

Extending Acoustic Beamformers to Unknown-Orientation Dipoles

Armin Goudarzi

`armin.goudarzi@dlr.de`

German Aerospace Center (DLR)

Daniel Ernst

German Aerospace Center (DLR)

Research Article

Keywords: acoustic beamforming, dipole sources, source orientation, microphone arrays, cross-spectral matrix, GPU computing

Posted Date: September 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11105567/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Extending Acoustic Beamformers to Unknown-Orientation Dipoles

Armin Goudarzi [ORCID](#)^{1*} and Daniel Ernst [ORCID](#)¹

^{1*}Department of Experimental Methods, German Aerospace Center (DLR), Bunsenstr. 10, 37073, Göttingen, Germany.

*Corresponding author(s). E-mail(s): armin.goudarzi@dlr.de;
Contributing authors: daniel.ernst@dlr.de;

Abstract

Many aeroacoustic sources radiate as dipoles, but their orientations are often unknown. Dipole beamforming nevertheless requires these orientations to construct the acoustic Green's functions. We represent an arbitrarily oriented dipole as a linear combination of three Cartesian dipoles. Substituting this representation into the normalized conventional beamforming criterion reduces local orientation selection to a real symmetric three-by-three generalized eigenproblem at each focus point. The resulting dipole transfer vectors can be used in established beamforming methods without prescribed orientations or surface-normal assumptions. A broadband extension couples orientation across frequencies while retaining separate source powers. In a synthetic four-dipole case, broadband conventional beamforming locates all four sources at every evaluated frequency, with mean absolute errors of 0.49 dB in level and **2.30°** in axis orientation, treating opposite axes as equivalent. With 64 microphones, ten frequencies, and 10,201 focus points, the beamforming computation takes 17.4 ms on a graphics processing unit. Applied to open-jet airframe-noise data, the method estimates different effective slat-track dipole axes in the 2- and 8-kHz octaves, consistent with previously reported frequency-scaling regimes. Orientation recovery depends on the array's spatial resolution, the beamforming method, and the quality of the estimated cross-spectral matrix.

Keywords: acoustic beamforming, dipole sources, source orientation, microphone arrays, cross-spectral matrix, GPU computing

1 Introduction

Unsteady aerodynamic loading on solid surfaces produces dipole sound, making dipole source models relevant to airframe components, fans, and propellers [1]. Using these models in beamforming requires the source orientation, which is generally unknown. Thus, array imaging often relies on a monopole model or a dipole axis inferred from the surface geometry. These choices avoid estimating the orientation, but monopole mismatch can bias both location and level [2, 3], and a surface-derived axis need not represent the effective acoustic dipole axis.

This paper proposes an extension of established beamformers that estimates unknown dipole orientations directly from microphone-array data. At each focus point, we calculate the microphone responses of three unit dipoles aligned with the x , y , and z axes. Any real dipole axis gives a linear combination of these responses. For conventional beamforming, the combination maximizing the normalized score is found through a 3×3 generalized eigenproblem formed from the responses and the cross-spectral matrix (CSM). Its largest eigenvalue gives the maximum score, and an associated eigenvector gives the axis. The corresponding transfer vector supplies the original beamforming method’s localization, subtraction, or source-power estimation. We call this analytical orientation selection orientation profiling.

The local orientation estimates can also be combined across frequencies. When a source can be assumed to have a common position and axis within a selected band, we extend broadband localization [4, 5] to share the axis while retaining frequency-dependent powers. Combining axes requires accounting for their sign ambiguity: an axis and its negative produce the same source covariance, but would cancel under direct vector averaging. We therefore average the sign-invariant outer products of narrowband axes.

An alternative to analytical orientation selection is to test discrete axes, as in the angular scan of Gao et al. [6]. Our solution removes this angular scan while retaining the spatial scan. Other directional formulations use spatial modes [7], monopole–multipole dictionaries [8], sparse Bayesian models [9], dedicated dipole arrays [10], or geometry constraints [11]. Gridless cross-spectral-matrix fitting and maximum-likelihood methods jointly optimize position, power, and orientation [5, 12], avoiding grid error but requiring coupled nonlinear optimization.

Small eigenproblems are already used to estimate dipole orientation in source imaging [13, 14]. We contribute a reusable orientation-selection step for established grid-based beamformers, including broadband coupling, pair weighting, and diagonal removal. We build on Sarradj’s normalizations [15] to distinguish the localization score from reported source power.

Conventional beamforming provides a direct test of orientation selection without an iterative reconstruction solver; other method extensions are verified in the appendices. Controlled-source, angular-grid, and randomized finite-snapshot tests assess location, power, and orientation, alongside central processing unit (CPU) and graphics processing unit (GPU) timings. An open-jet measurement with 512 microphones demonstrates the approach under wind-tunnel conditions and provides additional directional information for interpreting slat-track source mechanisms.

2 Theory

2.1 Cross-spectral matrix model

The microphone CSM contains auto-spectra on its diagonal and cross-spectra off the diagonal. We denote the Hermitian cross-spectral matrix supplied to the estimator by $\mathbf{C} \in \mathbb{C}^{M \times M}$. Indices $r = 1, \dots, F$, $m = 1, \dots, M$, $p = 1, \dots, P$, and $s = 1, \dots, S$ denote frequency bin, microphone, focus point, and source. For frequency f_r in hertz, $\omega_r = 2\pi f_r$ and $k_r = \omega_r/c$ are angular frequency and wavenumber. A calibrated pressure snapshot is $\mathbf{p}_r^{(\ell)} \in \mathbb{C}^M$, indexed by ℓ . We fix one bin and suppress r until [section 2.6](#), using f , ω , k , and $\mathbf{p}^{(\ell)}$. The input may be an ensemble CSM or a finite-snapshot estimate such as

$$\hat{\mathbf{C}}^{(K)} = \frac{1}{K} \sum_{\ell=1}^K \mathbf{p}^{(\ell)} (\mathbf{p}^{(\ell)})^H, \quad \mathbf{C} = \mathbf{C}^H. \quad (1)$$

The main cross-method comparison uses exact ensemble CSMs; the randomized stress test in [section 3.3](#) additionally uses finite-snapshot estimates. For S mutually incoherent compact sources, the standard population-model assumption is

$$\mathbf{C} = \sum_{s=1}^S q_s \mathbf{g}_s \mathbf{g}_s^H + \mathbf{N}, \quad q_s \geq 0, \quad (2)$$

Here q_s is the source auto-power, or source strength, in the chosen Green's function normalization. We use *source-coefficient power* when this normalization must be emphasized. For the chosen source model, the Green's function vector $\mathbf{g}_s \in \mathbb{C}^M$ collects the responses to a unit source at $\mathbf{y}_s$ evaluated at the microphone positions $\mathbf{x}_m$: $(\mathbf{g}_s)_m = g(\mathbf{x}_m, \mathbf{y}_s)$. It is the source-to-array transfer vector before beamforming normalization. The matrix $\mathbf{N} \in \mathbb{C}^{M \times M}$ contains noise and unmodeled contributions. Coherent sources require cross terms and are discussed as a limitation in [section 6.2](#).

An optional real symmetric pair-weight matrix $\mathbf{W} \in \mathbb{R}^{M \times M}$ defines

$$\mathbf{C}^{(W)} = \mathbf{W} \odot \mathbf{C} \in \mathbb{C}^{M \times M}. \quad (3)$$

Ordinary diagonal removal uses $W_{mm} = 0$ and $W_{mj} = 1$ for $m \neq j$. Pair weighting preserves Hermiticity but need not preserve positive semidefiniteness, even for nonnegative weights. We therefore distinguish the unweighted source model from a weighted imaging objective.

2.2 Dipole Green's functions

For illustration, let $d = \|\mathbf{x} - \mathbf{y}\|_2$ and $\hat{\mathbf{d}} = (\mathbf{x} - \mathbf{y})/d$. The free-field monopole, its source gradient, and a compact dipole with real unit axis $\mathbf{n} \in \mathbb{R}^3$ are

$$g^{(0)}(\mathbf{x}, \mathbf{y}) = \frac{\exp(-ikd)}{4\pi d}, \quad k = \frac{\omega}{c}, \quad (4)$$

$$\nabla_{\mathbf{y}} g^{(0)} = \left(ik + \frac{1}{d} \right) g^{(0)} \hat{\mathbf{d}}, \quad (5)$$

$$g^{(d)}(\mathbf{x}, \mathbf{y}, \mathbf{n}) = \chi \mathbf{n}^T \nabla_{\mathbf{y}} g^{(0)}, \quad \|\mathbf{n}\|_2 = 1, \quad (6)$$

Here χ fixes the source-coefficient normalization. Changing the derivative sign convention changes the global sign of the dipole axis but not its rank-one CSM. Each source is represented in each frequency bin by a real axis and a complex amplitude, describing a dipole oscillating along a fixed axis. Complex moments with noncollinear real and imaginary parts describe elliptically rotating dipole moments, for example from phase-shifted force components, and are excluded by this fixed-axis assumption. The axes are independent between bins unless one effective broadband axis is imposed in [section 2.6](#). Because a rank-one CSM is invariant to $\mathbf{n} \mapsto -\mathbf{n}$, only the unoriented axis in $\mathbb{S}^2/\{\pm 1\}$ is recoverable.

Let $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ be the Cartesian basis. At focus point p , the columns of $\mathbf{V}_p$ are the source-to-array transfer vectors for unit dipoles aligned with the x , y , and z axes. Collect them in the Cartesian dipole transfer matrix

$$\mathbf{V}_p = [\mathbf{v}_{p,1} \ \mathbf{v}_{p,2} \ \mathbf{v}_{p,3}] \in \mathbb{C}^{M \times 3}, \quad (\mathbf{v}_{p,a})_m = g^{(d)}(\mathbf{x}_m, \mathbf{y}_p, \mathbf{e}_a). \quad (7)$$

Linearity in orientation gives the standard Cartesian representation

$$\boxed{g_p^{(d)}(\mathbf{n}) = \mathbf{V}_p \mathbf{n} \in \mathbb{C}^M, \quad \mathbf{n} \in \mathbb{S}^2/\{\pm 1\}.} \quad (8)$$

The columns retain their physical Green's-function normalization: separately normalizing the columns would generally change the axis represented by $\mathbf{n}$. Beamforming weights are formed from the combined vector $\mathbf{V}_p \mathbf{n}$. The same representation applies to any propagation model whose dipole field is linear in $\mathbf{n}$, including differentiable connected, refracted, and image source Green's functions. The propagation model changes $\mathbf{V}_p$, not the orientation optimization developed below.

2.3 Orientation optimization

For the local derivation, fix one focus point and set $\mathbf{V} = \mathbf{V}_p \in \mathbb{C}^{M \times 3}$. Both frequency and focus-point indices remain suppressed through [section 2.5](#); every local statement applies independently when (r, p) is restored.

