

Article

Rank-Adaptive Bayesian Tensor Ring Completion for Low-Altitude 5D Radio Environment Map Construction

Ying Wang , Zhuo Sun *  and Hao Ma 

School of Information and Communication Engineering, Beijing University of Posts and Telecommunications, Beijing 100876, China; wying@bupt.edu.cn (Y.W.); haoma@bupt.edu.cn (H.M.)

* Correspondence: zhuosun@bupt.edu.cn

Abstract

The rapid development of the low-altitude economy demands comprehensive electromagnetic spectrum awareness. However, constructing a comprehensive radio environment map (REM) in this scenario is challenging, as spectrum sensing data collected by unmanned aerial vehicles (UAVs) in complex low-altitude environments is typically sparse, fragmented, and non-uniformly distributed across the high-dimensional space of time, frequency, and 3D space. To address these issues, this study proposes a rank-adaptive Bayesian tensor ring completion (Ra-BTRC) framework. The method models the low-altitude electromagnetic environment as a unified five-dimensional (5D) spectrum tensor. It then employs tensor ring (TR) decomposition to capture latent high-order correlations across all dimensions. To overcome the sensitivity of conventional TR methods to predefined ranks, Ra-BTRC introduces sparsity-inducing priors on the TR core factors, enabling variational Bayesian inference to learn observation uncertainty and infer effective TR ranks from sparse measurements without manually fixing the TR rank. Simulations demonstrate that Ra-BTRC significantly outperforms existing TR-based baselines, achieving more than 10 dB MMSE improvement at a 5% sampling rate while accurately recovering local spectrum structures and temporal dynamics. The proposed approach provides a robust and scalable solution for reliable global low-altitude spectrum cognition under stringent sensing budgets.

Keywords: low-altitude economy; radio environment map; tensor completion; tensor ring decomposition; Bayesian inference

1. Introduction

The low-altitude economy (LAE) is rapidly emerging as a major driver of a new wave of industrial upgrading and urban airspace development [1]. Intelligent low-altitude platforms, represented by unmanned aerial vehicles (UAVs) and electric vertical take-off and landing aircraft (eVTOLs), are being increasingly deployed in logistics delivery, emergency rescue, urban inspection, aerial transportation, and public safety [2]. Such aerial systems rely heavily on robust wireless communications for critical operations, including flight control, swarm coordination, and payload data transmission [3]. However, the inherently high mobility in low-altitude airspace, frequent channel fluctuations, and complex interference caused by heterogeneous services lead to continuous temporal and spatial variations in interference levels, channel conditions, and traffic demands. Consequently, this exacerbates spectrum congestion, reduces frequency reuse efficiency, and increases difficulty in resource management [4]. Furthermore, comprehensive electromagnetic spectrum awareness is



Academic Editor: Fabrizio Messina

Received: 14 April 2026

Revised: 12 June 2026

Accepted: 1 July 2026

Published: 3 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

foundational to various downstream wireless capabilities, including interference avoidance, adaptive link selection, interference-free trajectory or routing planning, localization, and environmental sensing [5]. Therefore, accurately estimating electromagnetic spectrum usage and its evolution patterns is a crucial prerequisite for enhancing the operational efficiency, reliability, and safety of the LAE [6].

The electromagnetic spectrum map, or radio environment map (REM), is a key technology for low-altitude spectrum cognition [7]. By fusing multidimensional sensing data into a digital twin of the electromagnetic environment, it provides global a priori information for 6G dynamic spectrum sharing and UAV path planning [8]. However, owing to the stringent size, weight, and power constraints of UAVs, individual sensing capabilities are restricted to limited communication range and specific service bands, making continuous wide-area coverage unattainable [9]. Additionally, the complex low-altitude environment characterized by severe shadowing from buildings and fast time-varying channels due to high mobility, causes the collected spectrum data to be highly sparse, fragmented, and non-uniformly distributed. Consequently, leveraging inherent spectrum sensing data correlations to accurately complete missing values and construct the spectrum situations of unobserved regions has emerged as the fundamental challenge in building practical low-altitude REMs [10].

For REM construction in low-altitude scenarios, existing studies mainly rely on UAV-borne measurements to infer the radio environment along flight trajectories. In particular, aerial spectrum surveying frameworks have been proposed to gather measurements online and continuously update radio map estimates during UAV missions [11]. To support integrated space-air-ground networking, 3D voxel-based REM construction has been investigated, where discretization accuracy, boundary effects, and flight costs are jointly considered and path planning is often incorporated to improve mapping efficiency [12]. In terms of modeling, a representative approach explicitly incorporates environmental geometry by reconstructing 3D obstacle semantics prior to radio mapping and further formulates air-to-ground channel gain prediction as a full-dimensional (6D) radio map construction problem, thereby effectively reducing the required measurement density [13]. Moreover, multi-UAV collaborative mapping in urban environments has also been studied to address the complexity of low-altitude propagation and to enable more systematic data collection and map construction pipelines [14].

Although these studies have advanced low-altitude REM from 2D to 3D, their primary outputs remain spatial maps constrained to specific frequency bands and limited time windows. Essentially, they employ slice-based or quasi-static processing in the time and frequency dimensions [11–14]. However, for low-altitude systems jointly shaped by mobility-induced temporal variation, frequency selectivity, and altitude-/blockage-dependent spatial heterogeneity, the spectrum situation can be more naturally viewed as a strongly coupled five-dimensional (5D) field over time, frequency, and three-dimensional space. Existing 3D methods inherently capture mere “snapshots” of this 5D system, failing to fully exploit the high-order correlations across time, frequency, and altitude layers. Conversely, naively extending low-dimensional construction frameworks to the full 5D setting under severely sparse and non-uniform observations leads to a highly ill-posed completion problem with severe sample-complexity and computational challenges, which hampers reconstruction accuracy and the construction of highly available and scalable global low-altitude REMs [15].

Low-rank tensor completion offers an effective paradigm for modeling high-dimensional spectrum data [16]. By capturing the latent low-rank structure of multidimensional observations, it transforms high-dimensional sparse sensing into a missing-entry recovery problem that leverages cross-domain correlations, thereby enabling high-dimensional REM

construction under limited sampling budgets [15–17]. Compared to traditional CANDECOMP/PARAFAC (CP) or Tucker decompositions [18], tensor ring (TR) decomposition represents high-order tensors via a circular cascade of core tensors, offering greater representational flexibility and a more balanced rank distribution [19]. Under controlled ranks, its parameter scale and computational overhead scale more gracefully with the tensor order, making it exceptionally well-suited for characterizing the tightly coupled 5D spectrum tensor structure [20]. Recent studies have introduced TR to REM completion. For instance, some works tensorize spatial–temporal spectrum data and employ TR low-rank factors to recover missing entries, further incorporating a dynamic window mechanism to reduce computational overhead and enhance real-time performance in time-varying scenarios [21]. However, existing TR completion methods (such as TR-ALS [22] and TR-WOPT [23] based on alternating optimization) typically require pre-defining or initializing the TR rank, followed by optimization under a fixed rank. A large rank causes overfitting and spurious structures under sparse observations, while a small rank leads to underfitting and lost spectral details [24]. This problem is particularly critical in low-altitude REM construction because the intrinsic rank of the spectrum tensor changes with emitter activity, propagation complexity, altitude-dependent blockage, and sensing sparsity. For example, smooth line-of-sight dominated regions may require only a compact rank structure, whereas multi-emitter interference, blockage-induced shadowing, or highly fragmented UAV observations may require higher model capacity. A manually fixed rank is therefore unlikely to remain suitable across different airspace regions, time slots, and sampling budgets.

To address the above challenges, we propose a Rank-Adaptive Bayesian Tensor Ring Completion (Ra-BTRC) method for constructing 5D REMs in low-altitude scenarios. The main contributions of this work are summarized as follows:

- We formulate the low-altitude electromagnetic environment as a unified 5D (T-F-X-Y-Z) dynamic spectrum tensor field. By introducing TR decomposition, our approach fully exploits the latent high-order correlations across spatial, temporal, and frequency domains, effectively mitigating the curse of dimensionality inherent in high-dimensional spectrum completion.
- To overcome the sensitivity of conventional TR algorithms to predefined ranks, we develop the Ra-BTRC model for sparse low-altitude REM construction. By imposing sparsity-inducing hierarchical priors on the TR core factors, the model automatically prunes redundant components during variational inference. This yields adaptive effective rank estimation without requiring prior rank specification.
- Extensive experiments demonstrate that Ra-BTRC significantly outperforms existing baselines at extremely low sampling rates (e.g., 1–5%). The proposed method accurately captures the strong local correlations and non-linear features of the 5D spectrum data, providing a highly reliable and scalable technical pathway for deploying 6G low-altitude cognitive networks under severe sampling constraints.

2. Related Works

2.1. REM Construction in Low-Altitude Scenarios

Recent years have witnessed increasing interest in constructing REMs for LAE networks, where UAVs/eVTOLs rely on reliable wireless links for command-and-control and payload data delivery. A common paradigm is to use UAV-borne measurements as the primary data source, since aerial platforms can flexibly probe the radio environment in 3D spaces and collect spectrum observations at locations that are difficult to instrument from the ground.

A representative line of work treats REM construction as an aerial surveying problem, a UAV autonomously collects measurements along its flight path and updates the radio

map estimate online to improve adaptability in dynamic environments. For example, Romero et al. [11] proposed an aerial spectrum surveying framework that performs online Bayesian learning to update the radio map estimate (and its uncertainty) whenever a new measurement is acquired, enabling real-time operation during UAV missions. To address the trade-off between sensing accuracy and operational costs, Wei et al. [12] investigated 3D voxel-based REM construction within integrated space–air–ground networks. Their work emphasizes the optimization of discretization accuracy and boundary errors through strategic path planning. Similarly, Ivanov et al. [25] conducted large-scale volumetric measurements to empirically validate the spatial variation of 3D aerial links, providing a data-driven basis for REM estimation under limited sampling budgets.

Beyond pure data-driven methods, recent studies have integrated environmental geometry to improve reconstruction efficiency. A representative approach proposed by Liu et al. [13] explicitly incorporates 3D obstacle semantics into the mapping process, formulating air-to-ground channel gain prediction as a full-dimensional (6D) radio map problem, which effectively reduces the required measurement density by leveraging physical priors. Furthermore, Lin et al. [14] explored multi-UAV collaborative mapping in urban environments, highlighting that the intricate multipath and shadowing effects in low-altitude channels prevent simplified analytical modeling and require more systematic data collection and completion pipelines.

Although the above studies have significantly advanced low-altitude REM construction, their primary outputs are still 2D maps at a fixed altitude or 3D spatial maps over a limited time window and/or a single (or a few) frequency bands. In other words, the time and frequency dimensions are often handled in a slice-wise manner, e.g., building separate maps per time window or per band, which effectively treats the underlying low-altitude radio environment as a sequence of quasi-static snapshots. This practice may lose exploitable high-order correlations across time–frequency–3D space, thereby limiting construction fidelity and scalability when measurements are extremely sparse.

