

Quantized Radio Map Estimation Using Tensor and Deep Generative Models

Subash Timilsina, Sagar Shrestha, and Xiao Fu

Abstract—*Spectrum cartography* (SC), also known as *radio map estimation* (RME), aims at crafting multi-domain (e.g., frequency and space) radio power propagation maps from limited sensor measurements. While early methods often lacked theoretical support, recent works have demonstrated that radio maps can be provably recovered using low-dimensional models—such as the *block-term tensor decomposition* (BTD) model and certain *deep generative models* (DGMs)—of the high-dimensional multi-domain radio signals. However, these existing provable SC approaches assume that sensors send real-valued (full-resolution) measurements to the fusion center, which is unrealistic. This work puts forth a *quantized* SC framework that generalizes the BTD and DGM-based SC to scenarios where heavily quantized sensor measurements are used. A maximum likelihood estimation (MLE)-based SC framework under a Gaussian quantizer is proposed. Recoverability of the radio map using the MLE criterion is characterized under realistic conditions, e.g., imperfect radio map modeling and noisy measurements. Simulations and real-data experiments are used to showcase the effectiveness of the proposed approach.

Index Terms—Radio map estimation, spectrum cartography, block-term tensor decomposition, deep generative model, Gaussian quantization.

I. INTRODUCTION

The *spectrum cartography* (SC) [also known as *radio map estimation* (RME)] technique was proposed to build multi-domain (e.g., space, frequency, and time) radio maps from limited sensor measurements that are sparsely acquired over a geographical area; see, e.g., [1]–[8]. The radio maps capture key characteristics of the radio frequency (RF) environment (e.g., interference propagation), and thus are critical for various tasks, such as opportunistic access, spectrum surveillance, beamforming, power allocation, and interference management; see [9]. From a signal processing viewpoint, estimating a multi-domain, high-dimensional and high-resolution radio map from limited sensor-acquired samples/measurements is an ill-posed inverse problem—where an infinite number of solutions exist in general, hindering the recoverability of the radio map.

Early SC/RME methods often assume that power propagation is smooth over the space. This made it possible to use various interpolation techniques for SC, e.g., the Kriging interpolation [3], thin plate splines [2], kernel methods [10], and Gaussian radial basis functions (RBF) [11]. Sparse representations of the radio maps in certain domains are also often leveraged for SC; see, e.g., [2], [4], [5], [12]. In recent years, low-rank matrix and tensor completion techniques were

The authors are with the School of EECS at Oregon State University. This work was supported in part by the National Science Foundation (NSF) under Projects NSF ECCS-2024058 and NSF CCF-2210004.

advocated for SC; see, e.g., [6], [13]–[15]. From a low-rank model completion viewpoint, some of these methods were shown to ensure the recovery of the radio map; see, e.g., [6], [15]. However, handcrafted priors such as sparsity and low rank do not always match with the reality, especially when heavy shadowing exists, e.g., in crowded urban or indoor environments; see discussions in [7].

To better model heavily shadowed environments, a number of deep learning-based methods were employed for SC. Compared to handcrafted prior-based approaches, deep learning-based methods need a training phase using oftentimes a large amount of training data, which creates extra workload. However, deep neural networks can represent complex scenarios as nonlinear functions in a succinct way, and thus can recover the radio map accurately under heavy shadowing. Earlier deep learning-based SC works in [8], [16]–[19] formulated the SC problem as image inpainting problems. The more recent work in [7] “embeds” a *deep generative model* (DGM)-learned prior into the *spatial loss fields* (SLFs) of the emitters. Then, a DGM-based data recovery problem was formulated. Compared to the inpainting-based methods, e.g., [8], [18], the embedded DGM-based method in [7] enjoys a lighter training burden and exhibits better generalizability. More importantly, it was shown to ensure recoverability of the radio map under realistic conditions—and such theoretical guarantees had been lacking in prior deep learning-based SC works.

Challenges. The recent developments of theory-backed SC (e.g., the works based on tensor [6] and DGM [7]) were built upon the premise that the sensors communicate with the fusion center using full-resolution real-valued measurements. This setup is unrealistic as real-world communication systems often designate limited bandwidth for signaling. The measurements sent to the fusion center by the sensors are usually heavily quantized. Nonetheless, SC using quantized measurements was much less studied in the literature. The line of “frugal sensing” work [20], [21] considered 1-bit power spectrum estimation, but did not consider the reconstruction of the spatial information. The work in [10] used kernel regression for SC with quantized sensor feedback. However, it assumed that the power spectral density (PSD) of each emitter is known to the fusion center—but estimating the PSDs itself is a highly non-trivial task [22], [23]. To our best knowledge, *blind* SC methods (i.e., SC without assuming knowing the emitters’ PSDs) using quantized measurements and the associated theoretical understanding have not been studied—for both handcrafted prior and DGM-based SC frameworks.

Contributions. In this work, we propose a quantized SC framework with provable recovery guarantees of the ground-

truth radio map. Our detailed contributions are as follows:

- **A Quantized SC Framework.** We propose a quantized SC framework that can flexibly work with various radio map models. We consider the setting where the measurements sent by the sensors are outputs of a Gaussian quantizer. Then, we formulate the quantized SC problem using the maximum-likelihood estimation (MLE) principle. We combine the quantization strategy with two radio map models, namely, *block-term tensor decomposition* (BTD) from [6] and DGM from [7], and design algorithms for tackling their respective MLE formulations. Both models are well-motivated and useful: The former is more economical to deploy (as no training data is needed) and the latter is more resilient to heavy shadowing.

- **Recoverability Analysis.** On the theory side, we show that the formulated MLE criteria under both the tensor model and the DGM of the radio maps ensure recovering the ground truth (up to bounded errors), if reasonable conditions are met. In particular, we show that our framework enjoys provable recovery even under challenging scenarios, e.g., when the low-rank tensor model or the DGM do not exactly match the ground truth. The results also reveal an intuitive trade-off between the model complexity and the sample complexity. Our analysis is a nontrivial integration of generalization error analysis and quantized data recovery, with careful consideration given to sensing paradigms in the context of SC.

Part of the work will appear at ICASSP 2023 [24], which introduced the basic idea and the DGM-based formulation. The journal version additionally includes 1) a tensor decomposition-based model that does not require off-line training; 2) detailed and unified recoverability analysis under both the DGM and BTD models; 3) more comprehensive simulations; and 4) real-data experiments.

Notation We use $x \in \mathbb{R}$, $\mathbf{x} \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{I \times J}$, $\underline{\mathbf{X}} \in \mathbb{R}^{I \times J \times K}$ to denote a scalar, vector, matrix, and tensor respectively. We adopt the Matlab notation $\mathbf{X}(i, :)$ and $\mathbf{X}(:, j)$ to represent the i th row and j th column of matrix X , respectively. $\underline{\mathbf{X}}(i, j, :)$ represents the (i, j) th tensor fiber. $\underline{\mathbf{X}}(i, j, k)$ represents the element of $\underline{\mathbf{X}}$ indexed by i, j, k . \circ and \circledast represent the outer product and the Hadamard product, respectively. The outer product between matrix \mathbf{U} and vector \mathbf{v} is defined as $\underline{\mathbf{X}} = \mathbf{U} \circ \mathbf{v}$ such that, $[\mathbf{U} \circ \mathbf{v}]_{i,j,k} = \mathbf{U}(i, j)\mathbf{v}(k)$. $\mathbf{0}$ and $\mathbf{1}$ represent the all-zero and all-one vectors, respectively. $\underline{\mathbf{0}}$ and $\underline{\mathbf{1}}$ represent the all-zero and all-one tensors, respectively. $\|\mathbf{x}\|_2$ and $\|\mathbf{X}\|_2$ represent the ℓ_2 norm and spectral norm, respectively. $\|\cdot\|_{\text{F}}$ denotes the Frobenius norm of matrices and tensors. We define $\|\underline{\mathbf{X}}\|_{\infty} = \max_{i,j,k} |\underline{\mathbf{X}}(i, j, k)|$. $|\mathcal{X}|$ denotes the cardinality of the set \mathcal{X} . For an integer I , the set $[I]$ denotes $\{1, \dots, I\}$. $\underline{\mathbf{L}} = \log(\underline{\mathbf{X}})$ is a tensor such that $\underline{\mathbf{L}}(i, j, k) = \log(\underline{\mathbf{X}}(i, j, k))$. “ \geq ” represents element-wise “larger than or equal to” (e.g., $\mathbf{X} \geq \mathbf{0}$ means every element of \mathbf{X} is nonnegative). $f'(x)$ is the first-order derivative of f at x .

II. BACKGROUND

A. Problem Setup

We focus on the SC scenario illustrated in Fig. 1 (a). The scenario is widely considered in the literature; see, e.g., [2],

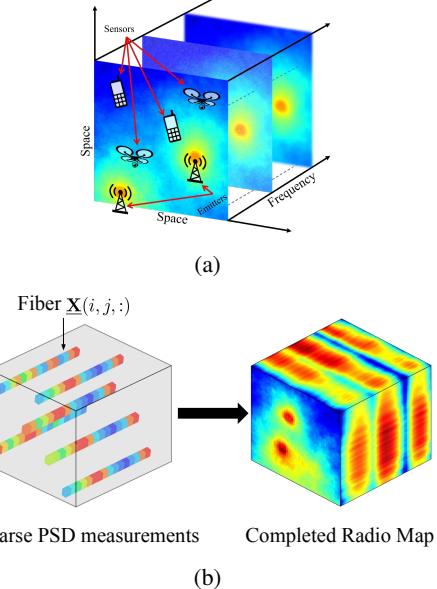


Fig. 1: (a) Scenario of interest. (b) Sampling pattern of the radio map; each sampled fiber is a PSD.

[6], [7], [10]. Specifically, we aim at recovering a *spatio-spectral* radio power density map induced by R emitters using measurements acquired by N sensors. The sensors are sparsely deployed over the region of interest. For the ease of exposure, we assume that the space domain is a 2D rectangle and is discretized into $I \times J$ grids (but the proposed design principles can be readily generalized to 3D cases). The frequency domain is also discretized into K frequency bins—and thus every spatial grid is associated with a K -dimensional *power spectral density* (PSD). In other words, the power propagation patterns of the R emitters over the K frequencies constitute an $I \times J \times K$ radio map tensor $\underline{\mathbf{X}} \in \mathbb{R}^{I \times J \times K}$, where the entry $\underline{\mathbf{X}}(i, j, k)$ is the PSD of the signal received at location (i, j) and frequency k . That is, every *fiber* [25] of the tensor, $\underline{\mathbf{X}}(i, j, :)$, represents the PSD of the received signal measured at the location (i, j) ; see Fig. 1 (b). We use the notation

$$\Omega = \{(i, j) | i \in [I], j \in [J]\}$$

to denote the set of sensor locations. Note that we often have

$$|\Omega| = N \ll IJ;$$

i.e., only a small number of sensors are available. We assume that every sensor acquires the full PSD $\underline{\mathbf{X}}(i_s, j_s, :)$ at its location (i_s, j_s) , where $s = 1, \dots, N$. If the sensors are able to transmit real-valued feedback to the fusion center, the goal of SC is to recover the full $\underline{\mathbf{X}}$ from the tensor fibers $\{\underline{\mathbf{X}}(i, j, :)\}_{(i, j) \in \Omega}$ at the fusion center.

B. Prior Works on Provable SC

Early developments of SC (see, e.g., [2]–[4], [12]–[14]) mostly focused on the methodology side but less considered the theoretical aspects, e.g., recoverability of the ground-truth radio map. More recently, there has been an increased research interest on theoretical understanding to SC. In this subsection,

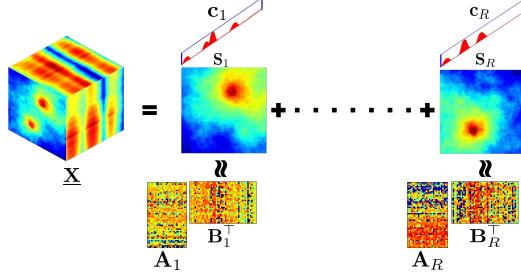


Fig. 2: The BTD-based model for radio maps used in [6].

we briefly review two SC models that were shown to guarantee recovery of the radio map.

1) Block-term Tensor Modeling and Provable Recovery:

The recent work [6] proposed a recoverability-guaranteed SC method from a block-term tensor completion viewpoint. The model in [6] starts by assuming that power propagation is coherent over the frequency band of interest, which in general holds when the ratio between the bandwidth of the frequencies of interest and its central frequency is not large [12], [22], [23], [26]. Under such circumstances, the radio map can be decomposed into the latent factors associated with spatial and spectral information, respectively [2], [6], [7], [10]:

$$\underline{X}(i, j, k) = \sum_{r=1}^R \underline{S}_r(i, j) \mathbf{c}_r(k) \iff \underline{X} = \sum_{r=1}^R \underline{S}_r \circ \mathbf{c}_r, \quad (1)$$

where $\underline{S}_r \in \mathbb{R}^{I \times J}$ is the *spatial loss field* (SLF) of emitter r , $\mathbf{c}_r \in \mathbb{R}^K$ is the PSD of emitter r , and \circ denotes the outer product. The SLF captures the spatial power propagation characteristics of an emitter, and the PSD reflects the emitter's spectral band occupancy.

The work [6] modeled each SLF \underline{S}_r in (1) as a low-rank matrix, i.e.,

$$\underline{S}_r = \mathbf{A}_r \mathbf{B}_r^\top,$$

where $\mathbf{A}_r \in \mathbb{R}^{I \times L}$ and $\mathbf{B}_r \in \mathbb{R}^{J \times L}$ with $\text{rank}(\mathbf{A}_r) = \text{rank}(\mathbf{B}_r) = L \ll \min\{I, J\}$. Fig. 2 illustrates the model. This low-rank model of \underline{S}_r connected (1) with a tensor model, namely, the *block-term tensor decomposition* (BTD) with multilinear rank- $(L, L, 1)$ model [27].

The model is well-motivated, as the individual SLFs often exhibit high correlations across the spatial domain, making \underline{S}_r 's approximately low-rank; see [6] for numerical evidence.

Using the BTD model of the radio map, the SC problem is formulated as follows:

$$\underset{\{\mathbf{c}_r, \mathbf{A}_r, \mathbf{B}_r\}_{r=1}^R}{\text{minimize}} \left\| \underline{M}_{\text{sens}} \circledast \left(\underline{X} - \sum_{r=1}^R (\mathbf{A}_r \mathbf{B}_r^\top) \circ \mathbf{c}_r \right) \right\|_{\text{F}}^2, \quad (2)$$

where $\underline{M}_{\text{sens}}$ is a sensing mask tensor such that $\underline{M}_{\text{sens}}(i, j, :) = \mathbf{1} \in \mathbb{R}^K$ if $(i, j) \in \Omega$ and $\underline{M}_{\text{sens}}(i, j, :) = \mathbf{0}$ otherwise. Notably, [6] leveraged the BTD model's essential uniqueness to show that the spatio-spectral radio map is provably recoverable under both random and *regular* sampling patterns.

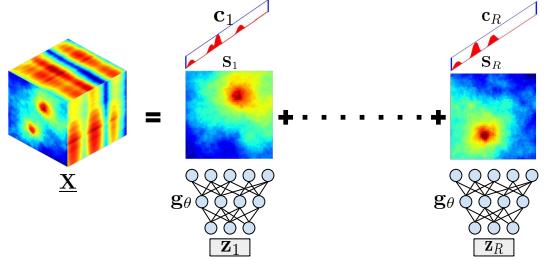


Fig. 3: Spectrum cartography using deep generative model from [7].

