Nucl. Phys. B 1026 (2026) 117453

Contents lists available at ScienceDirect

Nuclear Physics, Section B

journal homepage: www.elsevier.com/locate/nuclphysb

Bayesian polarimetric multi-source direction-of-arrival estimation for transient astronomy with sparse radio interferometric subarrays

Halim Zeghdoudi[a,*], Fatih Tank[b]

[a] LaPS Laboratory, Badji Mokhtar-Annaba University, Annaba, Algeria
[b] Department of Economics, Atilim University, Ankara, Turkey

ARTICLE INFO

Editor: Olalla Castro Alvaredo

Keywords:
Radio astronomy
Direction-of-arrival
Bayesian inference
Polarization
Fast radio bursts
RFI localization

ABSTRACT

Modern radio interferometers are increasingly challenged by fast transient events, complex radio-frequency interference (RFI), and observing conditions in which near-field and far-field emitters may coexist. Although classical direction-of-arrival (DOA) techniques can achieve high angular resolution, they are often developed for specific array geometries, tend to focus on single-source settings, and usually provide little information about uncertainty. Imaging-based methods, while powerful, are computationally demanding and can introduce delays that are not well suited to real-time transient astronomy.

In this work, we propose a Bayesian, polarization-aware framework for multi-source DOA estimation in arbitrary radio interferometric arrays. Starting from baseline-level covariance modeling and polarization-sensitive phase information, we build a probabilistic formulation that jointly infers source direction, polarization state, and, when relevant, source range. Because interferometric phase is inherently wrapped, uncertainty is modeled explicitly using circular statistical distributions, and posterior inference is carried out through a variational Bayesian scheme that remains computationally efficient.

Compared with deterministic or purely data-driven approaches, the proposed method offers a more physically grounded and statistically interpretable alternative. It incorporates array geometry, polarization structure, and prior astronomical knowledge directly into the inference process, while also delivering calibrated uncertainty estimates for source localization. Simulations using realistic LOFAR and SKA-Low configurations show robust multi-source separation, stable performance across wide bandwidths, and improved resilience in low signal-to-noise and near-field conditions.

Overall, the proposed framework enables imaging-free, uncertainty-aware localization of fast radio bursts, solar radio emission, and terrestrial RFI. It provides a statistically principled and computationally practical route toward real-time transient localization in next-generation radio observatories.

---

* Corresponding author.
E-mail address: halim.zeghdoudi@univ-annaba.dz (H. Zeghdoudi).

https://doi.org/10.1016/j.nuclphysb.2026.117453
Received 13 March 2026; Received in revised form 28 March 2026; Accepted 8 April 2026

1. Introduction

Radio interferometry is entering a new observational era. Modern arrays such as LOFAR [1] and the Square Kilometre Array (SKA) [2] combine large numbers of antennas, wide bandwidths, and high sensitivity, enabling the detection of fast radio bursts (FRBs), solar radio transients, pulsar emission, and a broad range of short-duration astrophysical phenomena. At the same time, these systems must operate in increasingly challenging radio-frequency interference (RFI) environments. In this context, rapid and reliable source localization has become a central requirement.

Classical direction-of-arrival (DOA) estimation methods such as MUSIC [3] and ESPRIT [4] achieve high angular resolution by exploiting covariance structure and signal subspaces. Subsequent developments, including EM/SAGE-based methods [5] and sparse reconstruction techniques [6,7], have extended these approaches to more challenging multi-source settings. More recently, deep learning methods have shown promising performance in controlled DOA estimation tasks [8–13].

However, most of these methods were developed for array geometries that are very different from those used in radio astronomy. Many assume uniform linear or otherwise highly structured layouts, and most are derived under far-field conditions. By contrast, interferometric arrays are designed primarily for Fourier-domain sampling and image formation rather than for direct DOA estimation. Imaging-based localization remains highly accurate, but it can also be computationally intensive and may introduce latency that is undesirable for rapid transient follow-up.

To reduce this gap, several imaging-free or visibility-domain localization strategies have been proposed for interferometric systems. Spatial filtering and subspace-projection techniques have been used for interference mitigation [14,15]. Visibility-domain optimization has also been applied to DOA estimation with arbitrary two-dimensional array layouts [16], while near-field interferometric corrections have been explored for altitude estimation of RFI emitters [17]. Particularly relevant to the present work is the recent baseline-based, polarization-sensitive interferometric DOA framework introduced in Yatawatta [18], which combines phase-difference projection (PDP) for phase unwrapping [19,20] with an ESPRIT-type estimator. That study showed convincingly that imaging-free DOA recovery is feasible from a selected subset of interferometric baselines.

The present work is closely related in spirit to that baseline-level formulation, but it is motivated by several limitations that remain open in current interferometric DOA methods. First, many existing approaches are most effective in the presence of a single dominant emitter, whereas practical radio observations often contain multiple simultaneous sources, such as overlapping solar burst components, multiple RFI emitters, or transient astronomical sources embedded in interference. Second, most available methods return deterministic point estimates without associated uncertainty measures. For astronomical applications, however, credible localization regions are often as important as the estimates themselves, especially for triggering multi-wavelength follow-up observations and for downstream statistical interpretation. Third, near-field and mixed-field effects are becoming increasingly relevant in modern observing environments, particularly in the presence of local interferers and propagation-induced distortions [21–24].

These challenges suggest the need for a probabilistic framework that is both physically grounded and computationally practical. Bayesian inference offers a natural way to incorporate prior information, model wrapped phase behavior through circular statistics, and return posterior distributions rather than single-point estimates. Such a formulation is attractive for interferometric localization because it can account explicitly for phase uncertainty, source multiplicity, polarization structure, and geometric constraints within a single inference procedure. More broadly, recent developments in statistical modeling for astronomy have shown the value of probabilistic methods in complex observational settings [25].

Motivated by these considerations, this paper develops a Bayesian, polarization-aware framework for multi-source direction-of-arrival estimation with sparse radio interferometric subarrays. Starting from baseline-level polarimetric statistics, we introduce a wrapped-phase likelihood based on circular distributions and derive a variational Bayesian approximation for joint inference of source directions, polarization parameters, and, when needed, source ranges. In the present study, the number of sources is assumed known in the controlled experiments, so that the focus remains on inference quality, uncertainty quantification, and robustness; automatic model-order selection is left for future work.

Compared with deterministic subspace-based estimators, the proposed approach is designed to support multiple sources, quantify uncertainty explicitly, and extend naturally to near-field settings. Compared with purely data-driven methods, it retains direct dependence on array geometry, interferometric phase structure, and physically interpretable prior assumptions. The result is an uncertainty-aware localization framework that is intended for rapid, imaging-light analysis of transient and interfered radio scenes.

