BeBeC-2026-D05

FUNDAMENTAL CRAMÉR-RAO BOUNDS FOR THE JOINT ESTIMATION OF WAVE SPEED AND SOURCE LOCATION

Ndubuisi Uchendu and Emiliano Rustighi
Department of Industrial Engineering, University of Trento
via Sommarive, 9 - 38123 Trento, Italy

ABSTRACT

The problem of joint source localisation and wave-speed estimation arises when the propagation speed is unknown. This paper characterises the fundamental accuracy limits via identifiability analysis and Cramér–Rao bounds for arbitrary sensor arrays, including degenerate planar and linear geometries, with specialisation to far-field and near-field models. Closed-form Fisher information expressions are derived for coloured noise via whitening, revealing an explicit, geometry-dependent coupling penalty between spatial parameters and wave speed. In the far field, joint identifiability of azimuth, elevation, and wave speed requires a genuinely three-dimensional array. Near-field estimation is generically identifiable due to wavefront curvature, with degeneracy arising only under specific source–array configurations. Numerical results validate the predicted scaling, coupling behaviour, and dependence on array geometry.

1 INTRODUCTION

Accurate localisation of acoustic sources using sensor arrays is fundamental in signal processing, with applications spanning aeroacoustics, sonar, seismology, medical ultrasound, and non-destructive testing. Classical direction-of-arrival (DOA) methods [13–15] and more recent localisation frameworks [4] assume a known propagation speed, reducing localisation to estimation of spatial parameters from inter-sensor phase or delay measurements. In many practical settings, this assumption is violated; wave speed depends on environmental conditions such as temperature, salinity, and material heterogeneity [5, 9, 16, 17]. Source location and wave speed must therefore be estimated jointly, yielding a nonlinear and intrinsically coupled inverse problem.

The Cramér–Rao bound (CRB) [2, 10] provides the fundamental performance benchmark for unbiased estimators. While CRBs for DOA with known wave speed are well established

[3, 6, 8, 12, 15, 19], results for joint estimation remain comparatively limited [20]. This paper makes two contributions: it establishes identifiability conditions for joint spatial and wave-speed estimation in far-field and near-field regimes, and derives closed-form CRBs for arbitrary array geometries under a stochastic Gaussian model with known coloured noise covariance. The analysis quantifies parameter coupling through the Schur complement and provides geometry-dependent design insights.

2 SIGNAL AND PARAMETER MODEL

Consider an array with M sensors located at positions $\mathbf{r}_m = [x_m, y_m, z_m]^\top$, $m = 1, \ldots, M$, centred at the origin, and a narrowband source signal $s(n)$. The source is characterised by spatial parameters $\boldsymbol{\xi} \in \mathbb{R}^p$ and an unknown wave speed $c > 0$, forming the parameter vector $\boldsymbol{\eta} = [\boldsymbol{\xi}^\top, c]^\top \in \mathbb{R}^{p+1}$. The array output is

$$\mathbf{y}(n) = s(n)\mathbf{a}(\boldsymbol{\eta}) + \mathbf{v}(n), \quad n = 1, \ldots, N, \tag{1}$$

where $s(n) \sim \mathcal{CN}(0, \sigma_s^2)$ is a zero-mean circularly symmetric complex Gaussian signal with variance σ_s^2, and $\mathbf{v}(n) \sim \mathcal{CN}(\mathbf{0}, \mathbf{R}_v)$ is spatially coloured noise with known positive-definite covariance $\mathbf{R}_v \succ 0$, independent of $s(n)$ and $\boldsymbol{\eta}$. The steering vector $\mathbf{a}(\boldsymbol{\eta})$ has m-th entry given by

$$a_m(\boldsymbol{\eta}) = \frac{1}{\beta_m(\boldsymbol{\xi})} \exp\big(-\mathrm{j}\omega\tau_m(\boldsymbol{\xi}, c)\big), \tag{2}$$

where $\tau_m(\boldsymbol{\xi}, c)$ denotes the effective propagation delay from the source to sensor m, and $\beta_m(\boldsymbol{\xi})$ accounts for deterministic amplitude attenuation. The snapshot covariance matrix is given by

$$\mathbf{R}(\boldsymbol{\eta}) = \sigma_s^2 \mathbf{a}(\boldsymbol{\eta})\mathbf{a}^{\mathsf{H}}(\boldsymbol{\eta}) + \mathbf{R}_v. \tag{3}$$

Far-field model. The far-field source direction is parametrised by azimuth ϕ and elevation θ, with the steering vector given by

$$a_m(\phi, \theta, c) = \exp\big(-\mathrm{j}\tfrac{\omega}{c}\, \mathbf{r}_m^\top \mathbf{u}(\phi, \theta)\big), \quad \mathbf{u}(\phi, \theta) = \begin{bmatrix} \cos\theta\cos\phi \\ \cos\theta\sin\phi \\ \sin\theta \end{bmatrix}. \tag{4}$$

Near-field model In the near field, the spatial parameter is the source position $\mathbf{p} = [x_s, y_s, z_s]^\top$, and

$$a_m(\mathbf{p}, c) = \rho_m(\mathbf{p})^{-\gamma} \exp\big(-\mathrm{j}\tfrac{\omega}{c}\, \rho_m(\mathbf{p})\big), \quad \rho_m(\mathbf{p}) = \|\mathbf{p} - \mathbf{r}_m\|. \tag{5}$$

where $\gamma = 0$ corresponds to the phase-only model and $\gamma = 1$ includes amplitude decay.