Let $\mathbf{A} \in \mathbb{C}^{M \times M}$ denote the Hermitian sensor-space matrix used in the beamforming objective. For conventional beamforming, $\mathbf{A} = \mathbf{C}$. The dependence on the dipole axis enters through the quadratic term $\mathbf{g}^H \mathbf{A} \mathbf{g}$. Substituting $\mathbf{g} = \mathbf{V} \mathbf{n}$ for a real axis $\mathbf{n}$ gives

$$\begin{aligned} \mathbf{g}^H \mathbf{A} \mathbf{g} &= \mathbf{n}^T (\mathbf{V}^H \mathbf{A} \mathbf{V}) \mathbf{n} \\ &= \mathbf{n}^T \text{Re}(\mathbf{V}^H \mathbf{A} \mathbf{V}) \mathbf{n}. \end{aligned} \quad (9)$$

The second equality holds because $\mathbf{V}^H \mathbf{A} \mathbf{V}$ is Hermitian and its skew-symmetric imaginary part has zero quadratic form for real $\mathbf{n}$. Thus the orientation dependence is

captured by a real symmetric 3×3 matrix. We abbreviate this matrix construction by

$$\boxed{\mathcal{K}[\mathbf{A}] = \text{Re}(\mathbf{V}^H \mathbf{A} \mathbf{V}) \in \mathbb{R}^{3 \times 3}.} \quad (10)$$

The quadratic term is therefore $\mathbf{n}^T \mathcal{K}[\mathbf{A}] \mathbf{n}$. Applying the same construction to the CSM, the identity matrix, and the pair-weighted CSM gives

$$\mathbf{S} = \mathcal{K}[\mathbf{C}], \quad \mathbf{G} = \mathcal{K}[\mathbf{I}], \quad (11)$$

$$\mathbf{S}_W = \text{Re}(\mathbf{V}^H (\mathbf{W} \odot \mathbf{C}) \mathbf{V}) = \mathcal{K}[\mathbf{C}^{(W)}] \in \mathbb{R}^{3 \times 3}. \quad (12)$$

They are real symmetric in exact arithmetic. Their six independent entries and numerical symmetrization are given in Appendix D.

Thus $\mathbf{S}$ is the CSM projected onto the three dipole transfer vectors, and $\mathbf{G}$ is their Cartesian transfer-vector Gram matrix. Both are 3×3 , independent of microphone count.

2.4 Generalized-eigenvalue solution

At each focus point, we maximize the localization score; source power is evaluated afterward (section 2.5). Let $\mathbf{K}_A = \mathcal{K}[\mathbf{A}]$ and let $\mathbf{K}_B \in \mathbb{R}^{3 \times 3}$ be the real symmetric positive-semidefinite normalization matrix, so that the denominator is $\mathbf{n}^T \mathbf{K}_B \mathbf{n}$.

When $\mathbf{K}_B$ is positive definite, whitening rescales its eigenmodes to unit normalization and converts the orientation problem into an ordinary real symmetric eigenproblem. Its largest eigenvalue gives the maximum response; transforming an associated eigenvector back gives the maximizing axis.

If the three Cartesian transfer vectors have a real linear dependence, the array cannot distinguish every dipole-axis component. We therefore solve only in their observable Cartesian subspace. The supported coefficient space of $\mathbf{K}_B$ and its projector are

$$\mathcal{U}_B = \text{range}(\mathbf{K}_B), \quad \mathbf{P}_B = \mathbf{K}_B^\dagger \mathbf{K}_B. \quad (13)$$

The quotient is solved on $\mathcal{U}_B$; null-response directions are excluded, and $\mathbf{P}_B \mathbf{n}$ is the canonical representative of a physical axis. The orientation that maximizes the normalized beamforming response is

$$\boxed{\begin{aligned} \rho(\mathbf{n}) &= \frac{\mathbf{n}^T \mathbf{K}_A \mathbf{n}}{\mathbf{n}^T \mathbf{K}_B \mathbf{n}}, \quad \mathbf{n} \in \mathcal{U}_B, \quad \mathbf{n}^T \mathbf{K}_B \mathbf{n} > 0, \\ \hat{\mathbf{n}} &\in \arg \max \rho(\mathbf{n}), \quad \rho^\star = \rho(\hat{\mathbf{n}}), \\ \mathbf{g}^\star &= \mathbf{V} \hat{\mathbf{n}} \in \mathbb{C}^M, \quad \|\hat{\mathbf{n}}\|_2 = 1, \end{aligned}} \quad (14)$$

On $\mathcal{U}_B$, $\mathbf{K}_B$ is positive definite, so the Rayleigh–Ritz theorem applies to the retained eigenmodes [16].

Stationarity on $\mathcal{U}_B$ gives the supported generalized eigenproblem. If $\ker(\mathbf{K}_B) \subseteq \ker(\mathbf{K}_A)$, it has the familiar form

$$\boxed{\mathbf{K}_A \hat{\mathbf{n}} = \rho^* \mathbf{K}_B \hat{\mathbf{n}}.} \quad (15)$$

Denote this possibly set-valued mapping by $\text{geig}_{\max}(\mathbf{K}_A, \mathbf{K}_B) = (\rho^*, \hat{\mathbf{n}})$; $\text{geig}_{\min}$ denotes the analogous minimum. If the largest supported eigenvalue is repeated, several axes share the same score and the orientation is nonunique. The eigengap tests uniqueness and numerical conditioning, not statistical confidence.

For every compatible dipole scan, $\mathbf{K}_B = \mathbf{G}$ and

$$\mathbf{n}^T \mathbf{G} \mathbf{n} = \|\mathbf{V} \mathbf{n}\|_2^2, \quad \ker(\mathbf{G}) = \ker(\mathbf{V}) \quad \text{for real axis vectors.} \quad (16)$$

Consequently, $\ker(\mathbf{G}) \subseteq \ker(\mathcal{K}[\mathbf{A}])$: an axis with zero array response contributes neither numerator nor denominator and is excluded from the solve. Physically, a dipole that produces no pressure at any microphone is unobservable; its source power cannot be inferred from these data. Excluding it prevents division by zero in the subsequent power normalization and does not imply zero source power.

The selected axis supplies the dipole transfer vector $\mathbf{g}^* = \mathbf{V} \hat{\mathbf{n}}$ for localization, subtraction, or source-power estimation. An *exact local profile* means analytical orientation maximization at one focus point, not exact localization, separation, or power estimation.

2.5 Beamforming localization score and source power

Peak finding and source-power reporting use different steering vectors and maps. We first consider an unweighted CSM. For $\mathbf{g} \in \mathbb{C}^M \setminus \{\mathbf{0}\}$, define the array-coupling factor $\beta(\mathbf{g}) = \mathbf{g}^H \mathbf{g}$. The steering vector $\mathbf{w}$ contains the beamforming weights derived from $\mathbf{g}$ using the selected steering formulation. Sarradj's unit response (UR) formulation III and constant-norm formulation IV give [15]

$$\begin{aligned} \mathbf{w}^{\text{III}} &= \frac{\mathbf{g}}{\beta(\mathbf{g})}, & (\mathbf{w}^{\text{III}})^H \mathbf{g} &= 1, \\ \mathbf{w}^{\text{IV}} &= \frac{\mathbf{g}}{\sqrt{M\beta(\mathbf{g})}}, & \|\mathbf{w}^{\text{IV}}\|_2^2 &= \frac{1}{M}. \end{aligned} \quad (17)$$

The source-power estimate and normalized beamforming score, hereafter called the localization score, are

$$\begin{aligned} \hat{q}^{\text{UR}}(\mathbf{g}) &= (\mathbf{w}^{\text{III}})^H \mathbf{C} \mathbf{w}^{\text{III}} = \frac{\mathbf{g}^H \mathbf{C} \mathbf{g}}{\beta(\mathbf{g})^2}, \\ \rho(\mathbf{g}) &= M(\mathbf{w}^{\text{IV}})^H \mathbf{C} \mathbf{w}^{\text{IV}} = \frac{\mathbf{g}^H \mathbf{C} \mathbf{g}}{\beta(\mathbf{g})} = \beta(\mathbf{g}) \hat{q}^{\text{UR}}(\mathbf{g}). \end{aligned} \quad (18)$$

The constant factor M scales the formulation-IV map without changing its maximizing location or axis. For the isolated model $\mathbf{C} = q\mathbf{g}\mathbf{g}^H$, $\hat{q}^{\text{UR}} = q$, whereas $\rho = q\beta(\mathbf{g})$.

Formulation III therefore compensates for the array coupling to report source power; formulation IV retains this factor for localization. Source power is expressed in the chosen Green's function normalization, not as received array power or a map normalized by its maximum.

For pair weighting, define the rank-one calibration denominator

$$\delta_W(\mathbf{g}) = \sum_{m=1}^M \sum_{j=1}^M W_{mj} |g_m|^2 |g_j|^2, \quad \hat{q}_W(\mathbf{g}) = \frac{\mathbf{g}^H \mathbf{C}^{(W)} \mathbf{g}}{\delta_W(\mathbf{g})}. \quad (19)$$

For the full CSM, $\delta_W = \beta(\mathbf{g})^2$; diagonal removal subtracts $\sum_m |g_m|^4$. The matched rank-one model remains calibrated whenever $\delta_W \neq 0$. General pair weighting need not preserve positive semidefiniteness, however. Beamforming scores may therefore be signed, and a nonnegative power interpretation additionally requires $\delta_W(\mathbf{g}) > 0$.

The corresponding localization score for peak finding is

$$\rho_W(\mathbf{g}) = \frac{\mathbf{g}^H \mathbf{C}^{(W)} \mathbf{g}}{\beta(\mathbf{g})} = \frac{\delta_W(\mathbf{g})}{\beta(\mathbf{g})} \hat{q}_W(\mathbf{g}). \quad (20)$$

The localization normalization is unchanged; the source-power estimate instead uses δ_W to compensate for the selected microphone pairs.

For a dipole, $\mathbf{g} = \mathbf{V}\mathbf{n}$ and $\beta(\mathbf{V}\mathbf{n}) = \mathbf{n}^T \mathbf{G}\mathbf{n}$, hence

$$\rho_W(\mathbf{n}) = \frac{\mathbf{n}^T \mathbf{S}_W \mathbf{n}}{\mathbf{n}^T \mathbf{G} \mathbf{n}}, \quad \hat{q}_W(\mathbf{n}) = \frac{\mathbf{n}^T \mathbf{S}_W \mathbf{n}}{\delta_W(\mathbf{V}\mathbf{n})}. \quad (21)$$

The score ρ_W is a generalized Rayleigh quotient. Power $\hat{q}_W$ is evaluated after axis selection because its denominator no longer gives a generalized Rayleigh quotient. Conventional beamforming (CBF) with ρ_W maximized over the axis independently at each focus point and frequency bin is termed *Profile-CBF*. It is the generic local dipole extension used throughout this work: it selects the axis that maximizes the chosen localization score whenever that score is well defined. Peak finding and position updates therefore use ρ_W ; spectra and level metrics use $\hat{q}_W$.

For every maximized map, let $\mathcal{P}_{\text{valid}}$ contain exactly the focus points with a valid orientation estimate, unique extremum, and admissible denominator. Spatial selection is

$$p^* \in \arg \max_{p \in \mathcal{P}_{\text{valid}}} \rho_p. \quad (22)$$

Invalid focus points are undefined and masked (or stored as $-\infty$), not set to zero; this distinction matters for signed pair-weighted maps.

2.6 Common broadband dipole axis

Profile-CBF estimates an axis independently in each frequency bin. Joint-frequency CBF (JCBF) first combines their sign-invariant outer products $\hat{\mathbf{n}}\hat{\mathbf{n}}^T$ into a common axis at each focus point. It then evaluates the shared localization score across the chosen band, retaining one source power per bin.