The remainder of the paper is organized as follows. Section 2 presents the multi-source polarimetric signal model and the associated baseline covariance formulation. Section 3 introduces the wrapped-phase likelihood and the Bayesian model. Section 4 develops the variational approximation, including the near-field extension and the main computational aspects. Section 5 reports simulation results for realistic LOFAR-like and SKA-Low-like array configurations. Section 6 discusses representative applications, including FRB localization, solar radio burst separation, and near-field RFI analysis. Finally, Section 7 summarizes the main conclusions, outlines the main limitations of the present study, and highlights directions for future work.

2. Signal model

We consider a radio interferometric array consisting of P dual-polarized receiving elements located at positions $\mathbf{x}_p \in \mathbb{R}^3$, for $p = 1, \ldots, P$. Each element records two orthogonal polarization components, so that the received signal at antenna p is naturally represented as a complex vector in $\mathbb{C}^2$. The interferometric geometry considered in this work is illustrated in Fig. 1, where the projection of a baseline onto the incident direction determines the corresponding phase difference.

Assume that the array observes K electromagnetic sources. For the k-th source, the unknown parameters are its direction of arrival (DOA), denoted by (θ_k, ϕ_k), its intrinsic polarization parameters (γ_k, η_k), and its narrowband complex baseband signal $s_k(t)$. In the present study, K is treated as known in the controlled experiments, so that the focus remains on multi-source inference, uncertainty quantification, and robustness. Automatic estimation of the number of active sources is left for future work.

The unit vector associated with the arrival direction of source k is defined as

$$\hat{\mathbf{s}}_k = \begin{bmatrix} \cos\theta_k \cos\phi_k \\ \cos\theta_k \sin\phi_k \\ \sin\theta_k \end{bmatrix}. \tag{1}$$

This baseline-level geometric formulation is closely related in spirit to recent interferometric DOA approaches that operate directly on visibility-domain phase information, notably the baseline-based polarimetric framework of [18]. The difference here is that the signal model is developed as the front end of a Bayesian multi-source inference scheme, rather than for a deterministic subspace estimator.

2.1. Polarimetric voltage model

The instantaneous voltage measured at antenna p is modeled as

$$\mathbf{v}_p(t) = \sum_{k=1}^{K} a_p(\theta_k, \phi_k)\, \mathbf{E}_p(\theta_k, \phi_k)\, \mathbf{p}_k\, s_k(t) + \mathbf{n}_p(t), \tag{2}$$

where $a_p(\theta_k, \phi_k)$ denotes the propagation factor, $\mathbf{E}_p(\theta_k, \phi_k) \in \mathbb{C}^{2\times 2}$ is the direction-dependent antenna response matrix, $\mathbf{p}_k$ is the polarization vector of the k-th source, and $\mathbf{n}_p(t)$ represents additive receiver noise.

Under the far-field assumption, the incident wavefront is approximately planar across the array, and the propagation factor takes the form

$$a_p(\theta_k, \phi_k) = \exp\left(j \frac{2\pi}{\lambda}\, \mathbf{x}_p^{\top} \hat{\mathbf{s}}_k \right), \tag{3}$$

where λ is the observing wavelength.

For a near-field source located at range R_k, the planar-wave approximation is no longer valid. In that case, the propagation term is written as

$$a_p^{\mathrm{NF}} = \exp\left(j \frac{2\pi}{\lambda} \left\| R_k \hat{\mathbf{s}}_k - \mathbf{x}_p \right\| \right), \tag{4}$$

which accounts for wavefront curvature across the array and introduces range-dependent phase variations across baselines. This extension is important for local emitters and radio-frequency interference, where near-field effects cannot be neglected.

2.2. Polarization representation

The intrinsic polarization state of source k is parameterized as

$$\mathbf{p}_k = \begin{bmatrix} \cos\gamma_k \\ \sin\gamma_k\, e^{j\eta_k} \end{bmatrix}, \tag{5}$$

where γ_k controls the relative amplitude of the two orthogonal polarization components, while η_k represents their relative phase.

This representation allows the direction and polarization state to be inferred jointly within a unified model. Such joint treatment is particularly useful when polarization provides an additional means of discriminating between partially overlapping sources.

2.3. Baseline covariance structure

For a baseline formed by antennas (p, q), we define the stacked voltage vector

$$\mathbf{v}_{pq}(t) = \begin{bmatrix} \mathbf{v}_p(t) \\ \mathbf{v}_q(t) \end{bmatrix} \in \mathbb{C}^4. \tag{6}$$

Its covariance matrix is then written as

$$\mathbf{R}_{pq} = \mathbb{E}\left\{ \mathbf{v}_{pq}(t) \mathbf{v}_{pq}^{H}(t) \right\}, \tag{7}$$

which captures the second-order statistical structure of the received signals on that baseline.

In the multi-source case, this covariance can be decomposed as

$$\mathbf{R}_{pq} = \sum_{k=1}^{K} \mathbf{R}_{pq}^{(k)} + \sigma^2 \mathbf{I}, \tag{8}$$

where $\mathbf{R}_{pq}^{(k)}$ denotes the contribution of the k-th source and $\sigma^2\mathbf{I}$ models additive receiver noise.

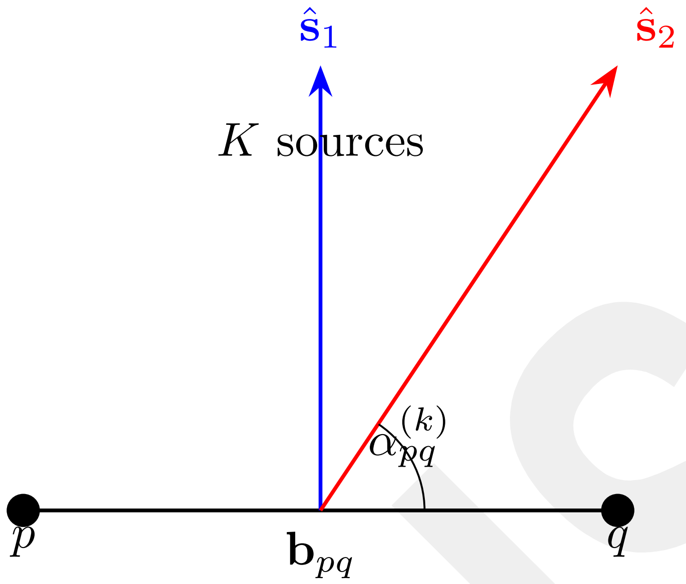

Fig. 1. Geometry of a dual-polarized interferometric array. Antennas p and q form baseline (p, q) and observe K incident sources. The projection of the baseline onto the source direction determines the corresponding interferometric phase difference.