3 IDENTIFIABILITY

Partial derivatives with respect to η_k are written $\dot{(\cdot)}_k$; all quantities are evaluated at the true parameter $\boldsymbol{\eta}_0$ unless stated otherwise.

3.1 General framework

The parameter $\boldsymbol{\eta}$ is *locally identifiable* at $\boldsymbol{\eta}_0$ if the only vector $\boldsymbol{\delta} \in \mathbb{R}^{p+1}$ satisfying $\sum_{k=1}^{p+1} \delta_k \frac{\partial \mathbf{R}}{\partial \eta_k} = \mathbf{0}$ is $\boldsymbol{\delta} = \mathbf{0}$. Under the single-source signal model, $\frac{\partial \mathbf{R}}{\partial \eta_k} = \sigma_s^2 \left(\dot{\mathbf{a}}_k \mathbf{a}^{\mathsf{H}} + \mathbf{a}\dot{\mathbf{a}}_k^{\mathsf{H}}\right)$, and each derivative admits the orthogonal decomposition

$$\dot{\mathbf{a}}_k = \alpha_k \mathbf{a} + \dot{\mathbf{a}}_{k,\perp}, \qquad \alpha_k = \frac{\mathbf{a}^{\mathsf{H}} \dot{\mathbf{a}}_k}{\|\mathbf{a}\|^2}, \qquad \dot{\mathbf{a}}_{k,\perp} = \left(\mathbf{I} - \frac{\mathbf{a}\mathbf{a}^{\mathsf{H}}}{\|\mathbf{a}\|^2}\right) \dot{\mathbf{a}}_k, \tag{6}$$

where $\alpha_k \mathbf{a}$ is the component of $\dot{\mathbf{a}}_k$ in the direction of $\mathbf{a}$, and $\dot{\mathbf{a}}_{k,\perp}$ is its orthogonal projection onto the complement of span($\mathbf{a}$). Since the rank-one component is common across all derivatives, parameter identifiability is governed by the orthogonal components, so the identifiability condition reduces to

$$\operatorname{rank}(\dot{\mathbf{A}}_\perp) = p + 1, \qquad \dot{\mathbf{A}}_\perp = [\dot{\mathbf{a}}_{1,\perp}, \ldots, \dot{\mathbf{a}}_{p+1,\perp}], \tag{7}$$

where rank$(\cdot)$ denotes matrix rank.

3.2 Far-field identifiability

Each steering derivative is $\dot{\mathbf{a}}_k = \mathrm{j}\mathbf{q}_k \odot \mathbf{a}_0$ with real sensor-projection vectors

$$q_{m,\phi} = -\tfrac{\omega}{c_0} (\partial_\phi \mathbf{u})^\top \mathbf{r}_m, \quad q_{m,\theta} = -\tfrac{\omega}{c_0} (\partial_\theta \mathbf{u})^\top \mathbf{r}_m, \quad q_{m,c} = \tfrac{\omega}{c_0^2} \mathbf{u}^\top \mathbf{r}_m,$$

where $\odot$ denotes the Hadamard (element-wise) product and $\mathbf{a}_0 = \mathbf{a}(\boldsymbol{\eta}_0)$. Since the array is centred, $\dot{\mathbf{a}}_{k,\perp} = \dot{\mathbf{a}}_k$. Diagonal weighting by $\mathbf{a}_0$ and scalar factors ω/c_0 and $1/c_0^2$ do not affect rank, so the projected derivative matrix satisfies

$$\operatorname{rank}(\dot{\mathbf{A}}_\perp) = \operatorname{rank}([\mathbf{q}_\phi, \mathbf{q}_\theta, \mathbf{q}_c]) = \operatorname{rank}(\mathbf{M}_r \mathbf{B}_0),$$

where $\mathbf{M}_r = [\mathbf{r}_1, \ldots, \mathbf{r}_M]^\top \in \mathbb{R}^{M \times 3}$ is the sensor-position matrix and $\mathbf{B}_0$ has columns proportional to $\{-\partial_\phi \mathbf{u}, -\partial_\theta \mathbf{u}, -\mathbf{u}\}$. For $\theta \neq \pm\pi/2$, $\mathbf{B}_0$ is full rank, giving:

Proposition 1 (Far-field identifiability)**.** Away from $\theta = \pm\pi/2$, where the azimuth ϕ is not uniquely defined due to a coordinate singularity, (ϕ, θ, c) are locally identifiable if and only if $\operatorname{rank}(\mathbf{M}_r) = 3$, i.e., the sensors span $\mathbb{R}^3$.

