

Article

High-Accuracy Off-Grid Sparse Bayesian Learning with Reliability-Guided Inference for Direction-of-Arrival Estimation

Wenchao He ^{1,2} , Haoran Wang ², Hongxi Zhao ²  and Yiran Shi ^{2,*} ¹ School of Mechanical and Electrical Engineering, Changchun Humanities and Sciences College, Changchun 130117, China; hewc23@mails.jlu.edu.cn² College of Communication Engineering, Jilin University, Changchun 130012, China; hrwang23@mails.jlu.edu.cn (H.W.); hxzhao24@mails.jlu.edu.cn (H.Z.)

* Correspondence: shiyiran@jlu.edu.cn

Abstract

Off-grid direction-of-arrival (DOA) estimation based on sparse Bayesian learning (SBL) can alleviate angular discretization mismatch, but its practical performance may be affected by unreliable posterior relevance statistics, sensitivity of effective error precision learning, and unstable offset correction. This paper proposes a reliability-guided stabilized off-grid SBL method for multisnapshot DOA estimation. The method is developed within the standard first-order multiple-measurement-vector Bayesian model and introduces three stabilization modules. First, a confidence-guided MAP-type shrinkage relevance update is introduced to suppress weak and non-dominant posterior components through reliability-controlled non-expansive shrinkage. Second, a posterior-concentration-guided damped noise update is introduced to stabilize scalar effective error precision learning when the sparse support is uncertain. Third, a trust-region cubic-regularized Newton refinement is formulated to obtain bounded active-support off-grid corrections from the posterior expected reconstruction error. Simulation results under off-grid deviation, varying SNRs, varying snapshot numbers, different source separations, and random-angle scenarios show that the proposed method achieves competitive and stable estimation performance compared with representative classical and sparse Bayesian baselines.

Keywords: direction-of-arrival estimation; off-grid sparse Bayesian learning; basis mismatch; multiple-measurement-vector model; reliability-guided inference

1. Introduction

Direction-of-arrival (DOA) estimation is a fundamental task in array signal processing, with applications in radar, sonar, wireless communications, and intelligent sensing. Since practical observations are often affected by noise, finite snapshots, and model mismatch, DOA estimators are expected to provide both accurate and stable localization results [1–4].

Classical DOA estimation methods include beamforming-based, subspace-based, and maximum-likelihood-based approaches. Capon proposed the minimum-variance distortionless response beamformer to improve spatial resolution by suppressing undesired interference [5]. MUSIC exploits the orthogonality between signal and noise subspaces for high-resolution estimation [6], while ESPRIT uses rotational invariance to avoid exhaustive spectral searching under suitable array structures [7]. Maximum-likelihood estimators can achieve high statistical efficiency, but usually require nonlinear multidimensional optimization [8]. Although these methods are well established, their performance depends



Received: 3 May 2026

Revised: 19 June 2026

Accepted: 9 July 2026

Published: 14 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.

strongly on reliable sample statistics and accurate array models. In low-SNR or limited-snapshot scenarios, the sample covariance may deviate from its ideal value, and the separation between the signal and noise subspaces becomes less reliable. These factors can lead to degraded resolution, missed sources, or unstable estimation results.

Sparse reconstruction provides another route for high-resolution DOA estimation. By discretizing the angular domain into an overcomplete dictionary, DOA estimation can be formulated as sparse spatial spectrum recovery. Malioutov et al. introduced this sparse reconstruction perspective for source localization [9]. Orthogonal matching pursuit (OMP) is a representative greedy sparse recovery method with low computational cost [10]. However, because OMP selects atoms sequentially, its estimation result can be sensitive to noise and early selection errors. More importantly, most fixed-grid sparse reconstruction methods suffer from basis mismatch when the true DOAs do not lie exactly on the predefined grid. This mismatch may spread the source energy over neighboring grid atoms and reduce the accuracy of the recovered sparse spectrum, especially when a coarse grid is used to control computational complexity.

To improve the stability of sparse recovery, sparse Bayesian learning (SBL) introduces a probabilistic inference framework for estimating sparse coefficients and relevance hyperparameters. Based on automatic relevance determination, SBL estimates the posterior distribution of sparse coefficients and learns relevance hyperparameters within a Bayesian framework [11,12]. Bayesian compressive sensing further showed that sparse signals and noise uncertainty can be jointly inferred from limited measurements [13]. Compared with greedy sparse methods, SBL exploits posterior mean, posterior covariance, and hyperparameter learning, making it attractive for multisnapshot DOA estimation with shared row sparsity. In the multiple-measurement-vector (MMV) setting, the common sparse support across snapshots can be used to improve estimation stability. However, conventional grid-based SBL still relies on a fixed angular dictionary and therefore remains affected by basis mismatch in off-grid scenarios.

To alleviate the basis-mismatch problem, Yang et al. proposed off-grid sparse Bayesian inference (OGSBI), in which first-order perturbation variables are introduced to characterize the deviations between true directions and their nearby grid points [14]. Dai et al. developed RootSBL, which improves continuous-angle refinement and computational efficiency through root-based off-grid estimation [15]. Subsequent studies have improved off-grid SBL from several representative aspects, including adaptive noise modeling, sparse prior design, continuous-angle refinement, and efficient Bayesian inference. Wen et al. introduced a variational sparse Bayesian framework for off-grid DOA estimation under nonuniform noise, showing that off-grid perturbations and noise uncertainty can be jointly inferred [16]. Yang and Yang incorporated hierarchical synthesis Lasso priors to enhance sparsity promotion and support discrimination for off-grid signals [17]. Shen et al. proposed an improved RootSBL method by further refining the root-based update strategy of the original RootSBL framework [18]. More recently, Li et al. developed an off-grid SBL estimator with three-stage hierarchical Laplace priors, and Tong et al. proposed an off-grid fast variational Bayesian inference method to reduce computational burden [19,20]. Deep-unfolding SBL networks have also been investigated to combine model-driven Bayesian inference with data-driven learning [21].

Although off-grid SBL alleviates angular basis mismatch by introducing continuous offset variables, first-order off-grid inference may still be affected by unreliable relevance learning, sensitive effective error precision updating, and unstable offset refinement. These issues can weaken support concentration and lead to less stable off-grid correction, especially under finite snapshots and noisy observations. Therefore, this paper focuses on stabilizing these sensitive inference steps within the first-order off-grid MMV-SBL framework.

To address these issues, this paper proposes a reliability-guided stabilized off-grid SBL method for multisnapshot DOA estimation. Within a first-order MMV Bayesian framework, CGBS-MAP, PCG-DNU, and TR-CRN are integrated to improve support concentration, stabilize effective error precision learning, and provide bounded off-grid refinement. The main contributions of this paper are summarized as follows:

1. An empirical-Bayes-inspired confidence-guided MAP-type shrinkage relevance update is proposed. By combining posterior confidence and local dominance, the update suppresses weak and non-dominant components. It improves sparse relevance concentration without using a hard support threshold.
2. A posterior-concentration-guided damped noise update is introduced. The residual-based effective error precision update is relaxed according to the current relevance concentration, which stabilizes precision learning during uncertain support formation while preserving the residual-based fixed point.
3. A trust-region cubic-regularized Newton refinement is developed for active-support off-grid correction. By imposing cubic damping and trust-region constraints, the refinement provides bounded offset updates and reduces overly aggressive correction within the local first-order approximation range.

The remainder of this paper is organized as follows. Section 2 introduces the signal model and problem formulation, Section 3 presents the proposed method, Section 4 reports the simulation results, and Section 5 concludes the paper.