Proposed common-axis coupling. Frequency and focus indices are now restored. For each $r = 1, \dots, F$ and focus point p , define

$$\mathcal{K}_{r,p}[\mathbf{A}] = \text{Re}(\mathbf{V}_{r,p}^H \mathbf{A} \mathbf{V}_{r,p}), \quad \mathbf{G}_{r,p} = \mathcal{K}_{r,p}[\mathbf{I}]. \quad (23)$$

If one effective source axis is assumed frequency independent over the processed band while its source power may vary, solve the frequency-by-frequency orientation problems

$$(\rho_{1,r,p}, \hat{\mathbf{n}}_{r,p}) = \text{eig}_{\max}(\mathcal{K}_{r,p}[\mathbf{C}_r], \mathbf{G}_{r,p}). \quad (24)$$

For pair-weighted imaging, $\mathbf{C}_r$ is replaced consistently by $\mathbf{W}_r \odot \mathbf{C}_r$. Let $\rho_{2,r,p}$ denote the second supported eigenvalue and let $\iota_{r,p}^{\text{uniq}}$ equal one only for a simple directional extremum. With nonnegative frequency weights α_r satisfying $\sum_r \alpha_r > 0$, define the sign-invariant orientation-averaging matrix

$$\boxed{\begin{aligned} \zeta_{r,p} &= \alpha_r \max(\rho_{1,r,p}, 0) \iota_{r,p}^{\text{uniq}}, \\ \mathbf{T}_p &= \sum_{r=1}^F \zeta_{r,p} \hat{\mathbf{n}}_{r,p} \hat{\mathbf{n}}_{r,p}^T \in \mathbb{R}^{3 \times 3}, \\ \hat{\mathbf{n}}_p^J &= \text{eigvec}_{\max}(\mathbf{T}_p). \end{aligned}} \quad (25)$$

The principal eigenvector defines the common axis when the largest eigenvalue is simple. It minimizes the weighted sum of $1 - (\mathbf{n}^T \hat{\mathbf{n}}_{r,p})^2$, the squared sine of the angle between unit axes, which is unchanged by sign reversal. The rank-deficient form and matched rank-one recovery proof are in Appendix A. Axis averaging remains distinct from spatial peak finding.

For a defined common axis, let $d_{r,p}^J = (\hat{\mathbf{n}}_p^J)^T \mathbf{G}_{r,p} \hat{\mathbf{n}}_p^J$. The broadband beamforming score uses the validity set

$$\mathcal{P}_{\text{valid}}^J = \{p : \hat{\mathbf{n}}_p^J \text{ is defined and } d_{r,p}^J > \tau_{G,r,p} \ \forall r \text{ with } \alpha_r > 0\}. \quad (26)$$

For $p \in \mathcal{P}_{\text{valid}}^J$, the fixed-axis beamforming score is

$$\bar{\rho}_p = \frac{1}{\sum_{r=1}^F \alpha_r} \sum_{r=1}^F \alpha_r \frac{(\hat{\mathbf{n}}_p^J)^T \mathcal{K}_{r,p}[\mathbf{C}_r] \hat{\mathbf{n}}_p^J}{d_{r,p}^J}. \quad (27)$$

The numerical guards in Appendix D exclude invalid focus points. For sources selected from this broadband score, power is still estimated separately in each frequency bin. At valid focus points, the unweighted unit-response estimate is

$$\hat{q}_{r,p}^{\text{JCBF}} = \max\left(\frac{(\hat{\mathbf{n}}_p^J)^T \mathcal{K}_{r,p}[\mathbf{C}_r] \hat{\mathbf{n}}_p^J}{(d_{r,p}^J)^2}, 0\right). \quad (28)$$

The squared denominator follows [equation \(18\)](#), whereas the localization score uses its first power. Pair-weighted data use the calibrated power expression in [equation \(21\)](#), evaluated at the same common axis. [Appendix A.2](#) distinguishes this score-weighted axis consensus from a source-covariance interpretation under additional assumptions.

2.7 Extension to other beamforming methods

The same Cartesian representation extends orientation estimation to other beamformers. The selected axis supplies a dipole Green’s function vector for localization, coherent subtraction, or inverse source-power estimation.

For CLEAN based on spatial source coherence (CLEAN-SC), the local eigenproblem is formed from the current residual CSM at each iteration. Its broadband extension (B-CLEAN-SC) also couples the axis across the selected band before coherent subtraction. Minimum-variance distortionless-response (MVDR) beamforming and multiple signal classification (MUSIC) instead minimize a quotient involving the loaded inverse CSM or the noise-subspace projector, respectively.

These frequency-local quadratic quotients admit exact orientation optimization by the local eigensolve. Deconvolution and covariance-fitting methods instead use an estimated axis to construct source responses for their power solver; this axis proposal does not generally optimize the global inverse objective. [Appendix B](#) gives the method-specific equations and summarizes the local operators in [Table 3](#).

3 Numerical and experimental evaluation

The four-source simulation verifies localization, source-power estimation, and axis recovery against known sources. Randomized simulations test sensitivity to separation, noise, and finite-snapshot CSM estimation. Wind-tunnel data illustrate how estimated axes aid source interpretation.

3.1 Four-source setup

[Figure 1](#) shows four mutually incoherent dipoles at $z/D = 0.5$ above a 64-sensor Fermat array with aperture $D = 1$ m and $c = 343$ m s⁻¹. The 101×101 focus grid spans $x/D, y/D \in [-0.5, 0.5]$ ($P = 10,201$). No beamformer receives a prescribed dipole axis or a surface-normal constraint. The Fermat array uses azimuth increment 137.508° and normalized radial law $\vartheta^{0.7}$. Exact noise-free ensemble CSMs are evaluated at $F = 10$ frequencies, $He_D = fD/c \in \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$ (343–3430 Hz). This range moves from poorly resolved sources to improved spatial resolution and increasingly coarse array sampling relative to the wavelength. With coupling $\kappa_{r,s} = M^{-1} \sum_m |g_{r,m,s}|^2$, each frequency-independent source-coefficient power is $q_s = [F^{-1} \sum_r \kappa_{r,s}]^{-1}$, giving equal band-mean received power rather than equal intrinsic coefficients. [Table 1](#) lists the source positions, axes, and intrinsic source-coefficient levels.




Fig. 1 Normalized four-dipole setup: (a) three-dimensional view and (b)–(d) xy , yz , and xz projections. Orange points and bidirectional axes denote ground truth; crosses mark axes normal to a projection. Numbers identify sources, and all projections use equal coordinate scaling

Table 1 Ground-truth source parameters.

$L_{q,s} = 10 \log_{10}(q_s/q_{\text{ref}})$ is the intrinsic source-coefficient level in the selected Green-function normalization ($q_{\text{ref}} = 1$); powers give equal band-mean received power at the array

Source	x/D	y/D	z/D	$\mathbf{n}_s$	$L_{q,s}$ (dB)
1	-0.30	-0.24	0.50	(1, 0, 0)	-6.3
2	-0.18	0.22	0.50	(0, 1, 0)	-5.9
3	0.16	-0.17	0.50	(0, 0, 1)	-12.6
4	0.30	0.27	0.50	$\frac{1}{\sqrt{3}}(1, 1, 1)$	-11.7

The complete Hermitian information is supplied through its packed upper triangle, and diagonal entries are retained. Monopole CBF provides the source-model baseline. Profile-CBF estimates dipole axes independently at each frequency, and JCBF couples the Profile-CBF axes across all ten bins, with frequency weights normalized by their initial score maxima. Appendix C lists the cross-method settings.

3.2 Evaluation metrics

At each frequency, the evaluation extracts at most $S = 4$ local maxima of the method-specific beamforming score that lie within 20 dB of its global maximum. A global linear assignment matches those candidates one-to-one to ground-truth positions, subject to the matching tolerance $r_{\text{match}} = 0.06D$. Unmatched candidates are false positives and unmatched sources are false negatives. With frequency-indexed true positives TP_r , false positives FP_r , and false negatives FN_r , define

$$F_{1,r} = \frac{2\text{TP}_r}{2\text{TP}_r + \text{FP}_r + \text{FN}_r}, \quad \bar{F}_1 = \frac{1}{F} \sum_{r=1}^F F_{1,r}. \quad (29)$$

The top- S protocol uses the known source count only in this controlled study and is not an operational detector. For the dense Conventional maps, assignment uses the beamforming-score maximum, and power and axis are read from the same unsmoothed focus point. Appendix C defines the method-specific reductions used only in the cross-method appendix.

Let $\mathcal{D}$ contain the frequency-source pairs assigned within r_{match} . Misses affect $\bar{F}_1$ but do not enter conditional errors. The conditional normalized position error is

$$e_{y,r,s} = \frac{\|\hat{\mathbf{y}}_{r,s} - \mathbf{y}_s\|_2}{D}, \quad (r, s) \in \mathcal{D}. \quad (30)$$

Here q denotes source-coefficient power in the chosen Green's function normalization, not radiated acoustic power in watts. The synthetic dipoles use $\chi = 1$. To compare the monopole baseline with their coefficients, we remove the leading wavenumber dependence:

$$\tilde{q}_{\text{mono}} = \frac{\hat{q}_{\text{mono}}}{k^2}, \quad k = \frac{2\pi f}{c} = \frac{2\pi H e_D}{D}. \quad (31)$$

This reference equates a dipole observed on its axis by a point receiver in the far field with the corresponding monopole estimate. Near-field terms and the different reception directions across an extended array remain as source-model mismatch. The correction uses no ground-truth axis and changes neither the localization score nor the detected positions.

Level error is conditional on successful localization and positive finite estimated and true powers. For the corresponding set $\mathcal{D}_q \subseteq \mathcal{D}$, its mean absolute error (MAE) is

$$\text{MAE}_L = \frac{1}{|\mathcal{D}_q|} \sum_{(r,s) \in \mathcal{D}_q} \left| 10 \log_{10} \frac{\hat{q}_{r,s}}{q_{r,s}} \right|. \quad (32)$$

For the monopole baseline, $\tilde{q}_{\text{mono}}$ replaces $\hat{q}$ in [equation \(32\)](#); tables identify this wavenumber-corrected comparison.

Dipole orientation is an axis rather than a directed vector. The source-wise orientation error in degrees therefore uses

$$\theta_{r,s}^\circ = \frac{180}{\pi} \arccos \left[\text{clip}_{[0,1]}(|\hat{\mathbf{n}}_{r,s}^T \mathbf{n}_s|) \right], \quad 0^\circ \leq \theta_{r,s}^\circ \leq 90^\circ, \quad (33)$$

which is sign invariant and evaluated only for detected pairs with a defined axis. Position, level, and axis errors are therefore conditional and must be read with $\bar{F}_1$.

3.3 Randomized finite-snapshot protocol

The second experiment examines how source separation, received source-power contrast, sensor-noise level, and snapshot count affect localization and axis recovery. Each realization contains two continuous off-grid positions at separation δ and two independent, uniformly distributed unoriented dipole axes. The array, 101×101 grid, focus plane, and ten Helmholtz numbers match [section 3.1](#).

We use frequency-by-frequency Profile-CBF with diagonal removal: $\mathbf{K}_A = \mathcal{K}[\hat{\mathbf{C}}^{(W)}]$, $\mathbf{K}_B = \mathbf{G}$, and $W_{mj} = 1 - \delta_{mj}$. This analytically maximizes the normalized localization score after diagonal removal. We report source recovery and sign-invariant axis error; Appendix [C.2](#) specifies the sampling.

Let $\kappa_{r,s} = M^{-1} \|\mathbf{g}_{r,s}\|_2^2$ be source-to-array coupling. The received source-power contrast is

$$\Delta L_{\text{rec},r} = 10 \log_{10} \frac{q_{r,1} \kappa_{r,1}}{q_{r,2} \kappa_{r,2}}, \quad q_{r,1} \kappa_{r,1} = 1, \quad q_{r,2} \kappa_{r,2} = 10^{-\Delta L_{\text{rec},r}/10}. \quad (34)$$

The strong received power is fixed while the weak source is reduced; noise is recalibrated from the resulting total signal power.