3.3 Near-field identifiability

Phase-only model ($\gamma = 0$). The projected derivative matrix reduces (up to rank-preserving scaling) to

$$\dot{\mathbf{A}}_{\perp,\mathrm{NF}} \propto [-\hat{\boldsymbol{\rho}}, \boldsymbol{\rho}] \in \mathbb{R}^{M \times 4}, \tag{8}$$

where $\hat{\boldsymbol{\rho}}$ stacks the per-sensor unit directions $\hat{\boldsymbol{\rho}}_m = (\mathbf{p} - \mathbf{r}_m)/\rho_m$, and $\boldsymbol{\rho} = [\rho_1, \ldots, \rho_M]^\top$ stacks the ranges. Local identifiability is therefore governed by the span of $\{\hat{\boldsymbol{\rho}}, \boldsymbol{\rho}\}$, leading to the condition $\operatorname{rank}([\hat{\boldsymbol{\rho}}, \boldsymbol{\rho}]) = 4$. Equivalently:

Proposition 2 (Near-field identifiability, $\gamma = 0$)**.** The parameters (x_s, y_s, z_s, c) are locally identifiable if and only if

$$\operatorname{rank}(\hat{\boldsymbol{\rho}}) = 3 \quad \text{and} \quad \boldsymbol{\rho} \notin \operatorname{span}(\hat{\boldsymbol{\rho}}).$$

Amplitude-and-phase model ($\gamma = 1$). For $\gamma = 1$, the spatial derivatives acquire an additional real component proportional to ρ_m^{-1}, so the projected derivative matrix has the rank-equivalent form

$$\dot{\mathbf{A}}_{\perp,\mathrm{NF}} \propto \left[-\hat{\boldsymbol{\rho}},\, \mathbf{1},\, \boldsymbol{\rho}\right], \tag{9}$$

with $\mathbf{1} = [1, \ldots, 1]^\top$. Hence, the local identifiability condition tightens to $\operatorname{rank}\left([\hat{\boldsymbol{\rho}},\, \mathbf{1},\, \boldsymbol{\rho}]\right) = 4$; equivalently:

Proposition 3 (Near-field identifiability, $\gamma = 1$)**.** The parameters (x_s, y_s, z_s, c) are locally identifiable if and only if

$$\operatorname{rank}(\hat{\boldsymbol{\rho}}) = 3 \quad \text{and} \quad \boldsymbol{\rho} \notin \operatorname{span}(\hat{\boldsymbol{\rho}}, \mathbf{1}),$$

i.e., the direction vectors span $\mathbb{R}^3$ and the range vector is not in the span of $\hat{\boldsymbol{\rho}}$ and the all-ones vector.

Table 1 summarises all three regimes. In contrast to the far field, where a three-dimensional (3-D) array is structurally required, near-field identifiability is generally configuration-dependent: for a non-degenerate array, violations of the rank conditions occur only on lower-dimensional subsets of the source-location space.

Table 1: Identifiability conditions for joint spatial and wave-speed estimation.

Regime	**Condition**	**Failure mechanisms**
Far field	$\operatorname{rank}(\mathbf{M}_r) = 3$	Insufficient spatial dimensionality (planar or linear arrays)
Near field, $\gamma = 0$	$\operatorname{rank}(\hat{\boldsymbol{\rho}}) = 3$, $\boldsymbol{\rho} \notin \operatorname{span}(\hat{\boldsymbol{\rho}})$	Direction vectors fail to span $\mathbb{R}^3$ or range vector aligns with the direction-vector subspace
Near field, $\gamma = 1$	$\operatorname{rank}(\hat{\boldsymbol{\rho}}) = 3$, $\boldsymbol{\rho} \notin \operatorname{span}(\hat{\boldsymbol{\rho}}, \mathbf{1})$	Direction vectors fail to span $\mathbb{R}^3$ or range vector aligns with the direction-vector subspace or equal-range configuration

4 CRAMÉR–RAO BOUND AND PARAMETER COUPLING

4.1 Fisher information matrix

Pre-whitening with $\mathbf{W} = \mathbf{R}_v^{-1/2}$ yields an equivalent white-noise model with steering vector $\tilde{\mathbf{a}} = \mathbf{W}\mathbf{a}$. Each whitened derivative decomposes as $\dot{\tilde{\mathbf{a}}}_k = (\tilde{\mathbf{a}}_0^{\mathsf{H}}\dot{\tilde{\mathbf{a}}}_k/\tilde{\alpha}_0)\tilde{\mathbf{a}}_0 + \dot{\tilde{\mathbf{a}}}_{k,\perp}$, where $\tilde{\alpha}_0 = \tilde{\mathbf{a}}_0^{\mathsf{H}}\tilde{\mathbf{a}}_0 = \mathbf{a}_0^{\mathsf{H}}\mathbf{R}_v^{-1}\mathbf{a}_0$. Applying the Slepian–Bangs formula [1, 11] with the Woodbury identity gives the Fisher information matrix (FIM):

$$\mathbf{J} = \frac{N\sigma_s^4}{(1+\sigma_s^2\tilde{\alpha}_0)^2}\dot{\tilde{\boldsymbol{\alpha}}}\dot{\tilde{\boldsymbol{\alpha}}}^\top + \frac{2N\sigma_s^4\tilde{\alpha}_0}{1+\sigma_s^2\tilde{\alpha}_0}\,\tilde{\mathbf{D}}, \tag{10}$$

where $\dot{\tilde{\boldsymbol{\alpha}}} = [\dot{\tilde{\alpha}}_1, \ldots, \dot{\tilde{\alpha}}_{p+1}]$ with $\dot{\tilde{\alpha}}_k = 2\,\mathrm{Re}\{\tilde{\mathbf{a}}_0^{\mathsf{H}} \dot{\tilde{\mathbf{a}}}_k\}$, $\tilde{\mathbf{D}} = \mathrm{Re}\{\dot{\tilde{\mathbf{A}}}_\perp^{\mathsf{H}} \dot{\tilde{\mathbf{A}}}_\perp\}$, and $\dot{\tilde{\mathbf{A}}}_\perp = [\dot{\tilde{\mathbf{a}}}_{1,\perp}, \ldots, \dot{\tilde{\mathbf{a}}}_{p+1,\perp}]$. For far-field and phase-only near-field models, $\dot{\tilde{\boldsymbol{\alpha}}} = \mathbf{0}$ and $\mathbf{J}$ reduces to the scaled Gram matrix $\tilde{\mathbf{D}}$.