2. Signal Model

2.1. Array Observation and Bayesian Sparse Model

Consider K far-field narrowband sources impinging on an M -element array from directions $\{\theta_k\}_{k=1}^K$. Let $\mathbf{x}(l) \in \mathbb{C}^M$ denote the received snapshot at time index l , and let L denote the number of snapshots. The array observation model is

$$\mathbf{x}(l) = \mathbf{A}(\boldsymbol{\theta})\mathbf{s}(l) + \mathbf{n}(l), \quad l = 1, 2, \dots, L \quad (1)$$

where $\mathbf{A}(\boldsymbol{\theta}) = [\mathbf{a}(\theta_1), \mathbf{a}(\theta_2), \dots, \mathbf{a}(\theta_K)] \in \mathbb{C}^{M \times K}$ is the steering matrix, $\mathbf{s}(l) \in \mathbb{C}^K$ is the source vector, and $\mathbf{n}(l) \in \mathbb{C}^M$ denotes additive noise. Stacking all snapshots gives

$$\mathbf{X} = [\mathbf{x}(1), \mathbf{x}(2), \dots, \mathbf{x}(L)] = \mathbf{A}(\boldsymbol{\theta})\mathbf{S} + \mathbf{N} \quad (2)$$

where $\mathbf{X} \in \mathbb{C}^{M \times L}$, $\mathbf{S} \in \mathbb{C}^{K \times L}$, and $\mathbf{N} \in \mathbb{C}^{M \times L}$.

For a linear array with sensor positions normalized by wavelength, the steering vector is written as

$$\mathbf{a}(\theta) = \left[e^{-j2\pi p_1 \sin \theta}, e^{-j2\pi p_2 \sin \theta}, \dots, e^{-j2\pi p_M \sin \theta} \right]^T \quad (3)$$

where $p_m = d_m/\lambda$ is the normalized position of the m th sensor, with d_m measured from the reference sensor. For a half-wavelength ULA, $p_m = (m-1)/2$.

After discretizing the angular search interval, the sparse source coefficients over the candidate grid are represented by a row-sparse matrix $\mathbf{W} \in \mathbb{C}^{N \times L}$. For each snapshot, the sparse coefficient vector is assigned the prior

$$\mathbf{w}(l) \sim \mathcal{CN}(\mathbf{0}, \text{diag}(\boldsymbol{\alpha})) \quad (4)$$

where $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \dots, \alpha_N]^T$ is a variance-type relevance hyperparameter vector. A larger α_n indicates a more active candidate direction, while a smaller α_n suppresses the corresponding grid component.

The residual after sparse grid modeling contains both observation noise and residual modeling error. It is denoted as an effective model error and modeled as

$$\mathbf{e}(l) \sim \mathcal{CN}(\mathbf{0}, \alpha_0^{-1} \mathbf{I}_M) \quad (5)$$

where α_0 is the scalar precision of the effective Gaussian error.

2.2. First-Order Off-Grid Sparse Model

The angular search interval is discretized into a grid $\{\tilde{\theta}_n\}_{n=1}^N$. Since true DOAs generally do not exactly coincide with predefined grid points, fixed-grid sparse modeling suffers from basis mismatch.

To compensate for this mismatch, a first-order off-grid approximation is introduced. For a direction around $\tilde{\theta}_n$, the steering vector is approximated as

$$\mathbf{a}(\theta) \approx \mathbf{a}(\tilde{\theta}_n) + \mathbf{b}(\tilde{\theta}_n)\beta_n \quad (6)$$

where, for an active grid point close to a true DOA, $\beta_n = \theta - \tilde{\theta}_n$ is the local off-grid offset, and

$$\mathbf{b}(\tilde{\theta}_n) = \left. \frac{\partial \mathbf{a}(\theta)}{\partial \theta} \right|_{\theta=\tilde{\theta}_n} \quad (7)$$

is the derivative of the steering vector with respect to the angular variable.

Figure 1 illustrates the first-order off-grid modeling principle. The off-grid offset β_n locally corrects the steering vector around the nearest grid direction, thereby reducing angular discretization mismatch.




Figure 1. Illustration of the first-order off-grid DOA model.

The grid dictionary and derivative dictionary are defined as

$$[\mathbf{A}_g]_{m,n} = e^{-j2\pi p_m \sin \tilde{\theta}_n} \quad (8)$$

$$[\mathbf{B}_g]_{m,n} = -j2\pi p_m \cos \tilde{\theta}_n e^{-j2\pi p_m \sin \tilde{\theta}_n} \quad (9)$$

where $\mathbf{A}_g \in \mathbb{C}^{M \times N}$ is the grid dictionary and $\mathbf{B}_g \in \mathbb{C}^{M \times N}$ is the derivative dictionary.

The effective first-order off-grid dictionary is then written as

$$\Phi(\beta) = \mathbf{A}_g + \mathbf{B}_g \text{diag}(\beta) \quad (10)$$

where $\beta = [\beta_1, \beta_2, \dots, \beta_N]^T$ is the real-valued off-grid offset vector. To keep the first-order approximation valid, each offset is bounded by

$$|\beta_n| \leq \frac{\Delta}{2} \quad (11)$$

where Δ is the angular grid interval.

Using the stacked observation matrix, the first-order off-grid sparse model becomes

$$\mathbf{X} = \Phi(\beta)\mathbf{W} + \mathbf{E} \quad (12)$$

where $\mathbf{W} \in \mathbb{C}^{N \times L}$ is the row-sparse source coefficient matrix and $\mathbf{E} \in \mathbb{C}^{M \times L}$ is the effective model error matrix. Unlike the physical noise matrix $\mathbf{N}$ in the original array model, $\mathbf{E}$ includes both observation noise and residual modeling error caused by angular discretization and first-order approximation.

3. Proposed Method

3.1. Overall Framework

The proposed method follows the first-order MMV off-grid Bayesian model in Section 2. Given $\mathbf{X}$, $\mathbf{A}_g$, $\mathbf{B}_g$, and the assumed source number K , it alternately updates the sparse relevance vector $\boldsymbol{\alpha}$, the scalar effective error precision α_0 , and the off-grid offset vector $\boldsymbol{\beta}$. In the following derivations, Φ denotes the current effective dictionary $\Phi(\boldsymbol{\beta})$.

At each iteration, MMV-SBL posterior inference is performed first. Then, CGBS-MAP, PCG-DNU, and TR-CRN update $\boldsymbol{\alpha}$, α_0 , and $\boldsymbol{\beta}$, respectively. After the stopping criterion is satisfied, the final DOAs are obtained from the dominant support and the refined offsets.

3.2. Bayesian Prior-Posterior Inference

Under the current effective dictionary Φ , the off-grid MMV observation model is written as

$$\mathbf{X} = \Phi \mathbf{W} + \mathbf{E} \quad (13)$$

where $\mathbf{W} \in \mathbb{C}^{N \times L}$ is the row-sparse source coefficient matrix and $\mathbf{E}$ denotes the effective model error matrix modeled as a circularly symmetric complex Gaussian term. Let $\mathbf{x}(l)$ and $\mathbf{w}(l)$ denote the l th columns of $\mathbf{X}$ and $\mathbf{W}$, respectively. For each snapshot, the likelihood is given by

$$p(\mathbf{x}(l) | \mathbf{w}(l), \alpha_0) = \mathcal{CN}(\mathbf{x}(l) | \Phi \mathbf{w}(l), \alpha_0^{-1} \mathbf{I}_M) \quad (14)$$

where α_0 is the scalar precision of the effective Gaussian error. Assuming conditional independence among snapshots, the likelihood of the whole observation matrix is

$$p(\mathbf{X} | \mathbf{W}, \alpha_0) = \prod_{l=1}^L p(\mathbf{x}(l) | \mathbf{w}(l), \alpha_0) \quad (15)$$

To impose common row sparsity across snapshots, each sparse coefficient vector $\mathbf{w}(l)$ is assigned a Gaussian prior with the shared covariance matrix:

$$p(\mathbf{w}(l) | \boldsymbol{\alpha}) = \mathcal{CN}(\mathbf{w}(l) | \mathbf{0}, \boldsymbol{\Gamma}), \quad \boldsymbol{\Gamma} = \text{diag}(\boldsymbol{\alpha}) \quad (16)$$

where $\boldsymbol{\alpha} = [\alpha_1, \alpha_2, \dots, \alpha_N]^T$ is the variance-type relevance hyperparameter vector. The same $\boldsymbol{\alpha}$ is shared by all snapshots, so that the rows of $\mathbf{W}$ are encouraged to have a common sparse support. The joint prior is

$$p(\mathbf{W} | \boldsymbol{\alpha}) = \prod_{l=1}^L p(\mathbf{w}(l) | \boldsymbol{\alpha}) \quad (17)$$

According to Bayes' rule, the posterior distribution is

$$p(\mathbf{W} | \mathbf{X}, \boldsymbol{\alpha}, \alpha_0) \propto p(\mathbf{X} | \mathbf{W}, \alpha_0) p(\mathbf{W} | \boldsymbol{\alpha}) \quad (18)$$