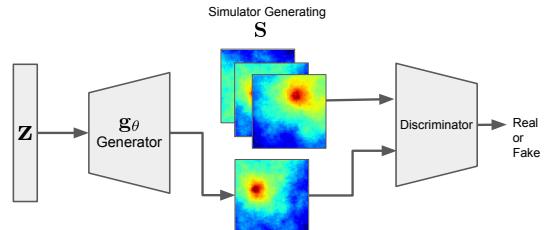


Fig. 4: Using simulated SLFs to learn a DGM of the SLFs via GAN.

2) Deep Generative Model (DGM)-Based Provable SC: In recent years, deep learning-based SC attracted much attention. The reason is that deep neural networks can represent very complex and heavily shadowed scenarios in a parsimonious way, and thus can boost performance of SC in challenging scenarios. However, most early attempts, e.g., those in [8], [18], formulated the SC problem as a supervised inpainting network learning problem. This formulation encountered training and generalization challenges, as the radio maps often have an extremely large latent state space. In addition, there is a lack of theoretical support to this line of work.

To exploit the representation power of deep neural networks with recoverability guarantees, the work [7] provided an alternative solution. The method can be regarded as an extension of the model in [6]. Realizing that the low-rank model for the SLFs in [6] may not always hold in practice, [7] modeled the SLF of every emitter using a data-driven DGM. To be specific, [7] proposed to learn a DGM of the SLF by simulating a large number of SLFs following a physical model. Then, every \underline{S}_r can be approximated as follows:

$$\underline{S}_r \approx \mathbf{g}_\theta(\mathbf{z}_r), \quad (3)$$

where $\mathbf{g}_\theta(\cdot) : \mathbb{R}^D \rightarrow \mathbb{R}^{I \times J}$ is a deep generative neural network that maps the “latent embedding” $\mathbf{z}_r \in \mathbb{R}^D$ such that $D \ll IJ$ to an SLF—see Fig. 3. The DGM is learned off-line, e.g., via autoencoder, variational autoencoder (VAE), or generative adversarial network (GAN); see Fig. 4. Using this learned DGM $\mathbf{g}_\theta(\mathbf{z}_r)$, the SC problem was cast as follows:

$$\underset{\{\mathbf{c}_r, \mathbf{z}_r\}_{r=1}^R}{\text{minimize}} \left\| \underline{M}_{\text{sens}} \circledast \left(\underline{X} - \sum_{r=1}^R \mathbf{g}_\theta(\mathbf{z}_r) \circ \mathbf{c}_r \right) \right\|_{\text{F}}^2. \quad (4)$$

The method used a DGM to describe each *individual* SLF (instead of the entire radio map as in [6], [18]). Consequently,

the DGM is much easier to learn in the training stage compared to the networks in [8], [18], which used a neural networks to learn the *aggregated* radio map from multiple emitters. This is because each SLF has a much smaller “state space” compared to that of the aggregated radio map; see more discussions in [7]. The work in [7] also showed that (4) guarantees recovering the radio map under reasonable conditions.

Remark 1 The BTD and DGM based methods have their respective advantages and disadvantages. The BTD model is simple to implement, but the *key limitation* lies in the designated low-rank prior for the SLFs. The low rank of S_r could be violated when the spatial correlation becomes weak—which typically happens when the environment has heavy shadowing effects, e.g., in urban areas. The DGM-based approach is capable of “encoding” complex, non-analytical prior information using the neural representation of $g_\theta(\cdot)$, which is suitable for modeling heavily shadowed cases. In addition, the DGM parameterization may substantially reduce the number of unknown parameters (if $\text{card}(z_r) = D \ll IJ$), which can greatly reduce the sample complexity of SC. However, the price to pay is that an extra offline training stage is needed. To generate training data for learning $g_\theta(\cdot)$, a certain level of awareness of the shadowing situation is required. These were not needed in the BTD work [6]. Hence, both models are useful and meaningful. The choice of the model depends on the scenario of interest.

C. Challenges - Lack of Quantized SC Approaches

Most of the previous SC works, including the BTD and DGM based methods, assumed that real-valued measurements (e.g., $\underline{\mathbf{X}}(i, j, :)$ for all $(i, j) \in \Omega$) can be transmitted to the fusion center (see, e.g., [2]–[7]). However, sending real-valued measurements is costly in terms of communication overhead. It is much more realistic that the sensors send quantized measurements. Quantized spectrum sensing was studied before (see, e.g., [28], [29] and the “frugal sensing” works [20], [21]) but the spatial information was not taken into consideration. Quantized SC was tackled in [10] using kernel regression. However, [10] assumed that the emitters’ PSDs are known, but estimating the PSDs itself is a hard problem [22], [23].

In addition to the lack of effective methodology, the theoretical aspects of quantized SC have yet to be studied. While the DGM-based framework in [7] and the BTD-based framework in [6] provided provable recoverability of the radio map with real-valued measurements, it remains an open question how to extend these results to handle quantized measurements without losing recoverability supports.

III. PROPOSED APPROACH

In this work, we propose a *quantized SC* framework where the sensors only feedback heavily quantized measurements. In particular, we will design a system that can flexibly work with the BTD and DGM models of the radio map in [6] and [7].

A. Sensing and Quantization

Following the setup in Sec. II-A, we assume that the sensor located at (i, j) can acquire the PSD of the received signal, i.e., $\underline{\mathbf{X}}(i, j, :) \in \mathbb{R}^K$. In addition, we consider the scenario where the sensors send a quantized version of $\underline{\mathbf{X}}(i, j, :)$. To be more specific, we employ a quantizer

$$\mathcal{Q}(\cdot) : \mathbb{R} \rightarrow \mathbb{Z}$$

that maps each element of the PSD received at (i, j) to an integer, i.e.,

$$\underline{\mathbf{Y}}(i, j, k) = \mathcal{Q}(\underline{\mathbf{X}}(i, j, k)), \quad k = 1, \dots, K$$

where $\underline{\mathbf{Y}}(i, j, :) \in \mathbb{Z}^K$ is the quantized PSD.

We employ the Gaussian quantization strategy that is widely in the literature [30]–[36]. To be specific, the quantization strategy is expressed as follows:

$$\begin{aligned} \underline{\mathbf{Y}}(i, j, k) &= \mathcal{Q}(h(\underline{\mathbf{X}}(i, j, k)) + \underline{\mathbf{V}}(i, j, k)) \quad (5) \\ \mathcal{Q}(x) &= q, \text{ if } b_{q-1} < x \leq b_q, \quad q \in [Q], \end{aligned}$$

where $\underline{\mathbf{V}}(i, j, k) \sim \mathcal{N}(0, \sigma^2)$ for all i, j, k are i.i.d. zero-mean Gaussian noise, and $\{b_q\}_{q=1}^Q$ are the pre-specified quantization bins (i.e., $[b_{q-1}, b_q]$ is the q th quantization interval). Here, the function $h(\cdot)$ is an invertible transformation function,

$$h(x) = \log(x + a), \quad (6)$$

where $a > 0$. For notation simplicity, we slightly abuse the notation by applying h to both entries and tensors, i.e., $h(\underline{\mathbf{X}})$ means taking the transform in (6) for each entry of $\underline{\mathbf{X}}$. Note that the function $h(\cdot)$ was not used in classic Gaussian quantization. In this work, we use this function to control the dynamic range of the data (also see Sec. IV-C for more discussions). Note that as h is invertible, it does not lose information of $\underline{\mathbf{X}}$; i.e., recovering $h(\underline{\mathbf{X}})$ from the quantized data recovers $\underline{\mathbf{X}}$.

Quantization with artificial noise like in (5) is called *dithering* in signal processing [37]. Dithering is known to be beneficial for retaining more information about the original signal in the quantized version by reducing the correlation between the quantization error and original signal [38].

B. Maximum Likelihood Estimation

Let $\underline{\mathbf{M}} = h(\underline{\mathbf{X}})$ be the transformed measurements of radio map. The quantized observations $\underline{\mathbf{Y}}$ has the following distribution:

$$\begin{aligned} \underline{\mathbf{Y}}(i, j, k) = q, \text{ with probability } f_q(\underline{\mathbf{M}}(i, j, k)), \\ \forall (i, j) \in \Omega, k \in [K] \end{aligned} \quad (7)$$

where f_q is defined as,

$$\begin{aligned} f_q(\underline{\mathbf{M}}(i, j, k)) &= \mathbb{P}(\underline{\mathbf{Y}}(i, j, k) = q \mid \underline{\mathbf{M}}(i, j, k)) \\ &= \Phi(b_q - \underline{\mathbf{M}}(i, j, k)) - \Phi(b_{q-1} - \underline{\mathbf{M}}(i, j, k)), \end{aligned} \quad (8)$$

in which $\Phi(x)$ is the cumulative distribution function (CDF) of the zero-mean and σ^2 -variance Gaussian variable (i.e., the CDF of the dithering noise in (5)). Using (8) and a certain parameterization of the radio map, one can formulate the SC problem as an MLE problem. In the next two subsections, we will use the BTD model in [6] and the DGM model in [7] to formulate two MLE criteria.

1) *BTD-Based Quantized SC*: As in [6], we first represent the radio map using the BTD model (see Fig. 2), i.e.,

$$\underline{\mathbf{X}}^{\text{BTD}} = \sum_{r=1}^R (\mathbf{A}_r \mathbf{B}_r^\top) \circ \mathbf{c}_r.$$

Using the described sensing and quantization procedure as in the previous section, the fusion center receives a quantized and incomplete tensor $\underline{\mathbf{Y}}(\Omega, :)$. To recover the ground-truth radio map, we consider the following MLE:

MLE via BTD

$$\underset{(\mathbf{A}, \mathbf{B}, \mathbf{C}) \in \mathcal{S}^{\text{BTD}}}{\text{minimize}} \quad F_{\Omega, \underline{\mathbf{Y}}}^{\text{BTD}}(\mathbf{A}, \mathbf{B}, \mathbf{C}), \quad (9)$$

where the objective function is defined as

$$F_{\Omega, \underline{\mathbf{Y}}}^{\text{BTD}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = - \sum_{(i,j) \in \Omega} \sum_{k=1}^K \sum_{q=1}^Q \mathbb{1}_{[\underline{\mathbf{Y}}(i,j,k)=q]} \log(f_q(h(\underline{\mathbf{X}}^{\text{BTD}}(i,j,k)))),$$

in which we have $\underline{\mathbf{X}}^{\text{BTD}}(i,j,k) = \left[\sum_{r=1}^R (\mathbf{A}_r \mathbf{B}_r^\top) \circ \mathbf{c}_r \right]_{i,j,k}$, $\mathbf{A}_r \in \mathbb{R}^{I \times L}$, $\mathbf{B}_r \in \mathbb{R}^{J \times L}$ and $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_R]$. In addition, we have used the notation $\mathbf{A} = [\mathbf{A}_1, \dots, \mathbf{A}_R]$ and $\mathbf{B} = [\mathbf{B}_1, \dots, \mathbf{B}_R]$. The constraint set is defined as follows:

$$\begin{aligned} \mathcal{S}^{\text{BTD}} = & \left\{ (\mathbf{A}, \mathbf{B}, \mathbf{C}) \mid \mathbf{A} \geq \mathbf{0}, \mathbf{B} \geq \mathbf{0}, \mathbf{C} \geq \mathbf{0} \right. \\ & \left. \|\mathbf{A}_r\|_{\text{F}} \leq \sqrt{\beta}, \|\mathbf{B}_r\|_{\text{F}} \leq \sqrt{\beta}, \|\mathbf{c}_r\|_2 \leq \kappa, \forall r \in [R] \right\}. \end{aligned} \quad (10)$$

The constraints in (10) mean that the SLFs and the PSDs are both nonnegative and bounded (with $\|\mathbf{S}_r\|_{\text{F}} \leq \beta$)—which are mild assumptions per their physical meaning. Particularly, we assume that the nonnegativity and boundedness of $\mathbf{S}_r = \mathbf{A}_r \mathbf{B}_r^\top$ are realized via having constraints on the latent factors of the BTD model. We also define the corresponding set:

$$\begin{aligned} \mathcal{X}^{\text{BTD}} = & \quad (11) \\ & \left\{ \underline{\mathbf{X}} \mid \underline{\mathbf{X}} = \sum_{r=1}^R (\mathbf{A}_r \mathbf{B}_r^\top) \circ \mathbf{c}_r, (\mathbf{A}, \mathbf{B}, \mathbf{C}) \in \mathcal{S}^{\text{BTD}} \right\}. \end{aligned}$$

2) *DGM-Based Quantized SC*: As mentioned, the BTD model-based approach may not work well when the spatial region exhibits heavy shadowing [7]. In such cases, it is natural to employ the DGM-based model in (3) to parameterize the radio map, i.e.,

$$\underline{\mathbf{X}}^{\text{DGM}} = \sum_{r=1}^R \mathbf{g}_\theta(\mathbf{z}_r) \circ \mathbf{c}_r, \quad (12)$$

where $\mathbf{g}_\theta(\cdot)$ is learned in an off-line manner as in [7]; see Fig. 3. The radio map model in (12) leads to the following MLE-based recovery criterion:

MLE via DGM

$$\underset{(\mathbf{C}, \underline{\mathbf{Z}}) \in \mathcal{S}^{\text{DGM}}}{\text{minimize}} \quad F_{\Omega, \underline{\mathbf{Y}}}^{\text{DGM}}(\mathbf{Z}, \mathbf{C}), \quad (13)$$

where we have

$$\begin{aligned} F_{\Omega, \underline{\mathbf{Y}}}^{\text{DGM}}(\mathbf{Z}, \mathbf{C}) = & \quad (14) \\ & - \sum_{(i,j) \in \Omega} \sum_{k=1}^K \sum_{q=1}^Q \mathbb{1}_{[\underline{\mathbf{Y}}(i,j,k)=q]} \log(f_q(h(\underline{\mathbf{X}}^{\text{DGM}}(i,j,k)))), \end{aligned}$$

in which $\underline{\mathbf{X}}^{\text{DGM}}(i,j,k) = [\sum_{r=1}^R \mathbf{g}_\theta(\mathbf{z}_r) \circ \mathbf{c}_r]_{i,j,k}$, $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_R]$, and $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_R]$. The constraint set is defined as follows:

$$\begin{aligned} \mathcal{S}^{\text{DGM}} = & \left\{ (\mathbf{C}, \mathbf{Z}) \mid \mathbf{g}_\theta(\mathbf{z}_r) \geq \mathbf{0}, \mathbf{C} \geq \mathbf{0}, \right. \\ & \left. \|\mathbf{g}_\theta(\mathbf{z}_r)\|_{\text{F}} \leq \beta, \|\mathbf{c}_r\|_2 \leq \kappa, \forall r \in [R] \right\}, \end{aligned} \quad (15)$$

where β and κ are positive constants. Similar as before, we define a corresponding set in the data domain:

$$\begin{aligned} \mathcal{X}^{\text{DGM}} = & \quad (16) \\ & \left\{ \underline{\mathbf{X}} \mid \underline{\mathbf{X}} = \sum_{r=1}^R \mathbf{g}_\theta(\mathbf{z}_r) \circ \mathbf{c}_r, (\mathbf{Z}, \mathbf{C}) \in \mathcal{S}^{\text{DGM}} \right\}. \end{aligned}$$

As in the continuous measurement case (cf. Remark 1), the BTD-based formulation in (9) is better suited for cases where training data is unavailable or the training cost is not affordable. The DGM-based formulation in (13) has a more expressive model to attain enhanced recovery accuracy, especially under challenging scenarios, e.g., heavy shadowing.

IV. OPTIMIZATION AND IMPLEMENTATION

In this section, we propose algorithms for handling the MLE problems in (9) and (13).