2.2. Tensor Decomposition and Completion in REM

REMs naturally exhibit multi-dimensional structures (e.g., space–frequency–time), which can be conveniently represented as high-order tensors. Under limited sensing budgets, recovering a complete REM from sparse observations is commonly formulated as a tensor completion problem by exploiting the latent low-rank structure of the underlying REM tensor. In such formulations, tensor decomposition provides a concrete parameterization of the low-rank prior by factorizing the REM tensor into a small set of latent factors, thereby reducing the degrees of freedom and enabling missing-entry recovery from incomplete measurements [18].

Early REM construction works primarily adopted CP and Tucker decompositions to capture multi-linear correlations. For instance, Tang et al. [16] utilized the HaLRTC formulation to recover missing measurements in multi-dimensional spectrum data, while Chen et al. [15] demonstrated that CP factorization can effectively exploit cross-domain correlations to reconstruct dynamic spatial–spectral–temporal REMs. While these foundational models offer concise representations, CP-rank determination is generally NP-hard and Tucker-based approaches often suffer from the curse of dimensionality as the tensor order increases, limiting their applicability in 5D REM construction.

To improve scalability, tensor train (TT) and TR decompositions have been introduced for high-order REM construction [21,26]. Although TT models offer linear storage complexity, their open-chain structure may fail to capture the cyclic or symmetric correlations inherent in circular UAV trajectories or periodic temporal patterns. In contrast, TR is particularly suitable for high-dimensional coupled REMs due to its circular struc-

ture and more balanced rank distribution. Recent advancements, such as the TRLRF and dynamic-window TR frameworks [21], have successfully employed TR low-rank factors to enhance reconstruction accuracy and computational efficiency in time-varying scenarios, demonstrating superior representational flexibility for 5D structures.

However, most TR completion methods (e.g., TR-ALS [22], TR-WOPT [23]) still require pre-defining the TR rank, which remains a primary bottleneck for autonomous REM construction. In dynamic low-altitude scenarios, the intrinsic rank of the spectrum tensor is typically unknown and fluctuates with environmental complexity. A rank set too high often leads to overfitting, where the model misinterprets measurement noise as signal structure, whereas a rank set too low results in underfitting, smoothing out critical spectral details such as localized interference peaks. This sensitivity to rank selection hinders the robustness and self-adaptive operation of TR-based frameworks in 6G cognitive networks, necessitating more advanced rank-adaptive inference mechanisms.

To clarify the research gap, Table 1 summarizes representative studies from the two lines of related work. Existing low-altitude REM studies mainly focus on 2D/3D spatial maps, while tensor completion-based REM methods exploit CP, Tucker, TT, or TR structures but generally do not address adaptive TR rank estimation for sparse 5D low-altitude REM reconstruction.

Table 1. Comparison of representative low-altitude REM construction and tensor completion-based REM methods.

Category	Works	Dimension	Strategy	Strength	Limitation
Low-altitude REM	Aerial surveying [11]; voxel REM [12]	X-Y/X-Y-Z	UAV sensing; online/voxel mapping	Supports UAV-based map update and spatial discretization	Limited time–frequency–3D coupling
Low-altitude REM	Geometry-assisted and multi-UAV mapping [13,14]	X-Y-Z or geometry-assisted map	Obstacle-aware and collaborative sensing	Improves urban mapping with environmental/UAV information	Not focused on sparse 5D spectrum tensor completion
Tensor REM	HaLRTC/Tucker; CP methods [15,16]	T-F-X/T-F-X-Y	Tucker/nuclear norm; CP factorization	Exploits multilinear low-rank structure	Rank selection and scalability are difficult in high order
Tensor REM	TT cartography [26]	X-Y-Z	Tensor train completion	Linear storage complexity	Weak direct boundary-mode coupling
Tensor REM	TRLRF and dynamic-window TR [21]	T-X-Y	Tensor ring completion	Flexible high-order representation	TR ranks require predefined or external selection
Proposed	Ra-BTRC	T-F-X-Y-Z	Bayesian tensor ring completion	Data-driven effective rank estimation	Currently validated mainly by controlled simulations

This comparison motivates the proposed Ra-BTRC framework, which focuses on rank-adaptive Bayesian TR completion for sparse 5D REM construction over time, frequency, and 3D space.

3. System Model and Problem Formulation

Consider a typical urban LAE network scenario, as illustrated in Figure 1. UAVs perform tasks such as cruising, inspection, or communication relay in low-altitude airspace while carrying radio-frequency sensing modules to measure the spectrum over target bands. The collected observations are then transmitted to the ground control center in real time or near real time via air-to-ground links. On the ground side, sparse measurements from UAVs

are fused to construct a complete REM, which supports low-altitude spectrum cognition and downstream tasks. The core objective considered in this paper is to recover the global spectrum situation from limited discrete measurements under constrained sensing budgets. Specifically, we aim to estimate and complete the spectrum situations across unmeasured spatial regions, frequency bands, and time slots, overcoming the challenge of sparse and non-uniformly missing observations. Based on this sensing architecture, this section formulates the low-altitude REM construction task as a sparse 5D tensor completion problem and then specifies the TR-based low-rank representation used in this work.




Figure 1. System model of low-altitude 5D REM construction. UAVs equipped with radio-frequency sensing modules collect sparse spectrum measurements during low-altitude missions and transmit them to the ground control center via air-to-ground links, where the observations are fused to complete the global spectrum situation.

3.1. Low-Altitude 5D Spectrum Tensor Model

We first discretize the space and establish a REM construction model. Let $\mathcal{A} \subset \mathbb{R}^3$ denote the physical domain of the target low-altitude airspace. After spatial discretization, $\mathcal{A}$ is represented by an $I_x \times I_y \times I_z$ voxel grid, where each voxel center is indexed by (x, y, z) with $x \in \{1, \dots, I_x\}$, $y \in \{1, \dots, I_y\}$, and $z \in \{1, \dots, I_z\}$. Similarly, the continuous observation time window is divided into I_t equal-length time slots, where each slot is sufficiently short to capture the instantaneous variation of spectrum occupancy (e.g., at the millisecond level). The target frequency band is partitioned into I_f orthogonal sub-channels or frequency bins to distinguish interference characteristics across heterogeneous service bands. Based on this discretization framework, we reasonably assume that the electromagnetic signal strength remains constant within each independent time-frequency-spatial voxel. Consequently, the global low-altitude electromagnetic spectrum situation can be modeled as a 5D spectrum tensor $\mathcal{X} \in \mathbb{R}^{I_t \times I_f \times I_x \times I_y \times I_z}$, where each entry $\mathcal{X}(t, f, x, y, z)$ represents the received spectrum power at the t -th time slot, the f -th frequency bin, and spatial location (x, y, z) , measured by the UAVs. In the proposed model, $\mathcal{X}$ is represented in normalized linear power units, because received powers from multiple emitters are physically additive in the linear domain.

According to the physical principles of electromagnetic wave propagation, $\mathcal{X}(t, f, x, y, z)$ can be approximately characterized by a limited number of dominant physical factors, such as transmitter/receiver characteristics, large-scale path loss, and shadow fading, over the 5D time–frequency–spatial domain [27]. Let $\mathbf{r} = (x, y, z)$ denote the voxel center coordinate.

The received spectrum power at time slot t , frequency bin f , and location $\mathbf{r}$ can be described by the following generic multi-transmitter model:

$$\mathcal{X}(t, f, \mathbf{r}) = \sum_{s \in \mathcal{S}} P_s(t, f) |h_s(t, f, \mathbf{r})|^2 + \sigma^2, \quad (1)$$

where $P_s(t, f)$ is the transmit power of the s -th signal source at (t, f) , σ^2 is the noise power. The channel frequency response $h_s(t, f, \mathbf{r})$ depends on space, time, and frequency, and can typically be further decomposed into large-scale and small-scale fading components, i.e.,

$$|h_s(t, f, \mathbf{r})|^2 \approx 10^{-\frac{PL_s(f, \mathbf{r})}{10}} \cdot 10^{-\frac{S_s(\mathbf{r})}{10}} \cdot |g_s(t, f, \mathbf{r})|^2, \quad (2)$$

where $PL_s(f, \mathbf{r})$ denotes the path loss associated with position $\mathbf{r}$ (e.g., geometric distance, blockage or Line-of-Sight conditions, and altitude) and frequency f ; $S_s(\mathbf{r})$ denotes the shadowing term, which is heavily dependent on environmental structures and varies continuously in space; and $g_s(t, f, \mathbf{r})$ captures the fast fading component (varying with t and f) induced by the Doppler effect and delay spread. Therefore, $\mathcal{X}$ is generated by the joint effects of $\mathbf{r}$, t , and f . Adjacent spatial voxels often exhibit similar $PL_s(f, \mathbf{r})$ and $S_s(\mathbf{r})$, neighboring frequency bins tend to share similar frequency responses, and both traffic occupancy and channel conditions are usually temporally correlated over short time slots. As a result, the spectrum tensor exhibits strong correlation and low-rank structure across the spatial, temporal, and frequency domains.

3.2. Sparse UAV Observation Model

Constrained by UAV trajectory coverage, scanning bandwidth, and sensing budgets, only partial received spectrum power observations can be obtained over the time-frequency-spatial voxel. Let $\mathcal{Y} \in \mathbb{R}^{I_t \times I_f \times I_x \times I_y \times I_z}$ denote the observed tensor, and define the observation set as $\Omega \subseteq [I_t] \times [I_f] \times [I_x] \times [I_y] \times [I_z]$, where $(t, f, x, y, z) \in \Omega$ indicates that a valid observation is available at the corresponding voxel point. We further define a binary mask tensor $\mathcal{M} \in \{0, 1\}^{I_t \times I_f \times I_x \times I_y \times I_z}$ as:

$$\mathcal{M}(t, f, x, y, z) = \begin{cases} 1, & (t, f, x, y, z) \in \Omega, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Then, the observation model can be formulated as

$$\mathcal{Y} = P_{\Omega}(\mathcal{X}) + \mathcal{E}, \quad (4)$$

where $\mathcal{E}$ represents the observation noise, which is modeled as zero-mean additive white Gaussian noise (AWGN); i.e., $\mathcal{E} \sim \mathcal{N}(0, \sigma^2)$. The Gaussian likelihood is imposed on the normalized linear-power observations. This assumption provides a tractable approximation for receiver thermal noise, sensing error, and residual modeling mismatch after normalization. More complicated correlated or heavy-tailed measurement errors can be incorporated by replacing the Gaussian likelihood. $P_{\Omega}(\cdot)$ denotes the sampling projection operator that retains only the entries within the observation set Ω , defined as $P_{\Omega}(\mathcal{X}) = \mathcal{M} \odot \mathcal{X}$, with $\odot$ representing the Hadamard product.