Each source contributes a structured low-rank term governed by its propagation phase, polarization state, and antenna response. The phase relationships encoded in $\mathbf{R}_{pq}$ form the basis of the proposed DOA inference procedure. In particular, the method operates directly on baseline-level polarimetric statistics rather than on an image-domain reconstruction, which makes it suitable for rapid interferometric localization. In practice, inference is carried out on a selected subset of informative baselines rather than on all possible antenna pairs; the exact configurations used for the LOFAR-like and SKA-Low-like experiments are specified later in Section 5.

3. Wrapped phase likelihood

The main interferometric observable used in this work is the phase difference measured on a baseline formed by antennas p and q. Working directly with baseline phases is closely related in spirit to recent visibility-domain formulations for interferometric DOA estimation, including the baseline-based polarimetric framework of Yatawatta [18]. In the present work, however, the objective is not limited to deterministic direction recovery. Instead, we seek a probabilistic formulation that can accommodate multiple simultaneous sources and provide uncertainty information together with the estimated directions.

3.1. Deterministic phase model

For source k, the ideal phase associated with baseline (p, q) is written as

$$\psi_{pq}^{(k)} = \frac{2\pi}{\lambda} b_{pq} \cos \alpha_{pq}^{(k)}, \tag{9}$$

where $b_{pq} = \|\mathbf{x}_p - \mathbf{x}_q\|$ denotes the baseline length, and $\alpha_{pq}^{(k)}$ is the angle between the baseline direction $\hat{\mathbf{b}}_{pq}$ and the source direction $\hat{\mathbf{s}}_k$.

In the single-source case, this relation provides a direct geometric mapping between the measured phase and the direction of arrival. In the multi-source setting, however, the observed visibility on a given baseline is a coherent superposition of contributions from all active emitters. We therefore write

$$V_{pq} = \sum_{k=1}^{K} \rho_k \exp\left(j\psi_{pq}^{(k)}\right), \tag{10}$$

where ρ_k represents the effective contribution of source k on that baseline. This term absorbs source strength together with the net effect of the antenna response and polarimetric weighting.

This formulation makes explicit why the multi-source problem is more difficult than the single-source case: the observed baseline phase is no longer tied to one emitter alone, but to the argument of a complex sum of source-dependent contributions.

3.2. Observed wrapped phase

The measured phase is taken to be the argument of the complex visibility, namely

$$\tilde{\psi}_{pq} = \arg(V_{pq}) + \epsilon_{pq}, \tag{11}$$

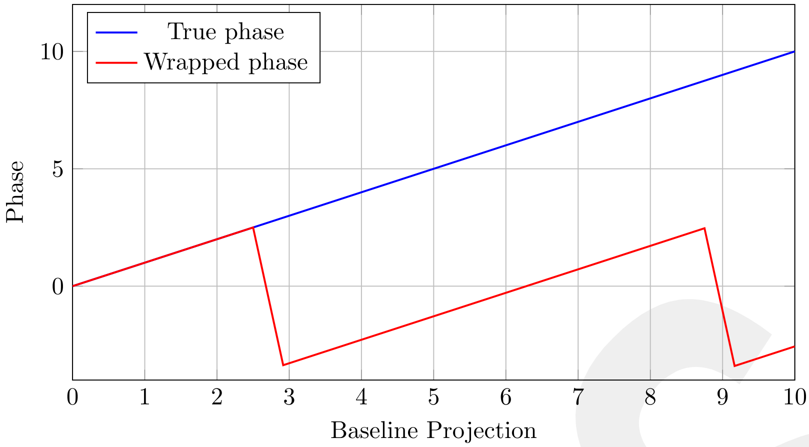

Fig. 2. Illustration of phase wrapping. The underlying phase evolves continuously with baseline projection, whereas the measured phase is restricted to the interval $(-\pi, \pi]$. The circular likelihood used in this work accounts for this periodicity explicitly.

where ϵ_{pq} accounts for thermal noise, residual calibration errors, and other perturbations affecting the phase estimate.

Since phase is intrinsically periodic, the observation satisfies $\tilde{\psi}_{pq} \in (-\pi, \pi]$. This periodicity is central to the estimation problem rather than a secondary implementation detail. When the physical phase extends beyond the principal interval, the observed quantity becomes wrapped, so that distinct geometric configurations may produce numerically similar phase values. The effect is illustrated in Fig. 2, which shows how a continuously varying phase is mapped back into the interval $(-\pi, \pi]$. In such situations, standard Euclidean error models become inappropriate, especially at low signal-to-noise ratio or when several sources contribute simultaneously to the same baseline.

3.3. Circular statistical likelihood

To account for the periodic nature of the phase, we model the wrapped observation using a Von Mises distribution,

$$p(\tilde{\psi}_{pq} \mid \Theta) = \frac{1}{2\pi I_0(\kappa)} \exp\left[\kappa \cos(\tilde{\psi}_{pq} - \mu_{pq})\right], \tag{12}$$

where

$$\mu_{pq} = \arg\left(\sum_{k=1}^{K} \rho_k e^{j\psi_{pq}^{(k)}}\right), \tag{13}$$

κ is the concentration parameter, and $I_0(\kappa)$ is the modified Bessel function of the first kind.

The Von Mises law may be viewed as the circular analogue of a Gaussian distribution. Its advantage in the present context is that it measures phase mismatch on the unit circle rather than on the real line, thereby handling phase wrapping in a statistically consistent way. This is one of the main points at which the proposed Bayesian formulation departs from purely deterministic phase-based estimators, including those that operate on the same baseline-level observables.

3.4. Multi-baseline likelihood

Let $\mathcal{B}$ denote the selected set of baselines used for inference. In practice, the method does not require all possible antenna pairs. Instead, it operates on a subset of informative baselines chosen to keep the inference procedure computationally tractable for sparse interferometric subarrays. The exact number of retained baselines for the LOFAR-like and SKA-Low-like configurations is given later in Section 5.

Assuming conditional independence across the selected baselines, the joint likelihood becomes

$$p(\{\tilde{\psi}_{pq}\}_{(p,q)\in\mathcal{B}} \mid \Theta) = \prod_{(p,q)\in\mathcal{B}} p(\tilde{\psi}_{pq} \mid \Theta). \tag{14}$$

This likelihood forms the probabilistic core of the Bayesian inference procedure developed in the next section. It also makes clear why baseline selection matters not only physically but also computationally: the number of retained baselines directly controls the cost of likelihood evaluation and therefore the cost of posterior approximation. This point is particularly relevant when comparing the proposed method with lighter deterministic approaches such as the baseline-level ESPRIT-type framework of[18].