A. Algorithm For BTD-Based MLE

To tackle Problem (9), we consider the following reformulated approximation:

$$\begin{aligned} \underset{\mathbf{A}, \mathbf{B}, \mathbf{C}}{\text{minimize}} \quad & \mathcal{L}_{\Omega, \underline{\mathbf{Y}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) \\ \text{subject to} \quad & \mathbf{A} \geq \mathbf{0}, \mathbf{B} \geq \mathbf{0}, \mathbf{C} \geq \mathbf{0}, \end{aligned} \quad (17)$$

where the objective function is expressed as follows:

$$\begin{aligned} \mathcal{L}_{\Omega, \underline{\mathbf{Y}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = & F_{\Omega, \underline{\mathbf{Y}}}^{\text{BTD}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + \lambda_1 r_1(\mathbf{A}) \\ & + \lambda_2 r_2(\mathbf{B}) + \lambda_3 r_3(\mathbf{C}), \end{aligned}$$

in which the regularization terms are as follows:

$$r_i(\cdot) = \|\cdot\|_{\text{F}}^2, \quad i = 1, 2, 3.$$

Note that promoting small $\|\mathbf{A}\|_{\text{F}}^2$ and $\|\mathbf{B}\|_{\text{F}}^2$ encourages small-energy SLFs $\|\mathbf{A}_r \mathbf{B}_r^\top\|_{\text{F}}$. Hence, such regularization terms *indirectly* enforce the constraints in (10), i.e., that the SLFs and the PSDs are bounded. Nonetheless, using the regularized reformulation instead of the constraints in (10) is more convenient for designing optimization algorithms.

We propose a block coordinate descent (BCD) procedure to handle (17). In the k th iteration, we first update \mathbf{C} as follows:

$$\begin{aligned} \mathbf{C}^{(k+1)} \leftarrow & \max(\mathbf{C}^{(k)} - \xi_3^{(k)} \nabla_{\mathbf{C}}(\mathcal{L}_{\Omega, \underline{\mathbf{Y}}}(\mathbf{A}^{(k)}, \mathbf{B}^{(k)}, \mathbf{C}^{(k)})) \\ & + \lambda_3 r_3(\mathbf{C})), \mathbf{0}), \end{aligned} \quad (18)$$

where $\bar{\nabla}_C(\cdot)$ is a gradient-related direction of $F_{\Omega, \underline{Y}}$ w.r.t C . Using the gradient w.r.t. C^1 , the direction $\bar{\nabla}_C(\cdot)$ is found by popular methods such as plain-vanilla gradient and momentum-assisted methods like Adagrad [40] and Adam [41]), and $\xi_3^{(k)}$ is the step size. The “ $\max(\cdot, 0)$ ” step projects the updated $C^{(k)}$ back to the nonnegative orthant.

Similarly, the update of A and B can be done as follows:

$$\begin{aligned} \mathbf{A}^{(k+1)} &\leftarrow \max(\mathbf{A}^{(k)} - \xi_1^{(k)} \bar{\nabla}_A(\mathcal{L}_{\Omega, \underline{Y}}(\mathbf{A}^{(k)}, \mathbf{B}^{(k)}, \mathbf{C}^{(k)}), \\ &\quad + \lambda_1 r_1(\mathbf{A})), \mathbf{0}), \end{aligned} \quad (19)$$

$$\begin{aligned} \mathbf{B}^{(k+1)} &\leftarrow \max(\mathbf{B}^{(k)} - \xi_2^{(k)} \bar{\nabla}_B(\mathcal{L}_{\Omega, \underline{Y}}(\mathbf{A}^{(k)}, \mathbf{B}^{(k)}, \mathbf{C}^{(k)}), \\ &\quad + \lambda_2 r_2(\mathbf{B})), \mathbf{0}), \end{aligned} \quad (20)$$

where $\xi_1^{(k)}$ and $\xi_2^{(k)}$ are the step sizes for A and B updates respectively. In practice, the step size parameters can be chosen using heuristics advocated in Adagrad or Adam, which are often effective. After the algorithm is terminated, the estimated radio map $\widehat{\mathbf{X}}$ is obtained using

$$\widehat{\mathbf{X}} = \sum_{r=1}^R (\widehat{\mathbf{A}}_r \widehat{\mathbf{B}}_r^\top) \circ \widehat{\mathbf{c}}_r,$$

where $\widehat{\mathbf{A}}_r$, $\widehat{\mathbf{B}}_r$ and $\widehat{\mathbf{c}}_r$ are the algorithm-found solutions.

Note that the procedure above is a typical inexact BCD approach, whose convergence properties were well studied and discussed in the literature; see [42]. The algorithm described in (18), (19) and (20) is referred to as *quantized spectrum cartography via BTD* (QuantSC-BTD).

B. Algorithm For DGM-Based MLE

Problem (13) can be tackled using similar ideas. To see this, we consider the following reformulation:

$$\begin{aligned} &\underset{\mathbf{C}, \mathbf{Z}}{\text{minimize}} \quad \mathcal{L}_{\Omega, \underline{Y}}(\mathbf{Z}, \mathbf{C}), \\ &\text{subject to } \mathbf{C} \geq \mathbf{0}, \end{aligned} \quad (21)$$

where

$$\mathcal{L}_{\Omega, \underline{Y}}(\mathbf{Z}, \mathbf{C}) = F_{\Omega, \underline{Y}}^{\text{DGM}}(\mathbf{Z}, \mathbf{C}) + \mu_1 q_1(\mathbf{Z}) + \mu_2 q_2(\mathbf{C}).$$

In practice, $q_1(\cdot) : \mathbb{R}^{R \times D} \rightarrow \mathbb{R}$ is used to regularize \mathbf{Z} (and thus also regularizing $\mathbf{S}_r = \mathbf{g}_\theta(\mathbf{z}_r)$). In this work, we set

$$q_1(\mathbf{Z}) = \sum_{r=1}^R \|\mathbf{z}_r\|_2^2 = \|\mathbf{Z}\|_{\text{F}}^2,$$

which would encourage the solution to reduce the Euclidean norm of the found SLFs, so that the specification in (15) is more likely to hold. Similar as before, we set $q_2(\mathbf{C}) = \|\mathbf{C}\|_{\text{F}}^2$. The nonnegativity constraint on \mathbf{C} is again added to reflect its physical meaning. Note that \mathbf{Z} needs not to be nonnegative and thus there is no such constraint on the latent embeddings. The SLFs are nonnegative and bounded. Hence, $\mathbf{g}_\theta(\cdot)$ uses the sigmoid activation functions at the output layer. Using sigmoid makes $\|\widehat{\mathbf{S}}_r\|_\infty \leq 1$. This does not lose generality or hurt the recoverability. To see this, assume that $\widehat{\mathbf{S}}_r = \mathbf{S}_r$ and $\widehat{\mathbf{c}}_r = \mathbf{c}_r$

¹Throughout this paper, the gradient computations are all implemented numerically by PyTorch [39].

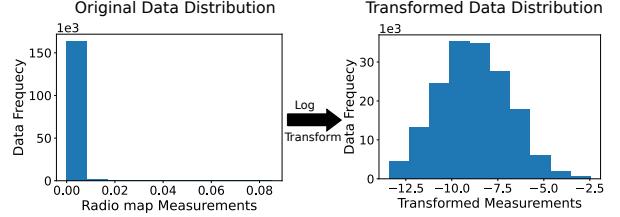


Fig. 5: Histograms of measurement values before (left) and after (right) the log-transform in (6). The radio map covers a $51 \times 51 \text{m}^2$ region with $R = 4$ emitters; $K = 64$ frequencies; shadowing parameters are $(X_c, \eta) = (50, 6)$ (see definitions in Sec. VI-A).

are estimated perfectly. One can always let $\widehat{\mathbf{S}}_r = \mathbf{S}_r / \|\mathbf{S}_r\|_\infty$ (so that $\|\widehat{\mathbf{S}}_r\|_\infty = 1$) and $\widehat{\mathbf{c}}_r = \|\mathbf{S}_r\|_\infty \mathbf{c}_r$. Such scaling/counter-scaling does not change the outer product; i.e., $\widehat{\mathbf{S}}_r \circ \widehat{\mathbf{c}}_r = \mathbf{S}_r \circ \mathbf{c}_r$ still holds—which means $\widehat{\mathbf{X}} = \underline{\mathbf{X}}$.

The BCD algorithm for tackling Problem (21) is similar as that for handling (17), except that only two blocks are updated. To be specific, in the k th iteration, we first update \mathbf{C} using gradient projection:

$$\begin{aligned} \mathbf{C}^{(k+1)} &\leftarrow \\ &\max(\mathbf{C}^{(k)} - v_2^{(k)} \bar{\nabla}_C(\mathcal{L}_{\Omega, \underline{Y}}(\mathbf{Z}^{(k)}, \mathbf{C}^{(k)}) + \mu_2 q_2(\mathbf{C})), \mathbf{0}). \end{aligned} \quad (22)$$

Then, the \mathbf{Z} -update is carried out as follows:

$$\begin{aligned} \mathbf{Z}^{(k+1)} &\leftarrow \mathbf{Z}^{(k)} - v_1^{(k)} \bar{\nabla}_Z(\mathcal{L}_{\Omega, \underline{Y}}(\mathbf{Z}^{(k)}, \mathbf{C}^{(k)})) \\ &\quad + \mu_1 q_1(\mathbf{Z})); \end{aligned} \quad (23)$$

In this case, as a DGM is involved, the gradient-related directions w.r.t. \mathbf{Z} can be found by back propagation-based methods, also using numerical gradient-finding tools. Again, when the algorithm stops, we obtain the estimated radio map $\widehat{\mathbf{X}}$ via

$$\widehat{\mathbf{X}} = \sum_{r=1}^R g_\theta(\widehat{\mathbf{z}}_r) \circ \widehat{\mathbf{c}}_r,$$

where $\widehat{\mathbf{z}}_r$ and $\widehat{\mathbf{c}}_r$ are the solutions found by the algorithm.

The algorithm in (22)-(23) is referred to as the *quantized SC via DGM* (QuantSC-DGM).

C. Quantization Interval Design

One challenge for quantizing radio maps is that the PSD values across space are extremely skewed; see Fig. 5. Directly applying simple quantizers (e.g., the uniform quantizer) does not work well. As mentioned, we propose to apply the log-transformation in (6) before quantization. Recall that the transformation is expressed as $h(x) = \log(x + a)$ with $a > 0$, which “condenses” the dynamic range of the radio map. Log-transformation is commonly used in quantization and data preprocessing for nonnegative data; see, e.g., [43].

After the log-transformation, the b_q ’s in (5) can be determined by making every interval roughly contain the same number of data points (i.e., $\underline{\mathbf{X}}(i, j, k)$). This strategy was employed in the matrix and tensor completion literature [44], [45] and was proven effective. The intuition is that we want

to make “maximum utilization” of all quantization levels. In an SC system, the $\{b_q\}$ should be determined by the system designers prior to deployment. Our suggested steps are as follows: 1) The system generates a set of simulated radio maps from a range of environment parameters that are considered reasonably reflecting the RF situation of the region of interest. 2) Using the simulated data, the system learns b_q for $q \in [Q]$ using the following principle:

$$b_q = \inf \left\{ m \left| \frac{q-1}{Q-1} \leq F(m), m = h(x) \right. \right\}, \forall q = 2, \dots, Q-1,$$

where $F(m)$ is the empirical CDF of the simulated \underline{M} ’s entries, and we set $b_0 = -\infty$ and $b_Q = +\infty$. We calculate bin boundaries from a large number (e.g., 1,000) of simulated radio maps and take average. 3) The fusion center broadcasts $\{b_1, \dots, b_Q\}$ to all the sensors. Then, the sensors use the b_q ’s to implement their quantizer.

V. RECOVERABILITY ANALYSIS

In this section, we present the recoverability analyses of the formulated problems in (9) and (13), respectively.

A. Technical Preparations

It is readily seen that the following fact holds:

Fact 1 For $\underline{X} \in \mathcal{X}^{\text{BTD}}$ and $\underline{X} \in \mathcal{X}^{\text{DGM}}$, there always exist $\alpha^{\text{BTD}} \in [0, \infty)$ and $\alpha^{\text{DGM}} \in [0, \infty)$ such that

$$\|\underline{X}\|_\infty \leq \alpha, \quad (24)$$

where $\alpha = \alpha^{\text{BTD}}$ if $\underline{X} \in \mathcal{X}^{\text{BTD}}$ and $\alpha = \alpha^{\text{DGM}}$ if $\underline{X} \in \mathcal{X}^{\text{DGM}}$.

Proof: The claim holds as the latent factors constituting \underline{X} in \mathcal{X}^{BTD} and \mathcal{X}^{DGM} are bounded; see (10) and (15). ■

As our framework uses a log-transformation to pre-process the sensor-acquired data, we will need the following lemma in our later analysis:

Lemma 1 Let $h(\underline{X})$ denote a tensor such that $[h(\underline{X})]_{i,j,k} = \log(\underline{X}(i, j, k) + a)$, where $\underline{X} \geq \mathbf{0}$ and $a > 0$. Suppose that $\|\underline{X}\|_\infty \leq \alpha$. Then, for $\underline{X} \neq \underline{X}'$, the following holds:

$$\frac{1}{\alpha + a} \leq \frac{\|h(\underline{X}) - h(\underline{X}')\|_{\text{F}}}{\|\underline{X} - \underline{X}'\|_{\text{F}}} \leq \frac{1}{a}. \quad (25)$$

The lemma says that $h(\cdot)$ exhibits continuity when the argument is nonnegative. The parameter $a > 0$ is used to prevent the mapped value goes to $-\infty$, which can be an inconsequential small value (i.e., a value that does not dominate the input signal) in practice. The proof is relegated to Appendix B. We will use the following fact:

Fact 2 Let $\mathcal{X} \in \{\mathcal{X}^{\text{BTD}}, \mathcal{X}^{\text{DGM}}\}$ and $\mathcal{M} = \{h(\underline{X}) \mid \forall \underline{X} \in \mathcal{X}\}$. Then, if $b_\ell \neq b_j$ for any $\ell \neq j \in [Q]$, there always exist bounded constants U_α , L_α and F_α such that

$$U_\alpha := \sup_m \max_{\ell \in [Q]} \log \left(\frac{1}{f_\ell(m)} \right) \quad (26a)$$

$$L_\alpha := \sup_m \max_{\ell \in [Q]} \left| \frac{\dot{f}_\ell(m)}{f_\ell(m)} \right| \quad (26b)$$

$$F_\alpha := \inf_m \max_{\ell \in [Q]} \frac{(f_\ell(m))^2}{\dot{f}_\ell(m)}, \quad (26c)$$

where α is as defined in Fact 1.

The proof of the Fact 2 is presented in Appendix C. Fact 2 means that if our radio map measurements are bounded, then $f_\ell(\cdot)$ does not change too sharply or (for at least one $\ell \in [Q]$) too slowly, reflected in L_α and F_α , respectively. The existence of these parameters makes the function locally analogous to functions that have gradient-Lipschitz continuity and strong convexity; see discussions in [30], [46].

B. Recoverability Under The BTD Model

We first analyze the recoverability of (17). Before we proceed, it should be mentioned that in the previous sections, for the simplicity of presentation, we slightly abused the notations by using \underline{X} , \underline{S}_r , \underline{c}_r , etc. to represent both the ground-truth parameters and the optimization variables (see, e.g., (17)). In the sequel, for the clarity of analysis, we will use \natural to denote the ground-truth parameters (e.g., \underline{C}^\natural), to distinguish them from the optimization variables (e.g., \underline{C} in (17)). Let us make the following assumption:

Assumption 1 There exists a constant $\nu^{\text{BTD}} \in [0, +\infty)$ such that

$$\min_{\underline{X} \in \mathcal{X}^{\text{BTD}}} \|\underline{X} - \underline{X}^\natural\|_\infty \leq \nu^{\text{BTD}}, \quad (27)$$

where \underline{X}^\natural represents the ground-truth radio map.