3.3. TR-Based Low-Rank Completion Formulation

Under the above model, low-altitude REM construction is formulated as a high-dimensional tensor completion problem. Specifically, given the partial observations $\mathcal{Y}$ over the index set Ω , the goal is to estimate the complete spectrum tensor $\mathcal{X}$ that matches

the observed entries while providing accurate predictions for unobserved ones [16]. The optimization objective can be formulated as:

$$\begin{aligned} \min_{\mathcal{X}} \quad & \|P_{\Omega}(\mathcal{X}) - \mathcal{Y}\|_F^2 \\ \text{s.t.} \quad & \mathcal{X} \in \mathcal{C}_{\text{low-rank}}, \end{aligned} \quad (5)$$

where $\mathcal{C}_{\text{low-rank}}$ denotes a generic low-rank tensor model class, whose concrete form depends on the adopted tensor decomposition. Therefore, a specific low-rank decomposition format is commonly employed to parameterize the low-rank structure into optimizable lower-order factors. Classical CP, Tucker, and TT decompositions can all be used to parameterize low-rank tensor completion models, but they exhibit different limitations for the low-altitude 5D spectrum tensor considered here. CP decomposition provides a compact factorization, but determining an appropriate CP rank is generally difficult and the required rank may become large in complex coupled scenarios [18]. Tucker decomposition can capture multi-linear correlations, but the size of its core tensor grows rapidly with the tensor order, leading to increased parameter count and computational overhead [18]. TT decomposition alleviates the curse of dimensionality through an open-chain structure, but the direct correlation between the first and last modes is not explicitly modeled [28]. In contrast, TR connects the boundary cores into a closed loop, providing a more balanced rank distribution and stronger flexibility for modeling coupled high-order structures.

To balance high-order representational capacity and scalable parameter complexity, this study adopts TR decomposition to parameterize the low-rankness of the spectrum tensor. By connecting the first and last core tensors of the TT format to form a closed loop, TR decomposition not only inherits the linear complexity advantage of TT (thus overcoming the curse of dimensionality) but also significantly enhances the representation of cyclic correlations in high-order tensors through its inherent rotational symmetry. This property is well aligned with the tightly coupled physical characteristics of the 5D time-frequency-spatial environment. Accordingly, this work fixes TR as the base tensor model and focuses on effective rank estimation within TR-based completion.

Based on TR decomposition theory, the 5D tensor $\mathcal{X}$ can be represented as the trace product of a sequence of third-order core factors $\{\mathcal{G}^{(n)}\}_{n=1}^5$; i.e.,

$$\mathcal{X}(t, f, x, y, z) = \text{Trace}\left(\mathcal{G}^{(1)}(:, t, :) \mathcal{G}^{(2)}(:, f, :) \mathcal{G}^{(3)}(:, x, :) \mathcal{G}^{(4)}(:, y, :) \mathcal{G}^{(5)}(:, z, :)\right), \quad (6)$$

where $\mathcal{G}^{(n)} \in \mathbb{R}^{R_n \times I_n \times R_{n+1}}$ ($n = 1, \dots, 5$) denotes the core factor corresponding to the n -th mode. The TR rank is defined as $\mathbf{R} = [R_1, \dots, R_5]$, which satisfies the circular constraint $R_6 = R_1$.

Substituting this formulation, the optimization problem (5) can be rewritten as

$$\begin{aligned} \min_{\{\mathcal{G}^{(n)}\}_{n=1}^5} \quad & \left\| P_{\Omega} \left(\text{Trace} \left(\prod_{n=1}^5 \mathcal{G}^{(n)} \right) \right) - \mathcal{Y} \right\|_F^2 + \lambda \cdot \mathcal{R}(\{\mathcal{G}^{(n)}\}) \\ \text{s.t.} \quad & \text{TR rank} = \mathbf{R}. \end{aligned} \quad (7)$$

This formulation corresponds to a deterministic TR completion model with a prescribed rank vector $\mathbf{R}$, where the first term enforces consistency between the reconstructed REM and the sampled UAV measurements at the sparse observation entries Ω . The second term serves as a regularizer, where $\lambda > 0$ is the balancing coefficient used to control the scale of the core factors, prevent overfitting, and implicitly promote a low-rank structure.

In practice, existing deterministic optimization algorithms for TR completion, such as alternating least squares (ALS) and the alternating direction method of multipliers (ADMM),

encounter a critical bottleneck: the TR rank vector $\mathbf{R}$ must be specified in advance [22,23,29]. However, in complex low-altitude electromagnetic environments, such prior knowledge is generally unavailable. The intrinsic tensor rank of the spectrum situation is jointly determined by the number of active transmitters, propagation complexity, traffic dynamics, and sensing sparsity, and may vary across time, frequency, and space. If the prescribed rank is too small, the model capacity is insufficient and important spectrum details may be over-smoothed; if it is too large, the model may fit noise and introduce spurious structures. Repeated trial-and-error searches for suitable ranks further incur substantial computational cost and are incompatible with the adaptive requirements of 6G low-altitude spectrum cognition. Therefore, adaptive TR rank estimation is necessary to adjust model complexity according to both environmental structure and observation sparsity, which motivates the Bayesian effective rank estimation mechanism developed in the next section.

4. Rank-Adaptive Bayesian Tensor Ring Completion

To address the aforementioned challenges, this section presents the proposed Ra-BTRC algorithm. Unlike deterministic TR completion methods that treat the TR rank as a fixed hyperparameter, Ra-BTRC starts from a moderately over-estimated TR rank and infers a compact effective rank through posterior shrinkage. Specifically, the method does not directly infer a discrete rank variable; instead, it introduces continuous sparsity hyperparameters associated with TR rank interfaces. These hyperparameters suppress unsupported rank components during variational Bayesian inference, thereby converting posterior shrinkage into data-driven effective rank estimation.

The key distinction from generic variational Bayesian tensor decomposition lies in the ring-structured coupling of the shrinkage prior. In many Bayesian tensor models, shrinkage is imposed on individual factors or components within a specific decomposition format. In contrast, a TR rank component is an interface shared by two adjacent core tensors. Therefore, the proposed prior ties each rank-interface hyperparameter to the corresponding slices of two neighboring TR cores, so that redundant components are pruned in a structurally consistent manner across the ring. This design is particularly important for 5D REM reconstruction, where model complexity must adapt to sparse observations while preserving coupled time–frequency–spatial correlations.

4.1. Probabilistic Generative Model

Based on the sparse observation model, we assume that each element of the observation noise $\mathcal{E}$ is independently drawn from a zero-mean Gaussian distribution. Let $\tau = \frac{1}{\sigma^2}$ denote the noise precision. Given the 5D low-altitude spectrum tensor $\mathcal{X}$ parameterized by its TR core factors $\{\mathcal{G}^{(n)}\}_{n=1}^5$, the likelihood function over the observed tensor $\mathcal{Y}$ is formulated as

$$p(\mathcal{Y}|\{\mathcal{G}^{(n)}\}_{n=1}^5, \tau) = \prod_{(i_1, \dots, i_5) \in \Omega} \mathcal{N}(\mathcal{Y}_{i_1, \dots, i_5} | \mathcal{X}_{i_1, \dots, i_5}, \tau^{-1}). \quad (8)$$

This likelihood characterizes the probability density of acquiring the specific UAV observations given the underlying true spectrum tensor and noise level. The learnable noise precision τ allows the model to adapt the residual uncertainty level from data. A small posterior precision corresponds to larger unexplained measurement uncertainty and therefore prevents the model from overfitting noisy sparse observations, whereas a large posterior precision indicates that the observed entries can be fitted with stronger confidence. In this sense, τ provides an uncertainty control mechanism for sparse UAV-borne spectrum measurements. Since the noise level is typically unknown in practical

low-altitude sensing tasks due to sensor conditions and environmental interference, we place a conjugate Gamma prior on τ to enable adaptive noise learning

$$p(\tau) = \text{Ga}(\tau|a, b) = \frac{b^a}{\Gamma(a)} \tau^{a-1} e^{-b\tau}, \quad (9)$$

where a and b are hyperparameters, typically set to small values to yield a weakly informative prior, so that posterior inference is mainly driven by the observed data.

To enable TR rank adaptation, we impose sparsity-inducing priors on the core factors at the level of TR rank interfaces rather than treating each core entry independently. In the circular topology of TR decomposition, adjacent core factors are connected via shared rank dimensions. Specifically, the third mode of $\mathcal{G}^{(n-1)}$ and the first mode of $\mathcal{G}^{(n)}$ are coupled through a common rank index. Therefore, when a rank component is deemed redundant, it should not vanish in only one core tensor; instead, it should be jointly suppressed across the corresponding adjacent cores. In other words, if the r -th rank component of $\mathcal{G}^{(n)}$ is heavily penalized along its first mode, the corresponding horizontal slice $\mathcal{G}^{(n)}(r, :, :)$ will be driven toward zero. Simultaneously, dictated by the shared rank index of the ring topology, this suppression forces the corresponding frontal slice $\mathcal{G}^{(n-1)}(:, :, r)$ along the third mode of $\mathcal{G}^{(n-1)}$ to shrink jointly. This coupled shrinkage mechanism effectively achieves the structural and adaptive pruning of redundant rank components.

Specifically, for the n -th core factor $\mathcal{G}^{(n)}$, we assume that each entry $g_{i,j,k}^{(n)}$ is independently drawn from a zero-mean Gaussian distribution, where $i \in [1, R_n]$, $j \in [1, I_n]$, and $k \in [1, R_{n+1}]$. The precision of this Gaussian distribution is not a constant, but jointly controlled by two independent hyperparameters, $\lambda_i^{(n)}$ and $\lambda_k^{(n+1)}$, formulated as:

$$p(g_{i,j,k}^{(n)} | \lambda_i^{(n)}, \lambda_k^{(n+1)}) = \mathcal{N}(g_{i,j,k}^{(n)} | 0, (\lambda_i^{(n)} \times \lambda_k^{(n+1)})^{-1}), \quad (10)$$

where $\lambda_i^{(n)}$ is the hyperparameter controlling the i -th component of the left interface rank R_n (i.e., corresponding to the horizontal slice of $\mathcal{G}^{(n)}$), whereas $\lambda_k^{(n+1)}$ controls the k -th component of the right interface rank R_{n+1} (i.e., corresponding to the frontal slice of $\mathcal{G}^{(n)}$). Due to the cyclic nature of the TR structure, we naturally have $\lambda_k^{(6)} \equiv \lambda_k^{(1)}$ when $n = 5$.

Let the hyperparameter vector associated with the n -th rank interface be denoted as $\lambda^{(n)} = [\lambda_1^{(n)}, \dots, \lambda_{R_n}^{(n)}]$. To facilitate the Bayesian inference of $\lambda^{(n)}$, we assign it a conjugate Gamma prior distribution

$$p(\lambda^{(n)} | \mathbf{c}_n, \mathbf{d}_n) = \prod_{r_n=1}^{R_n} \text{Ga}(\lambda_{r_n}^{(n)} | c_n^{r_n}, d_n^{r_n}), \quad (11)$$

where $\mathbf{c}_n$ and $\mathbf{d}_n$ are set to infinitesimally small values to construct a non-informative prior. This Gaussian–Gamma hierarchical prior induces the marginal distribution of the core factors with a heavy-tailed property [30]. Consequently, it achieves a delicate balance, enforcing strong shrinkage on the majority of redundant components while effectively preserving a minority of critical, informative rank components.