Since both the likelihood and the prior are complex Gaussian, the posterior distribution is also complex Gaussian:

$$p(\mathbf{W} | \mathbf{X}, \boldsymbol{\alpha}, \alpha_0) = \prod_{l=1}^L \mathcal{CN}(\mathbf{w}(l) | \boldsymbol{\mu}(l), \boldsymbol{\Sigma}) \quad (19)$$

The posterior covariance matrix shared by all snapshots and the posterior mean of the l th snapshot are given by

$$\Sigma = \left(\alpha_0 \Phi^H \Phi + \Gamma^{-1} \right)^{-1} \quad (20)$$

$$\mu(l) = \alpha_0 \Sigma \Phi^H \mathbf{x}(l) \quad (21)$$

For computational convenience, the equivalent form based on the matrix inversion lemma is used:

$$\Sigma = \Gamma - \Gamma \Phi^H \mathbf{C}^{-1} \Phi \Gamma \quad (22)$$

where

$$\mathbf{C} = \Phi \Gamma \Phi^H + \alpha_0^{-1} \mathbf{I}_M \quad (23)$$

The posterior mean matrix of all snapshots is then obtained as

$$\mathbf{M} = \Gamma \Phi^H \mathbf{C}^{-1} \mathbf{X} \quad (24)$$

The n th row $\mathbf{m}_n$ of $\mathbf{M}$ contains the posterior coefficient estimates of the n th candidate direction across all snapshots.

For sparse relevance learning, let

$$s_n = [\Sigma]_{n,n} \quad (25)$$

denote the posterior variance of the n th candidate direction. The posterior second-moment statistic is defined as

$$q_n = \frac{1}{L} \|\mathbf{m}_n\|_2^2 + s_n \quad (26)$$

Here, q_n combines the average posterior coefficient energy and the posterior uncertainty of the n th candidate direction.

In conventional SBL, the relevance parameter is usually updated directly as $\alpha_n = q_n$. However, in off-grid DOA estimation, finite snapshots, effective model error, and residual basis mismatch may produce weak and unreliable activations around the true support. Therefore, the proposed method introduces a confidence-guided MAP-type shrinkage update for α_n , as described in the next subsection.

3.3. Confidence-Guided MAP-Type Shrinkage Relevance Learning

To improve posterior spectral concentration, CGBS-MAP introduces a reliability-controlled shrinkage mechanism into the relevance update. This design aims to reduce weak and unreliable activations while retaining dominant posterior components. Compared with direct moment updating, the shrinkage form provides a smoother way to suppress low-confidence components. It also avoids introducing a hard support threshold before the posterior relevance profile becomes sufficiently concentrated. The following exponential form is introduced as an empirical-Bayes-inspired regularization term for α_n :

$$p(\alpha_n | \tau_n) \propto \exp(-\tau_n \alpha_n), \quad \alpha_n > 0 \quad (27)$$

Here, τ_n is not treated as a fixed prior parameter independent of the data. Instead, it is constructed adaptively from posterior reliability information and is used as a data-adaptive regularization coefficient. The resulting update is a reliability-guided regularized MAP-type shrinkage rule and does not correspond to a fully hierarchical Bayesian prior model.

The corresponding MAP-type relevance update is formulated as

$$\alpha_n^{\text{MAP}} = \arg \min_{\alpha_n > 0} \left(\log \alpha_n + \frac{q_n}{\alpha_n} + \tau_n \alpha_n \right) \quad (28)$$

Solving the stationary condition gives the unique positive solution

$$\alpha_n^{\text{MAP}} = \frac{2q_n}{1 + \sqrt{1 + 4\tau_n q_n}} \quad (29)$$

When $\tau_n = 0$, the update reduces to the ordinary SBL moment update $\alpha_n = q_n$. When $\tau_n > 0$, it provides a non-expansive shrinkage of q_n , thereby suppressing weak posterior components without hard thresholding.

The shrinkage strength is controlled by a posterior reliability score. For the scalar effective error precision considered in this paper, the effective error variance estimate is

$$\hat{\sigma}^2 = \frac{1}{\alpha_0} \quad (30)$$

The global confidence factor and local dominance factor are defined as

$$c_n = \frac{q_n}{q_n + \kappa_\sigma \hat{\sigma}^2 + \varepsilon} \quad (31)$$

and

$$d_n = \frac{q_n}{\max_{i \in \mathcal{N}_n} q_i + \varepsilon} \quad (32)$$

Here, $\mathcal{N}_n$ denotes a fixed local neighborhood including the n th grid point, ε avoids numerical division by zero, and κ_σ is a scale coefficient. In the experiments, $\mathcal{N}_n$ includes the n th grid point and its available immediate neighboring grid points. The factor c_n measures the error-relative significance of q_n , whereas d_n measures its local dominance.

The reliability score is constructed as

$$z_n = c_n^{\gamma_c} d_n^{\gamma_d} \quad (33)$$

where γ_c and γ_d are fixed shaping constants. The adaptive shrinkage rate is then assigned by

$$\tau_n = \text{clip}\left(\frac{\lambda}{z_n + \varepsilon}, \tau_{\min}, \tau_{\max}\right) \quad (34)$$

where $\text{clip}(x, a, b) = \min\{\max(x, a), b\}$, $\lambda > 0$ controls the overall shrinkage strength, and $\tau_{\min}$ and $\tau_{\max}$ are fixed bounds. The clipping operation keeps τ_n within a prescribed range and avoids either negligible or excessive shrinkage. A larger z_n leads to weaker shrinkage, while a smaller z_n leads to stronger suppression of unreliable components.

Since $\tau_n \geq 0$, the following non-expansive property holds:

$$0 < \alpha_n^{\text{MAP}} \leq q_n \quad (35)$$

Thus, CGBS-MAP does not amplify the ordinary posterior moment update, but selectively shrinks unreliable components according to their confidence and local dominance.

For numerical stability, a damped relevance update is used:

$$\alpha_n^{(t+1)} = \eta_\alpha \alpha_n^{\text{MAP}} + (1 - \eta_\alpha) \alpha_n^{(t)}, \quad 0 < \eta_\alpha \leq 1 \quad (36)$$

where η_α is the damping factor for relevance updating. Therefore, CGBS-MAP preserves the closed-form SBL relevance update structure while introducing reliability-aware shrinkage to improve posterior spectral concentration.

Figure 2 illustrates the reliability-controlled transformation from the posterior statistic q_n to the MAP-type shrinkage update, where reliable components are retained and weak non-dominant activations are suppressed.




Figure 2. Illustration of the CGBS-MAP relevance update.

3.4. Posterior-Concentration-Guided Damped Noise Update

The scalar effective error precision is updated from the posterior expected reconstruction error. Under the current effective dictionary Φ , define the posterior-mean residual matrix as

$$\mathbf{R} = \mathbf{X} - \Phi \mathbf{M} \quad (37)$$

For the m th observation dimension, the snapshot-averaged posterior residual energy is

$$E_m^2 = \frac{1}{L} \|\mathbf{r}_m\|_2^2 + [\Phi \Sigma \Phi^H]_{m,m} \quad (38)$$

where $\mathbf{r}_m$ is the m th row of $\mathbf{R}$. The first term is the posterior-mean residual energy, and the second term accounts for posterior uncertainty.

The use of $a_0 + M$ in the numerator is due to the snapshot-averaged normalization in (38). Specifically, if the posterior expected residual energy is written in its unnormalized form, the effective residual contribution is

$$\sum_{l=1}^L \mathbb{E} [\|\mathbf{x}(l) - \Phi \mathbf{w}(l)\|_2^2] = \|\mathbf{R}\|_F^2 + L \text{tr}(\Phi \Sigma \Phi^H) \quad (39)$$

Dividing this expression by L gives

$$\frac{1}{L} \|\mathbf{R}\|_F^2 + \text{tr}(\Phi \Sigma \Phi^H) = \sum_{m=1}^M E_m^2 \quad (40)$$

Thus, $\sum_{m=1}^M E_m^2$ represents a sensor-domain sum of snapshot-averaged posterior residual energies. Under this normalized formulation, the effective number of averaged residual terms is M . If the unnormalized likelihood form were used directly, an equivalent expression would contain an ML -dependent numerator and the corresponding unnormalized residual energy in the denominator.