The existence of $\nu^{\text{BTD}} > 0$ makes sense, which could be a result of sensing noise or just due to modeling errors.

Theorem 1 Suppose that Assumption 1 holds. Assume that Ω is uniformly sampled from $[I] \times [J]$ with replacement and that $N = |\Omega|$. Define $\underline{X}^* = \sum_{r=1}^R (\underline{A}_r^* \underline{B}_r^{*\top}) \circ \underline{c}_r^*$, where

$$(\underline{A}^*, \underline{B}^*, \underline{C}^*) = \arg \min_{(\underline{A}, \underline{B}, \underline{C}) \in \mathcal{S}^{\text{BTD}}} F_{\Omega, \underline{Y}}^{\text{BTD}}(\underline{A}, \underline{B}, \underline{C}), \quad (28)$$

is an optimal solution of the MLE via BTD criterion in (9). Then, with probability at least $1 - 2\delta$, we have

$$\begin{aligned} \frac{\|\underline{X}^* - \underline{X}^\natural\|_{\text{F}}^2}{IJK} &\leq \frac{8C_1C_2(1+\tau)}{K} \sqrt{\frac{R}{N}} + U_\alpha C_1 \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} \\ &\quad + U_\alpha C_1 \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + C_1 C_2 \nu^{\text{BTD}}, \end{aligned} \quad (29)$$

where, $\tau = 3 \left(\sqrt{((I+J)L+K) \log(3\sqrt{R}(\beta+\kappa))} + \sqrt{K \log(\kappa/2) + (I+J)L \log(\beta)} \right)$, and β, κ are defined in

(10). The constants C_1 and C_2 are given by $C_1 = 4(\alpha+a)^2/\mathbb{F}_\alpha$, and $C_2 = \mathbb{L}_\alpha/a$.

The proof is relegated to Appendix D-B. Theorem 1 shows that the *mean squared error* (MSE) of the *optimal* solution given by (9) is able to approach the ground-truth $\underline{\mathbf{X}}^\ddagger$, if $N = |\Omega|$ is sufficiently large and the tensor-based representation error (i.e., ν^{BTD}) is reasonably small. Our result also shows that $N = \Omega((I+J)L+K)R\log(3\sqrt{R}(\beta+\kappa))/K^2)$ is the sample complexity that suffices to ensure provable recovery, which is at the same order of the number of unknown parameters. This is consistent with the previous results from classic matrix/tensor completion literature under other low-rank models [30]–[34].

Remark 2 The recoverability of low-rank tensors under 1-bit quantization was studied in [33], but was under the canonical polyadic decomposition (CPD) model. Multi-level quantization based tensor completion was considered in [33], [34], [36], but again under tensor models such as the CPD and Tucker models. In this work, we formulate our problem under the BTD model, which is more suitable for the SC problem. Recoverability of radio map tensors under the BTD model was considered in [6] under a different setting. There, the sampling pattern was deterministic and no quantization was involved (a similar setting was used in the context of hyperspectral super-resolution [47]). Hence, their proof took a very different route via using the essential uniqueness of BTD, which is not needed in our case.

The random sampling and quantization settings in [33], [34], [36] are closer to ours, but the data acquisition in these works was done through independently sampling the entries of the tensor. However, in this work, we focus on the scenario where the sensors can sample the entire PSD $\underline{\mathbf{X}}(i, j, :)$, also known as the *fiber sampling* case [25]. Our proof extends those in [33], [34], [36] and carefully accommodates these differences.

C. Recoverability Under The DGM Model

The proof under the DGM of the radio map (13) follows a similar technical route, except that the DGM’s “complexity measure” needs to be carefully derived. As in the previous section, we make the following assumption:

Assumption 2 There exists a constant $\nu^{\text{DGM}} \in [0, +\infty)$ such that

$$\min_{\mathbf{X} \in \mathcal{X}^{\text{DGM}}} \|\mathbf{X} - \underline{\mathbf{X}}^\ddagger\|_\infty \leq \nu^{\text{DGM}}. \quad (30)$$

Ideally, if $\mathbf{g}_\theta(\cdot)$ is a universal function approximator, ν^{DGM} would have been zero (provided that other types of noise do not exist). However, in practice, $\mathbf{g}_\theta(\cdot)$ is represented by a neural network with finite depth/width, which means that it is never universal and thus a nonzero ν^{DGM} exists. Note that ν^{DGM} can be decreased if the neural network $\mathbf{g}_\theta(\cdot)$ is more expressive. A more expressive neural network often admits a deeper and wider neural architecture, which has a increased number of parameters.

We also make the following structural assumption:

Assumption 3 The generator network is an L -layer neural network

$$\mathbf{g}_\theta(\mathbf{z}) = \text{mat}(\zeta_L(\mathbf{A}_L(\dots\zeta_1(\mathbf{A}_1\mathbf{z})))), \quad (31)$$

where $\text{mat}(\cdot) : \mathbb{R}^{IJ} \rightarrow \mathbb{R}^{I \times J}$ reshapes a vector to a matrix (i.e., the inverse operation of vectorization), $\mathbf{A}_\ell \in \mathbb{R}^{D_\ell \times D_{\ell-1}}$ is the network weight in the ℓ th layer, in which $D_0 = D$ and $D_L = IJ$, $\zeta_\ell(\cdot) = [\zeta(\cdot), \dots, \zeta(\cdot)]^T : \mathbb{R}^{D_\ell} \rightarrow \mathbb{R}^{D_\ell}$ is a ϕ_ℓ -Lipschitz activation function with $\zeta(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$, and $P = \prod_{\ell=1}^L \phi_\ell \|\mathbf{A}_\ell\|_2 < \infty$. In addition, the latent embedding \mathbf{z} is from a bounded set, i.e., $\mathbf{z} \in \mathcal{Z} = \{\mathbf{z} \in \mathbb{R}^D : \|\mathbf{z}\|_2 \leq q\}$.

Assumption 3 follows the same setting in [7]. This assumption characterizes the structure of generative neural network. Note that many commonly used activation functions are Lipschitz continuous, e.g., ReLU, sigmoid, and tanh. In addition, (31) subsumes many deep network architectures as its special cases, e.g., the multi-layer perceptron (MLP) and the convolutional neural network (CNN).

With the assumptions, we show the following recoverability theorem:

Theorem 2 Suppose that Assumption 2 and 3 hold. Assume that Ω is uniformly sampled from $[I] \times [J]$ and that $|\Omega| = N$. Define $\underline{\mathbf{X}}^* = \sum_{r=1}^R \mathbf{g}_\theta(\mathbf{z}_r^*) \circ \mathbf{c}_r^*$, where

$$(\mathbf{C}^*, \mathbf{Z}^*) = \arg \min_{(\mathbf{C}, \mathbf{Z}) \in \mathcal{S}^{\text{DGM}}} F_{\Omega, \underline{\mathbf{Y}}}^{\text{DGM}}(\mathbf{C}, \mathbf{Z}), \quad (32)$$

is an optimal solution of the MLE via DGM criterion in (13). Then with probability at least $1 - 2\delta$, we have

$$\frac{\|\underline{\mathbf{X}}^* - \underline{\mathbf{X}}^\ddagger\|_\text{F}^2}{IJK} \leq \frac{8C_1C_2(1+\tau)}{K} \sqrt{\frac{R}{N}} + \mathbb{U}_\alpha C_1 \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} + \mathbb{U}_\alpha C_1 \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + C_1C_2\nu^{\text{DGM}}, \quad (33)$$

where

$$\tau = 3\sqrt{K\log(\frac{3}{2}\sqrt{R}\kappa(\beta+\kappa))} + D\log(3\sqrt{R}Pq(\beta+\kappa)),$$

and β, κ are defined in (15). The constants C_1 and C_2 are given by $C_1 = 4(\alpha+a)^2/\mathbb{F}_\alpha$, and $C_2 = \mathbb{L}_\alpha/a$.

The proof of Theorem 2 is presented in Appendix D-B. Theorem 2 shows that the optimization criterion in (13) ensures the recoverability of the ground-truth radio map $\underline{\mathbf{X}}^\ddagger$ —even when the radio map is parameterized by a learned DGM other than a classic low-rank model. There are two primary sources of errors in our recovery process. The first error occurs due to the limited fiber samples of radio maps measurement, also referred to as generalization (or statistical) error, which is given by the first term on the R.H.S. of (33). The second type of error is induced by the representation error of the DGM, i.e., ν^{DGM} , which is the last term in (33). One can see that the recovered radio map becomes closer to the ground truth with an increased sample size at the rate of $\mathcal{O}(1/\sqrt{N})$. When the DGM becomes more complex (e.g., when the depth/width of network is increased), the approximation error ν^{DGM} decreases. However, the increased network complexity

makes τ larger. This presents a reasonable trade-off between model and sample complexities. In short, one hopes to use a reasonably complex DGM so that ν^{DGM} is small, but not an overly expressive DGM that might lead to over-fitting effects.

VI. EXPERIMENTS

A. Synthetic Data Experiments

1) *Data Generation*: For data generation, we first consider the 2D geographical region of $50 \times 50 \text{ m}^2$ that has 51×51 grids. The region is discretized using indices $\{0, 1, \dots, 50\} \times \{0, 1, \dots, 50\}$ such that $I = J = 51$. The spectral domain has $K = 64$ subbands. We generate the radio maps by first generating the SLF of each emitters using the joint path loss model and spatial correlated log-normal shadowing model [48]. The two key parameters for this model are the decorrelation distance X_c and the shadowing variance η . Larger X_c and smaller η correspond to milder shadowing environment; see [7] for more details. The PSDs are also generated following the same method as in [6], [7].

2) *DGM Learning*: We learn the DGM using GAN. We set $D = 256$. The latent embedding z_r is sampled from the standard Gaussian distribution. The detailed architecture of the GAN is in Appendix A. To generate training samples for learning the DGM, we create 5,000 random SLFs using random emitter locations, a range of path loss coefficients, X_c and η in $[2, 2.5]$, $[30, 100]$ and $[3, 8]$, respectively. We then train our GAN using Adam optimizer with a batch size of 128 with 250 epochs. The initial learning rates for the discriminator and generator are set to be 0.0004 and 0.0001 respectively.

3) *Algorithm Settings*: We initialize \mathbf{A} and \mathbf{B} as zero matrices in the BTD model. For the DGM model, we initialize z_r using the standard normal distribution. For both models, we initialize the elements of \mathbf{C} using the uniform distribution between 0 and 1. Under both the DGM and BTD models, we stop the algorithms when the relative change of the cost function is smaller than 10^{-3} or when the BCD algorithms reach 300 iterations. For both of our reformulations in (21) and (17), the regularization terms are all set to be 10^{-3} . We set $L = 10$ for the BTD model following [6]. For optimization, we use the Adam optimizer for both DGM and BTD block updates. For both models, we set the initial learning rates for \mathbf{C} to be 0.003 and for the other variables to be 0.006.

4) *Quantization*: We set $a = 10^{-6}$ in $h(x) = \log(x + a)$. For quantization, we set the noise variance using a validation set consisting of 20 simulated radio maps. The validation set is generated with shadowing parameters $X_c \in [30, 100]$ and $\eta \in [3.0, 8.0]$. The recovering algorithms are applied to the validation set with various $\sigma^2 \in [0.5, 4.0]$ for different bits cases and the average best-performing $\sigma^2 = 1.7$ is selected. The quantization bins are constructed using 1,000 simulated radio maps data generated following the description in Sec. VI-A2. The b_q 's are constructed using the method in Sec. IV-C for $B \in [1, 8]$, where $B = \log_2(Q - 1)$ bits.

5) *Performance Metrics*: We use *log-domain normalized reconstruction error* (LNRE) as our metric:

$$\text{LNRE} = \frac{\|\widehat{\mathbf{M}} - \mathbf{M}^\dagger\|_F^2}{\|\mathbf{M}^\dagger\|_F^2}, \quad (34)$$

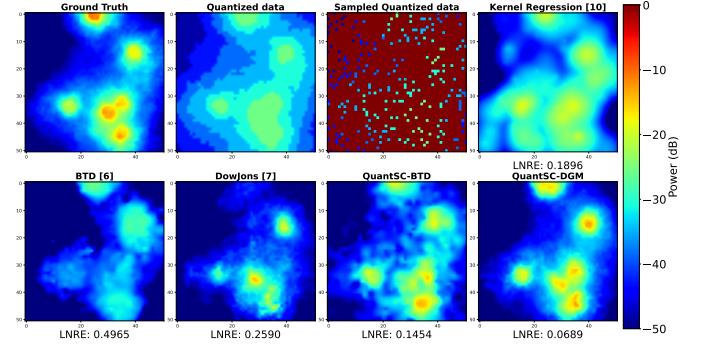


Fig. 6: Ground-truth and reconstructed radio maps by various methods at the 30th frequency bin; $\rho = 10\%$, $R = 6$, $X_c = 50$, $\eta = 6$ and $\sigma^2 = 1.7$, $B = 3$ bits.

where $\widehat{\mathbf{M}} = h(\widehat{\mathbf{X}})$ is the estimated log-transformed the radio map. The LNRE is an appropriate metric for skewed data as it prevents a few larger values in the \mathbf{X} from dominating the performance evaluation. All the LNREs are averaged over 10 random trials.

6) *Baselines*: We use the DGM-based SC method in [7], namely, “DowJons”, as it demonstrated state-of-the-art SC performance given that real-valued feedback can be used. We also use the BTD-based method for unquantized data in [6]. Additionally, we use the method in [10] based on kernel regression for quantized SC (referred to as “KR”). Note that [10] assumes that the PSDs of the emitters are known to the fusion center, but our method does not need this information.

7) *Results*: Fig. 6 shows an illustrative example where we have $R = 6$ emitters. For visualization, we show the recovery result at a single frequency bin. In this experiment, $\rho = 10\%$ of the 51×51 grids are sampled, where $\rho = |\Omega|/IJ \times 100\%$. We use $B = 3$ bits for quantizing every real-valued measurement. The shadowing parameters are set to be $X_c = 50$ and $\eta = 7$. The LNRE of the reconstructed radio maps over all frequencies are shown along with each images. One can see that the QuantSC-DGM offers the best recovery performance, both visually and in terms of LNRE. The QuantSC-BTD works reasonably, as it still recovers the positions of emitters well. The DowJons method clearly missed some emitters, which is not surprising—as it is not designed for quantized SC. The BTD-based method from [6] also performs poorly as it also does not consider quantized measurements. The KR method did not perform well and over-smoothing effects are observed.

Fig. 7 shows the result with heavier shadowing compared to the scenario in Fig. 6. Here, we use $R = 2$ and set X_c and η to be 40 and 8, respectively. All the other settings follow those in previous paragraph. Clearly, this is a much more challenging scenario as most of the methods fail. QuantSC-BTD does not work well in this case, as expected—since the low-rank assumption of the SLFs is grossly violated. The proposed QuantSC-DGM is the only one that produces a radio map close to the ground truth. This shows the power of DGM in terms of modeling complex environments.

Tables I and II show the performance of the methods under various η and X_c , respectively. Recall that the two parameters reflect the level of shadowing. In all cases, QuantSC-DGM

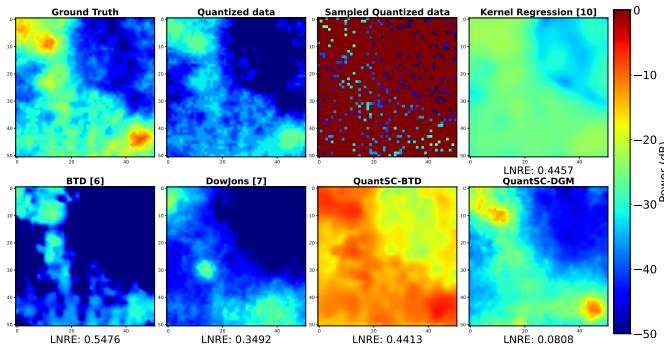


Fig. 7: Ground-truth and reconstructed radio maps under heavy shadowing by various methods at the 20th frequency bin; $\rho = 10\%$, $R = 2$, $X_c = 40$, $\eta = 8$ and $\sigma^2 = 1.7$, $B = 3$ bits.