The key mechanism of this prior is that, as the hyperparameter $\lambda_i^{(n)}$ approaches infinity ($\lambda_i^{(n)} \rightarrow \infty$) during Bayesian inference, the variance term $(\lambda_i^{(n)} \times \lambda_k^{(n+1)})^{-1}$ tends to zero. As a result, all core entries associated with index i are strongly shrunk toward zero, including the i -th horizontal slice $\mathcal{G}^{(n)}(i, :, :)$ of $\mathcal{G}^{(n)}$ and the i -th lateral slice $\mathcal{G}^{(n-1)}(:, :, i)$ of the previous core factor $\mathcal{G}^{(n-1)}$. Through this collaborative sparsity effect, the model automatically identifies and eliminates redundant rank components that contribute little

to spectrum tensor construction, thereby achieving structurally consistent pruning in the TR representation.

Accordingly, after posterior inference, the effective TR rank at the n -th interface can be estimated by counting the active rank components whose adjacent-core posterior energy remains above a small threshold ϵ :

$$\hat{R}_n = \left| \left\{ r : \sum_{j=1}^{I_n} \sum_{s=1}^{R_{n+1}} \mathbb{E}[(g_{r,j,s}^{(n)})^2] + \sum_{j=1}^{I_{n-1}} \sum_{u=1}^{R_{n-1}} \mathbb{E}[(g_{u,j,r}^{(n-1)})^2] > \epsilon \right\} \right|. \quad (12)$$

A rank component is regarded as inactive when the posterior energy of its two adjacent core slices falls below ϵ . This criterion makes the rank pruning mechanism explicit and converts posterior shrinkage into an effective TR rank estimate.

In summary, let $\Theta = \{\{\mathcal{G}^{(n)}\}_{n=1}^5, \{\lambda^{(n)}\}_{n=1}^5, \tau\}$ denote the set of all latent variables. The joint distribution of the observed spectrum tensor $\mathcal{Y}$ and Θ can be formulated as

$$p(\mathcal{Y}, \Theta) = p(\mathcal{Y} | \{\mathcal{G}^{(n)}\}_{n=1}^5, \tau) p(\tau | a, b) \prod_{n=1}^5 p(\mathcal{G}^{(n)} | \lambda^{(n)}, \lambda^{(n+1)}) p(\lambda^{(n)} | \mathbf{c}_n, \mathbf{d}_n). \quad (13)$$

The objective is to infer the posterior distribution $p(\Theta | \mathcal{Y})$ of the latent variables given the observations $\mathcal{Y}$. According to Bayes' theorem, this is expressed as:

$$p(\Theta | \mathcal{Y}) = \frac{p(\mathcal{Y}, \Theta)}{\int p(\mathcal{Y}, \Theta) d\Theta}. \quad (14)$$

Subsequently, the posterior expectations of the core factors are used to construct the global 5D REM $\mathcal{X}$, while the effective TR ranks characterize the compact posterior supported TR structure.

However, exact posterior inference is analytically intractable. This is because it requires integrating over the latent variable set Θ , while the TR model involves nonlinear multiplicative coupling among the core factors, and the sparsity-inducing priors introduce additional hierarchical dependencies, making the integral intractable in closed form. To obtain a computationally feasible approximation, we adopt variational Bayesian inference. Unlike deterministic TR optimization, which directly estimates a single set of core factors, variational Bayesian inference approximates posterior distributions over the core factors, sparsity hyperparameters, and noise precision. This enables the reconstructed REM to be obtained from posterior expectations while model uncertainty and rank relevance are updated jointly during inference. Specifically, an approximate posterior $q(\Theta)$ with a factorized form is introduced to approximate the true posterior $p(\Theta | \mathcal{Y})$. The approximation is obtained by maximizing the evidence lower bound (ELBO), or equivalently, by minimizing the Kullback–Leibler divergence $\text{KL}(q(\Theta) || p(\Theta | \mathcal{Y}))$, which leads to tractable iterative updates for the core factors and hyperparameters.

4.2. Variational Bayesian Inference with Ring-Structured Rank Shrinkage

The optimization principle follows the standard coordinate-ascent variational Bayesian framework, but the resulting updates are specialized to the proposed TR rank-interface shrinkage model. The key model-specific feature is that each sparsity hyperparameter is associated with a TR rank interface and is shared by two adjacent core tensors. Therefore, the posterior updates do not only shrink individual entries; they jointly regulate adjacent TR core slices and convert posterior shrinkage into structurally consistent effective rank pruning.

In variational inference, we introduce a family of tractable distributions $q(\Theta)$ to approximate the true posterior $p(\Theta \mid \mathcal{Y})$. Using the mean-field assumption, $q(\Theta)$ is factorized as

$$q(\Theta) = q(\tau) \prod_{n=1}^5 q(\mathcal{G}^{(n)}) \prod_{n=1}^5 q(\lambda^{(n)}). \quad (15)$$

The goal of variational inference is to minimize the KL divergence between $q(\Theta)$ and $p(\Theta \mid \mathcal{Y})$; i.e.,

$$\begin{aligned} q^*(\Theta) &= \arg \min_q \text{KL}(q(\Theta) \parallel p(\Theta \mid \mathcal{Y})) \\ &= \arg \min_q \int q(\Theta) \ln \frac{q(\Theta)}{p(\Theta \mid \mathcal{Y})} d\Theta \\ &= \arg \min_q \ln p(\mathcal{Y}) - \int q(\Theta) \ln \frac{p(\mathcal{Y}, \Theta)}{q(\Theta)} d\Theta. \end{aligned} \quad (16)$$

Since the first term is constant with respect to q , minimizing the KL divergence is equivalent to maximizing the ELBO:

$$\mathcal{L}(q) = \mathbb{E}_q[\ln p(\mathcal{Y}, \Theta)] - \mathbb{E}_q[\ln q(\Theta)]. \quad (17)$$

Under the mean-field assumption, the variational factors can be updated iteratively by coordinate ascent; i.e., one factor is updated while the others are fixed, which guarantees a non-decreasing ELBO. For any variable block $v \in \Theta$, the optimal variational factor satisfies

$$\ln q^*(v) = \mathbb{E}_{q(\Theta \setminus v)}[\ln p(\mathcal{Y}, \Theta)] + \text{const}, \quad (18)$$

where const collects all terms independent of v . Owing to the conjugate priors introduced in Section 3.1, the variational factors remain in standard distribution families, the posterior of each core tensor is Gaussian, while those of the precision hyperparameters and noise precision are Gamma distributions.

This variational procedure provides a tractable inference scheme for effective rank estimation, but it should not be interpreted as a strict guarantee of identifying the unique ground-truth TR rank. For sparse high-order spectrum tensor completion, exact identifiability generally depends on the sampling pattern, tensor incoherence, noise level, and intrinsic rank structure. Moreover, because the TR representation is non-convex and not unique up to gauge transformations, the coordinate-ascent variational updates are designed to monotonically improve the ELBO when each variational factor is updated exactly, but they do not guarantee global optimality of the inferred rank. Nevertheless, the coupled Gamma shrinkage prior provides an automatic complexity control mechanism: rank components that are weakly supported by the observations are assigned large posterior precision and are jointly suppressed in adjacent core tensors, yielding a compact posterior-supported TR structure.

The mean-field approximation also ignores posterior correlations among different core tensors and hyperparameters, which may make the inference sensitive to initialization when observations are extremely sparse. In our implementation, this risk is mitigated by using moderately over-estimated initial TR ranks, monitoring the ELBO and fitting error during iterations, and relying on the Gamma shrinkage priors to suppress unsupported rank components. Nevertheless, avoiding local optima under extremely low sampling rates remains an open issue for future work, where more advanced posterior approximations or multi-initialization strategies may be investigated.

For the n -th core factor $\mathcal{G}^{(n)}$, let its mode-2 lateral slice be denoted by $\mathbf{G}_j^{(n)} = \mathcal{G}^{(n)}(:, j, :) \in \mathbb{R}^{R_n \times R_{n+1}}$, $j = 1, \dots, I_n$, and define $\mathbf{g}_j^{(n)} = \text{vec}(\mathbf{G}_j^{(n)})$. When all other core factors and hyperparameters are fixed, the observation likelihood is linear Gaussian with respect to $\mathbf{g}_j^{(n)}$. Therefore, its variational posterior is multivariate Gaussian

$$q(\mathbf{g}_j^{(n)}) = \mathcal{N}(\mathbf{g}_j^{(n)} | \boldsymbol{\mu}_j^{(n)}, \boldsymbol{\Sigma}_j^{(n)}). \quad (19)$$

To derive $\boldsymbol{\mu}_j^{(n)}$ and $\boldsymbol{\Sigma}_j^{(n)}$, define the observation subset $\Omega_j^{(n)} = \{(i_1, \dots, i_5) \in \Omega \mid i_n = j\}$, which contains all observed entries whose n -th mode index equals j . For each observation index $\mathbf{i} = (i_1, \dots, i_5) \in \Omega_j^{(n)}$, contracting all core factors except $\mathcal{G}^{(n)}$ along the TR structure yields an effective design vector $\mathbf{h}_{\mathbf{i}}^{(n)}$, such that

$$\mathcal{X}_{\mathbf{i}} = (\mathbf{h}_{\mathbf{i}}^{(n)})^\top \mathbf{g}_j^{(n)}. \quad (20)$$

This leads to the standard Bayesian linear regression updates

$$(\boldsymbol{\Sigma}_j^{(n)})^{-1} = \mathbb{E}[\tau] \sum_{\mathbf{i} \in \Omega_j^{(n)}} \mathbb{E}[\mathbf{h}_{\mathbf{i}}^{(n)} (\mathbf{h}_{\mathbf{i}}^{(n)})^\top] + (\mathbb{E}[\boldsymbol{\Lambda}^{(n+1)}] \otimes \mathbb{E}[\boldsymbol{\Lambda}^{(n)}]), \quad (21)$$

$$\boldsymbol{\mu}_j^{(n)} = \mathbb{E}[\tau] \boldsymbol{\Sigma}_j^{(n)} \sum_{\mathbf{i} \in \Omega_j^{(n)}} y_{\mathbf{i}} \mathbb{E}[\mathbf{h}_{\mathbf{i}}^{(n)}], \quad (22)$$

where $\boldsymbol{\Lambda}^{(n)} = \text{diag}(\boldsymbol{\lambda}^{(n)})$. The precision term $\mathbb{E}[\boldsymbol{\Lambda}^{(n+1)}] \otimes \mathbb{E}[\boldsymbol{\Lambda}^{(n)}]$ is specific to the proposed ring-structured prior: it couples the left and right rank-interface precisions of the n -th TR core. Therefore, the posterior covariance of each core slice is controlled jointly by two neighboring rank interfaces, rather than by an independent element-wise regularizer. This structure is essential for transferring rank shrinkage across adjacent TR cores. Although the above expressions appear complex, the TR structure allows them to be implemented via a sequence of small-scale matrix operations, ensuring computational efficiency.