A normalized residual-based estimate of the scalar effective error precision is then written as

$$\alpha_{0,\text{raw}} = \frac{a_0 + M}{b_0 + \sum_{m=1}^M E_m^2} \quad (41)$$

where a_0 and b_0 are small positive hyperparameters. Since direct assignment may be unstable when the support is still diffuse, a damped precision update is used:

$$\alpha_0^{(t+1)} = (1 - \eta_0^{(t)}) \alpha_0^{(t)} + \eta_0^{(t)} \alpha_{0,\text{raw}} \quad (42)$$

The relaxation factor is determined by the concentration of the current relevance profile. Let $\mathcal{S}_K$ denote the indices of the K largest entries of α . The concentration index is defined as

$$\kappa_\alpha = \frac{\sum_{n \in \mathcal{S}_K} \alpha_n}{\sum_{n=1}^N \alpha_n + \varepsilon}, \quad \bar{\kappa}_\alpha = \text{clip}\left(\frac{\kappa_\alpha - K/N}{1 - K/N + \varepsilon}, 0, 1\right) \quad (43)$$

where $\bar{\kappa}_\alpha \in [0, 1]$ is the normalized concentration score. The damping factor is set as

$$\eta_0^{(t)} = \eta_{\min} + (\eta_{\max} - \eta_{\min})\bar{\kappa}_\alpha, \quad 0 \leq \eta_{\min} \leq \eta_{\max} \leq 1 \quad (44)$$

Here, $\eta_{\min}$ and $\eta_{\max}$ denote the lower and upper bounds of the precision-update relaxation factor.

As shown in Figure 3, PCG-DNU updates the precision conservatively when the relevance profile is diffuse and more rapidly when the support becomes concentrated.




Figure 3. Illustration of the PCG-DNU effect. A diffuse relevance profile leads to a small relaxation factor and a conservative precision update, whereas a concentrated profile leads to a larger relaxation factor and a faster approach to the raw empirical-Bayes target.

Since (42) is a convex combination of $\alpha_0^{(t)}$ and $\alpha_{0,\text{raw}}$, the updated precision remains bounded between them. When $\alpha_0^{(t)} = \alpha_{0,\text{raw}}$, the update also satisfies $\alpha_0^{(t+1)} = \alpha_0^{(t)}$. Therefore, PCG-DNU stabilizes effective error precision learning while preserving the residual-based fixed point.

3.5. Trust-Region Cubic-Regularized Off-Grid Refinement

After sparse relevance learning, off-grid refinement is performed on the active support to reduce residual basis mismatch. Let the active support be $\mathcal{I}$, whose size is $K_a \geq K$, selected from the K_a largest entries of α . For fixed posterior mean $\mathbf{M}$ and covariance Σ , the offset correction is based on the posterior expected reconstruction error:

$$\mathcal{L}_\beta(\beta) = \|\mathbf{X} - \Phi(\beta)\mathbf{M}\|_F^2 + L \text{tr}[\Phi(\beta)\Sigma\Phi^H(\beta)] \quad (45)$$

For clarity, the derivation of the active-support quadratic form is given below. Let the active offset vector be $\mathbf{u} = \beta_{\mathcal{I}}$ and write the dictionary perturbation on the active support as

$$\Phi(\mathbf{u}) = \mathbf{A}_g + \sum_{i \in \mathcal{I}} u_i \mathbf{G}_i, \quad \mathbf{G}_i = \mathbf{b}_i \mathbf{e}_i^T \quad (46)$$

where $\mathbf{b}_i$ is the i th column of $\mathbf{B}_g$ and $\mathbf{e}_i$ is the i th canonical basis vector. Then

$$\Phi(\mathbf{u})\mathbf{M} = \mathbf{A}_g\mathbf{M} + \sum_{i \in \mathcal{I}} u_i \mathbf{D}_i, \quad \mathbf{D}_i = \mathbf{b}_i \mathbf{m}_i \quad (47)$$

where $\mathbf{m}_i$ is the i th row of $\mathbf{M}$. Defining $\mathbf{R}_A = \mathbf{X} - \mathbf{A}_g \mathbf{M}$, the residual term becomes

$$\left\| \mathbf{R}_A - \sum_{i \in \mathcal{I}} u_i \mathbf{D}_i \right\|_F^2 \quad (48)$$

Because the offset variables u_i are real-valued while $\mathbf{D}_i$ and $\mathbf{G}_i$ are complex-valued, the real part of the corresponding complex inner products determines the real-valued gradient and Hessian terms. Expanding the residual term and the trace term in (45), and removing constants independent of $\mathbf{u}$, yields a real quadratic function of $\mathbf{u}$.

Let $\mathbf{u} = \beta_T$. Since the offsets are angular corrections, $\mathbf{u}$ is real-valued. Under the first-order off-grid model, (45) leads to

$$\min_{\mathbf{u}} \frac{1}{2} \mathbf{u}^T \mathbf{P} \mathbf{u} - \mathbf{v}^T \mathbf{u} \quad |u_i| \leq \frac{\Delta}{2}, \quad i \in \mathcal{I} \quad (49)$$

where Δ is the angular grid interval, $\mathbf{P} \in \mathbb{R}^{K_a \times K_a}$ is the real-valued quadratic coefficient matrix, and $\mathbf{v} \in \mathbb{R}^{K_a}$ is the real-valued linear coefficient vector. For $i, j \in \mathcal{I}$,

$$\begin{aligned} P_{ij} &= 2 \operatorname{Re} \left\{ \operatorname{tr} \left(\mathbf{D}_i^H \mathbf{D}_j \right) + L \operatorname{tr} \left(\mathbf{G}_i \Sigma \mathbf{G}_j^H \right) \right\}, \\ v_i &= 2 \operatorname{Re} \left\{ \operatorname{tr} \left(\mathbf{D}_i^H \mathbf{R}_A \right) - L \operatorname{tr} \left(\mathbf{G}_i \Sigma \mathbf{A}_g^H \right) \right\} \end{aligned} \quad (50)$$

Here, $\mathbf{R}_A = \mathbf{X} - \mathbf{A}_g \mathbf{M}$, $\mathbf{D}_i = \mathbf{b}_i \mathbf{m}_i$, and $\mathbf{G}_i = \mathbf{b}_i \mathbf{e}_i^T$. The operator $\operatorname{Re}\{\cdot\}$ ensures that $\mathbf{P}$ and $\mathbf{v}$ are real-valued quantities associated with the real offset vector $\mathbf{u}$. This treatment is equivalent to applying real-valued calculus to the real scalar cost function $\mathcal{L}_\beta$.

A direct Newton-type update may be unstable when $\mathbf{P}$ is ill-conditioned or when the offset step leaves the local valid range. Therefore, TR-CRN introduces cubic damping and trust-region constraints to stabilize active-support refinement. Figure 4 depicts the procedure of the trust-region cubic-regularized off-grid refinement.




Figure 4. Illustration of the trust-region cubic-regularized off-grid refinement.

Let $\mathbf{P}_{\text{reg}} = \mathbf{P} + \delta \mathbf{I}$, where $\delta > 0$ is a small diagonal loading factor, and define

$$g_i = (\mathbf{P}_{\text{reg}} \mathbf{u} - \mathbf{v})_i, \quad H_i = (\mathbf{P}_{\text{reg}})_{ii} \quad (51)$$

TR-CRN uses the one-dimensional cubic-regularized model

$$m_i(d) = g_i d + \frac{1}{2} H_i d^2 + \frac{\sigma_c}{3} |d|^3 \quad (52)$$

where $\sigma_c > 0$ is the cubic regularization coefficient. The stationary condition gives

$$d_i = -\frac{g_i}{H_i + \lambda_i}, \quad \lambda_i = \frac{-H_i + \sqrt{H_i^2 + 4\sigma_c|g_i|}}{2} \quad (53)$$

The term $\lambda_i \geq 0$ increases the effective curvature and suppresses oversized Newton steps.