TABLE I: LNRE under various η 's: $R = 6$, $X_c = 50$, $\rho = 10\%$, and $B = 3$ bits.

η	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
4.0	0.1351 ± 0.0092	0.2078 ± 0.0149	0.1352 ± 0.0082	0.0552 ± 0.0083
5.0	0.1284 ± 0.0153	0.2163 ± 0.0058	0.1384 ± 0.0142	0.0650 ± 0.0041
6.0	0.1949 ± 0.0102	0.2435 ± 0.0163	0.1327 ± 0.0092	0.0665 ± 0.0103
7.0	0.2310 ± 0.0210	0.2341 ± 0.0132	0.1461 ± 0.0097	0.0701 ± 0.0084
8.0	0.2705 ± 0.0076	0.2578 ± 0.0097	0.1593 ± 0.0085	0.0787 ± 0.0069

TABLE II: LNRE under various X_c 's: $R = 6$, $\eta = 6$, $\rho = 10\%$, and $B = 3$ bits.

X_c	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
30.0	0.2366 ± 0.0303	0.2116 ± 0.0142	0.1513 ± 0.0093	0.0699 ± 0.0068
50.0	0.1852 ± 0.0095	0.2192 ± 0.0102	0.1316 ± 0.0129	0.0625 ± 0.0041
70.0	0.1567 ± 0.0162	0.2218 ± 0.0092	0.1318 ± 0.0087	0.0620 ± 0.0080
90.0	0.1401 ± 0.0095	0.1783 ± 0.0105	0.1185 ± 0.0123	0.0579 ± 0.0091

TABLE III: LNRE under various R 's. $\eta = 6$, $X_c = 50$, $\rho = 10\%$, and $B = 3$ bits.

R	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
2	0.0818 ± 0.0193	0.1555 ± 0.0144	0.1125 ± 0.0241	0.0499 ± 0.0036
4	0.1300 ± 0.0324	0.2515 ± 0.0404	0.1301 ± 0.0101	0.0612 ± 0.0130
6	0.1730 ± 0.0078	0.2412 ± 0.0072	0.1442 ± 0.0129	0.0640 ± 0.0087
8	0.1902 ± 0.0309	0.2746 ± 0.0315	0.1621 ± 0.0137	0.0595 ± 0.0054
10	0.2295 ± 0.0639	0.3227 ± 0.0206	0.1860 ± 0.0160	0.0659 ± 0.0071
12	0.2197 ± 0.0127	0.3522 ± 0.0433	0.1872 ± 0.0111	0.0674 ± 0.0069
14	0.2181 ± 0.0099	0.3280 ± 0.0052	0.1985 ± 0.0181	0.0746 ± 0.0061
16	0.2149 ± 0.0083	0.3321 ± 0.0056	0.1951 ± 0.0031	0.0782 ± 0.0062
18	0.2296 ± 0.0167	0.3632 ± 0.0070	0.2148 ± 0.0076	0.0782 ± 0.0041
20	0.2356 ± 0.0174	0.3630 ± 0.0388	0.2117 ± 0.0075	0.0834 ± 0.0064

exhibits the lowest LNRE. The performance only deteriorates gracefully when X_c decreases and η increases, showing the method's robustness to heavy shadowing. The QuantSC-BTD approach offers the second-best LNRE performance in all cases. As expected, it works the best when $X_c = 90$ and $\eta = 6$, but suffers from a 40% LNRE increase when X_c reduces to 30. This suggests that the shadowing heavily affects the performance of the BTD-based approach, as mentioned in [7] as well.

Table III shows the LNREs of the reconstructed radio maps under various R 's, i.e., the number of emitters. All the methods see performance degradation when R increases, as more unknown parameters need to be estimated. However, the QuantSC-DGM and QuantSC-BTD still offer the best and

TABLE IV: LNRE under various \hat{R} 's. Ground-truth $R = 6$, $X_c = 40$, $\eta = 7.0$, $\rho = 10\%$, and $B = 3$ bits.

\hat{R}	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
2	-	0.3058 ± 0.0029	0.1790 ± 0.0098	0.0963 ± 0.0077
3	-	0.2931 ± 0.0047	0.1731 ± 0.0085	0.0852 ± 0.0068
4	-	0.2837 ± 0.0107	0.1706 ± 0.0097	0.0811 ± 0.0076
5	-	0.2447 ± 0.0617	0.1639 ± 0.0101	0.0695 ± 0.0049
6	0.1917 ± 0.0068	0.2398 ± 0.0090	0.1602 ± 0.0091	0.0682 ± 0.0083
7	-	0.2415 ± 0.0115	0.1612 ± 0.0105	0.0661 ± 0.0056
8	-	0.2439 ± 0.0102	0.1610 ± 0.0102	0.0673 ± 0.0083
9	-	0.2956 ± 0.0114	0.1736 ± 0.0098	0.0696 ± 0.0137
10	-	0.4474 ± 0.0220	0.1958 ± 0.0105	0.0716 ± 0.0093
11	-	0.4411 ± 0.0302	0.2017 ± 0.0107	0.0727 ± 0.0096

TABLE V: LNRE under various ρ 's. $R = 6$, $X_c = 50$, $\eta = 6$, and $B = 3$ bits.

ρ	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
3%	0.2104 ± 0.0452	0.2625 ± 0.0163	0.1980 ± 0.0074	0.0900 ± 0.0070
5%	0.1975 ± 0.0201	0.2506 ± 0.0141	0.1756 ± 0.0106	0.0761 ± 0.0091
10%	0.1722 ± 0.0124	0.2388 ± 0.0080	0.1429 ± 0.0116	0.0607 ± 0.0065
15%	0.1663 ± 0.0149	0.2220 ± 0.0079	0.1408 ± 0.0090	0.0557 ± 0.0083
20%	0.1416 ± 0.0075	0.1906 ± 0.0118	0.1315 ± 0.0071	0.0522 ± 0.0075

second-best LNRE performance.

The simulation in Table. IV tests the robustness to the wrongly estimated R 's. In the previous simulations, we assumed that the number of emitters R is accurately estimated by the algorithms—which may not always hold in practice. Table. IV shows the performance of all the methods when R is underestimated or overestimated—i.e., the performance under wrong R 's. For clarity, we use \hat{R} to denote the “estimated number of emitters” that is used by the algorithms—and our purpose is to observe how the algorithms behave when $\hat{R} \neq R$. The true R for this experiment setting is 6. The KR method needs to know the PSD of the emitters, and thus we use the $\hat{R} = R$ as its input. One can see that when $\hat{R} < R$, the algorithms have higher LNREs. The LNRE decreases when \hat{R} grows from 2 to 6, but increases again when $\hat{R} \geq 9$. This makes sense as an underestimated \hat{R} could not capture all the “principal components” of the radio map tensor, but an overestimated \hat{R} makes the complexity of the model (represented by τ in Theorems 1 and 2) higher. Hence, both seriously underestimated and largely overestimated \hat{R} could hurt the recovery accuracy. Similar phenomenon is also observed by changing the algorithm-used \hat{L} in the BTD (which may be different from the ground-truth L). The experiment using varying \hat{L} is not included due to page limitations, but the insights behind are the same.

Table V shows the performance of the algorithms under various ρ 's. Notably, the LNRE output by QuantSC-DGM using $\rho = 3\%$ is lower than the LNREs output by the other methods using $\rho = 20\%$ —again showing the expressive power of the DGM. However, the same DGM-empowered method DowJons performs worse compared to QuantSC-BTD under all ρ 's. This shows the effectiveness of our proposed quantized SC framework based on Gaussian quantization and MLE.

Table VI shows the performance under various numbers of quantization bits, i.e., B , used for each measurement. Compared to the baselines, the proposed methods QuantSC-DGM and QuantSC-BTD admit tangible margins when $B =$

TABLE VI: LNRE under various B 's. $R = 6$, $X_c = 50$, $\eta = 6$, and $\rho = 10\%$.

B	KR [10]	DowJons [7]	QuantSC-BTD	QuantSC-DGM
1	0.4670 ± 0.0059	0.7424 ± 0.0640	0.1815 ± 0.0112	0.0720 ± 0.0095
2	0.2885 ± 0.0260	0.5625 ± 0.0292	0.1548 ± 0.0159	0.0675 ± 0.0074
3	0.1781 ± 0.0082	0.2461 ± 0.0292	0.1434 ± 0.0140	0.0645 ± 0.0088
4	0.1547 ± 0.0114	0.1533 ± 0.0127	0.1385 ± 0.0132	0.0607 ± 0.0124
5	0.1384 ± 0.0057	0.0712 ± 0.0142	0.1339 ± 0.0168	0.0572 ± 0.0136
6	0.1128 ± 0.0100	0.0556 ± 0.0147	0.1242 ± 0.0078	0.0545 ± 0.0096
7	0.0970 ± 0.0079	0.0543 ± 0.0155	0.1028 ± 0.0119	0.0504 ± 0.0115
8	0.0743 ± 0.0090	0.0523 ± 0.0146	0.0972 ± 0.0109	0.0445 ± 0.0118

1, 2, 3, 4. The performance of the proposed methods are more than satisfactory when even only $B = 1$ bit is used. When $B \geq 5$, which means the quantization level reaches $Q = 31$, the performance of KR and DowJons catch up—as the quantization error gradually becomes negligible in these cases.

B. Real-Data Experiments

1) *Data Description*: The data was obtained within a $14 \times 34\text{m}^2$ indoor space on an office floor at the Mannheim University. The data was acquired across 9 distinct frequency bands centered at 2.412 GHz, 2.422 GHz, 2.427 GHz, 2.432 GHz, 2.437 GHz, 2.442 GHz, 2.447 GHz, 2.457 GHz, and 2.462 GHz, respectively [49]. The region is divided into $1 \times 1\text{m}^2$ grids, and 166 of these grids were installed with sensors at their centers which are placed throughout the hallway. The leftmost column in Fig. 8 displays the ground-truth radio map over the nine frequencies. The white blocks in the figures are rooms where measurements could not be taken. The measurements in our experiments are uniformly sampled from the hallway area. This type of mixed deterministic (rooms) and random (hallway) missing pattern is not exactly covered by the conditions made in our recoverability theorems, but the dataset can still be employed to test the algorithms. The results can indicate their usefulness in real-world settings and robustness to violation of the conditions in the theorems.

2) *Hyperparameter Settings*: The data is heavily skewed for the real data and the values are extremely small. The largest power measured is -40 dB. Hence, we take $a = 10^{-15}$ for our transform to avoid a dominating the output of $h(x + a)$. For quantization, we follow the same validation-based strategy as before: We set $\sigma^2 = 4.0$, which is tuned using a validation set consisting of 20 simulated radio maps. The validation set is generated with high shadowing parameter considering the indoor environment $X_c \in [20, 70]$ and $\eta \in [6.0, 9.0]$. The recovering algorithms are applied to the validation set with various $\sigma^2 \in [1.0, 5.0]$ and the best-performing σ^2 is selected. All the regularization parameter for BTD and DGM are set to be 10^{-4} . We set $R = 7$ and $L = 4$ following [7] and [6], respectively. We do not include KR as it needs the real PSDs of the emitters, which are unavailable.

3) *DGM Training*: To train GAN for real data, we first generate samples that have large η and small X_c to simulate heavy shadowing conditions—as we know that the real data was collected from an indoor environment. We select η ranging from 8 to 12 and X_c ranging from 5 to 50. We simulate the SLF within a $14 \times 34\text{m}^2$ region. We then generate 10,000

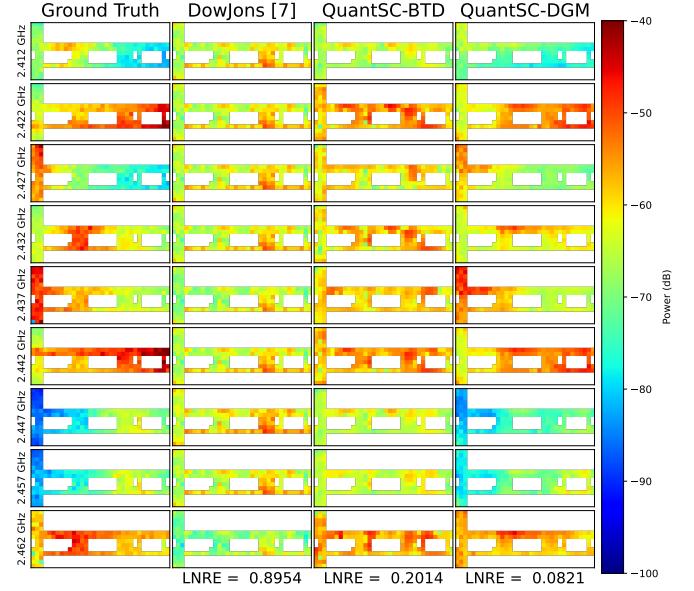


Fig. 8: Ground-truth and reconstructed radio maps on all the frequency bins; $\rho = 10\%$, $B = 3$.

TABLE VII: LNRE under various ρ 's for real data for $B = 3$.

ρ	DowJons [7]	QuantSC-BTD	QuantSC-DGM
5%	0.9474 ± 0.1017	0.2527 ± 0.0110	0.1144 ± 0.0084
10%	0.8944 ± 0.0790	0.2197 ± 0.0162	0.0866 ± 0.0079
15%	0.8761 ± 0.0524	0.2061 ± 0.0126	0.0809 ± 0.0068
20%	0.8452 ± 0.0303	0.1993 ± 0.0098	0.0782 ± 0.0102

samples of such SLFs and train the GAN to learn the DGM. The GAN is trained using the Adam algorithm and a batch size of 256, for a maximum of 500 epochs. The initial learning rates of the discriminator and the generator are 2×10^{-5} and 10^{-5} , respectively. The GAN architecture is in Appendix A.

4) *Result*: Fig. 8 shows the reconstructed radio maps over all 9 frequencies by the methods under 3-bit quantization. One can see that the DowJons method is incapable of accurately estimating the true radio map in such a challenging scenario. However, the proposed methods both output visually reasonable estimates, showing the effectiveness of our framework. As before, the estimated radio map by QuantSC-DGM is visually closer to the ground truth, but QuantSC-BTD offers satisfactory results without training a neural model off-line.

Tables VII-VIII show the quantitative results under various ρ 's and B 's. The performance of the algorithms are consistent with what we saw in the simulations. Remarkably, we have no access to the environment parameters of the real data. Hence, it is likely that our generated simulated SLFs for training the DGM have nontrivial model mismatches. However, both the qualitative and quantitative evaluations suggest that the proposed method work reasonably well, which shows the robustness of the proposed approach.