For the sparsity hyperparameters $\lambda_r^{(n)}$, each $\lambda_r^{(n)}$ controls the r -th component of the n -th TR rank. Owing to the conjugacy between the Gamma prior and the Gaussian precision, its variational posterior is also Gamma

$$q(\lambda_r^{(n)}) = \text{Ga}(\lambda_r^{(n)} | \tilde{c}_r^{(n)}, \tilde{d}_r^{(n)}). \quad (23)$$

Since $\lambda_r^{(n)}$ is shared by two adjacent core factors through the ring structure, its update incorporates information from both $\mathcal{G}^{(n)}$ and $\mathcal{G}^{(n-1)}$. A typical conjugate update is given by

$$\tilde{c}_r^{(n)} = c_r^{(n)} + \frac{1}{2}(I_n R_{n+1} + I_{n-1} R_{n-1}), \quad (24)$$

$$\tilde{d}_r^{(n)} = d_r^{(n)} + \frac{1}{2} \sum_{j=1}^{I_n} \sum_{s=1}^{R_{n+1}} \mathbb{E}[\lambda_s^{(n+1)}] \mathbb{E}[(g_{r,j,s}^{(n)})^2] + \frac{1}{2} \sum_{j=1}^{I_{n-1}} \sum_{u=1}^{R_{n-1}} \mathbb{E}[\lambda_u^{(n-1)}] \mathbb{E}[(g_{u,j,r}^{(n-1)})^2]. \quad (25)$$

When a rank component is simultaneously shrunk toward zero in the adjacent core factors, $\tilde{d}_r^{(n)}$ becomes small and $\mathbb{E}[\lambda_r^{(n)}]$ increases, which imposes stronger shrinkage on that component and eventually enables adaptive rank pruning. Different from element-wise Bayesian shrinkage, the update of $\lambda_r^{(n)}$ aggregates posterior energy from two adjacent core tensors. This design is necessary because a TR rank component is an interface shared by neighboring cores. Consequently, the posterior update of $\lambda_r^{(n)}$ determines whether the

corresponding interface component should remain active or be jointly suppressed in both adjacent cores, which is the main mechanism that differentiates Ra-BTRC from generic variational Bayesian tensor factorization.

The variational posterior of the noise precision τ also follows a Gamma distribution:

$$q(\tau) = \text{Ga}(\tau \mid \tilde{a}, \tilde{b}), \quad (26)$$

where

$$\tilde{a} = a_0 + \frac{|\Omega|}{2}, \quad (27)$$

$$\tilde{b} = b_0 + \frac{1}{2} \mathbb{E}[\|\mathcal{P}_\Omega(\mathcal{Y} - \mathcal{X})\|_F^2]. \quad (28)$$

Here, $|\Omega|$ denotes the number of observed samples. The expected squared error term can be expanded over the observation set as

$$\mathbb{E}[\|\mathcal{P}_\Omega(\mathcal{Y} - \mathcal{X})\|_F^2] = \sum_{i \in \Omega} (\mathcal{Y}_i^2 - 2\mathcal{Y}_i \mathbb{E}[\mathcal{X}_i] + \mathbb{E}[\mathcal{X}_i^2]), \quad (29)$$

where $\mathbb{E}[\mathcal{X}_i]$ and $\mathbb{E}[\mathcal{X}_i^2]$ can be computed from the posterior means and second-order moments of the core factors through TR contraction. The parameter $\tilde{b}$ therefore reflects the expected construction error of the current model. As the residual decreases, the estimated noise precision increases, corresponding to a smaller noise variance, which allows the model to balance underfitting and overfitting adaptively.

After obtaining the variational posterior $q(\Theta)$, the complete spectrum tensor is constructed using the posterior mean

$$\hat{\mathcal{X}} = \mathbb{E}_q[\text{TR}(\mathcal{G}^{(1)}, \dots, \mathcal{G}^{(5)})]. \quad (30)$$

In summary, the standard variational Bayesian framework provides the optimization backbone, while the distinctive part of Ra-BTRC lies in the ring-structured posterior updates induced by the coupled Gamma shrinkage prior. The core-factor update contains rank-interface precision coupling, the sparsity hyperparameter update aggregates posterior energy from adjacent cores, and the noise-precision update adapts the residual uncertainty level. Together, these model-specific updates enable data-driven effective rank estimation within the TR representation.

4.3. Computational Complexity Analysis

In the TR model, storing the complete 5D spectrum tensor only requires the core factors $\{\mathcal{G}^{(n)}\}_{n=1}^5$ and a small number of hyperparameters. The number of core parameters is $\sum_{n=1}^5 R_n I_n R_{n+1}$. Therefore, when the TR ranks are bounded, the storage complexity grows approximately linearly with the tensor order.

The main computational cost of the proposed algorithm lies in updating the posterior means and covariance parameters of the core factors $\{\mathcal{G}^{(n)}\}_{n=1}^5$ during variational inference. Let the current TR rank be R , the number of observed samples be $|\Omega|$ and, for simplicity, assume all modes of the 5D tensor have equal size I . In each variational update, to avoid the high cost of direct inversion (i.e., $O(I^3 R^6)$), the preconditioned conjugate gradient (PCG) method is used for iterative solving, reducing the cost of each linear system solve to $O(|\Omega| R^2 + I R^2)$. Let L_{cg} denote the average number of PCG iterations. Then, the complexity of updating a single core is approximately $O(L_{\text{cg}}(|\Omega| + I) R^2)$. Overall, the time complexity of one outer iteration is approximately $O(d \cdot L_{\text{CG}} \cdot |\Omega| R^2)$, where $d = 5$ is the tensor order.

This result shows that the computational complexity is mainly determined by the rank R and the number of observations $|\Omega|$, and scales linearly with the tensor order d and mode size I , thereby avoiding the exponential complexity growth typical of conventional high-dimensional tensor methods. In addition, owing to the proposed Rank-Adaptive pruning mechanism, the model can automatically drive R toward a compact intrinsic rank, which significantly reduces the practical runtime while maintaining completion accuracy. This makes the proposed method suitable for large-scale low-altitude sparse REM construction.

5. Numerical Experiments

In this section, extensive simulations are conducted to evaluate the effectiveness of the proposed Ra-BTRC method for low-altitude 5D REM construction under sparse observations. We first describe the simulation setup, including data generation, evaluation metrics, and baseline methods. We then compare the proposed method with representative TR-based approaches in terms of reconstruction accuracy, robustness under extremely low sampling rates, convergence behavior, and computational cost.

5.1. Simulation Setup

5.1.1. Data Generation

The simulation region is defined as a continuous 3D space and discretized with a horizontal spatial resolution of 1 m. A total of D dynamic emitters are deployed in the scene. Each emitter moves along a randomly initialized rectangular trajectory on the ground at a constant speed v , emulating realistic mobile communication services or airborne interference sources. The spectral characteristic of each emitter is described by a power spectral density (PSD) vector $\mathbf{c}_r(f)$. In this work, the PSD of the center frequency and its adjacent 1-3 frequency bins is generated by random superposition of sinc functions to mimic broadband occupied-band signatures [15]. Signal propagation follows a typical air-to-ground links channel model [31], and the path loss is defined as

$$L_{\text{dB}}(\mathbf{r}, t) = 32.4 + 20 \log_{10}(f_c) + 10 \cdot (A \cdot z^B) \cdot \log_{10}(d_{3\text{D}}(\mathbf{r}, t)) + \chi_\sigma(\mathbf{r}), \quad (31)$$

where f_c is the carrier frequency (GHz), $d_{3\text{D}}$ is the 3D Euclidean distance between the transmitter and receiver, and z denotes the UAV altitude. The parameters A and B are used to adjust the path-loss exponent at different altitude layers, thereby reflecting the non-uniform attenuation characteristics of low-altitude environments. The term χ_σ denotes zero-mean log-normal shadowing whose spatial correlation follows an exponential decay model,

$$\mathbb{E}[\xi(\mathbf{r}) \xi(\mathbf{r}')] = \eta \exp\left(-\frac{\|\mathbf{r} - \mathbf{r}'\|}{d_c}\right), \quad (32)$$

where η controls the shadowing strength and d_c is the correlation distance. Larger η and smaller d_c correspond to stronger shadowing and weaker spatial correlation, and can thus generate low-altitude electromagnetic environments of different complexity. The relevant parameter settings are listed in Table 2.

Finally, the global 5D spectrum tensor $\mathcal{X}(t, f, x, y, z)$ is generated by superposing the linear received-power fields of all emitters, $\tilde{S}_r(\mathbf{r}, t) = 10^{-L_{\text{dB}}(\mathbf{r}, t)/10}$, with their corresponding PSD vectors $\mathbf{c}_r(f)$:

$$\mathcal{X}(t, f, x, y, z) = \sum_{r=1}^D \tilde{S}_r(\mathbf{r}, t) \cdot \mathbf{c}_r(f). \quad (33)$$

The received power contribution of each emitter is first generated in the linear domain and then superposed across emitters. This avoids the physically invalid operation of

directly adding powers in dBm. For visualization, the reconstructed linear-power field can be converted to dBm by logarithmic scaling.

Table 2. Simulation parameter settings.

Parameter	Value
Length X	21
Width Y	21
Height Z	10
Frequency bin F	8
Time slot T	10
Speed v (grid/time slot)	1
Carrier frequency f_c (GHz)	2.4
Emitter number D	1, 3, 5
Path-loss exponent A	2
Path-loss exponent B	0
Shadowing strength η (dB)	8
Correlation distance d_c (m)	30

This generation process provides a controlled low-altitude REM benchmark with physically interpretable correlations and dynamic variations across the time, frequency, and 3D spatial domains. Specifically, the altitude-dependent path-loss term captures large-scale distance and height effects, while the spatially correlated log-normal shadowing term approximates the continuity and local variability of obstruction-induced attenuation. Although the simulator is not intended to be a full site-specific urban channel model, many practical imperfections in real UAV measurements can be interpreted in the tensor completion framework as observation perturbations or structured missing mechanisms. For example, hardware noise and calibration errors perturb the observed amplitudes, synchronization offsets may introduce local temporal mismatch, irregular trajectories and no-fly regions lead to non-uniform or block-structured sampling masks, and blockage or multipath may increase local non-stationarity and effective tensor rank. Ra-BTRC can partially mitigate these effects through Bayesian noise-precision learning, sparsity-induced rank regularization, and joint time–frequency–3D correlation modeling. Nevertheless, severe systematic bias, large-scale calibration drift, or highly structured non-random missing patterns may still degrade reconstruction accuracy. Future extensions may incorporate correlated or heteroscedastic likelihoods, calibration-aware preprocessing, trajectory-aware sampling models, environment-aware priors, and validation using real UAV spectrum measurements or public low-altitude datasets.