For statistically independent snapshots ℓ , synthetic pressure and the sample CSM are

$$\mathbf{p}_r^{(\ell)} = \sum_{s=1}^2 \mathbf{g}_{r,s} a_{r,s}^{(\ell)} + \mathbf{e}_r^{(\ell)}, \quad \hat{\mathbf{C}}_r^{(K)} = \frac{1}{K} \sum_{\ell=1}^K \mathbf{p}_r^{(\ell)} (\mathbf{p}_r^{(\ell)})^H, \quad (35)$$

where $a_{r,s}^{(\ell)} \sim \mathcal{CN}(0, q_{r,s})$ independently over r , s , and ℓ , and $\mathbf{e}_r^{(\ell)} \sim \mathcal{CN}(\mathbf{0}, \sigma_{e,r}^2 \mathbf{I})$ independently of every source coefficient. The sensor signal-to-noise ratio (SNR) is defined by

$$\text{SNR}_{e,r} = 10 \log_{10} \left(\frac{M^{-1} \text{tr} \mathbf{C}_{r,\text{sig}}}{\sigma_{e,r}^2} \right), \quad \mathbf{C}_{r,\text{sig}} = \sum_{s=1}^2 q_{r,s} \mathbf{g}_{r,s} \mathbf{g}_{r,s}^H. \quad (36)$$

Thus 0 dB gives equal mean signal and white-noise auto-power. Every scan receives the diagonal-removed (DR) CSM

$$\hat{\mathbf{C}}_r^{\text{DR}} = \hat{\mathbf{C}}_r^{(K)} - \text{diag}(\text{diag} \hat{\mathbf{C}}_r^{(K)}). \quad (37)$$

White sensor noise disappears from the ensemble off-diagonal CSM but not from finite-sample fluctuations, whose scale is $\mathcal{O}_p(K^{-1/2})$. Appendix C.2 gives the error decomposition.

The separation sweep uses all frequencies and the dimensionless proxy

$$\xi = He_D \|\mathbf{u}_{1,\parallel} - \mathbf{u}_{2,\parallel}\|_2, \quad \mathbf{u}_s = \frac{\mathbf{y}_s}{\|\mathbf{y}_s\|_2}, \quad (38)$$

where the subscript $\parallel$ denotes projection onto the array plane. Thus $2\pi\xi$ is a first-order aperture-scale phase-gradient separation, not a universal near-field resolution law. Appendix C.2 lists the separation, contrast, snapshot, and SNR levels.

Each condition uses 200 common-random-number replicates. The known two-source model selects two positive maxima and accepts assignments within $\min(0.06D, 0.49\delta)$. A pair is recovered only if both assignments succeed; axis error is conditional on assignment. Recovery uses Wilson intervals, and pooled or conditional quantities use a 2000-resample replicate-cluster bootstrap.

3.4 Open-jet wind-tunnel setup

The open-jet measurement uses a high-lift half-model previously investigated with monopole CLEAN-SC [17] (Fig. 5a), at $M_\infty = 0.2$ and $c = 344 \text{ m s}^{-1}$. A 512-microphone subset of the approximately $6 \times 3 \text{ m}$ array was about 3 m from the model; 70,821 surface-near focus points had 20 mm spacing. The 30 s, 48 kHz record yielded 1,404 Welch blocks (2,048-sample Hann windows, 50% overlap, $\Delta f = 23.44 \text{ Hz}$). B-CLEAN-SC with unknown dipole orientations (Appendix B.2) was applied separately to each octave, sharing one axis per reconstructed source only within that band. It used loop gain 0.25, at most 80 iterations, 20-dB dynamic range, diagonal removal, and Amiet shear-layer correction, without an orientation or surface-normal prior.

4 Computational implementation

For efficient CPU and GPU evaluation after eliminating the angular scan, we focus on the remaining microphone-pair contractions. The Hermitian CSM requires only

$$N_\Delta = \frac{M(M+1)}{2} \quad (39)$$

stored upper-triangular entries. The CSM-dependent contractions in equations (11) and (12) can therefore consume a packed upper triangle directly, and diagonal removal omits its M diagonal entries. The six independent entries of $\mathbf{S}_W$ are accumulated from packed pairs, whereas $\mathbf{G}$ requires an $\mathcal{O}(M)$ sensor reduction. Across the grid, the dominant work is $\mathcal{O}(FPN_\Delta)$ plus $\mathcal{O}(FPM)$ dipole transfer-vector calculation; the factorized angular-grid comparison is given in Appendix D.1.

Hardware-adaptive focus blocks avoid materializing the full $F \times 3 \times M \times P$ basis. CPU execution parallelizes frequencies and blocks; fused GPU kernels keep contractions, local solves, residuals, and outputs device resident. Appendix D gives the packed contraction, numerical guards, precision, and transfer policy.

Starting from a prepared CSM and focus coordinates, runtime measures the beam-former call through completed outputs, including Green’s function evaluation, local algebra, and any validation or device transfers within that call. Reusable pre-call setup, CSM generation, short-time Fourier transform (STFT), final host export, and serialization are excluded. Values are synchronized medians of nine runs after two warm-ups. CPU thread pools use at most 16 threads; gridded inverse profiles limit BLAS to one thread per frequency task.

5 Results

We compare spatial reconstructions, distinguish localization from power, and assess axis estimation, frequency coupling, finite-snapshot performance, and transfer to wind-tunnel data.

5.1 Spatial reconstruction and source power

Figure 2 compares monopole CBF and Profile-CBF. It changes both source model and selection normalization and does not isolate the source-model effect. Increasing frequency improves separation of overlapping responses. Across all ten frequencies, the monopole baseline yields $\bar{F}_1 = 0.55$, $0.014D$ mean position error, and a wavenumber-corrected conditional level MAE of 3.40 dB. Profile-CBF raises $\bar{F}_1$ to 0.80 without a prescribed axis (Table 6).




Fig. 2 Conventional source-selection maps: (a)–(c) monopole CBF (formulation III); (d)–(f) frequency-local Profile-CBF localization score. Columns show $He_D = 1, 2, 4$, respectively. Each map is normalized to its own maximum over 20 dB. Orange marks ground-truth positions and axes; arrowless black curves show the estimated unoriented dipole-axis field and carry no polarity

Figure 3 separates JCBF’s shared localization score from its frequency-dependent power. The common axis and broadband score combine information from the entire band. Source power additionally compensates for axis-dependent array coupling, so its maxima need not coincide with the source-selection maxima.




Fig. 3 JCBF: (a) source power at $He_D = 4$; (b) shared broadband localization score. The score and dipole-axis field are common to all ten frequency bins; power remains frequency dependent. Each scalar map is normalized separately over 20 dB, and streamlines are shown within that range. Peaks are selected from (b). Symbols follow Fig. 2

5.2 Axis estimation and frequency coupling

Table 2 compares the monopole baseline, frequency-local Profile-CBF, and JCBF. Profile-CBF and JCBF use the same localization score and fixed-axis power estimate (section 2.5); their difference is frequency coupling.

Table 2 Source-model and frequency-coupling comparison. E_L is conditional level MAE (dB; wavenumber-corrected for the monopole), E_θ is conditional projective-axis MAE (degrees), and E_y is mean position error normalized by D . CPU/GPU columns report median production runtimes excluding final export. A dash denotes an inapplicable quantity

Method	Coupling	CPU (ms)	GPU (ms)	$\bar{F}_1$	E_L	E_θ	E_y
Monopole CBF	Local	215.1	2.4	0.55	3.40	–	0.014
Profile-CBF	Local	1137.5	14.8	0.80	0.09	1.23	< 0.001
JCBF	Joint	2429.6	17.4	1.00	0.49	2.30	< 0.001

Coupling Profile-CBF into JCBF raises $\bar{F}_1$ from 0.80 to 1.00; on the common detected subset of 32 frequency–source pairs, their conditional sign-invariant dipole-axis MAEs are 1.23° and 2.30°. Thus the broadband gain comes from recovering more sources, not from reducing conditional axis error for sources already detected locally.

5.3 Randomized off-grid and finite-snapshot stress test

Figure 4 reports pair-recovery probability and conditional sign-invariant axis error for the complete Profile-CBF localization and axis-estimation pipeline.




Fig. 4 Randomized two-source performance of pair-weighted frequency-by-frequency Profile-CBF with diagonal removal. The upper row reports pair-recovery probability; the lower row reports conditional median sign-invariant axis error on a shared logarithmic scale. Panels (a, d) pool physical separations and frequencies in ξ bins; (b, e) vary received source-power contrast; and (c, f) vary snapshot count with no sensor noise and at SNRs of 0, -6, -12, and -18 dB. Panels (b,c,e,f) use $H e_D = 8$ and $\delta/D = 0.20$; the final snapshot marker is the exact ensemble CSM. Bars are 95% Wilson or replicate-cluster-bootstrap intervals as defined in [section 3.3](#).

Pair-recovery probability rises from 0.625 at $\xi = 1.25$ to 0.987 at $\xi = 2.25$ and exceeds 0.997 from $\xi = 2.75$. At $\xi = 2.25$, conditional median axis and position errors are 1.66° and $0.00444D$, the latter near the $0.00399D$ nearest-grid-point quantization reference. Received source-power contrast exposes the expected Conventional dynamic-range limit: pair-recovery probability decreases from 0.985 for equal powers to 0 at 20 dB, where only the strong source survives. The lower conditional axis error in panel (e) therefore reflects the increasing dominance of the strong source among detections, not improved pair recovery.

At 0-dB sensor SNR, pair-recovery probability increases from 0.175 at $K = 1$ to 0.960 at $K = 64$ and 0.985 at $K = 128$. With noise 18 dB above the signal, $K = 64$ gives a pair-recovery probability of 0.170 and 19.40° conditional axis error;

$K = 512$ reaches 0.925 and 7.04° . At $K = \infty$, all SNR cases converge to a pair-recovery probability of 0.995 and a conditional axis error of 1.00° , because diagonal removal removes spatially white noise exactly at the ensemble level. The nonzero limit reflects off-grid localization and the two-source, diagonal-removed imaging criterion, which remain even with an exact CSM. At low K or SNR, the conditional axis errors describe the subset of easier scenes that were successfully assigned.

5.4 Open-jet wind-tunnel demonstration

For the measured case in [section 3.4](#), [Fig. 5](#) extends the previously published monopole CLEAN-SC analysis [\[17\]](#) with effective dipole-axis estimates. At 2 kHz, the estimated slat-track axes are predominantly transverse to the freestream; at 8 kHz, they align with the slat leading-edge direction. The observed change in orientation complements the distinct low-frequency Helmholtz and high-frequency Strouhal scaling regimes of the slat-track contribution reported previously [\[18\]](#).




Fig. 5 Open-jet demonstration. (a) Full-resolution model (gray), 20-mm focus shell (green), 512 microphones (orange), and modeled shear layer for the Amiet correction (blue). Blue arrows indicate freestream direction. (b, c) Dipole B-CLEAN-SC in the 2- and 8-kHz octaves, each normalized to its maximum over 20 dB. Colored surfaces show Gaussian-smoothed levels; red arrows show estimated dipole axes, for which opposite directions are equivalent. Arrow length is fixed.

6 Discussion

6.1 Implications for dipole imaging

The results demonstrate that established beamformers can be extended to estimate dipole orientation together with source location and level. The dipole extension determines the axis from microphone measurements without prescribing a surface-normal orientation, adding directional information to the reconstructed sources.

In practice, dense Conventional maps should be read at positions selected by the localization score: the power and axis at those same points characterize the detected sources. Selecting power maxima instead can associate these quantities with a different location ([Fig. 3](#)).