4.2 Cramér–Rao bound

The CRB is obtained by inverting the FIM: $\mathrm{CRB}(\boldsymbol{\eta}) = \mathbf{J}^{-1}$. Writing $\mathbf{J} = \mathbf{J}_0 + \mathbf{z}\mathbf{z}^\top$ with $\mathbf{J}_0 = \frac{2N\sigma_s^4\tilde{\alpha}_0}{1+\sigma_s^2\tilde{\alpha}_0}\tilde{\mathbf{D}}$ and $\mathbf{z} = \frac{\sqrt{N}\sigma_s^2}{1+\sigma_s^2\tilde{\alpha}_0}\dot{\tilde{\boldsymbol{\alpha}}}$, the Sherman–Morrison formula gives

$$\mathrm{CRB}(\boldsymbol{\eta}) = \mathbf{J}_0^{-1} - \frac{\mathbf{J}_0^{-1}\mathbf{z}\mathbf{z}^\top\mathbf{J}_0^{-1}}{1 + \mathbf{z}^\top\mathbf{J}_0^{-1}\mathbf{z}}, \tag{11}$$

provided $\mathbf{J}_0 \succ 0$.

Since Eq. (10) is the sum of a rank-one term and a scaled positive-semidefinite matrix, $\mathbf{J} \succ 0 \iff \tilde{\mathbf{D}} \succ 0 \iff \mathrm{rank}(\dot{\tilde{\mathbf{A}}}_\perp) = p + 1$: identifiability and existence of a finite CRB are equivalent for the rank-one signal model. If $\mathbf{J}$ is rank-deficient, some parameters are locally non-identifiable and the classical CRB is undefined. In that case, the Moore–Penrose pseudoinverse $\mathbf{J}^\dagger$ may be used to characterise covariance bounds restricted to the identifiable subspace.

For white noise ($\mathbf{R}_v = \sigma_n^2\mathbf{I}$) and the phase-only model ($\dot{\boldsymbol{\alpha}} = \mathbf{0}$),

$$\mathrm{CRB}(\boldsymbol{\eta}) = \frac{1 + \alpha_0\,\mathrm{SNR}}{2N\alpha_0\,\mathrm{SNR}^2}\,\mathbf{D}^{-1}, \qquad \mathrm{SNR} = \sigma_s^2/\sigma_n^2, \quad \alpha_0 = \|\mathbf{a}_0\|^2, \tag{12}$$

with $\mathbf{D} = \mathrm{Re}\{\dot{\mathbf{A}}_\perp^{\mathsf{H}}\dot{\mathbf{A}}_\perp\}$. This reveals CRB $\propto \mathrm{SNR}^{-2}$ at low SNR, transitioning to SNR^{-1} at high SNR; geometry enters entirely through $\mathbf{D}^{-1}$.

4.3 Parameter coupling

Partitioning the FIM conformably with $(\boldsymbol{\xi}, c)$ gives

$$\mathbf{J} = \begin{bmatrix} \mathbf{J}_{\xi\xi} & \mathbf{j}_{\xi c} \\ \mathbf{j}_{\xi c}^\top & J_{cc} \end{bmatrix}. \tag{13}$$

The spatial CRB with unknown c follows from the Schur complement:

$$\mathrm{CRB}(\boldsymbol{\xi}) = \left(\mathbf{J}_{\xi\xi} - \mathbf{j}_{\xi c}J_{cc}^{-1}\mathbf{j}_{\xi c}^\top\right)^{-1} \succeq \mathrm{CRB}_0(\boldsymbol{\xi}), \tag{14}$$

where $\mathrm{CRB}_0(\boldsymbol{\xi}) = \mathbf{J}_{\xi\xi}^{-1}$ is the bound for known c. The coupling penalty $\boldsymbol{\Delta}_{\xi c} = \mathbf{j}_{\xi c}J_{cc}^{-1}\mathbf{j}_{\xi c}^\top \succeq \mathbf{0}$ inflates the spatial CRB relative to the known-c case. It vanishes if and only if $\mathbf{j}_{\xi c} = \mathbf{0}$, requiring both the amplitude coupling ($\dot{\tilde{\boldsymbol{\alpha}}}_\xi\dot{\tilde{\boldsymbol{\alpha}}}_c = 0$) and the geometric coupling ($\mathrm{Re}\{\dot{\tilde{\mathbf{A}}}_{\xi,\perp}^{\mathsf{H}}\dot{\tilde{\mathbf{a}}}_{c,\perp}\} = \mathbf{0}$) to vanish. For phase-only models ($\dot{\tilde{\boldsymbol{\alpha}}} = \mathbf{0}$), only the geometric condition remains. Geometrically, the coupling disappears when variations in c induce array phase changes that are orthogonal to those induced by spatial perturbations.