5.1.2. Metric

To evaluate the reconstruction performance under different UAV sensing constraints, three sampling patterns are considered in this work, as shown in Figure 2.

- **Random Sampling:** the observation set Ω is generated by random sampling over the global 5D voxel.
- **Trajectory-constrained Sampling:** the generation of the observation set Ω follows certain rules, simulating the fact that UAV measurements are concentrated on continuous flight trajectories and have limited local perception coverage, reflecting the reality that UAV-borne spectrum sensing is constrained by flight continuity and spatial accessibility.
- **Block Missing Sampling:** several spatial regions are assumed to be unavailable for sensing due to no-fly zones, building blockage, or task constraints, and the observation set Ω are randomly selected only from the remaining accessible regions.

These three patterns represent progressively more structured and non-uniform missing mechanisms compared with ideal random sampling. In this work, the sampling mask is treated as a given observation support rather than an explicit probabilistic missingness model. Therefore, the proposed completion framework can operate on arbitrary observation masks, but severe non-random missing patterns with large disconnected unobserved regions may still introduce systematic reconstruction bias.

For all sampling patterns, the sampling rate ρ is defined as the ratio between the number of observed samples and the total number of entries in the 5D spectrum tensor,

$$\rho = \frac{|\Omega|}{I_x \times I_y \times I_z \times I_t \times I_f} \times 100\%, \quad (34)$$

where $|\cdot|$ denotes set cardinality.




Figure 2. Visualization of different sampling patterns on a representative X-Y slice at a fixed time, frequency, and altitude. The first panel shows the complete ground-truth slice. The second panel shows random sampling, where observations are randomly selected from the slice. The third panel shows trajectory-constrained sampling, where observations are concentrated along continuous UAV flight paths. The fourth panel shows block missing sampling, where white regions denote unavailable areas that cannot be sampled.

To evaluate the construction error at each time slot, we compute the normalized mean square error (NMSE) as

$$\text{NMSE}(t) \text{ (dB)} = 10 \log_{10} \left(\frac{\|\hat{\mathcal{X}}_t - \mathcal{X}_t\|_F^2}{\|\mathcal{X}_t\|_F^2} \right), \quad (35)$$

where $\mathcal{X}_t \in \mathbb{R}^{I_f \times I_x \times I_y \times I_z}$ denotes the true spectrum tensor slice at time slot t , $\hat{\mathcal{X}}_t$ is the corresponding completion, and $\|\cdot\|_F$ denotes the Frobenius norm.

Furthermore, to measure the average completion error over the whole observation interval, we use the mean normalized mean square error (MMSE) as

$$\text{MMSE (dB)} = \frac{1}{I_t} \sum_{t=1}^{I_t} 10 \log_{10} \left(\frac{\|\hat{\mathcal{X}}_t - \mathcal{X}_t\|_F^2}{\|\mathcal{X}_t\|_F^2} \right). \quad (36)$$

This metric provides an overall measure of completion accuracy and stability in dynamic spectrum environments.

5.1.3. Baselines

As discussed in Sections 2.2 and 3, CP, Tucker, and TT-based tensor completion methods provide important alternatives for REM reconstruction, but they involve different tensor parameterizations and therefore address a broader model selection problem. Since the main objective of this work is to evaluate whether Bayesian learning improves sparse REM completion within the TR representation, we focus the experimental comparison on TR-family baselines. This choice allows us to isolate the effect of rank adaptivity from the representational differences among tensor formats. Specifically, we choose TRLRF [29] as

the main baseline because it is a representative and competitive TR completion method that imposes low-rank regularization on latent TR factors and is closely related to the rank-sensitivity issue addressed in this paper.

To fairly assess the effect of rank selection on completion performance, we consider two TRLRF settings.

- **Adaptive-rank setting (TRLRF-AR):** the effective TR rank inferred by Ra-BTRC, $\hat{\mathbf{R}} = [\hat{R}_1, \dots, \hat{R}_5]$, is used as the rank configuration for TRLRF. This setting evaluates the performance of a deterministic TR baseline when a favorable rank choice is provided.
- **Random-rank setting (TRLRF-RR):** the TR rank vector is randomly generated within a predefined range to simulate the rank-mismatch condition faced by fixed-rank TR methods when prior rank knowledge is unavailable. This setting reflects the sensitivity of TR completion to manual rank specification.

Thus, TRLRF-AR and TRLRF-RR together provide two controlled TR-family references: one with a favorable rank setting and one with rank mismatch. They are used to assess whether the proposed Bayesian rank-adaptive inference improves reconstruction beyond deterministic TR completion under comparable TR representations.

In addition, all methods are implemented under the same random mask, sampling rate, and hardware environment to ensure a fair comparison.

5.2. Overall Performance Comparison

Figure 3 compares the MMSE of different methods under varying sampling rates ρ . For TRLRF-RR, 8 random TR rank configurations are tested at each sampling rate, and the shaded region in the figure indicates the corresponding performance variation, reflecting the uncertainty caused by random rank selection. It can be observed that Ra-BTRC consistently achieves the lowest MMSE across all sampling rates, with a particularly clear advantage in the low-sampling regime. This is mainly because the variational Bayesian framework, together with sparsity-inducing priors, effectively exploits latent structural information in sparse observations and adaptively regularizes the model, thereby reducing the risk of overfitting under extremely limited data. In contrast, the shaded region of TRLRF-RR becomes significantly wider as the sampling rate decreases, indicating that its performance is highly sensitive to rank choice in sparse data scenarios. Inappropriate rank settings may lead to severe performance degradation. Although TRLRF-AR performs slightly better than TRLRF-RR by using a more suitable rank configuration, it still exhibits a clear gap from Ra-BTRC.

To examine the robustness of the proposed method under heterogeneous interference, we further evaluate multi-emitter scenarios with $D = 1$, $D = 3$, and $D = 5$. The emitters are assigned different initial positions and spectral signatures, resulting in more complex spatial frequency interference patterns. Figure 3 reports the MMSE under different sampling rates. As the number of emitters increases, the reconstruction task becomes more challenging because the spectrum tensor contains more localized peaks and overlapping interference components. Nevertheless, Ra-BTRC consistently achieves lower reconstruction error than the TRLRF-based baselines, indicating that the proposed adaptive Bayesian rank learning remains effective in more heterogeneous low-altitude interference environments. Although no simulation can guarantee performance in all real urban environments, the multi-emitter results provide additional evidence that Ra-BTRC remains effective when the spectrum field contains overlapping and heterogeneous interference components.






Figure 3. Comparison of MMSE versus sampling rate for the three methods under multi-emitter scenarios: (a) $D = 1$, (b) $D = 3$, (c) $D = 5$. For TRLRF-RR, 8 random rank settings are tested at each sampling rate; the shaded band indicates the MMSE range, and TRLRF-RR (mean) denotes the average over different rank settings.

To visually evaluate the structural recovery capability under extremely sparse observations, Figure 4 presents slice-based comparisons of the true and reconstructed 5D spectrum tensors at a sampling rate of 5%. The figure includes four types of 2D slices: two X-Y slices at different altitude levels, one X-Z slice to illustrate the 3D spatial structure and altitude-dependent variation, and one T-F slice at a fixed spatial voxel to show the temporal evolution and spectral structure. The first column shows the ground-truth slices, which represent the actual distribution of the low-altitude spectrum field. The second column shows the sparse observations, where substantial missing entries can be observed across time, space, and frequency, resulting in severe structural incompleteness. The third to fifth columns present the completion results of TRLRF-RR, TRLRF-AR, and the proposed Ra-BTRC, respectively. Overall, the completion results of Ra-BTRC are visually the closest to the ground truth in both spatial distribution and dynamic variation. It can accurately recover local hotspots and fine-grained structures, indicating that the proposed method more effectively exploits high-order time–frequency–spatial correlations under low sampling rates and thus achieves higher-quality and more stable 5D spectrum map completion. In contrast, although TRLRF-RR and TRLRF-AR can recover the overall trend, they still exhibit varying degrees of distortion in local details.




Figure 4. Visual comparison of 5D spectrum tensor slices at a sampling rate of $\rho = 5\%$. From left to right, the columns show the ground truth, sparse observations, and the completion results of TRLRF-RR, TRLRF-AR, and Ra-BTRC. From top to bottom, the rows correspond to two X – Y slices at different altitude levels, one X – Z slice, and one T – F slice at a fixed spatial location.

5.3. Performance Under Different Sampling Patterns

To further investigate the influence of practical UAV sensing constraints, Figure 5 compares the MMSE of TRLRF-RR, TRLRF-AR, and Ra-BTRC under three sampling patterns at a sampling rate of $\rho = 5\%$. It can be observed that the sampling mechanism has a clear impact on reconstruction accuracy. Random sampling achieves the best overall performance because the observations are more uniformly distributed over the 5D tensor, providing more balanced information across spatial, temporal, and frequency dimensions. In contrast, trajectory-constrained sampling leads to the largest reconstruction error, since the measurements are concentrated along limited UAV flight paths and leave large unobserved regions away from the trajectories. Block missing sampling shows intermediate

difficulty, although observations are randomly distributed in the accessible regions, the unavailable blocks introduce structured spatial gaps that cannot be directly sensed.

Among the compared methods, Ra-BTRC consistently achieves the lowest MMSE under all three sampling patterns. This indicates that the proposed rank-adaptive Bayesian tensor ring framework is more robust to non-random and structured missing mechanisms than the TRLRF-based baselines. The advantage is particularly meaningful under trajectory-constrained sampling, where the observation distribution deviates most strongly from the ideal random sampling assumption. Meanwhile, the degradation under trajectory-constrained sampling also indicates that strongly biased UAV coverage remains more challenging than random sampling.




Figure 5. MMSE comparison of the three algorithms under different sampling patterns when $\rho = 5\%$ and $D = 1$. For TRLRF-RR, the TR rank is randomly set to $\mathbf{R} = [5, 5, 5, 5, 5]$.

5.4. Performance Across Time, Frequency, and Altitude

Figure 6 shows the per-step NMSE of the three methods along the time, frequency, and altitude dimensions at a sampling rate of 7%. The right vertical axis in each subfigure reports the actual sampling ratio at each step, so as to exclude the effect of sampling imbalance. As shown, the sampling rates remain approximately uniform across all dimensions; hence, the performance differences mainly reflect the completion capability of the algorithms themselves. Along the time dimension in Figure 6a, as the time step increases, the completion errors of TRLRF-RR and TRLRF-AR show clear fluctuations, indicating that fixed-rank methods are less robust to local transient changes in dynamic spectrum fields and may accumulate errors over time. In contrast, the NMSE of Ra-BTRC remains consistently low and stable, demonstrating its robustness in temporally dynamic sensing tasks. Along the frequency dimension in Figure 6b, Ra-BTRC consistently achieves the lowest completion error and remains insensitive to band variation, confirming its strong ability to capture frequency-selective attenuation patterns. Notably, Figure 6c shows that, as the altitude z increases, the NMSE of Ra-BTRC exhibits a distinct decreasing trend. This agrees with the physical characteristics of low-altitude propagation, at higher altitudes, LoS propagation becomes more dominant, shadowing is reduced, and the spatial structure of the spectrum field becomes smoother. Ra-BTRC successfully captures this property and achieves more accurate completion in high-altitude regions. By contrast, the baseline methods, especially TRLRF-RR, show increasing errors at higher altitudes, suggesting that they fail to effectively exploit the signal correlations along the altitude dimension and may suffer from overfitting under sparse observations.