The usefulness of the recovered axis depends on spatial resolution and CSM quality. [Figure 4](#) shows that both affect localization and orientation. Where a beamformer

cannot separate overlapping sources, its estimated axis may describe their combined response rather than either source individually.

Frequency coupling is useful when a physical source retains its position and axis across a band: better-resolved frequencies can support localization in less-resolved bins. The ablation in Table 2 also shows why detection gains and axis accuracy must be assessed separately.

The wind-tunnel example illustrates why band selection matters: the estimated slat-track axes change with frequency (Fig. 5). For slender structures with nearby edges and interacting surfaces, these axes provide evidence for distinguishing source contributions without selecting a surface normal in advance. No independent orientation reference is available for this experiment, so the axes support interpretation rather than quantify orientation accuracy.

6.2 Scope and limitations

Validation across beamformer families is restricted to the controlled case (Appendix E). The randomized study tests only Profile-CBF and retains uncorrelated sources, matched propagation, and known source count. It therefore does not establish a cross-method robustness ranking.

The robustness curves evaluate the combined localization and orientation-estimation pipeline. Without a known-axis reference, losses cannot be attributed solely to orientation estimation. Conditional errors describe only assigned sources and can decrease when difficult sources cease to be detected.

The extension estimates a real, sign-invariant axis; more general complex-valued moments require a richer model. Poorly conditioned Gram matrices limit observable orientation components. Source coherence, calibration errors, and propagation mismatch can still bias the reconstruction; estimating orientation does not correct these model errors.

7 Conclusion

The proposed dipole extension enables established beamformers to reconstruct source location, power, and orientation without prescribed dipole axes or surface-normal assumptions. Expressing the dipole Green’s function in a three-component Cartesian basis replaces the angular scan with small local 3×3 eigenproblems. For sources with a common position and axis within a selected band, optional broadband coupling combines information across frequencies while retaining frequency-dependent source powers. The synthetic tests show how orientation recovery depends on spatial resolution and cross-spectral matrix quality, while the runtime measurements demonstrate efficient CPU and GPU implementations. In the open-jet experiment, estimated slat-track axes provide directional information beyond source location and level, supporting the interpretation of distinct aeroacoustic source mechanisms.

A Additional Cartesian results

A.1 Rank-deficient common-axis estimation

Some frequencies constrain only part of the dipole axis. Their observable projections can nevertheless determine a common axis. We extend the full-rank construction in [equation \(25\)](#) to this case. Let $\rho_{2,r,p}$ be the second supported generalized eigenvalue and define

$$\iota_{r,p}^{\text{uniq}} = \begin{cases} 1, & \rho_{1,r,p} > \rho_{2,r,p}, \\ 0, & \rho_{1,r,p} = \rho_{2,r,p}, \end{cases} \quad (40)$$

with value one for a valid one-dimensional supported space. A repeated extremum therefore contributes no directional evidence; in particular, spatially white sensor noise has a constant quotient over the observable space. Each local solution is represented in $\text{range}(\mathbf{G}_{r,p})$ with projector

$$\mathbf{P}_{r,p} = \mathbf{G}_{r,p}^\dagger \mathbf{G}_{r,p} = \text{proj}[\text{range}(\mathbf{G}_{r,p})]. \quad (41)$$

The construction combines complementary supported components from different frequency bins; it does not require the physical axis to lie in the intersection of all local ranges. Using the evidence weights $\zeta_{r,p}$ from [equation \(25\)](#), define

$$\boxed{\begin{aligned} \mathbf{A}_p^J &= \sum_{r=1}^F \zeta_{r,p} \mathbf{P}_{r,p} (\mathbf{I} - \hat{\mathbf{n}}_{r,p} \hat{\mathbf{n}}_{r,p}^T) \mathbf{P}_{r,p}, \\ \mathbf{B}_p^J &= \sum_{r=1}^F \zeta_{r,p} \mathbf{P}_{r,p}, \quad \mathbf{A}_p^J, \mathbf{B}_p^J \in \mathbb{R}^{3 \times 3}. \end{aligned}} \quad (42)$$

A unique common-axis estimate exists when $\mathbf{B}_p^J$ has a supported mode and the minimum supported generalized eigenvalue is simple. It is given by

$$(\lambda_p^J, \hat{\mathbf{n}}_p^J) = \text{geig}_{\min}(\mathbf{A}_p^J, \mathbf{B}_p^J), \quad (43)$$

which minimizes

$$\frac{\sum_r \zeta_{r,p} \|(\mathbf{I} - \hat{\mathbf{n}}_{r,p} \hat{\mathbf{n}}_{r,p}^T) \mathbf{P}_{r,p} \mathbf{n}\|_2^2}{\sum_r \zeta_{r,p} \|\mathbf{P}_{r,p} \mathbf{n}\|_2^2} \quad \text{over } \mathbf{n} \in \text{range}(\mathbf{B}_p^J). \quad (44)$$

If all $\mathbf{G}_{r,p}$ are positive definite, $\mathbf{P}_{r,p} = \mathbf{I}$ and

$$\mathbf{A}_p^J = \left(\sum_r \zeta_{r,p}\right) \mathbf{I} - \mathbf{T}_p, \quad \mathbf{B}_p^J = \left(\sum_r \zeta_{r,p}\right) \mathbf{I}, \quad (45)$$

so the minimum generalized eigenvector is exactly the principal eigenvector of $\mathbf{T}_p$ used in the main text.

Proposition 1 (exact joint recovery). At the true focus point, let the matched, unweighted, noise-free CSMs be

$$\mathbf{C}_r = q_r \mathbf{V}_{r,p} \mathbf{n}_p \mathbf{n}_p^T \mathbf{V}_{r,p}^H, \quad q_r \geq 0, \quad (46)$$

with $\alpha_r q_r > 0$ in at least one coupled bin, aggregate observability $\mathbf{n}_p \in \text{range}(\mathbf{B}_p^J)$, and projective identifiability: the combined observable projections determine a unique axis up to sign. Cauchy–Schwarz gives

$$\frac{q_r |(\mathbf{V}_{r,p} \mathbf{n})^H \mathbf{V}_{r,p} \mathbf{n}_p|^2}{\|\mathbf{V}_{r,p} \mathbf{n}\|_2^2} \leq q_r \|\mathbf{V}_{r,p} \mathbf{n}_p\|_2^2, \quad (47)$$

and, for every contributing bin ($\zeta_{r,p} > 0$), the local representative satisfies

$$\hat{\mathbf{n}}_{r,p} \hat{\mathbf{n}}_{r,p}^T = \frac{(\mathbf{P}_{r,p} \mathbf{n}_p)(\mathbf{P}_{r,p} \mathbf{n}_p)^T}{\|\mathbf{P}_{r,p} \mathbf{n}_p\|_2^2}. \quad (48)$$

Every summand of $\mathbf{A}_p^J$ consequently annihilates $\mathbf{n}_p$. The objective is nonnegative, and projective identifiability makes its zero minimum unique, proving $\hat{\mathbf{n}}_p^J = \pm \mathbf{n}_p$. Without aggregate observability, only the representative modulo the common null space is identifiable.

A.2 Conditional Cartesian covariance interpretation

The Cartesian matrices also admit a source-covariance interpretation under more restrictive assumptions than Profile-CBF. Suppose an isolated dipole produces $\mathbf{C} = q \mathbf{V} \mathbf{n} \mathbf{n}^T \mathbf{V}^H$, $q \geq 0$, without pair weighting, and $\mathbf{V}^H \mathbf{V}$ is real. Then $\mathbf{G} = \mathbf{V}^H \mathbf{V}$ and $\mathbf{S} = q \mathbf{G} \mathbf{n} \mathbf{n}^T \mathbf{G}$. Supported pseudoinversion gives

$$\hat{\mathbf{Q}}_{\text{cart}} = \mathbf{G}^\dagger \mathbf{S} \mathbf{G}^\dagger = q (\mathbf{P}_G \mathbf{n})(\mathbf{P}_G \mathbf{n})^T, \quad \mathbf{P}_G = \mathbf{G}^\dagger \mathbf{G}. \quad (49)$$

If $\mathbf{n} \in \text{range}(\mathbf{G})$, this is the covariance $q \mathbf{n} \mathbf{n}^T$ of the three Cartesian dipole coefficients; otherwise only the projected axis is recovered. For an observable axis and $q > 0$, its principal eigenvector gives the dipole axis and its trace gives the source-coefficient power. Summing such covariances across frequencies preserves a common axis and adds source powers. This differs from JCBF’s sum of unit-axis outer products weighted by localization scores, which is not a source-power covariance. The real-Gram condition holds for the homogeneous free-field derivative basis, whose sensor rows are complex scalars times real direction vectors. A general complex basis, overlapping sources, noise, or diagonal removal need not satisfy the identity. It is an interpretation of the isolated-source model, not the estimator used in the Profile-CBF comparisons.

B Orientation selection for other methods

Table 3 summarizes the local operators and how each method uses the selected axis. The following derivations specify the normalizations, residual updates, and inverse solvers.

Inverse methods estimate source powers by combining modeled responses from possible sources on the focus grid. These responses form a *dictionary*, and each response is an *atom*. In CSM fitting, an atom is the unit-power source contribution $\mathbf{g}\mathbf{g}^H$; in map deconvolution, it is the corresponding point-spread function. The orientation estimate supplies $\mathbf{g}_{r,p}^* = \mathbf{V}_{r,p}\hat{\mathbf{n}}_{r,p}$ to construct these responses. Examples include the deconvolution approach for mapping of acoustic sources (DAMAS) and multifrequency sparse Bayesian learning (MF-SBL). An axis estimated for an unresolved mixture represents an effective direction.

Table 3 Orientation selection across beamforming methods

Construction	Local operator or result	Local extremum	Use after the local step
Profile-CBF/JCBF	$\mathcal{K}_{r,p}[\mathbf{C}_r]$ or $\mathcal{K}_{r,p}[\mathbf{W}_r \odot \mathbf{C}_r]$	Maximum	Beamforming score; then fixed-axis source power
CLEAN-SC/B-CLEAN-SC	$\mathcal{K}_{r,p}[\mathbf{C}_r^{\text{res}}]$	Maximum	Standard coherent residual subtraction
MVDR	$\mathcal{K}_{r,p}[\mathbf{C}_{r,\lambda}^{-1}]$	Minimum	Pseudospectrum only
MUSIC Power Fit	$\mathcal{K}_{r,p}[\mathbf{E}_{r,\text{noise}}\mathbf{E}_{r,\text{noise}}^H]$	Minimum	Oriented peaks, then nonnegative covariance fitting
DAMAS/CSM Fit/MF-SBL	Dictionary-axis estimate	–	Oriented atom for the inverse solver

B.1 Adaptive and subspace scans

MVDR and MUSIC select the axis with the smallest inverse-CSM response or noise-subspace projection, respectively. Normalize each observable dipole response to unit norm:

$$\begin{aligned} \mathcal{D}_{r,p} &= \{\mathbf{n} \in \text{range}(\mathbf{G}_{r,p}) : \mathbf{n}^T \mathbf{G}_{r,p} \mathbf{n} > 0\}, \\ \bar{\mathbf{g}}_{r,p}(\mathbf{n}) &= \frac{\mathbf{V}_{r,p} \mathbf{n}}{\sqrt{\mathbf{n}^T \mathbf{G}_{r,p} \mathbf{n}}}, \quad \|\bar{\mathbf{g}}_{r,p}(\mathbf{n})\|_2 = 1. \end{aligned} \quad (50)$$

For MVDR, choose the loading in $\mathbf{C}_{r,\lambda} = \mathbf{C}_r + \lambda_r \mathbf{I}$ so that $\mathbf{C}_{r,\lambda} \succ 0$; a positive λ_r alone does not ensure this for an indefinite input. The orientation-adaptive profiles are

$$\eta_{r,p}^{\text{MVDR}} = \min_{\mathbf{n} \in \mathcal{D}_{r,p}} \frac{\mathbf{n}^T \mathcal{K}_{r,p}[\mathbf{C}_{r,\lambda}^{-1}] \mathbf{n}}{\mathbf{n}^T \mathbf{G}_{r,p} \mathbf{n}}, \quad \mathcal{P}_{r,p}^{\text{MVDR}} = 1/\eta_{r,p}^{\text{MVDR}}, \quad (51)$$

$$\eta_{r,p}^{\text{MUSIC}} = \min_{\mathbf{n} \in \mathcal{D}_{r,p}} \frac{\mathbf{n}^T \mathcal{K}_{r,p} [\mathbf{E}_{r,\text{noise}} \mathbf{E}_{r,\text{noise}}^H] \mathbf{n}}{\mathbf{n}^T \mathbf{G}_{r,p} \mathbf{n}}, \quad \mathcal{P}_{r,p}^{\text{MUSIC}} = 1/\eta_{r,p}^{\text{MUSIC}}. \quad (52)$$

Here $\mathbf{E}_{r,\text{noise}}$ contains the noise-subspace eigenvectors. Both quantities are normalized localization pseudospectra, not source-coefficient power; in particular, the MVDR expression is not a calibrated Capon source-power estimate. MUSIC applies a relative quotient floor of 10^{-7} , selects peaks from the oriented pseudospectrum, and obtains power from a nonnegative covariance fit.