4.4 Specialisations

Far field. With $\alpha_0 = M$ and $\mathbf{Q} = [\mathbf{q}_\phi, \mathbf{q}_\theta, \mathbf{q}_c]$, the white-noise CRB is

$$\mathrm{CRB}(\boldsymbol{\eta}) = \frac{1 + M\mathrm{SNR}}{2NM\mathrm{SNR}^2} \left(\mathbf{Q}^\top \mathbf{Q}\right)^{-1}. \tag{15}$$

For a centred planar array ($z_m = 0$), the elevation and wave-speed gradients satisfy

$$q_{m,c} = (\cot\theta_0/c_0)\, q_{m,\theta},$$

so the FIM is rank-deficient and individual estimation of θ or c from a single planar array is impossible. Only azimuth and the composite parameter

$$\tilde{\theta} = \theta + (\cot\theta_0/c_0)c$$

are identifiable. For a general planar array, the white-noise CRB for the identifiable parameters admits the factorisation

$$\mathrm{CRB}(\phi) = F_\phi W_\phi, \qquad \mathrm{CRB}(\tilde{\theta}) = F_\theta W_\theta,$$

where

$$F_\phi = \frac{1 + M\,\mathrm{SNR}}{2NM\,\mathrm{SNR}^2 k_0^2 \cos^2\theta_0}, \qquad F_\theta = \frac{1 + M\,\mathrm{SNR}}{2NM\,\mathrm{SNR}^2 k_0^2 \sin^2\theta_0}, \tag{16}$$

capture all SNR dependence, with $k_0 = \omega/c_0$, and

$$W_\phi = \frac{S_{hh}}{S_{xx}S_{yy} - S_{xy}^2}, \qquad W_\theta = \frac{S_{gg}}{S_{xx}S_{yy} - S_{xy}^2}, \tag{17}$$

are geometry weights expressed through the second-order sensor moments

$$S_{xx} = \sum_m x_m^2, \qquad S_{yy} = \sum_m y_m^2, \qquad S_{xy} = \sum_m x_m y_m,$$

with $h_m = x_m\cos\phi_0 + y_m\sin\phi_0$ and $g_m = -x_m\sin\phi_0 + y_m\cos\phi_0$. This shows that wave-speed coupling does not degrade azimuth accuracy, whereas elevation–wave-speed coupling is governed entirely by the planar geometry. The determinant $S_{xx}S_{yy} - S_{xy}^2$ measures the effective two-dimensional aperture and vanishes only for collinear arrays.

For a linear array along the x-axis, the phase-gradient vectors are all collinear and the model reduces to a single identifiable parameter $\nu = \frac{\cos\theta\cos\phi}{c}$, with

$$\mathrm{CRB}(\nu) = \frac{1 + M\mathrm{SNR}}{2NM\mathrm{SNR}^2\omega^2 S_{xx}}, \tag{18}$$

where $S_{xx} = d^2M(M^2-1)/12$ for a uniformly spaced array with inter-element spacing d. If c is known, the standard CRB expression for the identifiable parameter $u = \cos\theta\cos\phi$ (direction

cosine onto the array axis) is recovered [7, 18]:

$$\mathrm{CRB}(u) = c_0^2\,\mathrm{CRB}(v) = \frac{1 + M\mathrm{SNR}}{2NM\mathrm{SNR}^2\, k_0^2 S_{xx}}.$$

Near field. For the phase-only model under white noise,

$$\mathrm{CRB}(\boldsymbol{\eta}) = \frac{1 + M\mathrm{SNR}}{2NM\mathrm{SNR}^2}\left(\mathbf{G}^\top\mathbf{G}\right)^{-1}, \tag{19}$$

where $\mathbf{G} = k_0\mathcal{P}_\perp\,[-\hat{\boldsymbol{\rho}},\ \boldsymbol{\rho}/c_0]$, and $\mathcal{P}_\perp = \mathbf{I} - \frac{1}{M}\mathbf{1}\mathbf{1}^\top$. Here $\mathcal{P}_\perp$ centres each column and removes the component lying in the rank-one signal subspace.

Near-field identifiability depends jointly on the sensor geometry and source position through the rank of the direction- and range-induced structure. Degeneracy arises when the direction vectors fail to span $\mathbb{R}^3$, or when the range vector becomes aligned with the subspace generated by the direction field, resulting in a rank deficiency that renders spatial parameters and wave speed jointly unidentifiable. Unlike the far-field case, the near-field model remains locally identifiable for a planar array when the source lies off the array plane. If the source lies in the array plane, the source coordinate normal to the plane becomes unidentifiable, and the identifiable parameter set reduces to the two in-plane coordinates together with c. For a linear array along the x-axis, the identifiable parameters reduce to $(x_s, r_\perp, c)$, where $r_\perp = \sqrt{y_s^2 + z_s^2}$.

5 NUMERICAL VALIDATION

Unless stated otherwise, simulations use the phase-only white-noise model with $N = 1$ snapshot for CRB evaluation ($N = 64$ for Monte Carlo, 500 trials), $\sigma_s^2 = 1$, $c_0 = 1$, $\omega = 2\pi$, true DOA $(\phi_0, \theta_0) = (0.4, 0.6)$ rad, and $M = 16$ sensors. Results are reported in wavelength-normalised units. The 3-D array has sensor positions drawn from $\mathcal{N}(\mathbf{0}, 0.25\mathbf{I}_3)$ and then centred; the planar array is obtained by projecting the same layout onto $z_m = 0$. Near-field sources are specified as $\mathbf{p} = r\,\mathbf{u}(\phi_0, \theta_0)$ with fixed DOA and range r. The nominal configuration uses $r = 2$, giving $\mathbf{p} = [1.5, 0.6, 1.1]^\top$.