Figure 6. Per-dimension NMSE profiles at a sampling rate of $\rho = 7\%$: (a) time slot vs. NMSE, (b) frequency band vs. NMSE, and (c) height vs. NMSE. In each subfigure, the left axis denotes NMSE, while the right axis shows the corresponding sampling rate to illustrate sampling balance across dimensions.

To further evaluate the local completion capability under extremely sparse observations, Figure 7 shows the tracking curves of received spectrum power versus time, frequency, and altitude at a specific spatial location (11, 11) under a sampling rate of 5%. For visualization, the reconstructed linear-power values are converted to received signal strength indicator (RSSI) values in dBm. Regions with missing observations are explicitly marked, indicating that no samples are available at those positions along the corresponding dimension. Despite the presence of large contiguous missing intervals in

the time, frequency, and altitude dimensions, the completion curves of Ra-BTRC remain highly consistent with the ground truth (black curves). This indicates that the proposed method can effectively exploit high-order correlations in the tensor structure and borrow information from other dimensions to infer completely missing entries. In contrast, both TRLRF-RR and TRLRF-AR exhibit noticeable deviations in missing regions. In particular, TRLRF-RR shows substantial fluctuations around the true values and fails to accurately capture fine-grained signal variations, especially near peaks caused by rapid signal changes. These results further demonstrate the superior ability of Ra-BTRC to recover local spectrum features under very low sampling rates.




Figure 7. Single-point reconstruction profiles at a sampling rate of $\rho = 5\%$: (a) time slot vs. RSSI, (b) frequency band vs. RSSI, and (c) height vs. RSSI. The curves compare the ground truth and completed RSSI at a fixed spatial location, with missing regions explicitly marked.

5.5. Convergence and Rank Adaptivity Analysis

To evaluate the numerical stability, uncertainty learning, and adaptive rank estimation capability of the proposed Ra-BTRC algorithm, Figure 8 shows the evolution of key optimization variables and objective-related quantities during iteration at a sampling rate of 5%. Specifically, we record the ELBO, the data fitting error (err), and the posterior statistics of the core factors $\mathcal{G}$, the sparsity hyperparameters λ , and the noise precision τ .

It can be observed that the ELBO increases monotonically with the number of iterations and gradually converges, which is consistent with the theoretical property of variational Bayesian inference. Meanwhile, the data fitting error decreases rapidly in the initial stage and then stabilizes at a low level. This indicates that the model quickly captures the main structural information from sparse observations and further refines the completion through iterative updates of the core factors and hyperparameters. The evolution of $\mathbb{E}[\lambda]$ and related objective terms reveals the intrinsic mechanism of rank adaptation. At the early stage, the values of λ change significantly, reflecting the exploration of the relevance of different rank components. As the iterations proceed, some components are driven to very large values, corresponding to near-zero variances, so that redundant rank components are automatically pruned and the model gradually converges to a stable compact rank structure.

This process is entirely data-driven and requires no manual intervention, highlighting the full Bayesian adaptivity of the proposed method. In addition, the posterior statistics of the core tensors $\mathcal{G}$ and the noise precision τ exhibit synchronized convergence behavior, further confirming the numerical stability of the algorithm in joint multi-variable inference. The stable ELBO and posterior-shrinkage trends in Figure 8 indicate that the adopted initialization leads to a stable compact rank structure in the tested sparse setting, although this empirical convergence does not eliminate the possibility of local optima in general non-convex tensor inference. These posterior evolution results show that variational Bayesian inference contributes to sparse REM reconstruction through three coupled mechanisms: posterior uncertainty learning via τ , sparsity-induced regularization via λ , and TR-specific rank pruning through adjacent-core shrinkage.




Figure 8. Convergence and posterior shrinkage of Ra-BTRC at a sampling rate of $\rho = 5\%$. The figure shows the evolution of the ELBO, data fitting error, and posterior statistics of the core factors $\mathcal{G}$, the sparsity hyperparameters λ , and the noise precision τ during the iterative process.

Beyond numerical convergence, Figure 8 helps explain the roles of the main components of Ra-BTRC. The monotonic increase of the ELBO and the decrease of the data fitting error reflect the effect of variational Bayesian inference, which updates posterior distributions of both core tensors and hyperparameters. The evolution of $\mathbb{E}[\lambda]$ reflects the shrinkage effect induced by the Gamma sparsity priors. As some precision values become large, the corresponding adjacent TR core slices are suppressed, which provides direct evidence of adaptive rank pruning. Meanwhile, the evolution of the noise precision τ indicates that the model learns observation uncertainty from data and adaptively balances data fitting and regularization. These posterior-evolution results show that Bayesian inference, sparsity-induced shrinkage, adaptive rank pruning, and noise precision learning jointly contribute to stable sparse REM reconstruction within the unified Ra-BTRC framework.

To quantitatively characterize the rank adaptation behavior, Table 3 reports the initial TR rank, final effective TR rank, number of pruned components, and MMSE under different sampling rates. The experiment is conducted under the single-emitter setting ($D = 1$) with random sampling. The pruning threshold is fixed as $\epsilon = 10^{-4}$ for all sampling rates, and the maximum number of iterations is set to 200. The effective rank is computed using the posterior energy criterion defined in Section 4.1. The number of pruned components is calculated as $\sum_{n=1}^5 (R_n^{(0)} - \hat{R}_n)$ over the five TR rank interfaces.

Table 3. Quantitative rank adaptation results under different sampling rates ($D = 1$, random sampling, $\epsilon = 10^{-4}$).

ρ	$\mathbf{R}^{(0)}$	$\hat{\mathbf{R}}$	Pruned	MMSE (dB)
1%	[10, 10, 10, 10, 10]	[2, 3, 2, 5, 4]	34/50	−3.82
3%	[10, 10, 10, 10, 10]	[1, 6, 8, 2, 4]	29/50	−8.05
5%	[10, 10, 10, 10, 10]	[6, 3, 5, 6, 3]	27/50	−12.48
7%	[10, 10, 10, 10, 10]	[3, 6, 6, 8, 7]	20/50	−16.27
10%	[10, 10, 10, 10, 10]	[7, 9, 6, 5, 5]	18/50	−18.32

This observation is consistent with the motivation of adaptive rank estimation. The same initial TR rank is compressed into different final effective ranks under different sampling rates, indicating that the model complexity is mainly driven by data support and posterior shrinkage rather than being fixed by initialization. At lower sampling rates, stronger pruning is generally observed because sparse observations provide weaker support for high-rank components. As the sampling rate increases, more rank components are retained, allowing the model to preserve finer spectrum structures. These results make the adaptive rank learning process explicit and support the role of sparsity-induced effective rank estimation. This behavior is especially relevant to practical low-altitude deployment, where sensing density, interference complexity, and propagation conditions may change over time; a fixed rank selected offline is therefore unlikely to remain optimal across different operating conditions.

5.6. Complexity Analysis

To evaluate the practical computational cost of Ra-BTRC, we compare the end-to-end runtime and the number of iterations required for convergence under the same hardware conditions. At a sampling rate of 5%, Ra-BTRC exhibits noticeably faster convergence in terms of iterations. Benefiting from the effective exploitation of latent data structure by variational Bayesian inference, Ra-BTRC reaches a stable state and achieves its best completion performance within about 100 iterations. In contrast, the deterministic optimization methods TRLRF-AR and TRLRF-RR typically require about 300 iterations to converge or terminate, indicating a higher iteration burden.

All runtime results are measured under the same hardware environment with an NVIDIA RTX 3060 GPU (NVIDIA Corporation, Santa Clara, CA, USA). In terms of runtime, Ra-BTRC requires 183.6 s, while TRLRF-AR and TRLRF-RR require 141.1 s and 67.9 s, respectively. Thus, the runtime of Ra-BTRC is about 30.1% higher than that of TRLRF-AR and about 170.4% higher than that of TRLRF-RR. Although Ra-BTRC has a higher per-iteration computational cost, this additional time investment leads to substantially improved completion accuracy. Specifically, at the extreme sampling rate of 5%, Ra-BTRC achieves an MMSE of −12.48 dB, compared with −2.25 dB for TRLRF-AR and −2.20 dB for TRLRF-RR, corresponding to an accuracy gain of more than 10 dB.

Overall, Ra-BTRC attains a significantly better accuracy and complexity trade-off in sparse low-altitude REM completion. The moderate increase in computational cost is justified by the substantial gain in completion accuracy and robustness, making the proposed method more suitable for low-altitude spectrum cognition tasks with stringent reliability requirements.

The present experiments use a moderate tensor size to allow repeated evaluations under multiple sampling rates, emitter settings, and sampling patterns. Larger-scale tensor experiments and real-world deployment with higher spatial, temporal, and frequency resolutions will be further investigated in future work.

6. Discussion

The simulation results demonstrate that the proposed Ra-BTRC framework improves sparse 5D REM reconstruction not only in terms of numerical accuracy, but also in terms of deployment robustness and model adaptability. This section further discusses the implications of these findings from two perspectives: deployment-oriented insights for low-altitude spectrum management, and the role of Bayesian rank adaptation under uncertain and sparse observations.

6.1. Deployment-Oriented Insights from Sparse REM Reconstruction

The experimental results provide useful insights for 6G-enabled low-altitude spectrum management. As shown in Figure 3 and the complexity analysis, Ra-BTRC maintains reliable reconstruction performance under extremely sparse observations. This result is practically meaningful because dense spectrum sensing over the full time–frequency–3D space is difficult to realize in low-altitude operations. UAV-borne sensing is usually constrained by flight endurance, payload capacity, energy consumption, mission routes, and regulatory restrictions. Therefore, constructing a reliable 5D REM from limited observations can reduce the sensing burden while still providing global electromagnetic awareness for low-altitude networks.

The results under different sampling patterns further indicate the deployment relevance of the proposed framework. Figure 5 shows that Ra-BTRC consistently achieves the lowest MMSE under random, trajectory-constrained, and block-missing sampling patterns. This suggests that the method is effective not only under ideal random sensing, but also under more realistic conditions where observations are concentrated along UAV trajectories or where certain regions cannot be sampled due to no-fly zones, obstacles, or task constraints. From an operational perspective, the completed REM can support interference-aware trajectory planning, adaptive frequency selection, dynamic spectrum access, and communication resource allocation. From a management perspective, it can provide spectrum managers, network operators, and urban airspace regulators with a data-driven basis for identifying interference hotspots, coordinating heterogeneous low-altitude services, and improving the reliability and safety of UAV/eVTOL operations.