To keep the first-order approximation valid, both the offset and update step are bounded:

$$|\beta_i| \leq \frac{\Delta}{2}, \quad |d_i| \leq \rho\Delta, \quad 0 < \rho \leq \frac{1}{2} \quad (54)$$

The update is accepted only when $\mathcal{L}_\beta$ does not increase under fixed posterior statistics. Thus, TR-CRN provides bounded active-support off-grid offset correction rather than new support discovery.

3.6. DOA Extraction and Overall Iterative Procedure

After the stopping criterion is satisfied, the final DOA estimates are obtained from the dominant sparse support and the corresponding off-grid offsets. Since α is a variance-type relevance vector, its large entries indicate candidate grid points with stronger posterior source evidence. Let $\mathcal{I}_K = \{i_1, i_2, \dots, i_K\}$ denote the indices of the K largest entries of the final relevance vector α . The continuous-angle estimate is then obtained by adding the refined offset to the corresponding grid direction:

$$\hat{\theta}_k = \tilde{\theta}_{i_k} + \beta_{i_k}, \quad k = 1, 2, \dots, K \quad (55)$$

Here, $\tilde{\theta}_{i_k}$ is the selected grid point and β_{i_k} is the refined off-grid offset. The estimate $\hat{\theta}_k$ and the corresponding estimation error are measured in degrees. The estimated DOAs are sorted in ascending angular order for performance evaluation.

The iterative procedure stops when the relative change of the relevance vector becomes smaller than a prescribed tolerance, or when the maximum number of iterations is reached. The complete iterative procedure is summarized in Algorithm 1.

Algorithm 1 Overall procedure of the proposed reliability-guided off-grid SBL method

Input: Observation matrix $\mathbf{X}$, grid dictionary $\mathbf{A}_g$, derivative dictionary $\mathbf{B}_g$, assumed source number K , active support size K_a , maximum iteration number $T_{\max}$, and convergence tolerance ϵ_{tol} .

Output: Estimated DOAs $\{\hat{\theta}_k\}_{k=1}^K$ in egress.

1. Initialize $\alpha^{(0)}, \alpha_0^{(0)}, \beta^{(0)} = \mathbf{0}$, and set $t = 0$.
 2. Construct the effective dictionary $\Phi^{(t)} = \mathbf{A}_g + \mathbf{B}_g \text{diag}(\beta^{(t)})$.
 3. Perform MMV-SBL posterior inference with $\Phi^{(t)}$ to obtain $\mathbf{M}^{(t)}, \Sigma^{(t)}$, and $q_n^{(t)}$.
 4. Update the sparse relevance vector $\alpha^{(t+1)}$ using CGBS-MAP.
 5. Update the scalar effective error precision $\alpha_0^{(t+1)}$ using PCG-DNU.
 6. Select the active support from the K_a largest entries of $\alpha^{(t+1)}$, where $K_a \geq K$.
 7. Refine the active-support off-grid offsets using TR-CRN under the trust-region constraint and update $\beta^{(t+1)}$ accordingly.
 8. If the convergence criterion is satisfied or $t \geq T_{\max}$, stop; otherwise set $t = t + 1$ and return to Step 2.
 9. Select the final support $\mathcal{I}_K$ from the K largest entries of α and extract final DOAs by $\hat{\theta}_k = \tilde{\theta}_{i_k} + \beta_{i_k}$.
-

The proposed framework combines CGBS-MAP, PCG-DNU, and TR-CRN within a unified first-order MMV Bayesian inference procedure. These modules respectively improve relevance concentration, stabilize effective error precision learning, and provide bounded active-support off-grid offset correction. The final DOA estimates are obtained from the learned support and refined offsets.

It should be noted that the complete alternating procedure is not claimed to provide a strict monotonic decrease in a single global objective, because the relevance update and precision update include reliability-guided shrinkage and damping. The acceptance rule in TR-CRN only ensures that the offset-related posterior expected reconstruction error does not increase under fixed posterior statistics. Therefore, the proposed method is regarded as a stabilized iterative inference scheme. Accordingly, the experiments use the relative change of the relevance vector and the maximum iteration number as the stopping criteria.

3.7. Hyperparameter Determination and Reproducibility Settings

The key hyperparameters of the proposed method are γ_c , γ_d , $\tau_{\min}$, $\tau_{\max}$, and ρ . These parameters control the reliability shaping, adaptive shrinkage range, and trust-region offset refinement of the proposed method. To improve reproducibility and avoid case-specific tuning, the same default values are used in all simulation experiments unless otherwise stated. Their influence is further examined by the sensitivity analysis in Section 4.5.

For the CGBS-MAP module, γ_c and γ_d control the nonlinear shaping of the posterior confidence factor and the local dominance factor in (33), respectively. The bounds $\tau_{\min}$ and $\tau_{\max}$ in (34) restrict the adaptive shrinkage rate and prevent it from becoming either too weak or excessively strong. For the TR-CRN module, ρ in (54) controls the maximum offset update in one iteration. Since the first-order off-grid model is valid only in a local angular region, ρ is chosen to be smaller than 1/2. The key hyperparameters for the experiments are presented in Table 1.

Table 1. Default values of the key hyperparameters used in all simulation experiments.

Parameter	Default Value	Function
γ_c	0.8	Shape factor of posterior confidence in (33)
γ_d	1.2	Shape factor of local dominance in (33)
$\tau_{\min}$	0.25	Lower bound of the adaptive shrinkage rate in (34)
$\tau_{\max}$	4.0	Upper bound of the adaptive shrinkage rate in (34)
ρ	0.10	Trust-region step ratio in (54)

The values $\gamma_c = 0.8$ and $\gamma_d = 1.2$ provide moderate shaping of the confidence and local dominance factors. The bounds $\tau_{\min} = 0.25$ and $\tau_{\max} = 4.0$ restrict the adaptive shrinkage range so that unreliable components are suppressed without over-shrinking dominant components. The trust-region ratio $\rho = 0.10$ gives a conservative bounded offset step and helps keep the first-order approximation within its local valid region. The remaining implementation constants are fixed as $\lambda = 10^{-3}$, $\kappa_\sigma = 1$, $\eta_\alpha = 0.8$, $\eta_{\min} = 0.1$, $\eta_{\max} = 0.8$, $\delta = 10^{-6}$, $\sigma_c = 10^{-2}$, $a_0 = b_0 = 10^{-6}$, and $\varepsilon = 10^{-12}$. These constants are used for numerical scaling, damping, diagonal loading, cubic regularization, precision-update initialization, and numerical safeguards. They are kept unchanged in all experiments and are not tuned for individual cases.

3.8. Computational Complexity Analysis

Let T denote the average number of outer iterations. In each iteration, the main computational cost comes from MMV-SBL posterior inference. Constructing the data covariance matrix $\mathbf{C} = \Phi \Gamma \Phi^H + \alpha_0^{-1} \mathbf{I}_M$ requires approximately $O(M^2N)$ operations. Factorizing the

$M \times M$ matrix $\mathbf{C}$, or solving the corresponding linear systems, requires approximately $O(M^3)$ operations. Once the factorization of $\mathbf{C}$ is available, computing the posterior mean matrix $\mathbf{M} = \mathbf{\Gamma}\mathbf{\Phi}^H\mathbf{C}^{-1}\mathbf{X}$ requires approximately $O(M^2L + MNL)$ operations. Computing the diagonal posterior variances required by q_n requires approximately $O(MN)$ operations when only the diagonal entries of $\mathbf{\Sigma}$ are evaluated.

The CGBS-MAP relevance update mainly consists of element-wise operations over the N grid points and therefore introduces approximately $O(N)$ additional cost when the posterior statistics are already available. The PCG-DNU update uses the posterior residual energy and the relevance concentration score. Apart from residual evaluation, its additional damping operation is scalar and has low computational cost. For TR-CRN, the active support size is K_a . Forming the active-support quadratic terms requires approximately $O(K_a^2ML + K_a^2M)$ operations, and the coordinate-wise cubic-regularized update requires $O(K_a)$ operations. Since $K_a \ll N$ in the experiments, the TR-CRN overhead is limited to the selected active support rather than the full angular grid.

Combining the dominant terms, the approximate per-run complexity of the proposed method is

$$O\left(T\left(M^2N + M^3 + M^2L + MNL + K_a^2ML\right)\right). \quad (56)$$

Thus, the proposed method has a higher computational burden than greedy or root-based baselines, mainly because it retains iterative MMV Bayesian posterior inference and additionally performs reliability-guided stabilization and active-support off-grid refinement.