To express the covariance fit as a real least-squares problem without changing its Frobenius norm, vectorize a Hermitian matrix $\mathbf{X} \in \mathbb{C}^{M \times M}$ as

$$\begin{aligned} \text{vec}_R(\mathbf{X}) &= \text{col} \left((X_{mm})_{m=1}^M, (\sqrt{2} \text{Re } X_{mj})_{1 \leq m < j \leq M}^{\text{lex}}, (\sqrt{2} \text{Im } X_{mj})_{1 \leq m < j \leq M}^{\text{lex}} \right) \in \mathbb{R}^{M^2}, \\ \|\text{vec}_R(\mathbf{X})\|_2 &= \|\mathbf{X}\|_F. \end{aligned} \quad (53)$$

For oriented support $\mathcal{S}$, MUSIC Power Fit reports

$$\hat{\mathbf{q}}_{r,\mathcal{S}} \in \arg \min_{\mathbf{q}_{r,\mathcal{S}} \geq 0} \left\| \text{vec}_R(\mathbf{C}_r) - \sum_{p \in \mathcal{S}} q_{r,p} \text{vec}_R(\mathbf{g}_{r,p}^* (\mathbf{g}_{r,p}^*)^H) \right\|_2^2. \quad (54)$$

B.2 CLEAN-SC and B-CLEAN-SC

CLEAN-SC estimates an axis at each focus point from the current residual, selects the score maximum, and subtracts its coherent contribution. At each iteration, axis estimation and peak selection are

$$\begin{aligned} (\rho_{r,p}^*, \hat{\mathbf{n}}_{r,p}) &= \text{geig}_{\max}(\mathcal{K}_{r,p}[\mathbf{C}_r^{\text{res}}], \mathbf{G}_{r,p}), \\ p^* &\in \arg \max_{p \in \mathcal{P}_{\text{valid},r}} \rho_{r,p}^*. \end{aligned} \quad (55)$$

The unweighted source-coefficient increment uses loop gain μ and the selected axis's array coupling:

$$\beta_{r,p} = \hat{\mathbf{n}}_{r,p}^T \mathbf{G}_{r,p} \hat{\mathbf{n}}_{r,p} = \|\mathbf{V}_{r,p} \hat{\mathbf{n}}_{r,p}\|_2^2, \quad \Delta \hat{q}_{r,p^*} = \mu \frac{\rho_{r,p^*}^*}{\beta_{r,p^*}}, \quad (56)$$

Pair weighting replaces the increment by $\mu \rho_{r,p^*}^* \beta_{r,p^*} / \delta_{W_r}(\mathbf{g}_{r,p^*})$.

B-CLEAN-SC solves the same narrowband orientation problems and combines them into one common axis per focus point and iteration. It then selects the broadband score maximum using [equation \(27\)](#):

$$\begin{aligned} (\rho_{1,r,p}, \hat{\mathbf{n}}_{r,p}) &= \text{geig}_{\max}(\mathcal{K}_{r,p}[\mathbf{C}_r^{\text{res}}], \mathbf{G}_{r,p}), \\ p^* &\in \arg \max_{p \in \mathcal{P}_{\text{valid}}^J} \bar{\rho}_p. \end{aligned} \quad (57)$$

Per-frequency peak normalization limits domination by one energetic bin; powers and coherent subtractions remain frequency specific.

Let $\mathbf{g}_{r,p}^*$ be the selected source transfer vector, $\mathbf{u}_r = \mathbf{g}_{r,p}^* / \|\mathbf{g}_{r,p}^*\|_2$, $\pi_r = \mathbf{u}_r^H \mathbf{C}_r^{\text{res}} \mathbf{u}_r > 0$, and $\mathbf{d}_r = \mathbf{C}_r^{\text{res}} \mathbf{u}_r / \pi_r$. For a binary retained-pair mask $\mathbf{M}_r = \mathbf{M}_r^T \in \{0, 1\}^{M \times M}$, CLEAN-SC computes the coherent vector by a fixed-point inner iteration, followed by residual subtraction:

$$\begin{aligned} \mathcal{H}_r(\mathbf{h}) &= (\mathbf{1} - \mathbf{M}_r) \odot \mathbf{h} \mathbf{h}^H, \\ \mathbf{h}_r &= \frac{\mathbf{d}_r + \mathcal{H}_r(\mathbf{h}_r) \mathbf{u}_r}{\sqrt{1 + \mathbf{u}_r^H \mathcal{H}_r(\mathbf{h}_r) \mathbf{u}_r}}, \end{aligned} \quad (58)$$

$$\mathbf{C}_r^{\text{res}} \leftarrow \mathbf{C}_r^{\text{res}} - \mu \pi_r (\mathbf{M}_r \odot \mathbf{h}_r \mathbf{h}_r^H). \quad (59)$$

For the full mask, $\mathbf{h}_r = \mathbf{d}_r$ and the local refresh reduces to

$$\mathcal{K}_{r,p}[\mathbf{C}_r^{\text{res}}] \leftarrow \mathcal{K}_{r,p}[\mathbf{C}_r^{\text{res}}] - \mu \pi_r \text{Re}[\mathbf{V}_{r,p}^H \mathbf{h}_r \mathbf{h}_r^H \mathbf{V}_{r,p}]. \quad (60)$$

Arbitrary binary masks or continuous pair weights require the packed-pair contraction rather than this $\mathcal{O}(M)$ projection.

B.3 DAMAS and covariance-fitting methods

The evaluated inverse methods first estimate a shared axis for each candidate source by maximizing a ratio of sums (RoS):

$$\hat{\mathbf{n}}_p^{\text{RoS}} = \arg \max_{\substack{\mathbf{n} \in \text{range}(\mathbf{G}_p) \\ \mathbf{n}^T \mathbf{G}_p \mathbf{n} > 0}} \frac{\mathbf{n}^T (\sum_{r=1}^F \alpha_r \mathcal{K}_{r,p}[\mathbf{C}_r]) \mathbf{n}}{\mathbf{n}^T \mathbf{G}_p \mathbf{n}}, \quad \overline{\mathbf{G}}_p = \sum_{r=1}^F \alpha_r \mathbf{G}_{r,p}. \quad (61)$$

Frequency-dependent power and nonproportional Gram matrices can rotate the resulting axis. This step supplies oriented source responses to the inverse solver; it does not exactly optimize that solver's objective over orientation.

For oriented atoms, the formulation-III DAMAS system is

$$b_{r,i} = \frac{\mathbf{g}_{r,i}^H \mathbf{C}_r \mathbf{g}_{r,i}}{(\mathbf{g}_{r,i}^H \mathbf{g}_{r,i})^2}, \quad H_{r,ij} = \frac{|\mathbf{g}_{r,i}^H \mathbf{g}_{r,j}|^2}{(\mathbf{g}_{r,i}^H \mathbf{g}_{r,i})^2}, \quad (62)$$

$$\hat{\mathbf{q}}_r = \arg \min_{\mathbf{q}_r \geq 0} \|\mathbf{b}_r - \mathbf{H}_r \mathbf{q}_r\|_2^2. \quad (63)$$

The reported reduced-support variant first accepts S separated atoms by residual orthogonal matching pursuit (OMP). Its nonnegative least-squares (NNLS) runtime is not a full-grid DAMAS runtime.

Define the oriented covariance atom and residual vector

$$\mathbf{a}_{r,p}(\mathbf{n}) = \text{vec}_R[\mathbf{V}_{r,p} \mathbf{n} \mathbf{n}^T \mathbf{V}_{r,p}^H], \quad \mathbf{c}_r^{\text{res}} = \text{vec}_R(\mathbf{C}_r^{\text{res}}). \quad (64)$$

For a proposed shared axis, the exact nonnegative one-atom coefficient and bandwise residual decrease are

$$q_{r,p}^{\text{trial}} = \max\left(\frac{\mathbf{a}_{r,p}^T \mathbf{c}_r^{\text{res}}}{\mathbf{a}_{r,p}^T \mathbf{a}_{r,p}}, 0\right),$$

$$\Delta J_p = \sum_{r=1}^F \left[\|\mathbf{c}_r^{\text{res}}\|_2^2 - \|\mathbf{c}_r^{\text{res}} - q_{r,p}^{\text{trial}} \mathbf{a}_{r,p}\|_2^2 \right]. \quad (65)$$

CSM Fit NNLS accepts the largest positive decrease and solves on support $\mathcal{S}$

$$\hat{\mathbf{q}}_{\mathcal{S}} \in \arg \min_{\mathbf{q}_{\mathcal{S}} \geq 0} \sum_{r=1}^F \left\| \mathbf{C}_r - \sum_{p \in \mathcal{S}} q_{r,p} \mathbf{g}_{r,p}^* (\mathbf{g}_{r,p}^*)^H \right\|_F^2. \quad (66)$$

Projected MF-SBL applies automatic relevance determination (ARD) to N_a selected atoms collected in $\Phi_r \in \mathbb{R}^{M^2 \times N_a}$ and uses $\mathbf{\Gamma} = \text{diag}(\gamma)$ with $\gamma \in \mathbb{R}_{>0}^{N_a}$ shared over frequency. One projected reweighted-ridge step is

$$\hat{\mathbf{x}}_r = \Pi_+ [(\Phi_r^T \Phi_r + \sigma_r^2 \mathbf{\Gamma}^{-1})^{-1} \Phi_r^T \text{vec}_R(\mathbf{C}_r)], \quad (67)$$

$$\Sigma_r = \sigma_r^2 (\Phi_r^T \Phi_r + \sigma_r^2 \mathbf{\Gamma}^{-1})^{-1}, \quad \gamma_p \leftarrow \frac{1}{F} \sum_{r=1}^F (\hat{x}_{r,p}^2 + (\Sigma_r)_{pp}). \quad (68)$$

Here Π_+ projects coefficients onto the nonnegative orthant, and the fixed model-error variance sets the regularization scale:

$$\sigma_r^2 = \max\left(10^{-8} \frac{\|\text{vec}_R(\mathbf{C}_r)\|_2^2}{M^2}, 10^{-12}\right). \quad (69)$$

Rows at the pruning threshold are removed, so $\mathbf{\Gamma}$ remains positive definite on the active set. The source-selection quantity and reported coefficient are

$$\nu_p = \left(\sum_{r=1}^F \hat{x}_{r,p}^2 \right)^{1/2}, \quad \hat{q}_{r,p} = \hat{x}_{r,p}. \quad (70)$$

The projection and fixed model-error variance are computational approximations rather than an exact complex-Wishart likelihood.