Figure 1 confirms the SNR^{-2} scaling predicted by Eq. (12). In the far field and the near field, the unknown-c CRB is uniformly larger than the known-c case, reflecting the additional uncertainty introduced by joint estimation, while preserving the same asymptotic SNR slope. In the near field, all four parameters are identifiable for the chosen out-of-plane configuration. The CRB for c is of the same order as the spatial-parameter CRBs, indicating that wave-speed estimation has comparable statistical difficulty to direction finding in well-conditioned arrays.

Figure 2 isolates the geometric contribution to CRB inflation. Under array scaling (Fig. 2a), horizontal scaling and vertical compression increase the elevation penalty most strongly, while the azimuth penalty remains close to unity. The elevation sweep (Fig. 2b) shows a non-monotonic dependence on θ_0, with the strongest coupling at intermediate elevations. In the near field, CRB inflation grows with range in a coordinate-dependent way; for the configuration considered, x_s is most affected and z_s least.

Table 2 compares three far-field configurations: a planar uniform rectangular array (URA), a mildly perturbed (slightly non-planar) version of the URA with small random out-of-plane

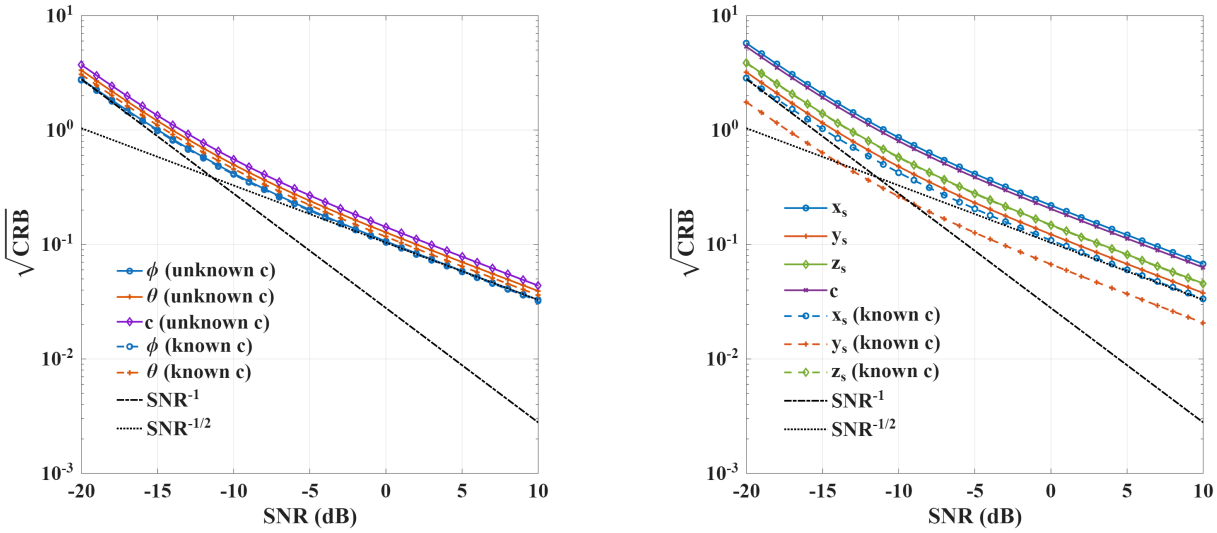

(a) Far-field 3-D array. *(b) Near-field planar array.*

Figure 1: Square-root CRB versus SNR. Solid curves correspond to unknown c and dashed curves to known c. In (b), the source location is $\mathbf{p} = [1.5, 0.6, 1]^\top$. Reference slopes proportional to $\mathrm{SNR}^{-1/2}$ and SNR^{-1} are shown for comparison. All values are wavelength-normalised.

