar{E}) \leq n_l$, we can loosen the bound to

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1, \bar{E}\right)$$

$$\leq 2\exp\left(-a(d_0 - d_{lb})^2 N_{ij}^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right).$$

Taking expectation over $\bar{E}$, for $\alpha = r_2(1 - \epsilon)^2$ and suitable positive constants c_1, c_2 and c, we have,

$$\Pr\left(d_{ij} \geq d_0 \Big| I_{ij} = 1\right)$$

$$\leq 2\exp\left(-a(d_0 - d_{lb})^2 \alpha^2 n^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right)$$

$$+ \Pr(N_{ij} < n\alpha)$$

$$\leq 2\exp\left(-c_1 n^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right) + \exp(-c_2 n) \tag{49}$$

$$\leq c' \exp\left(-cn^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right), \tag{50}$$

where in (49) we have used (20). (50) is obtained by a similar argument as used to obtain (43) using Cauchy Schwarz inequality.

Rows in Different Clusters: Next consider the case $I_{ij} = 0$, i.e. rows i and j are in different clusters. The bounding technique is similar to the case when $I_{ij} = 1$. We wish to evaluate $\Pr(d_{ij} \leq d_0 | I_{ij} = 0, \bar{E}, \mathbf{X})$. Let $s_{ij}(p, q)$ be the number of columns k such that $\mathbf{Y}_{ik} \neq e, \mathbf{Y}_{jk} \neq e, \mathbf{X}_{ik} = p$ and $\mathbf{X}_{jk} = q$. Then for a $I_{ij} = 0$ and fixed $\bar{E}, \mathbf{X}$, the random variable $N_{ij} d_{ij}$ is given by

$$N_{ij} d_{ij} = \sum_{p,q \in \mathsf{A}} \sum_{\substack{\mathbf{Y}_{ik}, \mathbf{Y}_{jk} \neq e \\ \mathbf{X}_{ik} = p, \mathbf{X}_{jk} = q}} 1\left(\mathbf{Y}_{ik} \neq \mathbf{Y}_{jk}\right).$$

The above summation has $s_{ij}(p, q)$ i.i.d Bernoulli random variables of mean μ_{pq}, for each $(p, q) \in \mathsf{A}^2$, all the random variables being independent. Using the Chernoff Bound, we may then write

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right)$$
$$\leq \frac{\prod_{p,q} (1 - \mu_{pq} + \mu_{pq} e^\theta)^{s_{ij}(p,q)}}{e^{d_0 N_{ij} \theta}} \, \theta \leq 0.$$

By using the inequality $1 + x \leq e^x$ we obtain

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right)$$
$$\leq \exp\left(N_{ij} \beta_{ij}(e^\theta - 1) - N_{ij} d_0 \theta\right), \quad \text{for any } \theta \leq 0,$$

where

$$\beta_{ij} = \frac{\sum_{p,q \in \mathsf{A}} \mu_{pq} s_{ij}(p, q)}{N_{ij}}. \tag{51}$$

Using $\theta = \min\{\ln(d_0 / \beta_{ij}), 0\}$, we obtain

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}, \mathbf{X}\right)$$
$$\leq \begin{cases} \exp\left(N_{ij}(d_0 - \beta_{ij}) + N_{ij} d_0 \ln\left(\frac{\beta_{ij}}{d_0}\right)\right) & \text{if } \beta_{ij} \geq d_0 \\ 1 & \text{if } \beta_{ij} < d_0. \end{cases} \tag{52}$$

But $s_{ij}(p, q) \leq N_{ij}$, and so

$$\frac{\beta_{ij}}{d_0} \leq \frac{\sum_{p,q} \mu_{pq}}{d_0} = 1 + s,$$

where s is defined as

$$s := \frac{\sum_{p,q} \mu_{pq}}{d_0} - 1.$$

So for $\beta_{ij} \geq d_0$, we have $0 \leq (\beta_{ij} - d_0)/d_0 \leq s$. But for any $0 < c < 1/(2(1+s))$, the function $f(x) = \ln(1+x) - x + cx^2$ is a decreasing function on $[0, s]$. So we have the following

$$\ln\left(\frac{\beta_{ij}}{d_0}\right) \leq \frac{\beta_{ij} - d_0}{d_0} - c\left(\frac{\beta_{ij} - d_0}{d_0}\right)^2.$$