4. Simulation Results and Discussion

This section evaluates the proposed method through eight controlled experiments: off-grid deviation evaluation, SNR sweep, snapshot sweep, hyperparameter sensitivity analysis, source-separation evaluation, random-angle generalization, ablation study, and runtime comparison. These experiments examine off-grid compensation accuracy, noise-level sensitivity, multisnapshot performance, parameter robustness, scenario generalization, module contribution, and computational burden under a consistent simulation protocol.

4.1. Performance Metrics and Experimental Settings

The main experimental settings are summarized in Table 2. For each experimental setting, all compared methods are evaluated under the same Monte Carlo protocol.

Table 2. Main experimental settings.

Item	Setting
Array geometry	Half-wavelength-spaced ULA, $M = 8$
Angular range/grid	$[-90^\circ, 90^\circ], 1^\circ$
Grid points	$N = 181$
Source scenario	Two uncorrelated far-field narrowband sources
True DOAs/source number	$[-15.24^\circ, 15.12^\circ], K = 2$
Noise model	Circularly symmetric complex Gaussian
Monte Carlo trials	$P = 100$
Maximum iterations/tolerance of iterative methods	100, 10^{-4}
Active support size	$K_a = 2K$

Unless otherwise specified, the compared baselines include MUSIC, Root-MUSIC, ESPRIT, OMP, OGSBL, RootSBL, and OGFVBI, while several supplementary experiments report a representative subset for clarity. The SNR is defined as

$$\text{SNR} = 10 \log_{10} \frac{\|\mathbf{A}(\boldsymbol{\theta})\mathbf{S}\|_F^2}{\|\mathbf{N}\|_F^2}. \quad (57)$$

For each target SNR, the complex Gaussian noise variance is adjusted according to (57). Unless otherwise stated, the sources have equal powers and mutually independent waveforms.

RMSE and success rate are used to evaluate estimation accuracy and localization reliability. Since the estimated DOAs may not be ordered consistently with the true DOAs, the estimates are first matched to the true directions by the minimum total angular error in each Monte Carlo trial. All DOA estimates and angular errors are measured in degrees. The RMSE is defined as

$$\text{RMSE} = \sqrt{\frac{1}{KP} \sum_{r=1}^P \sum_{k=1}^K (\hat{\theta}_k^{(r)} - \theta_k)^2} \quad (58)$$

where K is the number of sources, P is the number of Monte Carlo trials, θ_k is the true DOA, and $\hat{\theta}_k^{(r)}$ is the matched estimate in the r th trial.

A trial is regarded as successful only when all matched DOA errors are smaller than a prescribed tolerance τ . The success rate is defined as

$$\text{Success Rate} = \frac{1}{P} \sum_{r=1}^P \mathbb{I} \left(\max_{1 \leq k \leq K} |\hat{\theta}_k^{(r)} - \theta_k| < \tau \right) \times 100\% \quad (59)$$

where $\mathbb{I}(\cdot)$ is the indicator function. Unless otherwise stated, $\tau = 0.5^\circ$, which is stricter than the 1° grid interval. For the single-source off-grid deviation experiment, the absolute error $e_\theta = |\hat{\theta} - \theta_{\text{true}}|$ is reported in degrees.

4.2. Off-Grid Deviation Evaluation

To directly evaluate the off-grid compensation capability of different methods, a single-source off-grid deviation experiment is conducted. The same $M = 8$ half-wavelength ULA and 1° angular grid are used. The nearest grid point is fixed at $\theta_{\text{grid}} = 0^\circ$, and the true DOA is moved away from this grid point by different fractions of the grid interval. The normalized off-grid deviation ratio is defined as

$$\rho_{\text{off}} = \frac{|\theta_{\text{true}} - \theta_{\text{grid}}|}{\Delta_g} \quad (60)$$

where $\Delta_g = 1^\circ$ is the grid interval. In this experiment, ρ_{off} is set to 0, 0.1, 0.2, 0.3, and 0.4, corresponding to true DOAs of 0° , 0.1° , 0.2° , 0.3° , and 0.4° , respectively. The number of snapshots is fixed at $L = 100$, the SNR is fixed at 5 dB, and $K = 1$ is used only for this deviation test.

For each value of ρ_{off} , Table 3 reports the absolute estimation error

$$e_\theta = |\hat{\theta} - \theta_{\text{true}}| \quad (61)$$

The left column denotes the normalized off-grid deviation ratio, and the remaining columns give the estimation errors of different algorithms in degrees.

Table 3 shows that the proposed method achieves the lowest average error among the compared methods in this illustrative single-source setting, while RootSBL also provides very close average performance. It should also be noted that the proposed method is not uniformly best at every individual deviation ratio; RootSBL and OGFVBI obtain smaller errors at some points. Therefore, this experiment is interpreted as a direct demonstration of off-grid correction behavior rather than a complete proof of general off-grid superiority. In contrast, OMP is restricted to the fixed angular grid, so its error increases as the true DOA

moves away from the grid point. The subsequent source-separation and random-angle experiments further complement this single-source off-grid test by examining multi-source and randomly generated angular scenarios.

Table 3. Absolute estimation errors in degrees under different normalized off-grid deviation ratios.

Normalized Off-Grid Deviation ρ_{off}	Proposed	OGSBL	RootSBL	OGFVBI	OMP
0.00	0.03422	0.07218	0.04170	0.01137	0.00000
0.10	0.04723	0.09947	0.03941	0.06613	0.10000
0.20	0.04286	0.09319	0.03706	0.05304	0.20000
0.30	0.04045	0.07718	0.03466	0.02446	0.30000
0.40	0.01875	0.06118	0.03224	0.09474	0.40000
Average	0.03670	0.08064	0.03701	0.04995	0.20000

4.3. Performance Under Varying SNR

The SNR experiment evaluates the estimation accuracy and localization reliability of different methods under varying noise levels. Unless otherwise stated, the common settings in Table 2 are used. In this experiment, the number of snapshots is fixed at $L = 200$, while the SNR is varied from -10 dB to 10 dB. Each point is averaged over $P = 100$ Monte Carlo trials, and the success threshold is set to 0.5° . This experiment examines whether each method can maintain accurate and reliable off-grid localization as the observation noise level changes.

Figure 5a shows that the RMSE of most methods decreases as the SNR increases. OMP exhibits an error floor because it is restricted to the fixed angular grid. OGSBL, RootSBL, and OGFVBI improve the estimation accuracy through off-grid refinement or Bayesian inference, while MUSIC, Root-MUSIC, and ESPRIT provide classical subspace-based references. The CNN baseline also shows improved accuracy as the SNR increases, but its performance is limited by the grid-based output and is less stable in the low-SNR region. The proposed method achieves lower RMSE at several SNR points, especially in the low-to-medium SNR region.

Figure 5b shows that the success rates generally increase with SNR. The CNN baseline has a relatively low success rate at low SNRs but improves significantly as the SNR increases. The proposed method obtains a higher success rate at most low-to-medium SNR points, while several methods approach high success rates as the SNR increases.




Figure 5. Performance comparison under varying SNR. (a) RMSE versus SNR. (b) Success rate versus SNR. The number of snapshots is fixed at $L = 200$, and each point is averaged over 100 Monte Carlo trials. A trial is counted as successful only when all matched DOA errors are smaller than 0.5° .

This behavior is consistent with the design of the reliability-guided inference chain: CGBS-MAP improves relevance concentration, PCG-DNU stabilizes effective error precision learning, and TR-CRN provides bounded active-support offset correction.

4.4. Performance Under Varying Snapshot Numbers

This experiment investigates the influence of the number of snapshots on estimation accuracy and localization reliability. The common settings in Table 2 are used, with the SNR fixed at 5 dB. The number of snapshots is varied from 50 to 500, and each point is averaged over $P = 100$ Monte Carlo trials. This setting evaluates whether each method can maintain reliable off-grid localization under different amounts of available temporal observations.

Figure 6a shows that the RMSE values of most methods decrease as the number of snapshots increases, indicating that additional multisnapshot data improve estimation reliability. OMP remains limited by fixed-grid mismatch, whereas the off-grid Bayesian methods benefit from continuous-angle refinement. The CNN baseline also benefits from the increase in snapshots because a larger number of snapshots provides a more stable sample covariance matrix. The proposed method maintains low RMSE over the tested snapshot range, showing stable estimation accuracy under different data lengths.