Table 1
Main simulation settings.

Parameter	Value / Range
Array geometries	LOFAR-like, SKA-Low-like
Number of sources K	1–5
Frequency range	50–200 MHz
Snapshots per realization	1000
SNR	-10 to 20 dB
Relative source power	0 to -6 dB
Near-field fraction	30%
Near-field range	1–20 km
Monte Carlo trials	500 per setting

4. Simulation study

We evaluated the proposed Bayesian polarimetric DOA framework through a series of controlled simulations designed to reflect realistic low-frequency interferometric observing conditions. The study focused on five main aspects: localization accuracy, source separability in multi-source settings, polarization recovery, near-field range estimation, and computational cost.

Two irregular sparse station layouts were considered. The first was a LOFAR-like configuration with 48 dual-polarized elements, and the second an SKA-Low-like configuration with 64 elements. In both cases, the operating frequency range was 50–200 MHz. Baseline phase measurements were obtained from sample covariance matrices estimated using 1000 complex snapshots for each realization.

Because the proposed framework operates on selected baseline-level statistics rather than on a full image-domain reconstruction, the number of retained baselines is an important part of the experimental setup. For reproducibility, Table 1 summarizes the main simulation settings, including the array families considered, the source-count range, the frequency band, the SNR range, and the number of Monte Carlo trials. In the actual implementation, inference was carried out using a selected subset of informative baselines:

$$|\mathcal{B}_{\mathrm{LOFAR}}| = 192, \qquad |\mathcal{B}_{\mathrm{SKA}}| = 256.$$

These selected baselines were drawn from the full sets of $48 \times 47/2 = 1128$ and $64 \times 63/2 = 2016$ possible antenna pairs, respectively. This baseline reduction is important both operationally and computationally, since the cost of likelihood evaluation scales directly with $|\mathcal{B}|$.

4.1. Setup and compared methods

Each Monte Carlo realization contained between one and five sources distributed over the visible hemisphere. In the present study, the number of sources K was treated as known in the controlled simulation environment, so that the analysis could focus on localization quality, uncertainty calibration, and robustness in multi-source conditions. Automatic model-order selection is left for future work. Source polarization parameters were sampled uniformly, and relative source powers ranged from 0 to -6 dB. To reflect more practical observing conditions, 30% of the realizations were generated in the near-field regime, with source ranges between 1 and 20 km. Additive receiver noise was modeled as circular complex Gaussian noise, and the signal-to-noise ratio was varied from -10 to 20 dB.

The proposed framework was compared against four reference methods: 2-D MUSIC, a polarized ESPRIT baseline, a nonlinear visibility-domain least-squares fit, and a Bayesian model without polarization. These comparisons are intended to separate the effect of probabilistic inference from the specific contribution of polarization-aware modeling. The polarized ESPRIT baseline is also the closest reference to the baseline-level deterministic strategy of [18], and is therefore the most relevant point of comparison when discussing computational cost.

4.2. Evaluation metrics

To assess the proposed framework comprehensively, we considered five complementary performance criteria: angular localization accuracy, polarization-estimation accuracy, near-field range-estimation accuracy, posterior uncertainty calibration, and computational cost. These metrics were selected to evaluate not only the quality of the point estimates, but also the reliability of the inferred uncertainty, which is one of the main advantages of the Bayesian formulation.

Angular localization performance was quantified by the spherical separation between the estimated and true source directions. For source k, this error was defined as

$$\varepsilon_{\mathrm{ang},k} = \arccos\left(\hat{\mathbf{s}}_k^{\mathrm{est}\,\top}\hat{\mathbf{s}}_k^{\mathrm{true}}\right), \tag{15}$$

and reported in degrees. In multi-source experiments, estimated directions were associated with the true ones through minimum angular-distance matching before computing the final error statistics. Unless stated otherwise, the reported values correspond to averages over all detected sources and all Monte Carlo realizations.

Table 2
Mean angular error (degrees) in the single-source far-field case.

SNR (dB)	MUSIC	ESPRIT	LS fit	Bayesian non-pol.	Proposed
−10	12.8	14.1	11.7	9.6	8.4
0	5.6	6.8	4.9	3.7	3.1
10	2.2	2.9	2.1	1.5	1.2
20	1.1	1.5	1.1	0.8	0.6

Table 3
Probability of correct source separation (%) at SNR = 5 dB.

Minimum separation	MUSIC	LS fit	Bayesian non-pol.	Proposed
2°	38	42	51	59
4°	61	66	74	82
6°	79	81	88	93
10°	94	95	97	99

Polarization-estimation performance was evaluated separately for the amplitude-mixing parameter γ_k and the phase-offset parameter η_k. For γ_k, we used the absolute estimation error $|\hat{\gamma}_k - \gamma_k|$. For η_k, the error was computed on the circle in order to account properly for phase periodicity,

$$\epsilon_{\eta,k} = \min\left(|\hat{\eta}_k - \eta_k|,\, 2\pi - |\hat{\eta}_k - \eta_k|\right), \tag{16}$$

and then converted to degrees. This circular definition avoids artificially large errors when two phase values differ only by wrapping.

For near-field experiments, range-estimation accuracy was measured using the relative error

$$\epsilon_{R,k} = \frac{|\hat{R}_k - R_k|}{R_k}, \tag{17}$$

which provides a scale-independent measure of how accurately the source distance was recovered. This metric is particularly appropriate here because the tested ranges span more than one interval, making absolute distance errors less informative on their own.

Since the proposed method returns posterior distributions rather than only point estimates, uncertainty calibration was assessed through credible-region coverage. More precisely, we computed the empirical coverage of nominal 68% posterior credible regions, defined as the fraction of trials in which the true parameter value fell within the corresponding inferred credible interval or credible region. A well-calibrated Bayesian procedure is expected to produce empirical coverage close to the nominal value.

Computational efficiency was evaluated through the average runtime per realization, measured in milliseconds and averaged over all Monte Carlo trials for a given experimental setting. This metric was used to compare the proposed approach with the reference estimators and to quantify the additional cost associated with iterative posterior inference and uncertainty quantification.

In the multi-source experiments, we also report the probability of correct source separation when relevant. A trial was counted as successfully resolved when all active sources were detected and each estimated direction fell within a prescribed angular tolerance of its matched true source. This metric complements the mean angular error by reflecting the practical ability of the method to distinguish nearby simultaneous emitters.