B.4 Gridless comparator

Gridless CSM Fit jointly optimizes position and axis and does not use the local orientation-selection step:

$$\min_{\mathbf{y}_s, \|\mathbf{n}_s\|_2=1, q_{r,s} \geq 0} \sum_{r=1}^F \left\| \mathbf{C}_r - \sum_{s=1}^S q_{r,s} \mathbf{V}_r(\mathbf{y}_s) \mathbf{n}_s \mathbf{n}_s^T \mathbf{V}_r(\mathbf{y}_s)^H \right\|_F^2. \quad (71)$$

It requires a source count and initialization policy and provides the continuous-position reference in the controlled comparison.

C Numerical and robustness details

Table 4 lists the result-defining settings for the controlled comparison.

Table 4 Numerical settings used in the controlled cross-method comparison

Method	Settings
Conventional	Monopole: formulation III. Profile-CBF: formulation IV score maximized locally, then formulation III power. JCBF: the same score and power at a common axis; all ten bins coupled with weights normalized by initial score maxima
CLEAN-SC	Loop gain $\mu = 1$; at most 10 extracted components and 10 inner source-component iterations; 20 dB stopping range
B-CLEAN-SC	Loop gain $\mu = 0.5$; at most 20 extracted components and 10 inner source-component iterations; 20 dB stopping range; all ten bins coupled
MUSIC Power Fit	Signal-subspace rank and retained support 4; relative floor 10^{-7} ; -18 dB support threshold; at most 300 NNLS iterations, tolerance 10^{-8}
Reduced-support DAMAS-NNLS and CSM Fit NNLS	Residual OMP with $K_{\text{OMP}} = S = 4$; reduced NNLS cap 300 and tolerance 10^{-8}
Projected MF-SBL	$K_{\text{OMP}} = 2S = 8$ seeds in a 2048-point dirty-score pool; 32-atom candidate set; 4 retained rows; at most 80 ARD iterations; tolerance and relative noise floor 10^{-8} ; damping 0.35
Gridless CSM Fit	Four sources; one optimizer start within $0.035D$ in the focus plane; analytic derivatives; at most 120 optimizer iterations; L-BFGS-B on the CPU and Adam on the GPU

C.1 Cross-method output reductions

Ground-truth-centered regions of interest (ROIs) of radius $0.06D$ are used only to summarize successfully detected cross-method outputs. Let $\mathcal{I}_s$ contain the focus points in the ROI associated with assigned source s . Dense maps use the beamforming-score-selected focus point; sparse maps sum raw coefficients and form a sign-invariant

effective axis:

$$\hat{q}_{r,s}^{\text{ROI}} = \sum_{p \in \mathcal{I}_s} \hat{q}_{r,p}, \quad \hat{\mathbf{n}}_{r,s}^{\text{ROI}} = \text{eigvec}_{\max} \left(\sum_{p \in \mathcal{I}_s} \hat{q}_{r,p} \hat{\mathbf{n}}_{r,p} \hat{\mathbf{n}}_{r,p}^T \right). \quad (72)$$

These reductions define the cross-method level and axis errors, not the wind-tunnel display smoothing. MUSIC uses fitted source power, CLEAN methods sum accepted coefficient increments, and gridless results are matched at their native positions.

C.2 Random-scene construction and uncertainty

For a requested separation δ , draw $\varphi \sim \mathcal{U}(0, 2\pi)$ and set $\mathbf{h} = (\delta/2)(\cos \varphi, \sin \varphi)$. The pair center is uniform over the rectangle that keeps $\mathbf{c} \pm \mathbf{h}$ at least $0.025D$ inside the $[-0.5D, 0.5D]^2$ focus domain, and

$$\mathbf{y}_{1,2} = (\mathbf{c} \mp \mathbf{h}, 0.5D). \quad (73)$$

For each source, $\mathbf{z}_s \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3)$ and $\mathbf{n}_s = \mathbf{z}_s / \|\mathbf{z}_s\|_2$, with sign canonicalized only for storage. This samples the projective axis uniformly because all reported quantities are invariant under $\mathbf{n}_s \mapsto -\mathbf{n}_s$. Deterministic substreams from master seed 20260815 reuse the same positions, axes, and standard-normal draws across levels of each one-factor sweep.

The separation sweep uses $\delta/D \in \{0.025, 0.05, 0.075, 0.10, 0.15, 0.20, 0.30, 0.40\}$, pooled in ξ bins of width 0.5, at equal received powers, $K = 128$, and 0-dB sensor SNR. Both remaining sweeps fix $He_D = 8$ and $\delta/D = 0.20$. The contrast sweep uses $\{0, 3, 6, 10, 20\}$ dB at $K = 128$ and 0-dB sensor SNR. The snapshot sweep uses equal received source powers and $K \in \{1, 2, 4, 8, 16, 32, 64, 128, 256, 512, \infty\}$, with no sensor noise or sensor SNRs of 0, -6, -12, and -18 dB.

Let $\mathbf{G}_r^{\text{src}} = [\mathbf{g}_{r,1}, \mathbf{g}_{r,2}] \in \mathbb{C}^{M \times 2}$ collect the two source responses, let $\mathbf{Q}_r = \text{diag}(q_{r,1}, q_{r,2})$, and denote snapshot covariances by hats. The finite-sample perturbation before diagonal removal can be decomposed as

$$\begin{aligned} \hat{\mathbf{C}}_r - \mathbf{C}_r &= \mathbf{G}_r^{\text{src}} (\hat{\mathbf{Q}}_r - \mathbf{Q}_r) (\mathbf{G}_r^{\text{src}})^H \\ &\quad + \mathbf{G}_r^{\text{src}} \hat{\mathbf{R}}_{ae,r} + \hat{\mathbf{R}}_{ae,r}^H (\mathbf{G}_r^{\text{src}})^H + \hat{\mathbf{R}}_{ee,r} - \sigma_{e,r}^2 \mathbf{I}, \end{aligned} \quad (74)$$

where $\hat{\mathbf{R}}_{ae,r} = K^{-1} \sum_{\ell} \mathbf{a}_r^{(\ell)} (\mathbf{e}_r^{(\ell)})^H$ and $\hat{\mathbf{R}}_{ee,r} = K^{-1} \sum_{\ell} \mathbf{e}_r^{(\ell)} (\mathbf{e}_r^{(\ell)})^H$. Every centered term is $\mathcal{O}_p(K^{-1/2})$. Diagonal removal cancels the ensemble term $\sigma_{e,r}^2 \mathbf{I}$ but retains random off-diagonal fluctuations. A -6-dB SNR step multiplies $\sigma_{e,r}^2$ by approximately four; when the noise-noise term dominates, approximately 16 times as many snapshots are therefore required for the same stochastic scale.

Each condition begins with 200 independent base replicates. The pair-recovery indicator is Bernoulli and receives a Wilson interval. For pooled separation bins, the resampling unit is the complete base replicate across its contributing frequencies and separations; conditional axis and position statistics resample only replicate clusters with defined values. Reported percentile intervals use 2000 resamples. Artifact

manifests retain contributing row counts, surviving replicate counts, and the common-random-number identities, so conditional uncertainty is not computed by treating frequency–source rows as independent.

D Angular-grid and computational details

The local matrices can be assembled directly from packed Hermitian CSMs, without expanding the full sensor matrix. For $1 \leq a \leq b \leq 3$, the weighted contraction is

$$(S_W)_{ab} = \text{Re} \left(\sum_{m=1}^M V_{ma}^* W_{mm} C_{mm} V_{mb} + \sum_{m < j} W_{mj} [V_{ma}^* C_{mj} V_{jb} + V_{ja}^* C_{mj}^* V_{mb}] \right). \quad (75)$$

Only the stored upper triangle is traversed; diagonal removal sets $W_{mm} = 0$. Numerical tolerances distinguish observable directions from round-off and nearly tied orientation scores. We use $\epsilon_{32} = 2^{-23} \simeq 1.19 \times 10^{-7}$, the relative spacing near unity in single precision, and a safety factor of 32. This factor is a numerical margin, not a bit count. Both backends symmetrize the local matrices and retain a metric mode only when

$$\lambda_i(\mathbf{K}_B) > \tau_B, \quad \tau_B = 32\epsilon_{32} \max_j \lambda_j(\mathbf{K}_B). \quad (76)$$

The threshold τ_B scales with the largest metric eigenvalue; its dipole counterpart $\tau_{G,r,p}$ rejects axes with negligible array coupling:

$$\tau_{G,r,p} = 32\epsilon_{32} \lambda_{\max}(\mathbf{G}_{r,p}), \quad (77)$$

and the exact uniqueness test is replaced in finite precision by

$$\Delta\rho_{r,p} = \rho_{1,r,p} - \rho_{2,r,p} > \tau_{\text{gap}}^{\text{abs}} + \tau_{\text{gap}}^{\text{rel}} \max(|\rho_{1,r,p}|, |\rho_{2,r,p}|), \quad \tau_{\text{gap}}^{\text{abs}} = 0, \quad \tau_{\text{gap}}^{\text{rel}} = 32\epsilon_{32}. \quad (78)$$

Focus points without a supported mode, finite quotient, unique extremum, or valid denominator are excluded from $\mathcal{P}_{\text{valid}}$ and represented by a mask or $-\infty$ during maximization. Signed valid scores remain distinct from invalid focus points.

All $10 \times 10,201 = 102,010$ controlled-case Gram matrices are positive definite under this rule; the minimum observed ratio $\lambda_{\min}(\mathbf{G}_{r,p})/\lambda_{\max}(\mathbf{G}_{r,p})$ is 3.40×10^{-2} . The production implementation nevertheless forms the general projectors. Its GPU kernel fuses projector construction, frequency reduction, and the supported 3×3 solve without a frequency–focus tensor of projectors, so the measurements exercise the rank-deficient-capable path.

D.1 Factorized orientation-grid comparator

The Gao-style baseline evaluates the same normalized beamforming score $\rho_W(\mathbf{n})$ as Profile-CBF. For integer $L \geq 1$, its complete projective grid contains the north pole, rings $\theta_j = j\pi/(2L)$, $j = 1, \dots, L-1$, with $4L$ azimuths, and one equator with $2L$

azimuths. Antipodal duplicates are omitted, giving

$$N_{\Omega,L} = 1 + 2L(2L - 1), \quad r_{\text{cov},L} \leq \bar{r}_{\text{cov},L} := \arccos \left[\cos^2 \left(\frac{\pi}{4L} \right) \right]. \quad (79)$$

The actual projective covering radius is

$$r_{\text{cov},L} = \sup_{\mathbf{n} \in \mathbb{S}^2} \min_{\mathbf{d}_j \in \mathcal{G}_L} \arccos |\mathbf{n}^T \mathbf{d}_j|. \quad (80)$$

The displayed bound is a conservative half-cell bound, not necessarily the attained radius. The evaluated $L \in \{4, 8, 16\}$ yields 57, 241, and 993 axes. Because the Cartesian contraction is shared, factorization changes the grid cost from $\mathcal{O}(FPN_{\Omega}N_{\Delta})$ to $\mathcal{O}(FPN_{\Delta}) + \mathcal{O}(FPN_{\Omega})$; the analytic solution removes the second term and its covering error.