Using this in (52), for $\beta_{ij} \geq d_0$, we have

$$\Pr\left(d_{ij} \leq d_0 \big| I_{ij} = 0, \bar{E}, \mathbf{X}\right) \leq \exp\left(-cN_{ij}\frac{(\beta_{ij} - d_0)^2}{d_0}\right). \tag{53}$$

Taking expectation over $\mathbf{X}$, we obtain

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0, \bar{E}\right)$$
$$\leq \Pr\left(\beta_{ij} < d_0 \Big| I_{ij} = 0, \bar{E}\right)$$
$$+ E\left[\exp\left(-cN_{ij}\frac{(\beta_{ij} - d_0)^2}{d_0}\right) \Big| I_{ij} = 0, \bar{E}\right]. \tag{54}$$

Then we follow the same line of arguments as in the case when $I_{ij} = 1$. Note that now

$$N_{ij}\beta_{ij} = \sum_{p \in \mathsf{A}} \mu_{pp} \sum_{l=1}^{t} n_l(\bar{E}) 1(\mathbf{X}_{i\{l\}} = p)$$
$$= \sum_{l=1}^{t} n_l(\bar{E}) \sum_{p \in \mathsf{A}} \mu_{pp} 1(\mathbf{X}_{i\{l\}} = p), \tag{55}$$

where the random variable

$$Z_l = \sum_{p,q \in \mathsf{A}} \mu_{pq} 1(\mathbf{X}_{i\{l\}} = p) 1(\mathbf{X}_{j\{l\}} = q)$$

takes the value μ_{pq} with probability $1/|\mathsf{A}|^2$. The mean of Z_l is $d_{ub} = \sum_{p,q} \mu_{pq}/|\mathsf{A}|^2$. Further, Z_l's are i.i.d Applying Lemma 3 and (20) as in the case of $I_{ij} = 1$, we again have

$$\Pr\left(d_{ij} \leq d_0 \Big| I_{ij} = 0\right)$$
$$\leq c' \exp\left(-cn^2/(n_1^2 + n_2^2 + \ldots + n_t^2)\right).$$

Since there are at most $m(m-1)/2$ pairs of rows, the result follows by the union bound. ∎

REFERENCES

[1] http://www.netflixprize.com/

[2] S.T. Aditya, O. Dabeer, B.K. Dey, "A Channel Coding Perspective of Recommendation Systems," *ISIT*, 2009, Seoul, Korea.

[3] G. Adomavicius, A. Tuzhilin, "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Tran. Knowledge and Data Engineering*, vol. 17, no. 6, pp. 734-749, June 2005.

[4] E. J. Candes, Y. Plan, "Matrix completion with noise," *Arxiv Preprint arXiv:0903.3131*, 2009

[5] E. J. Candes, B. Recht, "Exact Matrix Completion via Convex Optimization," *Arxiv Preprint arXiv:0805.4471*, 2008

[6] E. J. Candes, T. Tao, "The Power of Convex Relaxation: Near-Optimal Matrix Completion," *Arxiv Preprint arXiv:0903.1476*, 2009

[7] S. Chakrabarti, "Mining the Web," Morgan Kaufmann Publishers, San Fransisco, 2003

[8] T.M. Cover, J.A. Thomas, *Elements of Information Theory,* Second Edition, John Wiley & Sons, 2006

[9] A. Felfernig, G. Friedrich, L. Schmidt-Thieme, "Guest Editor's Introduction: Recommender Systems," *IEEE Intelligent Systems*, vol. 22 no. 3, pp. 18-21, May 2007

[10] W. Hoeffding, "Probability Inequalities for Sums of Bounded Random Variables," *Journal of the American Statistical Association*, vol. 58, no. 301, pp. 13-30, Mar 1963.

[11] R. Keshavan, A. Montanari, S. Oh, "Learning low rank matrices from O(n) entries," *Allerton* 2008.

[12] R. Keshavan, A. Montanari, S. Oh, "Matrix Completion from a Few Entries," *ISIT* 2009, Seoul, Korea.

[13] Y. Koren, "Factorization Meets the Neighborhood: a Multifaceted Collaborative Filtering Model," *ACM Int. Conference on Knowledge Discovery and Data Mining (KDD'08)*, 2008

[14] K. Lee and Y. Bresler, "ADMiRA: Atomic Decomposition for Minimum Rank Approximation," *Arxiv Preprint arXiv:0905.0044*, 2009

[15] M. Mitzenmacher, E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis,* Cambridge University Press, 2005

[16] B. Recht, M. Fazel, P. A. Parrilo, "Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization," submitted to *SIAM Review*, 2007