6.2. Bayesian Rank Adaptation Under Uncertain Observations

Variational Bayesian inference is particularly suitable for practical low-altitude REM construction because it provides a probabilistic mechanism for handling uncertain and sparse sensing data and, within the TR-based completion framework, further alleviates the rank-selection problem. On the measurement side, real UAV spectrum observations inevitably contain uncertainty caused by receiver noise, calibration errors, synchronization mismatch, platform mobility, shadowing, and multipath propagation. By estimating posterior distributions of the TR core factors and learning the noise precision from data, the Bayesian formulation balances data fitting and uncertainty-aware regularization instead of producing only a deterministic point estimate. This mechanism is consistent with the convergence behavior in Figure 8, where the posterior statistics and noise precision gradually stabilize during the iterative inference process.

On the observation side, low-altitude sensing data are inherently sparse and non-uniform because UAV trajectories can cover only a small subset of the full 5D tensor. The sparsity-inducing priors in Ra-BTRC regularize the completion process under limited observations and enable information sharing across time, frequency, and 3D space. This explains why Ra-BTRC can still recover local spectrum structures and temporal–frequency–altitude variations under very low sampling rates, as illustrated in Figures 4 and 7. Therefore, the role of variational Bayesian inference is not limited to numerical optimization; it provides a

principled way to stabilize high-dimensional REM completion when the available UAV measurements are sparse, fragmented, and non-uniformly distributed.

On the model side, TR-based completion requires an appropriate rank configuration, whereas the intrinsic rank of a real low-altitude spectrum tensor is generally unknown and may vary with emitter activity, altitude-dependent propagation, blockage, interference density, and sensing coverage. As reported in Table 3, the same initial TR rank is automatically compressed into different effective ranks under different sampling rates, demonstrating that Ra-BTRC can adapt model complexity according to data support. This automatic rank adaptation is important for real deployment because network operators cannot reliably predefine a single TR rank for different regions, time periods, sampling budgets, and service conditions.

7. Conclusions

This study investigated 5D REM construction for LAE scenarios under two major challenges: the highly dynamic, high-dimensional electromagnetic environment and the extreme sparsity of UAV-collected measurements. To address these issues, we proposed the Ra-BTRC framework. By modeling the low-altitude spectrum situation as a unified 5D tensor over time, frequency, and 3D space (T - F - X - Y - Z), the proposed method exploits high-order correlations across coupled domains. Moreover, by introducing sparsity-inducing hierarchical priors and variational Bayesian inference, it estimates effective TR ranks in a data-driven manner and adaptively controls model complexity under sparse observations.

The simulation results showed that Ra-BTRC consistently outperformed representative baselines and achieved accurate and robust completion, even at very low sampling rates. These results support the main premise of this work; namely, that jointly modeling time, frequency, and 3D spatial dependencies, together with adaptive Bayesian rank learning, is crucial for reliable low-altitude REM construction under severe observation sparsity. More broadly, the proposed framework provides a practical basis for spectrum cognition in AI and 6G-enabled low-altitude networks, with potential value for dynamic spectrum access, interference-aware trajectory planning, and airspace management. Future work will focus on measurement-driven validation and practical robustness enhancement of the proposed framework in real low-altitude UAV sensing environments.

Author Contributions: Conceptualization, Y.W. and Z.S.; methodology, Y.W. and Z.S.; software, Y.W. and H.M.; validation, Y.W. and H.M.; formal analysis, Y.W. and Z.S.; investigation, Y.W.; resources, Z.S.; data curation, Y.W. and H.M.; writing—original draft preparation, Y.W.; writing—review and editing, Y.W., Z.S. and H.M.; visualization, Y.W.; supervision, Z.S.; project administration, Z.S.; funding acquisition, Z.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China under Grants 62171063.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The simulation data presented in this study can be reproduced using the parameters and methodology detailed in Section 5.1.1. The generated datasets are also available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Zhang, Y.; Xu, R.; Hu, B.; Nie, K.; Zhao, H.; Chen, B.; Kong, Q. Enabling Scalable and Efficient Low-Altitude Airspace Utilization for Dense Urban Operations. *Aerospace* **2026**, *13*, 294.
2. Xu, J.; Guan, C.; Wang, Y.; Zhuang, J.; Gan, W. A Systematic Review of Urban Air Mobility Development: eVTOL Drones' Technological Challenges and Low-Altitude Policies of Shenzhen. *Drones* **2025**, *9*, 842.
3. Zeng, Y.; Wu, Q.; Zhang, R. Accessing From the Sky: A Tutorial on UAV Communications for 5G and Beyond. *Proc. IEEE* **2019**, *107*, 2327–2375.
4. Wang, Y.; Sun, G.; Sun, Z.; Wang, J.; Li, J.; Zhao, C.; Wu, J.; Liang, S.; Yin, M.; Wang, P.; et al. Toward realization of low-altitude economy networks: Core architecture, integrated technologies, and future directions. *IEEE Trans. Cogn. Commun. Netw.* **2025**, *11*, 2788–2820.
5. Zhang, S.; Liu, H.; Liu, Q.; Chen, K.; Ma, S.; Niyato, D. Large AI model in the air: Ubiquitous intelligent networks for low altitude economy. *IEEE Netw.* **2026**, *40*, 160–167.
6. Li, B.; Chen, J. Radio map-assisted routing and predictive resource allocation over dynamic low-altitude networks. *IEEE Trans. Wirel. Commun.* **2026**, *25*, 9955–9970.
7. Romero, D.; Kim, S.-J. Radio map estimation: A data-driven approach to spectrum cartography. *IEEE Signal Process. Mag.* **2022**, *39*, 53–72.
8. Wang, J.; Zhang, J.; Zhang, Y.; Sun, Y.; Nie, G.; Shi, L.; Zhang, P.; Liu, G. Radio environment knowledge pool for 6G digital twin channel. *IEEE Commun. Mag.* **2025**, *63*, 158–164.
9. Jasim, M.A.; Hassan, H.U.; Ali, M. A Survey on Spectrum Management for Unmanned Aerial Vehicles (UAVs). *IEEE Access* **2021**, *10*, 11443–11499.
10. Shrestha, S.; Fu, X.; Hong, M. Deep spectrum cartography: Completing radio map tensors using learned neural models. *IEEE Trans. Signal Process.* **2022**, *70*, 1170–1184.
11. Romero, D.; Shrestha, R.; Teganya, Y.; Chepuri, S.P. Aerial spectrum surveying: Radio map estimation with autonomous UAVs. In *Proceedings of the 2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, Espoo, Finland, 21–24 September 2020; IEEE: New York, NY, USA, 2020; pp. 1–6.
12. Wei, Z.; Yao, R.; Kang, J.; Chen, X.; Wu, H. Three-dimensional spectrum occupancy measurement using UAV: Performance analysis and algorithm design. *IEEE Sens. J.* **2022**, *22*, 9146–9157.
13. Liu, W.; Chen, J. UAV-aided radio map construction exploiting environment semantics. *IEEE Trans. Wireless Commun.* **2023**, *22*, 6341–6355.
14. Lin, J.; Wang, H.; Wu, T.; Shen, Z.; Jiang, R.; Fan, X. Unmanned Aerial Vehicle-Enabled Aerial Radio Environment Map Construction: A Multi-Stage Approach to Data Sampling and Path Planning. *Drones* **2025**, *9*, 81.
15. Chen, X.; Wang, J.; Huang, Q. Dynamic spectrum cartography: Reconstructing spatial-spectral-temporal radio frequency map via tensor completion. *IEEE Trans. Signal Process.* **2025**, *73*, 1184–1199.
16. Tang, M.; Ding, G.; Wu, Q.; Xue, Z.; Tsiftsis, T.A. A Joint Tensor Completion and Prediction Scheme for Multi-Dimensional Spectrum Map Construction. *IEEE Access* **2016**, *4*, 8044–8052.
17. Yin, K.; Fang, S.; Chu, F.; Fan, Y. Compressed tensor completion: Approach for UAV-aided 3-D radio map construction. *IEEE Internet Things J.* **2024**, *11*, 40516–40531.
18. Kolda, T.G.; Bader, B.W. Tensor decompositions and applications. *SIAM Rev.* **2009**, *51*, 455–500.
19. Zhao, Q.; Zhou, G.; Xie, S.; Zhang, L.; Cichocki, A. Tensor ring decomposition. *arXiv* **2016**, arXiv:1606.05535.
20. Long, Z.; Zhu, C.; Liu, J.; Liu, Y. Bayesian low rank tensor ring for image recovery. *IEEE Trans. Image Process.* **2021**, *30*, 3568–3580.
21. Hu, Y.; Zhang, J.; Liu, Y.; Liu, G.; Chen, G. Spectrum map construction method based on dynamic window size tensor ring low-rank factors. *Radioengineering* **2024**, *33*, 163–172.
22. Wang, W.; Aggarwal, V.; Aeron, S. Efficient low rank tensor ring completion. In *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, 22–29 October 2017; IEEE: New York, NY, USA, 2017; pp. 5697–5705.
23. Yuan, L.; Zhao, Q.; Cao, J. Higher-dimension tensor completion via low-rank tensor ring decomposition. In *Proceedings of the 2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, Honolulu, HI, USA, 12–15 November 2018; IEEE: New York, NY, USA, 2018; pp. 1–6.
24. Yu, J.; Zhou, G.; Sun, W.; Xie, S. Robust to rank selection: Low-rank sparse tensor-ring completion. *IEEE Trans. Neural Netw. Learn. Syst.* **2021**, *34*, 2451–2465.
25. Ivanov, A.; Muhammad, B.; Tonchev, K.; Mihovska, A.; Poulkov, V. UAV-based volumetric measurements toward radio environment map construction and analysis. *Sensors* **2022**, *22*, 9705.
26. Wang, X.; Shen, B.; Huang, X. Unbalanced Spectrum Cartography Algorithm Based on Overlapping Ket Augmentation and Tensor Train. *Acta Electronica Sinica* **2024**, *52*, 2468–2476. <https://doi.org/10.12263/DZXB.20240004>.
27. Lee, D.; Kim, S.J.; Giannakis, G.B. Channel gain cartography for cognitive radios leveraging low rank and sparsity. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 5953–5966.

28. Oseledets, I.V. Tensor-train decomposition. *SIAM J. Sci. Comput.* **2011**, *33*, 2295–2317.
29. Yuan, L.; Li, C.; Mandic, D.; Cao, J.; Zhao, Q. Tensor ring decomposition with rank minimization on latent space: An efficient approach for tensor completion. *Proc. AAAI Conf. Artif. Intell.* **2019**, *33*, 9151–9158.
30. Tipping, M.E. Sparse Bayesian learning and the relevance vector machine. *J. Mach. Learn. Res.* **2001**, *1*, 211–244.
31. Zhu, Q.; Yao, M.; Bai, F.; Chen, X.; Zhong, W.; Hua, B.; Ye, X. A general altitude-dependent path loss model for UAV-to-ground millimeter-wave communications. *Front. Inf. Technol. Electron. Eng.* **2021**, *22*, 767–776.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.