VII. CONCLUSION

A novel framework has been introduced for SC that operates effectively with heavily quantized sensor measurements. Unlike previous approaches to provable SC, which assumed

TABLE VIII: LNRE under various bits for real data, $\rho = 10\%$

Bits	DowJons [7]	QuantSC-BTD	QuantSC-DGM
1	0.9534 \pm 0.0624	0.2842 \pm 0.0162	0.0998 \pm 0.0111
2	0.9264 \pm 0.0233	0.2431 \pm 0.0100	0.0897 \pm 0.0076
3	0.8032 \pm 0.0581	0.2045 \pm 0.0095	0.0879 \pm 0.0082
4	0.7177 \pm 0.0310	0.1934 \pm 0.0128	0.0862 \pm 0.0073
5	0.5350 \pm 0.0186	0.1887 \pm 0.0101	0.0812 \pm 0.0092
6	0.2350 \pm 0.0270	0.1810 \pm 0.0137	0.0797 \pm 0.0057
7	0.1809 \pm 0.0148	0.1811 \pm 0.0083	0.0791 \pm 0.0087
8	0.1414 \pm 0.0159	0.1803 \pm 0.0070	0.0786 \pm 0.0072

full-resolution real-valued measurements, the new maximum likelihood estimation (MLE)-based framework uses measurements discretized by a Gaussian quantizer. It is compatible with both BTD and DGM-based radio map representations, with the former being training-free, and the latter showing higher resilience to heavy shadowing. The recoverability of the framework has been characterized under realistic conditions, such as imperfect radio map modeling. Simulations and real-data experiments have demonstrated the effectiveness of the proposed approach. The proposed quantized SC framework offers a practical and realistic solution to radio map estimation, with provable guarantees.

REFERENCES

- [1] G. Mateos, J.-A. Bazerque, and G. B. Giannakis, "Spline-based spectrum cartography for cognitive radios," in *Proc. IEEE Conf. Rec. Asilomar Conf. Signals Syst.*, 2009, pp. 1025–1029.
- [2] J. A. Bazerque, G. Mateos, and G. B. Giannakis, "Group-lasso on splines for spectrum cartography," *IEEE Trans. Signal Process.*, vol. 59, no. 10, pp. 4648–4663, 2011.
- [3] G. Boccolini, G. Hernandez-Penalosa, and B. Beferull-Lozano, "Wireless sensor network for spectrum cartography based on kriging interpolation," in *Proc. IEEE PIMRC*, 2012, pp. 1565–1570.
- [4] S.-J. Kim and G. B. Giannakis, "Cognitive radio spectrum prediction using dictionary learning," in *Proc. IEEE GLOBECOM*, 2013, pp. 3206–3211.
- [5] B. A. Jaywickrama, E. Dutkiewicz, I. Oppermann, G. Fang, and J. Ding, "Improved performance of spectrum cartography based on compressive sensing in cognitive radio networks," in *Proc. IEEE ICC*, 2013, pp. 5657–5661.
- [6] G. Zhang, X. Fu, J. Wang, X.-L. Zhao, and M. Hong, "Spectrum cartography via coupled block-term tensor decomposition," *IEEE Trans. Signal Process.*, vol. 68, pp. 3660–3675, 2020.
- [7] S. Shrestha, X. Fu, and M. Hong, "Deep spectrum cartography: Completing radio map tensors using learned neural models," *IEEE Trans. Signal Process.*, vol. 70, pp. 1170–1184, 2022.
- [8] Y. Teganya and D. Romero, "Data-driven spectrum cartography via deep completion autoencoders," in *Proc. IEEE ICC*, 2020, pp. 1–7.
- [9] S. Bi, J. Lyu, Z. Ding, and R. Zhang, "Engineering radio maps for wireless resource management," *IEEE Wirel. Commun.*, vol. 26, no. 2, pp. 133–141, 2019.
- [10] D. Romero, S.-J. Kim, G. B. Giannakis, and R. López-Valcarce, "Learning power spectrum maps from quantized power measurements," *IEEE Trans. Signal Process.*, vol. 65, no. 10, pp. 2547–2560, 2017.
- [11] M. Hamid and B. Beferull-Lozano, "Non-parametric spectrum cartography using adaptive radial basis functions," in *Proc. IEEE ICASSP*, 2017, pp. 3599–3603.
- [12] J. A. Bazerque and G. B. Giannakis, "Distributed spectrum sensing for cognitive radio networks by exploiting sparsity," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1847–1862, 2009.
- [13] B. Khalfi, B. Hamdaoui, and M. Guizani, "Airmap: Scalable spectrum occupancy recovery using local low-rank matrixapproximation," in *Proc. IEEE GLOBECOM*, 2018, pp. 206–212.
- [14] D. Schäufele, R. L. Cavalcante, and S. Stanczak, "Tensor completion for radio map reconstruction using low rank and smoothness," in *Proc. IEEE SPAWC*, 2019, pp. 1–5.
- [15] H. Sun and J. Chen, "Propagation map reconstruction via interpolation assisted matrix completion," *IEEE Trans. Signal Process.*, vol. 70, pp. 6154–6169, 2022.
- [16] Q. Niu, Y. Nie, S. He, N. Liu, and X. Luo, "Recnet: A convolutional network for efficient radiomap reconstruction," in *Proc. IEEE ICC*, 2018, pp. 1–7.
- [17] V. V. Ratnam, H. Chen, S. Pawar, B. Zhang, C. J. Zhang, Y.-J. Kim, S. Lee, M. Cho, and S.-R. Yoon, "Fadenet: Deep learning-based mm-wave large-scale channel fading prediction and its applications," *IEEE Access*, vol. 9, pp. 3278–3290, 2020.
- [18] X. Han, L. Xue, F. Shao, and Y. Xu, "A power spectrum maps estimation algorithm based on generative adversarial networks for underlay cognitive radio networks," *Sensors*, vol. 20, no. 1, p. 311, 2020.
- [19] E. Kriještorac, S. Hanna, and D. Cabric, "Spatial signal strength prediction using 3d maps and deep learning," in *Proc. IEEE ICC*, 2021, pp. 1–6.
- [20] O. Mehanna and N. D. Sidiropoulos, "Frugal sensing: Wideband power spectrum sensing from few bits," *IEEE Trans. Signal Process.*, vol. 61, no. 10, pp. 2693–2703, 2013.
- [21] A. Konar, N. D. Sidiropoulos, and O. Mehanna, "Parametric frugal sensing of power spectra for moving average models," *IEEE Trans. Signal Process.*, vol. 63, no. 5, pp. 1073–1085, 2014.
- [22] X. Fu, N. D. Sidiropoulos, J. H. Tranter, and W.-K. Ma, "A factor analysis framework for power spectra separation and multiple emitter localization," *IEEE Trans. Signal Process.*, vol. 63, no. 24, pp. 6581–6594, 2015.
- [23] X. Fu, N. D. Sidiropoulos, and W.-K. Ma, "Power spectra separation via structured matrix factorization," *IEEE Trans. Signal Process.*, vol. 64, no. 17, pp. 4592–4605, 2016.
- [24] S. Timilsina, S. Shrestha, and X. Fu, "Deep spectrum cartography using quantized measurements," in *Proc. IEEE ICASSP*, 2023.
- [25] X. Fu, S. Ibrahim, H.-T. Wai, C. Gao, and K. Huang, "Block-randomized stochastic proximal gradient for low-rank tensor factorization," *IEEE Trans. Signal Process.*, vol. 68, pp. 2170–2185, 2020.
- [26] Y. L. Polo, Y. Wang, A. Pandharipande, and G. Leus, "Compressive wide-band spectrum sensing," in *Proc. IEEE ICASSP*, 2009, pp. 2337–2340.
- [27] L. De Lathauwer, "Decompositions of a higher-order tensor in block terms—part ii: Definitions and uniqueness," *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 3, pp. 1033–1066, 2008.
- [28] W. Han, J. Li, Z. Tian, and Y. Zhang, "Efficient cooperative spectrum sensing with minimum overhead in cognitive radio," *IEEE Trans. Wireless Commun.*, vol. 9, no. 10, pp. 3006–3011, 2010.
- [29] D. Lee, T. Sasaki, T. Yamada, K. Akabane, Y. Yamaguchi, and K. Uehara, "Spectrum sensing for networked system using 1-bit compressed sensing with partial random circulant measurement matrices," in *Proc. IEEE VTC*, 2012, pp. 1–5.
- [30] M. A. Davenport, Y. Plan, E. Van Den Berg, and M. Wootters, "1-bit matrix completion," *Information and Inference: A Journal of the IMA*, vol. 3, no. 3, pp. 189–223, 2014.
- [31] S. A. Bhaskar, "Probabilistic low-rank matrix completion from quantized measurements," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2131–2164, 2016.
- [32] Y. Cao and Y. Xie, "Categorical matrix completion," in *Proc. IEEE CAMSAP*, 2015, pp. 369–372.
- [33] N. Ghadermarzy, Y. Plan, and O. Yilmaz, "Learning tensors from partial binary measurements," *IEEE Trans. Signal Process.*, vol. 67, no. 1, pp. 29–40, 2018.
- [34] B. Li, X. Zhang, X. Li, and H. Lu, "Tensor completion from one-bit observations," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 170–180, 2018.
- [35] P. McCullagh, "Regression models for ordinal data," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 42, no. 2, pp. 109–127, 1980.
- [36] C. Lee and M. Wang, "Tensor denoising and completion based on ordinal observations," in *Proc. ICML*. PMLR, 2020, pp. 5778–5788.
- [37] L. Schuchman, "Dither signals and their effect on quantization noise," *IEEE Trans. Commun. Technol.*, vol. 12, no. 4, pp. 162–165, 1964.
- [38] S. P. Lipshitz, R. A. Wannamaker, and J. Vanderkooy, "Quantization and dither: A theoretical survey," *Journal of the audio engineering society*, vol. 40, no. 5, pp. 355–375, 1992.
- [39] Pytorch, "torch.autograd," 2022. [Online]. Available: <https://pytorch.org/docs/stable/autograd.html>
- [40] J. Duchi, E. Hazan, and Y. Singer, "Adaptive subgradient methods for online learning and stochastic optimization," *Journal of Machine Learning Research*, vol. 12, no. 7, 2011.
- [41] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.

[42] M. Hong, M. Razaviyayn, Z.-Q. Luo, and J.-S. Pang, “A unified algorithmic framework for block-structured optimization involving big data: With applications in machine learning and signal processing,” *IEEE Signal Process. Mag.*, vol. 33, no. 1, pp. 57–77, 2015.

[43] E. E. Papalexakis, N. D. Sidiropoulos, and R. Bro, “From k-means to higher-way co-clustering: Multilinear decomposition with sparse latent factors,” *IEEE Trans. Signal Process.*, vol. 61, no. 2, pp. 493–506, 2012.

[44] A. S. Lan, C. Studer, and R. G. Baraniuk, “Matrix recovery from quantized and corrupted measurements,” in *Proc. IEEE ICASSP*, 2014, pp. 4973–4977.

[45] A. S. Lan, C. Studer, A. E. Waters, and R. G. Baraniuk, “Tag-aware ordinal sparse factor analysis for learning and content analytics,” in *Proc. International Conference on Educational Data Mining, EDM*, 2013.

[46] T. Cai and W.-X. Zhou, “A max-norm constrained minimization approach to 1-bit matrix completion,” *Journal of Machine Learning Research*, vol. 14, pp. 3619–3647, 2013.

[47] M. Ding, X. Fu, T.-Z. Huang, J. Wang, and X.-L. Zhao, “Hyperspectral super-resolution via interpretable block-term tensor modeling,” *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 3, pp. 641–656, 2020.

[48] A. Goldsmith, *Wireless communications*. Cambridge University Press, 2005.

[49] T. King, S. Kopf, T. Haenselmann, C. Lubberger, and W. Effelsberg. Crawdad dataset mannheim/compass (v. 2008-04-11). [Online]. Available: <https://crawdad.org/mannheim/compass/20080411>

[50] S. Shalev-Shwartz and S. Ben-David, *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.

[51] D. Pollard, “Empirical processes: Theory and applications,” in *NSF-CBMS Regional Conference Series in Probability and Statistics*. JSTOR, 1990, pp. i–86.

[52] M. Ding, X. Fu, and X.-L. Zhao, “Fast and structured block-term tensor decomposition for hyperspectral unmixing,” *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.*, 2023.

[53] W. Hoeffding, “Probability inequalities for sums of bounded random variables,” *Journal of the American statistical association*, vol. 58, no. 301, pp. 13–30, 1963.

[54] M. Mohri, A. Rostamizadeh, and A. Talwalkar, *Foundations of machine learning*. MIT press, 2018.

[55] P. L. Bartlett, D. J. Foster, and M. J. Telgarsky, “Spectrally-normalized margin bounds for neural networks,” in *Proc. NIPS*, vol. 30. Curran Associates, Inc., 2017.

APPENDIX A GAN ARCHITECTURE

Table IX [left] presents the detailed architecture of the GAN used in our synthetic data experiments. The “Deconv” term represents the block of operation stacked with transposed convolutional operations, followed by batch normalization and ReLU activations. Similarly, the “Conv” term represents the block stacked with convolutional operations, followed by batch normalization and leaky ReLU activations. The “Conv2d” refers to the convolutional operation followed by sigmoid activations. As mentioned, in order to ensure the nonnegativity of the generated SLFs, we use the sigmoid activation in the output layer of the generator. #C, #S, and #P denote the number of channels, stride size, and pad size, respectively. Table IX [right] presents the neural network structures used in the real data experiments.

APPENDIX B PROOF OF LEMMA 1

Let us assume that there are positive ℓ_{\min} and ℓ_{\max} such that the following holds:

$$\ell_{\min} \leq \frac{|h(x) - h(y)|}{|x - y|} \leq \ell_{\max} \quad (35)$$

$$\Rightarrow \ell_{\min}|x - y| \leq |h(x) - h(y)| \leq \ell_{\max}|x - y|. \quad (36)$$

TABLE IX: GAN architecture in simulations. [Left] For synthetic data experiments. [Right] For real data experiments

Generator					
Layer	Filter	#C	#S	#P	
Deconv	3×3	128	1	0	
Deconv	4×4	64	2	1	
Deconv	4×4	32	2	1	
Deconv	4×4	16	2	0	
Deconv	4×4	2	2	0	
Conv2d	4×4	1	1	0	
Discriminator					
Conv	4×4	16	2	1	
Conv	4×4	32	2	1	
Conv	4×4	64	2	1	
Conv	4×4	128	2	1	
Conv2d	3×3	1	1	0	

Generator					
Layer	Filter	#C	#S	#P	
Deconv	3×3	128	1	0	
Deconv	3×3	64	2	1	
Deconv	3×3	32	(1,2)	(0,1)	
Deconv	3×3	16	2	1	
Deconv	2×2	8	(1,2)	0	
Deconv	2×2	4	1	0	
Conv2d	2×2	1	1	0	
Discriminator					
Conv	2×2	16	2	1	
Conv	4×4	32	1	1	
Conv	4×4	64	1	(1,0)	
Conv	3×3	128	2	1	
Conv	3×3	256	2	(1,0)	
Conv2d	2×2	1	1	0	

Applying the above to the corresponding entries of the two tensors $\underline{\mathbf{X}}$ and $\underline{\mathbf{X}'}$, and then squaring and summing over all the entries, we have

$$\begin{aligned} \ell_{\min}^2 \|\underline{\mathbf{X}} - \underline{\mathbf{X}'}\|_F^2 &\leq \|h(\underline{\mathbf{X}}) - h(\underline{\mathbf{X}'})\|_F^2 \leq \ell_{\max}^2 \|\underline{\mathbf{X}} - \underline{\mathbf{X}'}\|_F^2 \\ \Rightarrow \ell_{\min} &\leq \frac{\|h(\underline{\mathbf{X}}) - h(\underline{\mathbf{X}'})\|_F}{\|\underline{\mathbf{X}} - \underline{\mathbf{X}'}\|_F} \leq \ell_{\max}. \end{aligned} \quad (37)$$

To find out ℓ_{\min} and ℓ_{\max} , consider the following:

$$\begin{aligned} &\min_{x,y \in [0,\alpha]} \frac{|\log(x+a) - \log(y+a)|}{|x-y|} \\ &\leq \frac{\|h(\underline{\mathbf{X}}) - h(\underline{\mathbf{X}'})\|_F}{\|\underline{\mathbf{X}} - \underline{\mathbf{X}'}\|_F} \leq \max_{x,y \in [0,\alpha]} \frac{|\log(x+a) - \log(y+a)|}{|x-y|} \\ &\stackrel{(b)}{\Rightarrow} \min_{x,y \in [0,\alpha]} \frac{\log(x+a) - \log(y+a)}{|x-y|} \\ &\leq \frac{\|h(\underline{\mathbf{X}}) - h(\underline{\mathbf{X}'})\|_F}{\|\underline{\mathbf{X}} - \underline{\mathbf{X}'}\|_F} \leq \max_{x,y \in [0,\alpha]} \frac{\log(x+a) - \log(y+a)}{x-y}, \end{aligned} \quad (38)$$

where x and y represent two arbitrary entries of $\underline{\mathbf{X}}$ and $\underline{\mathbf{X}'}$, respectively, (b) holds because the numerator and the denominator have the same sign.