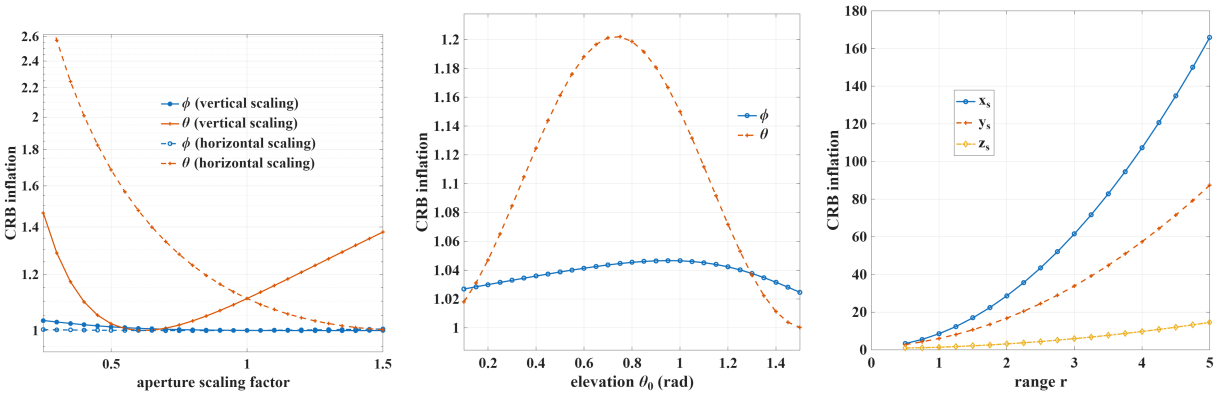

(a) Array scaling. *(b) Elevation sweep.* *(c) Near-field range.*

Figure 2: Coupling-induced CRB inflation at 0 dB *SNR. Inflation is defined as* $\mathrm{CRB}(\cdot \mid \text{unknown}, c) \big/ \mathrm{CRB}(\cdot \mid \text{known}, c)$

perturbations of the element positions (with standard deviation $\sigma_z = 0.05$, i.e., 10% of the inter-element spacing), and a 3-D random array. For the perturbed URA, the reported values are medians over 500 random perturbation realisations, reflecting its stochastic geometry-dependent conditioning; for the 3-D array, a representative realisation is used, as its statistics are effectively invariant across realisations. The planar URA remains rank deficient: $\mathrm{CRB}(\phi)$ is finite, whereas $\mathrm{CRB}(\theta)$ and $\mathrm{CRB}(c)$ diverge due to the non-identifiability of θ and c. Introducing out-of-plane perturbations restores full rank and yields finite CRBs, but the FIM remains moderately ill conditioned, with substantially larger uncertainty in θ and c than in the 3-D case. The fully 3-D array is well conditioned, confirming that a genuinely three-dimensional aperture is required for

Table 2: Far-field identifiability by array geometry at 0 dB *SNR. Planar URA values correspond to a single deterministic geometry; perturbed URA values are medians over 500 random realisations; 3-D array values correspond to a representative realisation (statistics are invariant across realisations).* $\lambda_{\min}(\mathbf{J})$ *denotes the smallest FIM eigenvalue and* $\kappa(\mathbf{J})$ *the condition number. CRB values are wavelength-normalised.*

Geometry	$\mathrm{rank}(\mathbf{J})$	$\lambda_{\min}(\mathbf{J})$	$\kappa(\mathbf{J})$	$\sqrt{\mathrm{CRB}(\phi)}$	$\sqrt{\mathrm{CRB}(\theta)}$	$\sqrt{\mathrm{CRB}(c)}$	**Interpretation**
Planar URA	2	0	∞	0.06	unbounded	unbounded	Rank-deficient; θ and c are not identifiable
Perturbed URA	3	≈ 2.3	$\approx 1.6 \times 10^2$	≈ 0.064	≈ 0.54	≈ 0.38	Full rank, but moderately ill-conditioned
3-D random	3	≈ 38	≈ 3.2	≈ 0.127	≈ 0.118	≈ 0.127	Full rank and well conditioned

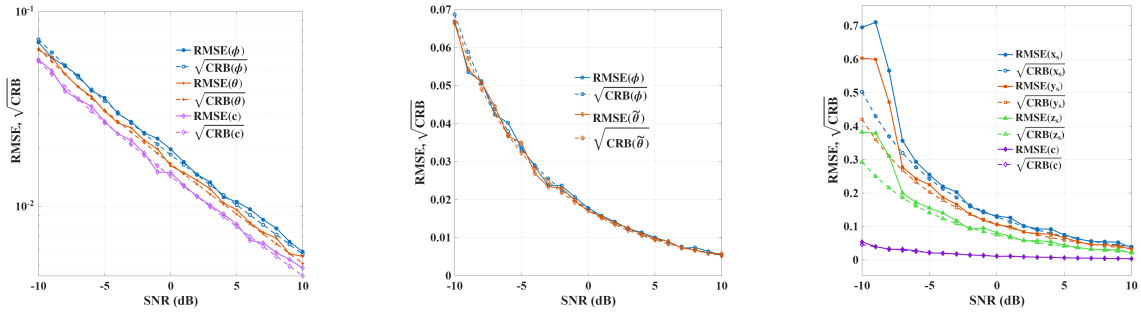

(a) Far-field 3-D array. *(b) Planar far-field configuration (degenerate case).* *(c) Near-field planar configuration with an out-of-plane source.*

Figure 3: Monte Carlo validation (500 trials; solid lines: RMSE, dashed lines: $\sqrt{\mathrm{CRB}}$*). In (b), only* ϕ *and the identifiable combination* $\tilde{\theta}$ *have finite CRB.*

far-field joint identifiability when $\theta_0 \neq \pm\pi/2$.

Finally, Fig. 3 validates the analytical CRB against Monte Carlo simulation (joint maximum-likelihood estimation using grid search) for three configurations: a far-field 3-D random array, a far-field planar array, and a near-field planar array with an out-of-plane source. In all identifiable cases, the empirical root-mean-square error (RMSE) approaches $\sqrt{\mathrm{CRB}}$ at moderate-to-high SNR. In the planar far-field case, only ϕ and the identifiable combination $\tilde{\theta}$ converge to finite values. Taken together, the Monte Carlo results corroborate the analytical CRB across identifiable and singular regimes.