The timing comparison also fixes the unit-response power calculation; only the axis-selection step differs.

Table 5 Continuous 3×3 orientation optimization and projective angular grids. N_{Ω} is the number of axes, $\bar{r}_{\text{cov}}$ an upper bound on the projective covering radius, $\Delta\rho_{\text{max}}$ the maximum relative score deficit from the analytic optimum within its 20-dB range, and $E_{\theta,\text{eig}}$ the mean projective-axis difference. Times show median $\pm$ median absolute deviation in milliseconds from five runs after two warm-ups; quality values use GPU results after CPU–GPU parity verification

Estimator	N_{Ω}	$\bar{r}_{\text{cov}}$ ($^{\circ}$)	CPU (ms)	GPU (ms)	$\Delta\rho_{\text{max}}$ (%)	$E_{\theta,\text{eig}}$ ($^{\circ}$)	$\bar{F}_1$
Profile-CBF	—	—	1121.5 ± 0.7	13.4 ± 0.0	—	—	0.80
Gao-style angular grid	57	15.86	1083.0 ± 29.1	16.8 ± 0.4	30.55	9.33	0.80
Gao-style angular grid	241	7.95	1229.6 ± 11.3	14.8 ± 0.0	12.43	4.75	0.80
Gao-style angular grid	993	3.98	1998.3 ± 27.9	36.2 ± 0.0	3.64	2.38	0.80

Table 5 shows decreasing score deficits and axis differences as the angular grid is refined. Profile-CBF reaches the continuous optimum without a covering-resolution parameter or the additional $\mathcal{O}(FPN_{\Omega})$ term. This diagnostic path materializes the local matrices and three outputs; these diagnostic runtimes are kept separate from the production timings in Table 6.

D.2 Precision and benchmark setup

The controlled NVIDIA CUDA scans store CSMs and dipole transfer matrices as complex64, accumulate in float32, and promote the 3×3 eigensolve and stored axes to float64. No float16 or tensor-core arithmetic is used.

The one-factor campaign contains 13 unique cases around the baseline $P = 51^2$, $F = 10$, and $M = 64$: $P \in \{11^2, 51^2, 101^2, 501^2, 1001^2\}$, $F \in \{1, 5, 10, 50, 100\}$, and $M \in \{16, 32, 64, 128, 256\}$. Each point is a median of five synchronized runs after one warm-up. The compact comparison uses monopole CBF, Profile-CBF, and JCBF. Their dominant packed-Hermitian contraction is represented by the work coordinate $FP(M+1)/2$.

D.3 One-factor computational scaling

Figure 6 reports focus-grid scaling and all one-factor sweeps against the packed-work coordinate. The asymptotic trend follows the dominant contraction; GPU launch and synchronization overhead has its greatest relative effect at the smallest workloads.




Fig. 6 Runtime scaling of the measured conventional profiles on the CPU (a, b) and GPU (c, d). Panels (a, c) vary focus count P ; panels (b, d) combine the one-factor P , F , and M sweeps against packed work $FP M(M+1)/2$. Error bars show one median absolute deviation, and black dashed lines indicate linear scaling. The legend identifies the included profiles; every point is measured rather than extrapolated

E Cross-method verification

The broad controlled comparison verifies the Cartesian interface across estimators with different source-selection rules and power estimates; it is not a universal ranking of those estimators.

Table 6 reports runtime excluding final export with 4096-point focus blocks. CPU–GPU parity is checked for the gridded profiles. Gridless fitting uses different CPU and

GPU optimizers; its reported quality is the GPU result. $\bar{F}_1$ uses each method’s source-selection quantity; conditional errors include only correctly assigned sources. Timings were measured on AMD Ryzen 9 5950X 16-Core Processor and NVIDIA GeForce RTX 2070 SUPER (125.7 GiB system memory, 7.6 GiB device memory). The software environment used Python 3.13.12, NumPy 2.4.6, CuPy 14.1.1, CUDA runtime 13.2. The processor provides 32 logical threads.

Table 6 Runtime and GPU quality over ten frequencies and four sources. Median synchronized runtimes are in milliseconds from nine runs after two warm-ups and exclude final export. E_L is conditional level MAE in decibels (wavenumber-corrected for the monopole baseline, [equation \(31\)](#)), E_θ is conditional sign-invariant dipole-axis MAE in degrees, and E_y is mean position error normalized by D . A dash denotes an inapplicable quantity

Method	Source	Coupling	CPU (ms)	GPU (ms)	$\bar{F}_1$	E_L	E_θ	E_y
Conventional	Monopole	Local	215.1	2.4	0.55	3.40	–	0.014
Profile-CBF	Dipole	Local	1137.5	14.8	0.80	0.09	1.23	< 0.001
JCBF	Dipole	Joint	2429.6	17.4	1.00	0.49	2.30	< 0.001
CLEAN-SC	Dipole	Local	1628.0	106.2	0.82	0.14	1.31	< 0.001
B-CLEAN-SC	Dipole	Joint	18782.1	338.4	1.00	0.32	0.51	< 0.001
MUSIC Power Fit	Dipole	Local	450.6	61.6	1.00	0.00	0.00	< 0.001
Reduced-support								
DAMAS-NNLS	Dipole	Joint	1733.6	63.6	0.85	0.03	1.18	0.003
CSM Fit NNLS	Dipole	Joint	2591.2	397.5	1.00	0.01	1.05	< 0.001
Projected								
MF-SBL	Dipole	Joint	2252.3	418.6	1.00	0.01	1.05	< 0.001
Gridless CSM Fit	Dipole	Joint	11199.6	405.2	1.00	0.07	3.34	0.008

Statements and Declarations

Author contributions. Armin Goudarzi developed the methodology and software, conducted the numerical investigation and analysis, prepared the visualizations, and wrote the original draft. Daniel Ernst contributed to methodology, interpretation, and critical revision of the manuscript. Both authors read and approved the final manuscript.

Competing interests. The authors have no competing interests to declare that are relevant to the content of this article.

Data availability. The controlled data are generated deterministically from the reported equations and parameters. The randomized data use the fixed master seed 20260815 and the protocol in [section 3.3](#). Evaluated tables, scene-level numerical arrays, and generation scripts are available from the corresponding author on reasonable request. The wind-tunnel data underlying Fig. 5 are available from the corresponding author on reasonable request.

Code availability. A readable, unoptimized NumPy reference implementation of frequency-by-frequency Profile-CBF accompanies this article as Supplementary Code 1

(`reference_conventional.py`). It is provided with the submission and, upon publication, will be accessible through the electronic supplementary material linked to the online article. It reproduces the deterministic four-source geometry and generates the frequency-by-frequency beamforming-score and fixed-axis source-power maps for $He_D = 1, 2, 4$. It is distinct from the optimized Profile-CBF and JCBF implementations used for the reported performance measurements, which are not publicly distributed. Further implementation details are available from the corresponding author on reasonable request.

Use of generative artificial intelligence (AI). OpenAI Codex was used to assist with method implementation and manuscript proofreading. The authors take full responsibility for the content of the manuscript.

References

- [1] Curle, N.: The influence of solid boundaries upon aerodynamic sound. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences **231**(1187), 505–514 (1955) <https://doi.org/10.1098/rspa.1955.0191>
- [2] Jordan, P., Fitzpatrick, J.A., Valiere, J.-C.: Measurement of an aeroacoustic dipole using a linear microphone array. The Journal of the Acoustical Society of America **111**(3), 1267–1273 (2002) <https://doi.org/10.1121/1.1446052>
- [3] Liu, Y., Quayle, A.R., Dowling, A.P., Sijtsma, P.: Beamforming correction for dipole measurement using two-dimensional microphone arrays. The Journal of the Acoustical Society of America **124**(1), 182–191 (2008) <https://doi.org/10.1121/1.2931950>
- [4] Goudarzi, A.: B-CLEAN-SC: CLEAN-SC for broadband sources. JASA Express Letters **3**(9), 094804 (2023) <https://doi.org/10.1121/10.0020992>
- [5] Goudarzi, A.: Global, and local optimization beamforming for broadband sources. The Journal of the Acoustical Society of America **155**(1), 262–273 (2024) <https://doi.org/10.1121/10.0024247>
- [6] Gao, J., Wu, H., Jiang, W.: Dipole-based beamforming method for locating dipole sources with unknown orientations in three-dimensional domains. The Journal of the Acoustical Society of America **147**(1), 125–136 (2020) <https://doi.org/10.1121/10.0000491>
- [7] Bouchard, C., Havelock, D.I., Bouchard, M.: Beamforming with microphone arrays for directional sources. The Journal of the Acoustical Society of America **125**(4), 2098–2104 (2009) <https://doi.org/10.1121/1.3089221>
- [8] Suzuki, T.: L1 generalized inverse beam-forming algorithm resolving coherent/incoherent, distributed and multipole sources. Journal of Sound and Vibration **330**(24), 5835–5851 (2011) <https://doi.org/10.1016/j.jsv.2011.05.021>

- [9] Pan, W., Wei, L., Feng, D., Shi, Y., Chen, Y., Li, M.: Multipole transfer matrix model-based sparse bayesian learning approach for sound source identification. *Applied Acoustics* **221**, 109987 (2024) <https://doi.org/10.1016/j.apacoust.2024.109987>
- [10] Porteous, R., Prime, Z., Doolan, C.J., Moreau, D.J., Valeau, V.: Three-dimensional beamforming of dipolar aeroacoustic sources. *Journal of Sound and Vibration* **355**, 117–134 (2015) <https://doi.org/10.1016/j.jsv.2015.06.030>
- [11] Chiocchetti, L., Marx, D., Valeau, V., Ollivier, F., Marchiano, R.: Three-dimensional microphone array for the reconstruction of compact dipole aeroacoustic sources with spatially varying orientation. *Applied Acoustics* **246**, 111243 (2026) <https://doi.org/10.1016/j.apacoust.2026.111243>
- [12] Li, J., Wang, X.: Super-resolution localization and orientation estimation of multiple dipole sound sources: From a maximum likelihood framework to wind tunnel validation. *Journal of Sound and Vibration* **595**, 118764 (2025) <https://doi.org/10.1016/j.jsv.2024.118764>
- [13] Van Veen, B.D., Drongelen, W., Yuchtman, M., Suzuki, A.: Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on Biomedical Engineering* **44**(9), 867–880 (1997) <https://doi.org/10.1109/10.623056>
- [14] Sekihara, K., Nagarajan, S.S., Poeppel, D., Marantz, A.: Asymptotic SNR of scalar and vector minimum-variance beamformers for neuromagnetic source reconstruction. *IEEE Transactions on Biomedical Engineering* **51**(10), 1726–1734 (2004) <https://doi.org/10.1109/TBME.2004.827926>
- [15] Sarradj, E.: Three-dimensional acoustic source mapping with different beamforming steering vector formulations. *Advances in Acoustics and Vibration* **2012**, 292695 (2012) <https://doi.org/10.1155/2012/292695>
- [16] Golub, G.H., Van Loan, C.F.: *Matrix Computations*, 4th edn. Johns Hopkins University Press, Baltimore (2013)
- [17] Goudarzi, A., Ernst, D., Ahlefeldt, T., Philipp, F., Spehr, C.: Far-field level estimation in open wind tunnels with beamforming and CLEAN based on source coherence. *CEAS Aeronautical Journal* (2026) <https://doi.org/10.1007/s13272-026-00986-8>
- [18] Goudarzi, A., Spehr, C., Herbold, S.: Expert decision support system for aeroacoustic source type identification using clustering. *The Journal of the Acoustical Society of America* **151**(2), 1259–1276 (2022) <https://doi.org/10.1121/10.0009322>