Using the mean value theorem, there exists $z \in [x, y]$ such that

$$h(z) = \frac{\log(x+a) - \log(y+a)}{x-y},$$

where $\dot{h}(z)$ is the derivative of h at z . We know that $\dot{h}(z) = \frac{1}{z+a}$. Since $\dot{h}(z)$ is monotonically increasing with z , the minimum and maximum of (38) can be bounded by the minimum and maximum value of $\dot{h}(z)$. Hence,

$$\min_{x,y \in [0,\alpha]} \frac{\log(x+a) - \log(y+a)}{x-y} = \min_{z \in [0,\alpha]} \frac{1}{z+a} = \frac{1}{\alpha+a}.$$

Similarly, we have $\max_{x,y \in [0,\alpha]} \frac{\log(x+a) - \log(y+a)}{x-y} = \frac{1}{a}$. Combining the above with (38) concludes the proof.

APPENDIX C PROOF OF FACT 2

Note that $\underline{\mathbf{M}}$ is bounded if $\|\underline{\mathbf{X}}\|_\infty \leq \alpha$:

$$\begin{aligned} \|\underline{\mathbf{M}} - \mathbf{0}\|_\infty &= \|h(\underline{\mathbf{X}}) - h((1-a)\mathbf{1})\|_\infty \\ \implies \|\underline{\mathbf{M}}\|_\infty &\stackrel{(a)}{\leq} \max_{i,j,k} \frac{1}{a} |\underline{\mathbf{X}}(i,j,k) - (1-a)| \\ &\stackrel{(b)}{\leq} \frac{1}{a} \max_{i,j,k} (|\underline{\mathbf{X}}(i,j,k)| + |(1-a)|) \leq \frac{\alpha + |1-a|}{a}, \end{aligned}$$

where (a) used (25) and (b) used the triangle inequality.

Note that $|f_\ell(m)| = |\dot{\Phi}(b_\ell - m) - \dot{\Phi}(b_{\ell-1} - m)| < \infty$ since $\dot{\Phi}(\cdot)$ is upper bounded for a given variance. Moreover, the denominator

$$\begin{aligned} f_\ell(m) &= \Phi(b_\ell - m) - \Phi(b_{\ell-1} - m) \\ &= \int_{b_{\ell-1}}^{b_\ell} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(\frac{-t^2}{\sigma^2}\right) dt > 0, \end{aligned}$$

for given $b_\ell > b_{\ell-1}, \ell = \{2, \dots, Q\}$. Hence the supremum and infimum defined in (26) exist, provided that $\underline{\mathbf{X}}(i,j,k)$ is upper bounded by α .

APPENDIX D COVERING NUMBERS AND PROOF OF THEOREMS

A. Covering Numbers

We will use the following definition repeatedly:

Definition 1 (Covering Number [50]) *The covering number $N(\mathcal{W}, \epsilon)$ of a set \mathcal{W} with parameter $\epsilon > 0$ is the smallest cardinality of any $\bar{\mathcal{W}} \subseteq \mathcal{W}$ such that for any $\mathbf{w} \in \mathcal{W}$, there exists a $\bar{\mathbf{w}} \in \bar{\mathcal{W}} \subset \mathcal{W}$ satisfying $\|\bar{\mathbf{w}} - \mathbf{w}\| \leq \epsilon$. The discretized set $\bar{\mathcal{W}}$ is called the ϵ -net of \mathcal{W} .*

Roughly speaking, the continuous set \mathcal{X} can be approximated by a discrete set whose cardinality is $N(\mathcal{X}, \epsilon)$.

Lemma 2 (Covering Number of \mathcal{X}^{BTD}) *The covering number of \mathcal{X}^{BTD} , i.e., $N(\mathcal{X}^{\text{BTD}}, \epsilon)$, is upper bounded by*

$$\left(\frac{3(\kappa + \beta)R}{\epsilon}\right)^{(I+J)L+K} \beta^{(I+J)L} \left(\frac{\kappa}{2}\right)^{RK}. \quad (39)$$

Proof: Let $\bar{\mathcal{A}}_r$ denote an $\frac{\epsilon}{2\sqrt{\beta}}$ -net of $\mathcal{A}_r = \{\mathbf{A}_r \in \mathbb{R}^{I \times L} \mid \|\mathbf{A}_r\|_\text{F} \leq \sqrt{\beta}, \mathbf{A}_r \geq 0\}$. Similarly, let $\bar{\mathcal{B}}_r$ denote an $\frac{\epsilon}{2\sqrt{\beta}}$ -net of $\mathcal{B}_r = \{\mathbf{B}_r \in \mathbb{R}^{J \times L} \mid \|\mathbf{B}_r\|_\text{F} \leq \sqrt{\beta}, \mathbf{B}_r \geq 0\}$. According to [51, Lemma 4.1], the covering number of an ϵ -net of D -dimensional Euclidean ball, i.e., $\{\mathbf{x} \in \mathbb{R}^D \mid \|\mathbf{x}\|_2 \leq \delta\}$ is upper bounded by $\left(\frac{3\delta}{\epsilon}\right)^D$. Therefore,

$$|\bar{\mathcal{A}}_r| \leq \left(\frac{6\beta}{\epsilon}\right)^{IL} \frac{1}{2^{IL}}, \quad |\bar{\mathcal{B}}_r| \leq \left(\frac{6\beta}{\epsilon}\right)^{JL} \frac{1}{2^{JL}}, \quad (40)$$

where the constant factors $\frac{1}{2^{IL}}$ and $\frac{1}{2^{JL}}$ are due to non-negativity constraints in the sets \mathcal{A}_r and \mathcal{B}_r , respectively. Next, consider the set $\mathcal{F}_r = \{\mathbf{S}_r = \mathbf{A}_r \mathbf{B}_r^\top \in \mathbb{R}^{I \times J} \mid \text{rank}(\mathbf{S}_r) \leq$

$L, \|\mathbf{S}_r\|_\text{F} \leq \beta, \mathbf{S}_r \geq 0, \mathbf{A}_r \in \mathcal{A}_r, \mathbf{B}_r \in \mathcal{B}_r\}$. We can construct an ϵ -net of \mathcal{F}_r such that the following holds:

$$\begin{aligned} \|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_\text{F} &= \|\mathbf{A}_r \mathbf{B}_r^\top - \bar{\mathbf{A}}_r \bar{\mathbf{B}}_r^\top\|_\text{F} \\ &\leq \|\mathbf{A}_r(\mathbf{B}_r - \bar{\mathbf{B}}_r)^\top\|_\text{F} + \|(\mathbf{A}_r - \bar{\mathbf{A}}_r)\bar{\mathbf{B}}_r^\top\|_\text{F} \end{aligned} \quad (41a)$$

$$\begin{aligned} &\leq \|\mathbf{A}_r\|_\text{F} \|\mathbf{B}_r^\top - \bar{\mathbf{B}}_r^\top\|_\text{F} + \|\mathbf{A}_r - \bar{\mathbf{A}}_r\|_\text{F} \|\bar{\mathbf{B}}_r^\top\|_\text{F} \quad (41b) \\ &\leq \sqrt{\beta} \left(\frac{\epsilon}{2\sqrt{\beta}} + \frac{\epsilon}{2\sqrt{\beta}} \right) \leq \epsilon, \end{aligned}$$

where (41a) uses the triangle inequality and (41b) follows from the Cauchy-Schwarz inequality. The cardinality of an ϵ -net of \mathcal{F}_r is equal to the product of cardinalities of the sets $\bar{\mathcal{A}}_r$ and $\bar{\mathcal{B}}_r$. Let $\bar{\mathcal{F}}_r$ be the $\frac{\epsilon}{R(\kappa + \beta)}$ -net of \mathcal{F}_r , then

$$|\bar{\mathcal{F}}_r| \leq \left(\frac{3\beta R(\kappa + \beta)}{\epsilon}\right)^{(I+J)L}. \quad (42)$$

Now, let $\bar{\mathcal{C}}_r$ denote an $\frac{\epsilon}{R(\kappa + \beta)}$ -net of $\mathcal{C}_r = \{\mathbf{c}_r \in \mathbb{R}^K \mid \|\mathbf{c}_r\|_2 \leq \kappa, \mathbf{c}_r \geq 0\}$. Then, using [51, Lemma 4.1], we get

$$|\bar{\mathcal{C}}_r| \leq \left(\frac{3\kappa R(\kappa + \beta)}{\epsilon}\right)^K \frac{1}{2^K}. \quad (43)$$

Finally, consider a discrete set $\bar{\mathcal{X}}^{\text{BTD}}$. We hope to construct this set such that for every $\underline{\mathbf{X}} \in \mathcal{X}^{\text{BTD}}$, there exists an $\bar{\mathbf{X}} \in \bar{\mathcal{X}}^{\text{BTD}}$ such that $\|\underline{\mathbf{X}} - \bar{\mathbf{X}}\|_\text{F} \leq \epsilon$. To this end, we let

$$\bar{\mathcal{X}}^{\text{BTD}} = \left\{ \bar{\mathbf{X}} \mid \bar{\mathbf{X}} = \sum_{r=1}^R \bar{\mathbf{S}}_r \circ \mathbf{c}_r, \bar{\mathbf{S}}_r \in \bar{\mathcal{F}}_r, \bar{\mathbf{c}}_r \in \bar{\mathcal{C}}_r \right\}.$$

To see that the above construction is an ϵ -net of \mathcal{X}^{BTD} , we show that there exists an $\bar{\mathbf{X}} \in \bar{\mathcal{X}}^{\text{BTD}}$ that satisfies the following chain of inequalities for any $\underline{\mathbf{X}} \in \mathcal{X}^{\text{BTD}}$:

$$\begin{aligned} \|\underline{\mathbf{X}} - \bar{\mathbf{X}}\|_\text{F} &= \left\| \sum_{r=1}^R \mathbf{S}_r \circ \mathbf{c}_r - \sum_{r=1}^R \bar{\mathbf{S}}_r \circ \bar{\mathbf{c}}_r \right\|_\text{F} \\ &\stackrel{(a)}{=} \|\mathbf{C}\mathbf{S} - \bar{\mathbf{C}}\bar{\mathbf{S}}\|_\text{F} \leq \|\mathbf{C}(\mathbf{S} - \bar{\mathbf{S}})\|_\text{F} + \|(\mathbf{C} - \bar{\mathbf{C}})\bar{\mathbf{S}}\|_\text{F} \\ &\leq \|\mathbf{C}\|_\text{F} \|\mathbf{S} - \bar{\mathbf{S}}\|_\text{F} + \|\mathbf{C} - \bar{\mathbf{C}}\|_\text{F} \|\bar{\mathbf{S}}\|_\text{F} \\ &= \sqrt{R}\kappa \left(\sum_{r=1}^R \|\mathbf{S}_r - \bar{\mathbf{S}}_r\|_\text{F}^2 \right)^{\frac{1}{2}} + \sqrt{R}\beta \left(\sum_{r=1}^R \|\mathbf{c}_r - \bar{\mathbf{c}}_r\|_2^2 \right)^{\frac{1}{2}} \\ &\stackrel{(b)}{\leq} \sqrt{R}\kappa \left(R \left(\frac{\epsilon}{R(\kappa + \beta)} \right)^2 \right)^{\frac{1}{2}} + \sqrt{R}\beta \left(R \left(\frac{\epsilon}{R(\kappa + \beta)} \right)^2 \right)^{\frac{1}{2}} \\ &\leq \epsilon \end{aligned}$$

where $\mathbf{C} = [\mathbf{c}_1, \dots, \mathbf{c}_R]$ and $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_R]^\top$ with $\mathbf{s}_r = \text{vec}(\mathbf{S}_r)$ for $r \in [R]$ and (a) is based on the matrix representation of the outer product; see [7] [52]. The inequality (b) holds by picking the $\bar{\mathbf{S}}_r$ from $\bar{\mathcal{F}}_r$ that has the smallest distance to \mathbf{S}_r . The same applies to the pick of $\bar{\mathbf{c}}_r$.

The cardinality of $\bar{\mathcal{X}}^{\text{BTD}}$ is equal to the product of the cardinality of R sets of $\bar{\mathcal{F}}_r$ and $\bar{\mathcal{C}}_r$, respectively. Hence, the covering number of set \mathcal{X}^{BTD} is upper bounded by (39). ■

Under Assumption 3, the covering number of \mathcal{X}^{DGM} is bounded as well [7]:

Fact 3 (Covering Number of \mathcal{X}^{DGM} [7]) Under Assumption 3, the covering number of \mathcal{X}^{DGM} satisfies

$$N(\mathcal{X}^{\text{DGM}}, \epsilon) \leq (3R(\beta+\kappa)/\epsilon)^{R(K+D)} (\kappa/2)^{RK} (Pq)^{RD}. \quad (45)$$

B. Proof of Theorems 1 and 2

Let us define $\mathcal{T} = \{(0, 0), \dots, (I, J)\}$. We take N samples denoted by $\omega_1, \dots, \omega_N$ from \mathcal{T} uniformly with replacement and collect the indices in $\Omega = \{\omega_1, \omega_2, \dots, \omega_N\}$. Then, for every $\omega_i \in \Omega, k \in [K]$, $\underline{Y}_{\omega_i k}$ is observed with the probability mass function defined by $\{f_\ell(\underline{M}_{\omega_i k})\}_{\ell=1}^Q$. Here, we define a shorthand notation by letting $\underline{Y}_{\omega k} = \underline{Y}(i, j, k)$ and $\underline{Y}_{\omega:} = \underline{Y}(i, j, :)$, where $\omega \in \Omega$. We use the similar notation for \underline{M} . For any two fibers, $\underline{M}_{\omega:}$ and $\underline{Y}_{\omega:}, \forall \omega \in \Omega$ the function $g: \mathbb{R}^K \times \mathbb{R}^K \rightarrow \mathbb{R}$ is defined as follows:

$$g(\underline{M}_{\omega:}, \underline{Y}_{\omega:}) = \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} \log \left(\frac{1}{f_\ell(\underline{M}_{\omega k})} \right). \quad (46)$$

Recall that $\underline{M} = h(\underline{X})$. Then we define

$$\begin{aligned} F_{\Omega, \underline{Y}}(\underline{M}) &= \frac{1}{N} \sum_{(i, j) \in \Omega} g(\underline{M}(i, j, :), \underline{Y}(i, j, :)) \\ &= \frac{1}{N} \sum_{(i, j) \in \Omega} \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} \log \left(\frac{1}{f_\ell(\underline{M}_{ijk})} \right). \end{aligned}$$

Let, $\underline{M}^\natural = h(\underline{X}^\natural)$ be the ground truth tensor and $\underline{M}^* = h(\underline{X}^*)$ be the optimal solution of the optimization problem (5).

$$\begin{aligned} &\mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:}) - g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})] \\ &= \sum_{(i, j) \in \mathcal{T}} \frac{1}{IJ} \mathbb{E}_{\underline{Y}|\omega=(i, j)} (g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:}) - g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})) \\ &= \frac{1}{IJ} \left(\sum_{(i, j) \in \mathcal{T}} \sum_{k=1}^K \sum_{\ell=1}^Q P(\underline{Y}_{\omega k} = \ell) \log \left(\frac{f_\ell(\underline{M}_{\omega k}^*)}{f_\ell(\underline{M}_{\omega k}^\natural)} \right) \right) \\ &= \frac{K}{IJK} \sum_{(i, j, k) \in [I] \times [J] \times [K]} \sum_{\ell=1}^Q f_\ell(\underline{M}_{\omega k}^\natural) \log \left(\frac{f_\ell(\underline{M}_{\omega k}^*)}{f_\ell(\underline{M}_{\omega k}^\natural)} \right) \\ &= -K \cdot \text{KL}(\underline{M}^* || \underline{M}^\natural), \end{aligned} \quad (47)$$

where $\text{KL}(\underline{M}^* || \underline{M}^\natural) = \frac{1}{IJK} \sum_{i, j, k} \text{KL}(\underline{M}_{ijk}^* || \underline{M}_{ijk}^\natural)$ and $\text{KL}(\underline{M}_{ijk}^* || \underline{M}_{ijk}^\natural) = \sum_{\ell=1}^Q f_\ell(\underline{M}_{ijk}^\natural) \log \frac{f_\ell(\underline{M}_{ijk}^*)}{f_\ell(\underline{M}_{ijk}^\natural)}$, which follow the definitions in [33].