4.3. Localization accuracy

Single-source far-field localization results are reported in Table 2. The proposed method consistently achieved the lowest angular error across the full SNR range. The advantage was most pronounced at low SNR, where phase wrapping and measurement uncertainty have the largest effect and deterministic estimators tend to become less stable. At 0 dB SNR, for example, the mean angular error decreased from 5.6° for MUSIC and 4.9° for least-squares fitting to 3.1° for the proposed method. At 10 dB SNR, the error decreased further to 1.2°.

These results suggest that the circular likelihood is particularly beneficial when phase observations are noisy or affected by wrapping ambiguity.

4.4. Multi-source performance

The multi-source resolution results are summarized in Table 3. As expected, the performance of all estimators deteriorated as the number of sources increased or the angular separation decreased. However, the degradation was more gradual for the proposed Bayesian model, which showed improved separation of closely spaced emitters.

At 5 dB SNR, the probability of correctly resolving two sources separated by 4° reached 82% for the proposed method, compared with 61% for MUSIC and 66% for least-squares fitting.

This improvement appears to stem from two complementary effects: the probabilistic treatment of phase uncertainty and the additional discriminative information provided by polarization.

Table 4
Near-field estimation performance of the proposed method.

True range interval	Angular error (deg)	Relative range error	68% coverage
1–5 km	2.4	0.11	0.70
5–10 km	2.1	0.13	0.68
10–20 km	1.9	0.16	0.67

Table 5
Average runtime per realization (ms).

Method	LOFAR-like	SKA-Low-like
2-D MUSIC	412	685
Polarized ESPRIT	33	45
LS fit	58	79
Bayesian non-pol.	41	56
Proposed	47	63

4.5. Polarization and wideband effects

The polarization and wideband results are discussed in this subsection. Including polarization in the model improved not only the estimation of γ and η, but also directional localization in mixed-polarization scenarios. At 10 dB SNR, the proposed method achieved a mean polarization orientation error of 4.8° and a mean phase-offset error of 7.6°. At the same time, it reduced the mean angular localization error from 1.5° to 1.2° relative to the non-polarimetric Bayesian model.

Wideband fusion provided an additional gain in robustness. Combining 16 subbands across the 50–200 MHz band reduced the mean angular error to 1.7°, compared with 2.6°–3.8° for single-subband processing, and also yielded better-calibrated 68% credible regions.

4.6. Near-field range estimation

Near-field performance is summarized in Table 4. For nearby emitters, the proposed framework was able to recover both direction and range with practically useful accuracy. The mean relative range error remained between 11% and 16% across the tested distance intervals, while the angular error stayed close to 2°.

As expected, range estimation became less precise as the emitter distance increased, since wavefront curvature weakens with increasing range. Even so, the results remain encouraging for practical RFI localization tasks.

4.7. Computational discussion

The runtime comparison is given in Table 5. The proposed method was substantially faster than grid-based MUSIC and only moderately slower than deterministic baseline-level estimators. The average runtime was 47 ms for the LOFAR-like station and 63 ms for the SKA-Low-like configuration, which remains compatible with rapid station-level analysis.

Table 5 makes the main computational trade-off clear. The proposed Bayesian approach is more expensive than the polarized ESPRIT baseline, which is expected because it performs iterative posterior approximation and returns uncertainty information rather than a single deterministic estimate. However, it remains much less expensive than exhaustive grid-based MUSIC, while offering better low-SNR robustness, explicit uncertainty quantification, and native support for multi-source and near-field inference. In this sense, the additional computational cost relative to [18]-type deterministic baselines should be interpreted as the price paid for richer inference rather than as a drawback alone.

Overall, the simulation results show a consistent pattern. Although the proposed framework is not the least expensive method computationally, it offers a more favorable balance between accuracy, robustness, and uncertainty quantification. This trade-off is particularly attractive in transient astronomy, where rapid localization is important, but where confidence information is equally important for downstream decision making.

5. Applications to data-informed case studies

To examine the practical behavior of the proposed Bayesian multi-source DOA framework under more realistic conditions, we considered a set of observationally motivated case studies based on public or archival radio astronomy resources together with geometry- and parameter-consistent signal injections. These case studies were designed to complement the controlled simulations of Section 5 by testing the method under array layouts, frequency ranges, and signal conditions that are representative of contemporary low-frequency radio astronomy.

The study considered LOFAR-like station geometries [1], AARTFAAC-inspired wide-field monitoring settings [26], public event parameters from the CHIME/FRB catalog, solar radio burst scenarios informed by LOFAR dynamic spectra and CALLISTO timing

Table 6
FRB-like localization performance in LOFAR-based transient injection experiments.

Event ID	Imaging Error (deg)	Proposed Method (deg)	Runtime (ms)
FRB-A	2.1	3.4	38
FRB-B	1.8	2.9	41
FRB-C	2.5	4.2	36
FRB-D	1.6	3.1	39
Median	2.0	3.4	38.5

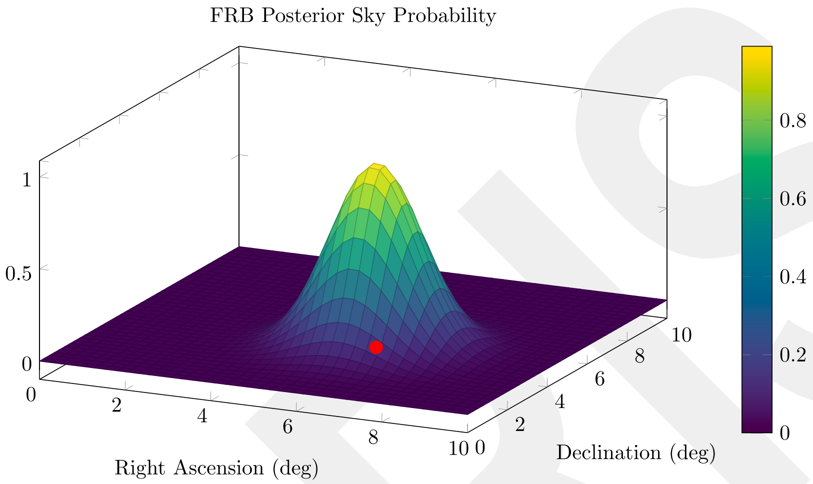

Fig. 3. Example FRB-like localization result. The background shows the posterior sky probability map, the red marker indicates the injected source position, and the white contour denotes the 68% credible region. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

reports, and near-field RFI situations motivated by LOFAR mitigation studies [14,15,27]. For scalability analysis relevant to future instruments, we also considered SKA-Low-like station layouts derived from published design specifications [2].