Figure 6b presents the corresponding success-rate curves. The CNN baseline shows an increasing success rate as the number of snapshots increases, but its reliability is lower when the number of snapshots is small. The proposed method maintains high localization reliability over the tested snapshot range and remains competitive with RootSBL, Root-MUSIC, and OGFVBI. This behavior is consistent with the reliability-guided inference chain: CGBS-MAP improves support concentration, PCG-DNU stabilizes effective error precision learning, and TR-CRN refines the offsets within a bounded local region.




Figure 6. Performance comparison under varying snapshot numbers. (a) RMSE versus the number of snapshots. (b) Success rate versus the number of snapshots. The SNR is fixed at 5 dB, and each point is averaged over 100 Monte Carlo trials. The success threshold is 0.5° .

4.5. Hyperparameter Sensitivity Analysis

To examine the influence of the key hyperparameters, a one-factor-at-a-time sensitivity analysis is conducted. The basic setting is $M = 8$, $K = 2$, $\theta = [-15.24^\circ, 15.12^\circ]$, $L = 200$, Gaussian noise, and $P = 200$ Monte Carlo trials. For each test, only one parameter is varied while the other parameters are fixed at the default values in Table 1. The tested hyperparameters include γ_c , γ_d , $\tau_{\min}$, $\tau_{\max}$, and ρ . For each hyperparameter value, the SNR is varied from -10 dB to 10 dB.

Figure 7 shows the sensitivity analysis of the key hyperparameters. In each subfigure, one parameter is varied while the remaining parameters are fixed at their default values. The RMSE curves show similar decreasing trends as the SNR increases, indicating that the proposed method does not exhibit severe sensitivity to moderate variations of a single hyperparameter. The default parameter settings generally provide stable and competitive estimation performance over the tested SNR range, without requiring case-specific tuning.




Figure 7. Hyperparameter sensitivity analysis of the proposed method. (a) γ_c ; (b) γ_d ; (c) $\tau_{\min}$; (d) $\tau_{\max}$; (e) ρ .

4.6. Performance Under Different Source Separations

To further evaluate the resolution capability of the compared methods, a two-source experiment with different angular separations is conducted. In this experiment, the SNR is fixed at 5 dB, the number of snapshots is fixed at $L = 200$, the grid interval is 1° , and each point is averaged over $P = 200$ Monte Carlo trials.

Figure 8 shows the estimation performance under different source separations. When the two sources are closely spaced, the steering vectors become highly correlated, and the DOA estimation problem becomes more challenging. As the angular separation increases, the source distinguishability improves, and most methods achieve lower estimation errors. The proposed method maintains stable performance over the tested separation range, indicating that the reliability-guided sparse Bayesian inference and bounded off-grid refinement can effectively support multi-source localization under different angular configurations.




Figure 8. Performance comparison under different source separations. The SNR is fixed at 5 dB, the number of snapshots is fixed at $L = 200$, and each point is averaged over 200 Monte Carlo trials.

4.7. Random-Angle Generalization Experiment

To avoid evaluating the algorithms only at a fixed DOA pair, an additional random-angle experiment is conducted. In each Monte Carlo trial, the two true DOAs are independently generated from two separated angular intervals:

$$\theta_1 \sim \mathcal{U}(-30^\circ, -10^\circ), \quad \theta_2 \sim \mathcal{U}(10^\circ, 30^\circ) \quad (62)$$

The SNR is fixed at 5 dB, the number of snapshots is fixed at $L = 200$, and $P = 100$ Monte Carlo trials are used.

Table 4 shows that the proposed method achieves a competitive RMSE of 0.1204° with a 100% success rate in the random-angle scenario. Root-MUSIC and RootSBL achieve slightly lower RMSEs because the sources are well separated and the ULA model is exactly satisfied. Nevertheless, the proposed method remains very close to these strong baselines and achieves lower RMSE than OGSBL, ESPRIT, and OGFVBI in this test. This result suggests that the proposed method is not only effective for a fixed DOA pair but also remains stable under randomly generated angular scenarios.

Table 4. Random-angle generalization results. The two DOAs are randomly generated from $[-30^\circ, -10^\circ]$ and $[10^\circ, 30^\circ]$ in each Monte Carlo trial.

Method	RMSE (Deg)	Success Rate (%)
Proposed	0.1204	100.0
OGSBL	0.1657	100.0
Root-MUSIC	0.1190	100.0
ESPRIT	0.1527	100.0
OGFVBI	0.1367	99.5
RootSBL	0.1176	100.0

4.8. Ablation Study

To examine the contribution of each proposed inference module, an ablation study is conducted under the same basic setting as the SNR experiment. The SNR is varied from -10 dB to 10 dB, the number of snapshots is fixed at $L = 200$, and each point is averaged over $P = 100$ Monte Carlo trials. Four variants are compared: the complete method, denoted as Proposed-full; the method without CGBS-MAP, denoted as w/o CGBS-MAP; the method without PCG-DNU, denoted as w/o PCG-DNU; and the method without TR-CRN, denoted as w/o TR-CRN.

In the w/o CGBS-MAP variant, the sparse relevance update is replaced by the standard posterior second-moment update. In the w/o PCG-DNU variant, the scalar effective error precision is directly updated by the residual-based estimate $\alpha_{0,\text{raw}}$ without concentration-guided damping. In the w/o TR-CRN variant, the cubic-regularized trust-region refinement is removed and the offset correction is performed by a regularized direct update.

Figure 9a shows that removing any proposed module generally increases the RMSE at several SNR points. Without CGBS-MAP, weak neighboring activations are less effectively suppressed, which may affect support selection. Without PCG-DNU, the effective error precision update becomes more sensitive to uncertain support formation. Without TR-CRN, the offset correction loses the bounded cubic-regularized refinement and may become less stable. These results indicate that the three modules contribute to different stages of the inference chain.

Figure 9b shows that the complete method maintains a higher or comparable success rate at most SNR points. This means that the proposed modules improve not only the average estimation error, but also the repeatability of successful localization. Overall, the ablation results suggest that the observed performance gain is related to the combined effect

of reliability-guided relevance shrinkage, concentration-guided effective error precision relaxation, and bounded off-grid refinement.




Figure 9. Ablation study under varying SNR. (a) RMSE versus SNR. (b) Success rate versus SNR. The complete method is compared with variants that remove CGBS-MAP, PCG-DNU, or TR-CRN. The success threshold is 0.5° .

4.9. Computational Runtime Comparison

All simulations were performed on a computer with a 13th Gen Intel(R) Core(TM) i7-13650HX CPU and 24.0 GB RAM using MATLAB R2023b under the same framework without parallel computation. The running times are listed in Table 5.

Table 5. Average runtime under the SNR sweep.

Method	Average Runtime (s)
RootSBL	0.1207
OGSBL	0.3200
OGFVBI	0.3908
Proposed	0.4152

The proposed method has the largest average runtime among the compared methods because it retains iterative MMV Bayesian posterior inference and additionally performs reliability-guided relevance shrinkage (CGBS-MAP), damped effective error precision updating (PCG-DNU), and active-support off-grid refinement (TR-CRN). Compared with other SBL-based baselines, the proposed method uses a larger iteration budget, which is reflected in the computational cost.

For the iterative SBL-based methods, the observed iteration records indicate that most algorithms reached their prescribed maximum iteration numbers under the strict convergence tolerance. Therefore, the runtime values primarily reflect the preset iteration budgets rather than natural early-stopping convergence. The main computational burden comes from the MMV-SBL posterior inference step, including covariance construction, matrix factorization, and posterior mean computation. Since TR-CRN is applied only to the selected active support and PCG-DNU is a scalar relaxation step, their additional overhead is relatively limited.