Let us denote $\widetilde{\underline{M}} = \arg \min_{\underline{M} \in \mathcal{M}} \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})]$, where $\mathcal{M} \in \{\mathcal{M}^{\text{BTD}}, \mathcal{M}^{\text{DGM}}\}$, depending on the model that we use. That is, $\widetilde{\underline{M}}$ is the solution that we can obtain from the expected version of our MLE.

Also, we have $F_{\Omega, \underline{Y}}(\underline{M}^*) \leq F_{\Omega, \underline{Y}}(\widetilde{\underline{M}})$, which implies the following chain of inequalities:

$$\begin{aligned} 0 &\leq F_{\Omega, \underline{Y}}(\widetilde{\underline{M}}) - F_{\Omega, \underline{Y}}(\underline{M}^\natural) + F_{\Omega, \underline{Y}}(\underline{M}^\natural) - F_{\Omega, \underline{Y}}(\underline{M}^*) \\ &\leq \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:})] - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})] + F_{\Omega, \underline{Y}}(\underline{M}^\natural) \\ &\quad - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:})] + \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})] - F_{\Omega, \underline{Y}}(\underline{M}^*) \\ &\quad + F_{\Omega, \underline{Y}}(\widetilde{\underline{M}}) - F_{\Omega, \underline{Y}}(\underline{M}^\natural) \end{aligned} \quad (48)$$

$$\begin{aligned} &\stackrel{(a)}{\leq} \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:}) - g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})] \\ &\quad + |F_{\Omega, \underline{Y}}(\underline{M}^\natural) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:})]| \\ &\quad + |F_{\Omega, \underline{Y}}(\underline{M}^*) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^*, \underline{Y}_{\omega:})]| + |F_{\Omega, \underline{Y}}(\widetilde{\underline{M}}) - F_{\Omega, \underline{Y}}(\underline{M}^\natural)| \\ &\stackrel{(b)}{\Rightarrow} K \cdot \text{KL}(\underline{M}^* || \underline{M}^\natural) \leq |F_{\Omega, \underline{Y}}(\widetilde{\underline{M}}) - F_{\Omega, \underline{Y}}(\underline{M}^\natural)| \\ &\quad + |F_{\Omega, \underline{Y}}(\underline{M}^\natural) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:})]| \\ &\quad + \sup_{\underline{M} \in \mathcal{M}} |F_{\Omega, \underline{Y}}(\underline{M}) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})]|, \end{aligned} \quad (49)$$

where (a) is by the triangle inequality and (b) follows (47).

Recall that $\min_{\widetilde{\underline{X}} \in \mathcal{X}_{R, \theta}} \|\widetilde{\underline{X}} - \underline{X}^\natural\|_\infty \leq \nu$, where $\nu \in \{\nu^{\text{BTD}}, \nu^{\text{DGM}}\}$ and $\nu \geq 0$ is a constant. The first term on the R.H.S. of (49) can be bounded as follows

$$\begin{aligned} &|F_{\Omega, \underline{Y}}(\widetilde{\underline{M}}) - F_{\Omega, \underline{Y}}(\underline{M}^\natural)| \\ &= \left| \frac{1}{N} \sum_{(i, j) \in \Omega} \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} (\log(f_\ell(\underline{M}_{ijk}^\natural)) - \log(f_\ell(\widetilde{\underline{M}}_{ijk}))) \right| \\ &\stackrel{(a)}{\leq} \frac{1}{N} \sum_{(i, j) \in \Omega} \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} |\log(f_\ell(\underline{M}_{ijk}^\natural)) - \log(f_\ell(\widetilde{\underline{M}}_{ijk}))| \\ &\stackrel{(b)}{\leq} \frac{1}{N} \sum_{(i, j) \in \Omega} \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} \mathsf{L}_\alpha \frac{1}{a} \|\underline{X}^\natural - \widetilde{\underline{X}}\|_\infty \\ &\stackrel{(c)}{\leq} \frac{\mathsf{L}_\alpha \nu}{aN} \sum_{(i, j) \in \Omega} \sum_{k=1}^K \sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} \leq \frac{\mathsf{L}_\alpha \nu}{aN} \sum_{(i, j) \in \Omega} \sum_{k=1}^K 1 \leq \frac{K \mathsf{L}_\alpha \nu}{a}. \end{aligned}$$

In the above, (a) used the triangle inequality. In addition, (b) used the property of log-concavity, Fact 2, and (35). The (c) part used the fact that $\sum_{\ell=1}^Q \mathbb{1}_{[\underline{Y}_{ijk} = \ell]} = 1$.

The second term in (49) can be bounded using Hoeffding's inequality [53] on independent random variable $Z_\omega = g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})$ and $Z_\omega \in [0, KU_\alpha]$. Then with probability greater than $1 - \delta$, the following holds:

$$|F_{\Omega, \underline{Y}}(\underline{M}^\natural) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}^\natural, \underline{Y}_{\omega:})]| \leq \sqrt{(KU_\alpha)^2 \log(\frac{1}{\delta})/2N}. \quad (51)$$

For the third term in (49), observe that $F_{\Omega, \underline{Y}}(\underline{M})$ is a sample average of the samples of joint random variable (ω, \underline{Y}) drawn independently from their joint distribution. And the corresponding expected value is $\mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})]$. Hence, the third term can be bounded via Rademacher Complexity-based arguments [50]. Specifically, with probability at least $1 - \delta$, the following holds

$$\begin{aligned} &\sup_{\underline{M} \in \mathcal{M}} |F_{\Omega, \underline{Y}}(\underline{M}) - \mathbb{E}_{\omega, \underline{Y}}[g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})]| \\ &\leq 2\widehat{\mathcal{R}}(\mathcal{G}) + KU_\alpha \sqrt{\frac{8 \log(\frac{2}{\delta})}{N}}, \end{aligned} \quad (52)$$

where $\widehat{\mathcal{R}}(\mathcal{G})$ is the empirical Rademacher complexity of the function class \mathcal{G} defined as follows:

$$\mathcal{G} = \{(\omega, \underline{Y}) \mapsto g(\underline{M}_{\omega:}, \underline{Y}_{\omega:}) | \underline{M} = h(\underline{X}), \underline{X} \in \mathcal{X}\}.$$

Using Talagrand Lemma [54, Lemma 4.2] along with Lemma 1 and Lipschitz property of $g(\underline{M}_{\omega:}, \underline{Y}_{\omega:})$ in terms of f_ℓ in Fact 2,

$$\widehat{\mathcal{R}}(\mathcal{G}) \leq \frac{\mathsf{L}_\alpha}{a} \widehat{\mathcal{R}}(\mathcal{X}),$$

where $\mathcal{X} \in \{\mathcal{X}^{\text{BTD}}, \mathcal{X}^{\text{DGM}}\}$. Using Dudley's entropy integral [55, Lemma A.5], one can bound the empirical Rademacher complexity by the covering number (cf. Appendix D-A) of the set as follows:

$$\widehat{\mathcal{R}}(\mathcal{X}) \leq \zeta(\mathcal{X}, \mu), \quad (53)$$

where

$$\zeta(\mathcal{X}, \mu) = \inf_{\mu > 0} \left(\frac{4\mu}{\sqrt{N}} + \frac{12}{N} \int_{\mu}^{\sqrt{N}} \sqrt{\log N(\mathcal{X}, \epsilon)} d\epsilon \right). \quad (54)$$

Combining (49), (50), (51), (52), and (53), we get

$$\begin{aligned} \text{KL}(\underline{\mathbf{M}}^* \parallel \underline{\mathbf{M}}^\natural) &\leq \quad (55) \\ \frac{2\mathsf{L}_\alpha \zeta(\mathcal{X}, \mu)}{aK} + \frac{K\mathsf{U}_\alpha}{K} \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} + \frac{\mathsf{U}_\alpha K}{K} \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + \frac{K\mathsf{L}_\alpha \nu}{aK}. \end{aligned}$$

The relation between Hellinger distance and KL divergence [30] is as follows:

$$d_H^2(\underline{\mathbf{M}}^*, \underline{\mathbf{M}}^\natural) \leq \text{KL}(\underline{\mathbf{M}}^* \parallel \underline{\mathbf{M}}^\natural). \quad (56)$$

From [32, Lemma 2] and Fact 2, we can bound Hellinger distance using the following:

$$\frac{\mathsf{F}_\alpha}{4} \frac{\|\underline{\mathbf{M}}^* - \underline{\mathbf{M}}^\natural\|_{\text{F}}^2}{IJK} \leq d_H^2(\underline{\mathbf{M}}^*, \underline{\mathbf{M}}^\natural). \quad (57)$$

Combining (56), (57) and Lemma 1, we get

$$\begin{aligned} \frac{\|\underline{\mathbf{M}}^* - \underline{\mathbf{M}}^\natural\|_{\text{F}}^2}{IJK} &\leq \frac{4}{\mathsf{F}_\alpha} \text{KL}(\underline{\mathbf{M}}^* \parallel \underline{\mathbf{M}}^\natural) \\ \implies \frac{\|\underline{\mathbf{X}}^* - \underline{\mathbf{X}}^\natural\|_{\text{F}}^2}{IJK} &\leq \frac{4(\alpha + a)^2}{\mathsf{F}_\alpha} \text{KL}(\underline{\mathbf{M}}^* \parallel \underline{\mathbf{M}}^\natural). \end{aligned} \quad (58)$$

Let us define the following constants,

$$C_1 = \frac{4(\alpha + a)^2}{\mathsf{F}_\alpha}, \text{ and } C_2 = \frac{\mathsf{L}_\alpha}{a}. \quad (59)$$

Combining (55), (59) and (58),

$$\begin{aligned} \frac{\|\underline{\mathbf{X}}^* - \underline{\mathbf{X}}^\natural\|_{\text{F}}^2}{IJK} &\leq \quad (60) \\ C_1 \left(\frac{2C_2 \zeta(\mathcal{X}, \mu)}{K} + \mathsf{U}_\alpha \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} + \mathsf{U}_\alpha \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + C_2 \nu \right). \end{aligned}$$

Since $(\log N(\mathcal{X}, \epsilon))^{1/2}$ increases with the decrease in ϵ , the integral in (54) can be upper bounded as,

$$\int_{\mu}^{\sqrt{N}} (\log N(\mathcal{X}, \epsilon))^{1/2} d\epsilon \leq \sqrt{N} (\log N(\mathcal{X}, \mu))^{1/2}. \quad (61)$$

In the next two subsections, we will prove Theorems 1 and 2, respectively, using the above technical preparations.

1) *Proof of Theorem 1* : Setting $\mu = \sqrt{R}$ and $\mathcal{X} = \mathcal{X}^{\text{BTD}}$ in $\zeta(\mathcal{X}, \mu)$, and using Lemma 2 and (61),

$$\begin{aligned} &\int_{\mu}^{\sqrt{N}} (\log N(\mathcal{X}^{\text{BTD}}, \epsilon))^{1/2} d\epsilon \\ &= \sqrt{NR} \left(\sqrt{((I+J)L+K) \log(3\sqrt{R}(\beta+\kappa))} \right. \\ &\quad \left. + \sqrt{K \log(\kappa/2)} + (I+J)L \log(\beta) \right) = \frac{\sqrt{NR}}{3} \tau, \end{aligned} \quad (62)$$

where we have

$$\begin{aligned} \tau &= 3 \left(\sqrt{((I+J)L+K) \log(3\sqrt{R}(\beta+\kappa))} \right. \\ &\quad \left. + \sqrt{K \log(\kappa/2)} + (I+J)L \log(\beta) \right). \end{aligned}$$

Combining (62) and (54),

$$\zeta(\mathcal{X}^{\text{BTD}}, \sqrt{R}) = 4(1+\tau) \sqrt{\frac{R}{N}}. \quad (63)$$

Combining (60) and (63),

$$\begin{aligned} \frac{\|\underline{\mathbf{X}}^* - \underline{\mathbf{X}}^\natural\|_{\text{F}}^2}{IJK} &\leq \frac{8C_1 C_2 (1+\tau)}{K} \sqrt{\frac{R}{N}} + \mathsf{U}_\alpha C_1 \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} \\ &\quad + \mathsf{U}_\alpha C_1 \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + C_1 C_2 \nu^{\text{BTD}}. \end{aligned} \quad (64)$$

Note that as both the second and third terms in (49) are upper bounded with probability at least $1 - \delta$, the over all upper bound in (64) holds with probability at least $1 - 2\delta$. This completes the proof of Theorem 1.

2) *Proof of Theorem 2* : Setting $\mu = \sqrt{R}$ and $\mathcal{X} = \mathcal{X}^{\text{DGM}}$ in $\zeta(\mathcal{X}, \mu)$, and using Fact 3 and (61),

$$\begin{aligned} &\int_{\mu}^{\sqrt{N}} (\log N(\mathcal{X}^{\text{DGM}}, \epsilon))^{1/2} d\epsilon \\ &= \sqrt{NR} \sqrt{(K+D) \log(3\sqrt{R}(\beta+\kappa)) + K \log(\kappa/2) + D \log(Pq)} \\ &= \sqrt{NR} \sqrt{K \log(\frac{3}{2} \sqrt{R} \kappa(\beta+\kappa)) + D \log(3\sqrt{R} Pq(\beta+\kappa))} \\ &= \frac{\sqrt{NR}}{3} \tau, \end{aligned} \quad (65)$$

where we have

$$\tau = 3 \sqrt{K \log(\frac{3}{2} \sqrt{R} \kappa(\beta+\kappa)) + D \log(3\sqrt{R} Pq(\beta+\kappa))}.$$

Combining (65) and (54),

$$\zeta(\mathcal{X}^{\text{DGM}}, \sqrt{R}) = 4(1+\tau) \sqrt{\frac{R}{N}}. \quad (66)$$

Combining (60) and (66), we have the following holds with probability at least $1 - 2\delta$:

$$\begin{aligned} \frac{\|\underline{\mathbf{X}}^* - \underline{\mathbf{X}}^\natural\|_{\text{F}}^2}{IJK} &\leq \frac{8C_1 C_2 (1+\tau)}{K} \sqrt{\frac{R}{N}} + \mathsf{U}_\alpha C_1 \sqrt{\frac{\log(\frac{1}{\delta})}{2N}} \\ &\quad + \mathsf{U}_\alpha C_1 \sqrt{\frac{8\log(\frac{2}{\delta})}{N}} + C_1 C_2 \nu^{\text{DGM}}. \end{aligned} \quad (67)$$

This completes the proof of Theorem 2.