6 CONCLUSION

This paper derived Cramér–Rao bounds for joint source localisation and wave-speed estimation, together with a complete identifiability characterisation in far-field and near-field regimes. Far-field joint estimation requires a genuinely three-dimensional array; planar geometries are structurally rank-deficient and admit only identifiable parameter combinations. Near-field localisation is generically identifiable due to wavefront curvature, with parameter coupling governed by array geometry. The Schur-complement formulation quantifies the wave-speed uncertainty penalty on spatial estimation, showing that geometric diversity, rather than aperture alone, governs robustness to coupling and model degeneracy. These results provide practical design guidance: three-dimensional apertures are essential for full far-field identifiability, while near-field arrays should maximise directional diversity to suppress coupling between spatial and wave-speed parameters. Natural extensions include wideband formulations, moving sources, and data-driven array optimisation under the identifiability constraints derived here.

ACKNOWEDGEMENTS

This work was financially supported by "Fondazione CARITRO" in the framework of the project "Seismic camera and digital-twin for water loss detection", call "Bando ricerca e sviluppo 2023/2024".

REFERENCES

[1] W. J. Bangs. *Array Processing with Generalized Beamformers*. Ph.D. thesis, Yale University, 1971.

[2] H. Cramér. *Mathematical methods of statistics*. Princeton University Press, Princeton, NJ, 1946.

[3] A. Dogandžić and A. Nehorai. "Cramér–Rao bounds for estimating range, velocity, and direction with an active array." *IEEE Transactions on Signal Processing*, 49(6), 1122–1137, 2001. doi:10.1109/78.923295.

[4] E. Grinstein, E. Tengan, B. Çakmak, T. Dietzen, L. Nunes, T. van Waterschoot, M. Brookes, and P. A. Naylor. "Steered response power for sound source localization: a tutorial review." *EURASIP Journal on Audio, Speech, and Music Processing*, 2024(1), 59, 2024. doi:10.1186/s13636-024-00377-z.

[5] K. C. Ho, X. Lu, and L. Kovavisaruch. "Source localisation using TDOA and FDOA measurements in the presence of receiver location errors: Analysis and solution." *IEEE Transactions on Signal Processing*, 55(2), 684–696, 2007. doi:10.1109/TSP.2006.885744.

[6] Y. D. Huang and M. Barkat. "Near-field multiple source localization by passive sensor array." *IEEE Transactions on Antennas and Propagation*, 39(7), 968–975, 1991. doi: 10.1109/8.86917.

[7] D. H. Johnson and D. E. Dudgeon. *Array Signal Processing: Concepts and Techniques*. Pearson, Englewood Cliffs, NJ, 1st edition, 1993.

[8] Y. Liang, W. Liu, Q. Shen, W. Cui, and S. Wu. "A review of closed-form cramér-rao bounds for doa estimation in the presence of gaussian noise under a unified framework." *IEEE Access*, 8, 175101–175124, 2020. doi:10.1109/ACCESS.2020.3026203.

[9] G. Nolet. *A Breviary of Seismic Tomography*. Cambridge University Press, Cambridge, 2008. doi:10.1017/CBO9780511984709.

[10] C. R. Rao. "Information and the accuracy attainable in the estimation of statistical parameters." *Bulletin of the Calcutta Mathematical Society*, 37, 81–89, 1945.

[11] D. Slepian. "Estimation of signal parameters in the presence of noise." Technical report, 1954. Reprinted in IRE Trans. Inf. Theory, vol. 3, pp. 68–89, 1954.

[12] D. Starer and A. Nehorai. "Passive localization of near-field sources by path following." *IEEE Transactions on Signal Processing*, 42(3), 677–680, 1994.

[13] P. Stoica and A. Nehorai. "MUSIC, maximum likelihood, and Cramér–Rao bound." *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(5), 720–741, 1989. doi: 10.1109/29.17564.

[14] P. Stoica and A. Nehorai. "MUSIC, maximum likelihood, and Cramér–Rao bound: Further results and comparisons." *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(12), 2140–2150, 1990. doi:10.1109/29.61539.

[15] P. Stoica and A. Nehorai. "Performance study of conditional and unconditional direction-of-arrival estimation." *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(10), 1783–1795, 1990. doi:10.1109/29.60109.

[16] A. M. Thode. "Source ranging with minimal environmental information using a virtual receiver and waveguide invariant theory." *Journal of the Acoustical Society of America*, 108(4), 1582–1594, 2000. doi:10.1121/1.1289362.

[17] D. Tollefsen, S. E. Dosso, and B. Khodabandeloo. "Matched-field source tracking with environmental sequential monte carlo filtering." *Journal of the Acoustical Society of America*, 141(5), 3526–3535, 2017. doi:10.1121/1.4983806.

[18] H. L. Van Trees. *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory: 4*. Wiley-Interscience, New York, NY, 1st edition, 2002.

[19] A. J. Weiss and B. Friedlander. "On the Cramér-Rao bound for direction finding of correlated signals." *IEEE Transactions on Signal Processing*, 41(1), 495, 1993. doi: 10.1109/TSP.1993.193187.

[20] S. Xu, M. Yang, C. Li, B. Tang, Y. Yang, L. Chen, and Y. Sun. "Joint source localization and propagation speed estimation using tdoa with hypothesized propagation speed." *Digital Signal Processing*, 159, 104934, 2025. doi:10.1016/j.dsp.2024.104934.