4.10. Discussion

The simulation results indicate that the proposed method can provide competitive and stable off-grid DOA estimation in the tested controlled scenarios by jointly enhancing support formation, effective error precision learning, and offset refinement. The off-grid deviation, SNR, and snapshot experiments show stable off-grid localization behavior under

the considered mismatch, noise-level, and data-length conditions. The hyperparameter sensitivity analysis shows that the selected default parameter set provides a stable operating point, while moderate parameter variations do not lead to algorithmic breakdown. The source-separation and random-angle experiments supplement the fixed-angle tests by evaluating different angular configurations and randomly generated DOA pairs. The ablation study indicates that CGBS-MAP, PCG-DNU, and TR-CRN contribute to different stages of the inference chain, so the observed performance improvement is related to their coordinated effect rather than to a single isolated modification. In particular, improved relevance concentration provides a more reliable active support for subsequent bounded offset correction.

The improved accuracy and reliability are obtained at the cost of increased computational burden. This is mainly because the proposed method retains iterative MMV-SBL posterior inference and adds reliability-guided relevance shrinkage, damped effective error precision learning, and active-support offset refinement. However, the additional cost is mainly associated with improved inference stability and continuous-angle correction. Therefore, the proposed method is more suitable for scenarios where off-grid estimation accuracy and estimation stability are more important than minimum runtime.

5. Conclusions

This paper proposed a reliability-guided stabilized off-grid SBL method for multisnapshot DOA estimation within a first-order MMV Bayesian model. The method stabilizes three key inference steps: sparse relevance learning, scalar effective error precision updating, and active-support off-grid refinement. Specifically, CGBS-MAP enhances sparse relevance concentration by suppressing weak and non-dominant posterior components, PCG-DNU stabilizes effective error precision learning through concentration-guided damping, and TR-CRN provides bounded continuous-angle offset correction with cubic regularization and trust-region constraints. The proposed method is evaluated through off-grid deviation, SNR, snapshot, hyperparameter sensitivity, source-separation, random-angle, ablation, and runtime experiments. Simulation results show that the proposed method achieves competitive DOA estimation performance compared with representative baselines in the tested settings. The ablation results support the contribution of the three modules, while the runtime and theoretical complexity analyses indicate an accuracy–complexity tradeoff with increased computational burden. The current study is still limited to controlled simulation scenarios with known source number, uncorrelated sources, equal source powers, and ideal array calibration. Future work will consider extensions to more challenging scenarios, such as coherent sources, unequal source powers, unknown source number estimation, array model imperfections, and more systematic learning-based comparisons.

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

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data reported in this study are generated from numerical simulations and can be reproduced using the model parameters and experimental settings described in the manuscript. Additional simulation data are available from the corresponding author upon reasonable request.

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

Abbreviations

DOA	Direction of Arrival
SBL	Sparse Bayesian Learning
MMV	Multiple Measurement Vector
MAP	Maximum a Posteriori
OMP	Orthogonal Matching Pursuit
OGFVBI	Off-Grid Fast Variational Bayesian Inference
SNR	Signal-to-Noise Ratio
RMSE	Root Mean Square Error
CGBS-MAP	Confidence-Guided MAP-Type Shrinkage Relevance Learning
PCG-DNU	Posterior-Concentration-Guided Damped Noise Update
TR-CRN	Trust-Region Cubic-Regularized Newton
CRN	Cubic-Regularized Newton

References

1. Wang, H.; Bai, Y.; Ren, J.; Wang, P.; Xu, T.; Zhang, W.; Zhang, G. DOA Estimation Method for Vector Hydrophones Based on Sparse Bayesian Learning. *Sensors* **2024**, *24*, 6439. [\[CrossRef\]](#) [\[PubMed\]](#)
2. Li, N.; Zhang, X.; Lv, F.; Zong, B. Direction-of-Arrival Estimation via Sparse Bayesian Learning Exploiting Hierarchical Priors with Low Complexity. *Sensors* **2024**, *24*, 2336. [\[CrossRef\]](#) [\[PubMed\]](#)
3. Mecklenbräuker, C.F.; Gerstoft, P.; Ollila, E.; Park, Y. Robust and Sparse M-Estimation of DOA. *Signal Process.* **2024**, *220*, 109461. [\[CrossRef\]](#)
4. Huang, Y.; Zhang, Y.; Tao, J.; Wen, C.; Liao, G.; Hong, W. Off-Grid DOA Estimation via a Deep Learning Framework. *Sci. China Inf. Sci.* **2023**, *66*, 222305. [\[CrossRef\]](#)
5. Capon, J. High-Resolution Frequency-Wavenumber Spectrum Analysis. *Proc. IEEE* **1969**, *57*, 1408–1418. [\[CrossRef\]](#)
6. Schmidt, R.O. Multiple Emitter Location and Signal Parameter Estimation. *IEEE Trans. Antennas Propag.* **1986**, *34*, 276–280. [\[CrossRef\]](#)
7. Roy, R.; Kailath, T. ESPRIT—Estimation of Signal Parameters via Rotational Invariance Techniques. *IEEE Trans. Acoust. Speech Signal Process.* **1989**, *37*, 984–995. [\[CrossRef\]](#)
8. Stoica, P.; Nehorai, A. MUSIC, Maximum Likelihood, and Cramer-Rao Bound. *IEEE Trans. Acoust. Speech Signal Process.* **1989**, *37*, 720–741. [\[CrossRef\]](#)
9. Malioutov, D.; Cetin, M.; Willsky, A.S. A Sparse Signal Reconstruction Perspective for Source Localization with Sensor Arrays. *IEEE Trans. Signal Process.* **2005**, *53*, 3010–3022. [\[CrossRef\]](#)
10. Tropp, J.A.; Gilbert, A.C. Signal Recovery from Random Measurements via Orthogonal Matching Pursuit. *IEEE Trans. Inf. Theory* **2007**, *53*, 4655–4666. [\[CrossRef\]](#)
11. Tipping, M.E. Sparse Bayesian Learning and the Relevance Vector Machine. *J. Mach. Learn. Res.* **2001**, *1*, 211–244.
12. Wipf, D.P.; Rao, B.D. Sparse Bayesian Learning for Basis Selection. *IEEE Trans. Signal Process.* **2004**, *52*, 2153–2164. [\[CrossRef\]](#)
13. Ji, S.; Xue, Y.; Carin, L. Bayesian Compressive Sensing. *IEEE Trans. Signal Process.* **2008**, *56*, 2346–2356. [\[CrossRef\]](#)
14. Yang, Z.; Xie, L.; Zhang, C. Off-Grid Direction of Arrival Estimation Using Sparse Bayesian Inference. *IEEE Trans. Signal Process.* **2013**, *61*, 38–43. [\[CrossRef\]](#)
15. Dai, J.; Bao, X.; Xu, W.; Chang, C. Root Sparse Bayesian Learning for Off-Grid DOA Estimation. *IEEE Signal Process. Lett.* **2017**, *24*, 46–50. [\[CrossRef\]](#)
16. Wen, C.; Shi, G.; Zhang, X. Off-Grid DOA Estimation Under Nonuniform Noise via Variational Sparse Bayesian Inference. *Signal Process.* **2017**, *137*, 69–79. [\[CrossRef\]](#)
17. Yang, J.; Yang, Y. Sparse Bayesian DOA Estimation Using Hierarchical Synthesis Lasso Priors for Off-Grid Signals. *IEEE Trans. Signal Process.* **2020**, *68*, 872–884. [\[CrossRef\]](#)
18. Shen, J.; Liu, T.; Guo, X.; Wang, Y. Off-Grid DOA Estimation Using Improved Root Sparse Bayesian Learning. *EURASIP J. Adv. Signal Process.* **2023**, *2023*, 68. [\[CrossRef\]](#)
19. Li, N.; Zhang, X.; Zong, B.; Lv, F.; Xu, J.; Wang, Z. An Off-Grid Direction-of-Arrival Estimator Based on Sparse Bayesian Learning with Three-Stage Hierarchical Laplace Priors. *Signal Process.* **2024**, *218*, 109371. [\[CrossRef\]](#)

20. Tong, X.; Chen, L.; Liu, Y.; Zhang, J. An Off-Grid DOA Estimation Method via Fast Variational Bayesian Inference. *Electronics* **2025**, *14*, 2781. [[CrossRef](#)]
21. Gong, Z.; Su, X.; Hu, P.; Liu, S.; Liu, Z. Deep Unfolding Sparse Bayesian Learning Network for Off-Grid DOA Estimation with Nested Array. *Remote Sens.* **2023**, *15*, 5320. [[CrossRef](#)]

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.