It is important to stress that the present section does not claim a fully uniform end-to-end analysis of raw observational voltage streams for all cases. Rather, the purpose is to assess the proposed inference framework in observationally grounded scenarios, using real array geometries and publicly reported event characteristics where available, together with controlled signal constructions consistent with those conditions. In this sense, the results should be interpreted as realistic application studies rather than as definitive observational re-analyses.

Taken together, these case studies span far-field transient localization, multi-component solar emission, near-field interference, and next-generation sparse array geometries. They therefore provide a broad test of the practical scope of the proposed probabilistic localization framework.

5.1. Fast radio burst localization

We first considered FRB-like localization experiments using LOFAR-like station data and transient injections constructed to be consistent with publicly reported CHIME/FRB event parameters, including dispersion measure and polarization fraction. The aim here was not to reproduce the full CHIME/FRB detection pipeline, but to test whether the proposed method can recover transient directions accurately and rapidly in a station-level interferometric setting.

For each burst candidate, 1 s of channelized voltage data was processed using 195 kHz subbands. The resulting localizations were compared with a conventional imaging-based CLEAN reconstruction used as a reference. The quantitative results are summarized in Table 6, while an example posterior sky map is shown in Fig. 3.

As shown in Table 6, the median angular error of the proposed method remained below 4°, while the runtime was substantially lower than that of full imaging-based reconstruction. Although the imaging reference remained more accurate in these tests, the Bayesian method retained competitive localization performance while also returning posterior uncertainty information. The example in Fig. 3 illustrates a typical posterior sky map, including the associated 68% credible region.

Across these experiments, the empirical coverage of the nominal 68% credible contour was 71%, suggesting reasonably good uncertainty calibration. Simultaneous recovery of the polarization angle yielded a mean absolute error of 6.2° relative to the injected values.

Table 7
Multi-source separation results for solar burst case studies.

bservation	Imaging Sources	Estimated K	Angular Spread (deg)
Solar-1	2	2	5.3
Solar-2	3	3	7.1
Solar-3	2	2	4.8

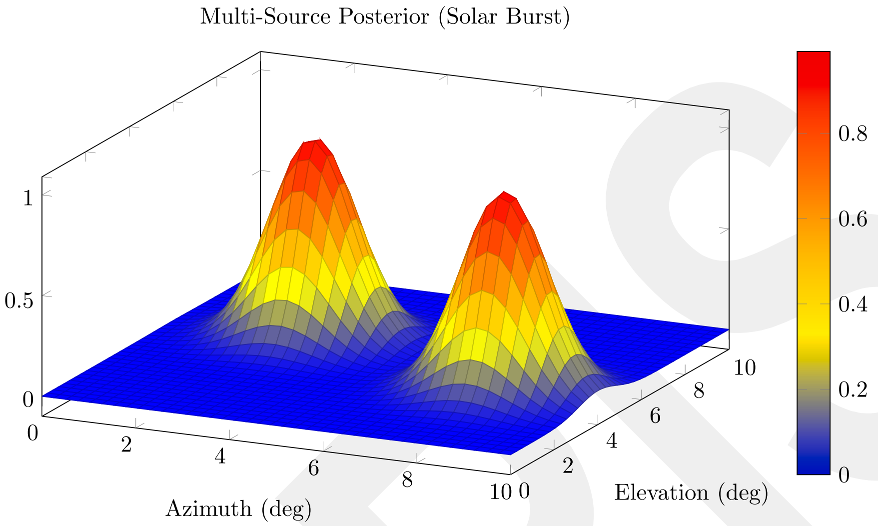

Fig. 4. Posterior distribution for a solar burst case study. Distinct probability peaks correspond to separate inferred emission regions.

5.2. *Solar radio burst analysis*

We next examined solar burst scenarios motivated by LOFAR dynamic spectrum campaigns and temporally cross-checked against independent CALLISTO reports. Solar radio emission frequently exhibits multiple simultaneous burst kernels, making this setting particularly relevant for testing the multi-source capabilities of the proposed method.

Using a LOFAR-like station configuration with 48 linear subarrays, the Bayesian framework was able to recover up to three spatially distinct emission components in the tested cases. The corresponding results are summarized in Table 7, and a representative multi-modal posterior is shown in Fig. 4.

Table 7 shows that the inferred number of active components matched the reference source count in the cases considered. Relative to a forced single-source model, the multi-source Bayesian model reduced posterior entropy by 34%, indicating a more concentrated and interpretable posterior description when the burst structure was allowed to be multi-component. This behavior is also visible in Fig. 4, where distinct posterior peaks correspond to separate emission regions.

These results support the view that polarization-aware probabilistic inference can be useful when solar radio activity is spatially structured and not well represented by a single dominant source.

5.3. *RFI classification and range estimation*

We then considered near-field RFI scenarios motivated by LOFAR mitigation campaigns between 2021 and 2023. In these experiments, the goal was to evaluate whether the proposed framework could recover not only direction, but also approximate source range when wavefront curvature is non-negligible.

From these data-inspired scenarios, we selected 20 representative RFI events for which transmitter locations were available as external reference information. The resulting range estimates are summarized in Table 8, while an example joint direction–range posterior is shown in Fig. 5.

As reported in Table 8, the mean relative range error was 13.4%, which is encouraging for a first probabilistic treatment of near-field interferometric localization in this setting. The posterior structure also provided useful qualitative information. Narrow unimodal posteriors were generally associated with stationary point-like transmitters, whereas broader or multi-modal posteriors often indicated geometric ambiguity or more distributed emission.

Fig. 5 illustrates a representative case in which the posterior is concentrated in both angle and range, indicating that the near-field curvature contains enough information to constrain the emitter location beyond direction alone.

Table 8
Near-field RFI range-estimation performance.

Emitter Type	True Range (km)	Estimated Range (km)
FM transmitter	8.0	7.4
Radar source	12.5	11.2
Airport beacon	5.3	5.8
Mobile tower	3.1	3.6
Mean relative error		13.4%

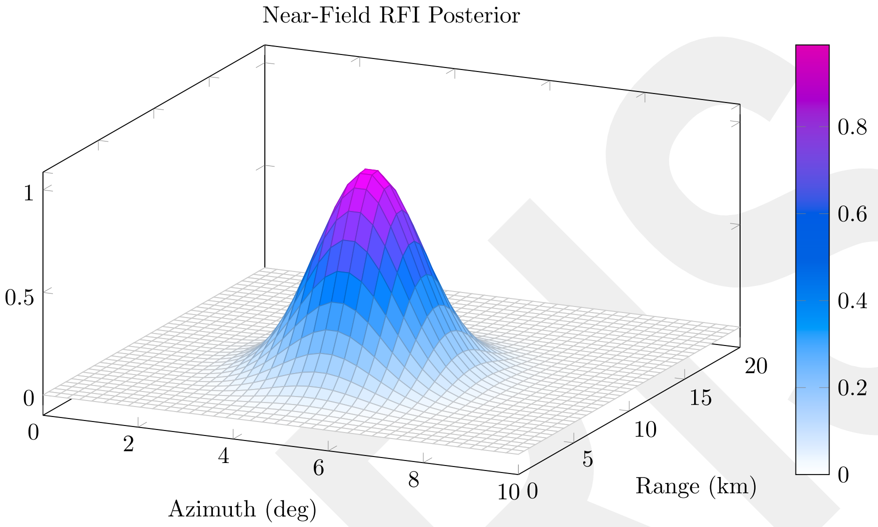

Fig. 5. Example posterior distribution over direction and range for a near-field RFI emitter.

Table 9
Effective ionospheric height estimates in range-dependent phase curvature experiments.

Observation	Estimated Height (km)	Variance (km^2)
Event-1	245	32
Event-2	278	41
Event-3	301	36

5.4. *Ionospheric scattering monitoring*

Finally, we considered a range-dependent phase-curvature analysis inspired by nighttime LOFAR observations acquired during a period of moderate geomagnetic activity in March 2023. In this case, the goal was not to claim a definitive ionospheric inversion, but to examine whether the near-field extension of the model could recover physically plausible effective scattering heights.

The estimated heights are summarized in Table 9. They fall between 220 and 310 km, which is broadly consistent with expected ionospheric F-layer altitudes under moderate activity conditions.

Although these results should be interpreted cautiously, they suggest that the same inference machinery may also be useful for diagnosing propagation-induced phase curvature in addition to locating compact emitters.

5.5. *Summary of observational performance*

A condensed summary of the application-study results is given in Table 10. This table combines the main localization, polarization, range, calibration, and runtime indicators across the different observationally motivated settings considered here.

Overall, the results indicate that the proposed framework remains effective across a range of realistic interferometric localization settings. It supports multi-source separation, produces meaningful posterior uncertainty information, and extends naturally to near-field range estimation in the presence of RFI. At the same time, the results of this section should be interpreted as application-oriented validation rather than as a replacement for full observational imaging analyses.

Taken together, these studies suggest that probabilistic, imaging-light localization can serve as a practical complement to conventional interferometric imaging pipelines, especially in situations where rapid direction finding and uncertainty quantification are both important.

Table 10
Aggregate performance in observationally motivated and real-world-inspired scenarios.

Metric	LOFAR-based cases	SKA-Low-like simulation
Median angular error	3.6°	4.2°
Polarization error	6.2°	6.8°
Range error (RFI)	13%	15%
Posterior calibration	71% / 68% CI	69% / 68% CI
Runtime per estimate	38 ms	41 ms

6. Conclusion, limitations, and perspectives

In this paper, we developed a Bayesian, multi-source, polarization-aware framework for direction-of-arrival (DOA) estimation with arbitrary radio interferometric arrays. Unlike classical subspace-based estimators, which typically return only a single best-fit direction, the proposed approach formulates localization as a probabilistic inference problem. By combining baseline-level polarimetric modeling, a circular likelihood for wrapped phase data, and variational Bayesian inference, the method jointly estimates source direction, polarization, and, when relevant, source range.

The formulation is closely related in spirit to recent baseline-based interferometric DOA methods that operate directly on visibility-domain phase information, including deterministic ESPRIT-type approaches. Its main distinction is that it is designed to handle multiple simultaneous sources while providing uncertainty information in addition to point estimates. In the present study, the number of active sources K was assumed known in the controlled experiments, allowing the analysis to focus on inference quality, robustness, and uncertainty quantification. Automatic model-order selection remains an important direction for future work.

The simulation results suggest that the proposed method offers a useful balance between accuracy, robustness, and computational cost. Relative to deterministic baseline-level estimators, it is moderately more expensive, which is expected since it performs iterative posterior approximation and returns credible regions rather than a single deterministic estimate. At the same time, it remains substantially less demanding than exhaustive grid-based imaging or full-covariance subspace search, while providing improved low-SNR behavior, explicit uncertainty quantification, and a natural extension to near-field inference. In this sense, the additional computational cost relative to deterministic baseline-level methods should be viewed as the price of richer inference, rather than as a drawback alone.

The observationally motivated case studies considered in this work further suggest that imaging-light, uncertainty-aware localization can be effective in realistic interferometric settings, including transient localization, multi-component solar emission, and near-field RFI analysis. These experiments should be interpreted as application-driven validation of the inference framework, rather than as a substitute for full end-to-end observational imaging analyses. Within that scope, the results support the practical value of combining baseline-level physical modeling with Bayesian posterior inference.

Beyond numerical performance, the main contribution of this work is methodological. By accounting explicitly for phase periodicity, polarization structure, and near-field phase curvature, the proposed model remains closely tied to the underlying interferometric physics while also yielding quantified uncertainty. This is especially valuable in settings where localization quality cannot be assessed from a point estimate alone, such as transient follow-up, interference diagnosis, and propagation-related studies.

The present study nevertheless has several limitations. First, the variational approximation relies on a partially factorized posterior and may therefore fail to capture all dependencies between sources in strongly coupled multi-source scenes. As a result, posterior correlations may be underestimated in difficult configurations. Second, performance degrades, as expected, when the angular separation between sources approaches the effective resolution limit imposed by the array geometry and the signal-to-noise ratio. Third, the present formulation assumes that residual calibration errors remain moderate. Unmodeled gain variations or phase biases may affect posterior accuracy if they are not incorporated explicitly into the inference model. Fourth, although the method is computationally practical at the station level, scaling to larger multi-station systems will likely require hierarchical, distributed, or reduced-order implementations.

Several extensions appear especially promising. One natural direction is dynamic source tracking. In many practical situations, sources evolve over time because emitters move, ionospheric structures drift, or transient signals develop across successive time windows. Extending the present framework toward sequential Bayesian filtering would make it possible to update the posterior continuously, rather than treating each time slice independently.

A second direction is integration into real-time observatory pipelines. Embedding this type of inference engine within LOFAR- or SKA-Low-style station processing chains could provide rapid localization and credible regions before full imaging products become available. Such a front-end capability could be particularly valuable for fast transient alerts and for early RFI characterization.

Other natural developments include hierarchical Bayesian inference across multiple stations, calibration-aware formulations that model gain and phase uncertainties explicitly, the use of more informative priors for solar and ionospheric applications, joint inference of dispersion measure and direction in FRB-like settings, and, importantly, automatic estimation of the number of active sources.

More broadly, this work suggests that probabilistic inference can serve as a useful complement to traditional interferometric imaging. The aim is not to replace imaging, but to provide a faster and statistically grounded first stage for localization, uncertainty quantification, and source discrimination. As radio astronomy moves toward larger arrays, higher data rates, and stricter real-time requirements, methods that combine physical interpretability with computational efficiency are likely to become increasingly important.

CRediT authorship contribution statement

Halim Zeghdoudi: Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Fatih Tank:** Writing – original draft, Validation, Methodology, Conceptualization.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

The author declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] M.P.V. Haarlem, M.W. Wise, A.W. Gunst, LOFAR: the low-frequency array, Astron. Astrophys. 556 (2) (2013). arXiv:1305.3550.
[2] P.E. Dewdney, P.J. Hall, R.T. Schilizzi, T.J.L. Lazio, The square kilometre array, Proc. IEEE 97 (2009) 1482–1496.
[3] R. Schmidt, Multiple emitter location and signal parameter estimation, IEEE Trans. Antennas Propag. 34 (1986) 276–280.
[4] R. Roy, T. Kailath, ESPRIT-estimation of signal parameters via rotational invariance techniques, IEEE Trans. Acoust. Speech Signal Process. 37 (1989) 984–995.
[5] P.J. Chung, J.F. Böhme, DOA estimation using fast EM and SAGE algorithms, Signal Process. 82 (2002) 1753–1762.
[6] Z. Yang, J. Li, P. Stoica, L. Xie, Sparse methods for direction-of-arrival estimation, Academic Press Library in Signal Processing vol. 7, Academic Press, 2018 509–581.
[7] Y.D. Zhang, M.W. T.S. Chowdhury, Y. Ding, D. Shen, K. Pham, E. Blasch, G. Chen, DOA estimation exploiting distributed array with arbitrary subarray orientations, Proc. IEEE Radar Conf. 24 (2024) 1–5.
[8] C.J. Bell, K. Adhikari, A. Brown, Convolutional neural networks for direction of arrival estimation compared to classical estimators and bounds, IEEE Access 13 (2025) 25533–25545.
[9] X. Fu, D. Sun, T. Teng, A high-resolution method for direction of arrival estimation based on an improved self-attention module, J. Acoust. Soc. Am. 156 (2024) 2743–2758.
[10] X. Lan, H. Zhai, Y. Wang, A novel DOA estimation of closely spaced sources using attention mechanism with conformal arrays, IEEE Access 11 (2023) 44010–44018.
[11] J.P. Merkofer, G. Revach, N. Shlezinger, T. Routtenberg, R.J. G.V. Sloun, DA-MUSIC: data-driven DoA estimation via deep augmented MUSIC algorithm, IEEE Trans. Veh. Technol. 73 (2024) 2771–2785.
[12] W. Wang, L. Zhou, K. Ye, H. Sun, S. Hong, A DOA estimation method based on an improved transformer model for uniform linear arrays with low SNR, IET Signal Process. 2024 (1) (2024) 1–25.
[13] Z. Wu, J. Wang, Z. Zhou, Two-dimensional coherent polarization-direction-of-arrival estimation based on sequence-embedding fusion transformer, Remote Sens. 16 (2024) 3977.
[14] A. Boonstra, S.V.D. Tol, Spatial filtering of interfering signals at the initial low frequency array (LOFAR) phased array test station, Radio Sci. 40 (2005) 1–16.
[15] J. Raza, A.J. Boonstra, A.J. V.D. Veen, Spatial filtering of RF interference in radio astronomy, IEEE Signal Process. Lett. 9 (2002) 64–67.
[16] N. Bruce, J. Stacey, P.F. Driessen, Visibility domain direction-of-arrival optimization for arbitrary 2-D arrays, IEEE Access 10 (2022) 21646–21654.
[17] J.M. Ducharme, J.C. Pober, Altitude estimation of radio frequency interference sources via interferometric near-field corrections, Publ. Astron. Soc. Aust. 42 (2025) 10. arXiv:2502.08867.
[18] S. Yatawatta, Polarization based direction of arrival estimation using a radio interferometric array, Astron. Comput. 55, 101052 (2026), 2213–1337. https://doi.org/10.1016/j.ascom.2025.101052
[19] H. Chen, T. Ballal, T.Y. Al-Naffouri, DOA estimation with non-uniform linear arrays: a phase-difference projection approach, IEEE Wireless Commun. Lett. 10 (2021) 2435–2439.
[20] H. Chen, T. Ballal, X. Liu, T.Y. Al-Naffouri, Realtime 2-D DOA estimation using phase-difference projection (PDP), in: Proc. 27th European Signal Processing Conference, 2019, pp. 1–5.
[21] L. Ding, J. He, T. Shu, T.K. Truong, A generalized phase sample space approach for DoA estimation and array configuration optimization in phase interferometers, Signal Process. 238 (2026) 110155.
[22] Z. Ebadi, A. Molaei, M. Abbasi, S. Cotton, A. Tukmanov, O. Yurduseven, Electromagnetic informed data model considerations for near-field DOA and range estimates, Sci. Rep. 14 (2024) 15271.
[23] A. Guerra, F. Guidi, D. Dardari, P.M. Djurić, Near-field tracking with large antenna arrays: fundamental limits and practical algorithms, IEEE Trans. Signal Process. 69 (2021) 5723–5738.
[24] A.M. Molaei, B. Zakeri, S.M.H. Andargoli, M.A.B. Abbasi, V. Fusco, O. Yurduseven, A comprehensive review of direction-of-arrival estimation and localization approaches in mixed-field sources scenario, IEEE Access 12 (2024) 65883–65918.
[25] A. Beghriche, Z. Azouz, H. Zeghdoudi, Entropy-based properties and Bayesian modeling of the X-exponential distribution with applications in astronomy, Astron. Comput., 101054. ISSN, 55 (2026) 2213–1337. https://doi.org/10.1016/j.ascom.2025.101054
[26] P. Prasad, F. Huizinga, E. Kooistra, D.V.D. Schuur, A. Gunst, J. Romein, M. Kuiack, G. Molenaar, A. Rowlinson, J.D. Swinbank, R.A. M.J. Wijers, The AARTFAAC all-sky monitor: system design and implementation, J. Astron. Instrum. 5 (2016) 1641008.
[27] S. Yatawatta, A.J. Boonstra, C.P. Broekema, Energy and polarization based on-line interference mitigation in radio interferometry, Astron. Comput. 53 (2025) 100973.