Branch-Aware Phase-to-Bearing Diagnostics in an Idealized Helical DAS Phase Array

Ning Hu, Yuanyuan Cao, Chengtao Huang, and Kunjie Shi

***Abstract*—Circular phase error is invariant to integer multiples of 2π and can therefore conceal a wrong phase branch. This study evaluates a branch-preserving phase-to-bearing protocol in a locked, idealized 60-channel helical phase array with normalized multiple-signal-classification azimuth estimation. The protocol separates local phase error, samplewise branch-label mismatch, target-frequency spatial consistency, and signed recovery, nominal-manifold, and total bearing errors. One counterexample combines a circular RMSE of 0.0224 rad with an absolute RMSE of 3.77 rad and a branch-mismatch density of 0.36. Under a fixed compound-geometry stress, accurate recovery leaves an approximately 13° nominal-manifold bias while the actual-manifold oracle remains near zero. Another checkpoint yields a small total error through cancellation of two large opposite-signed components. The first-order diagnostic also fails under the reported compound stress. Phase- and task-failure labels disagree in both directions across locked design cells. The results apply to the tested point-channel simulator; transfer to measured DAS has not been established.**

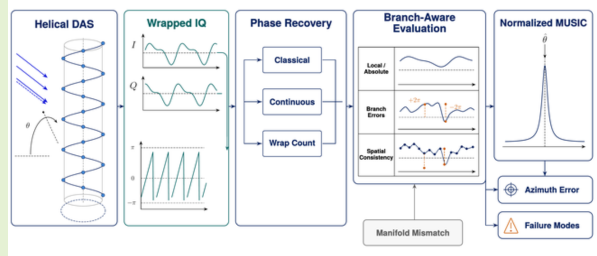

***Index Terms*—Branch ambiguity, cycle slip, direction-of-arrival estimation, distributed acoustic sensing, phase recovery, sensor evaluation, sensor model analysis.**

I. Introduction

DISTRIBUTED acoustic sensing (DAS) converts an optical fiber into a dense vibration-sensing aperture and is increasingly investigated for marine acoustics and array processing [1], [2]. Helically winding a fiber around a cylindrical host can supply sensor-dependent directional response and propagation delay within a compact aperture. Before those spatial cues can be used for azimuth estimation, however, coherent in-phase/quadrature observations must be converted from principal phase to a continuous phase field.

That conversion contains a discrete ambiguity. If ϕ is continuous phase and $\phi_w^{(0)} = \mathcal{W}(\phi) \in [-\pi, \pi)$ is its noiseless principal value, then

$$\phi = \phi_w^{(0)} + 2\pi k^{(0)}, \qquad k^{(0)} \in \mathbb{Z}. \tag{1}$$

A method can reproduce $\exp(j\phi)$ with a wrong integer count, which circular RMSE cannot reveal because it is invariant to 2π shifts. Absolute phase RMSE does not distinguish diffuse error from a few branch-wrong samples, and neither metric captures the target-frequency spatial pattern seen by an array estimator. Retaining branch information through the task is therefore necessary to avoid false confidence from low circular phase error or small total bearing error and to separate recovery, manifold, and cancellation effects.

This study examines how branch errors and nominal-manifold mismatch propagate through a locked, idealized helical phase-array–normalized-MUSIC azimuth pipeline. Simulated truth supplies integer-branch labels, paired actual/nominal manifolds, and record-level signed diagnostics. Representative classical and learned instances probe the pipeline; their unequal inputs, capacities, and optimization budgets preclude an overall method ranking.

The study asks four questions. **RQ1:** What diagnostic information is added by branch-aware metrics when circular or pointwise metrics conceal an integer error? **RQ2:** Under which tested SNR, frequency, amplitude, and mismatch conditions do locked instances exhibit transitions or distinct failure regimes? **RQ3:** When does a first-order tangent decomposition explain signed recovery and nominal-manifold contributions to normalized-MUSIC error, and where does that approximation fail? **RQ4:** How do phase failure, azimuth error, and task failure vary across record-length-qualified easy, hard, and mismatch pipeline settings?

The contribution is a record-level diagnostic protocol that distinguishes recovery-limited, nominal-manifold-limited, cancellation-dominated, and nonlinear/unresolved regimes in the locked pipeline. It retains branch labels through the array task and checks signed recovery, nominal, and total bearing

Received 19 July 2026. This work was supported by the State Key Laboratory of Ocean Engineering, Shanghai Jiao Tong University, under Grant GKZD010089, and by the State Key Laboratory of Acoustics and Marine Information, Chinese Academy of Sciences, under Grant SKLA202406. (Ning Hu and Yuanyuan Cao contributed equally to this work.) (Corresponding author: Ning Hu.)

Ning Hu is with the School of Mechanical Engineering, Hangzhou Dianzi University, Hangzhou 310018, China (e-mail: jluhooning@163.com).

Yuanyuan Cao is with the School of Communication Engineering, Hangzhou Dianzi University, Hangzhou 310018, China (e-mail: 25080206@hdu.edu.cn).

Chengtao Huang and Kunjie Shi are with the School of Mechanical Engineering, Hangzhou Dianzi University, Hangzhou 310018, China.

components against paired MUSIC outcomes. The constituent phase metrics, MUSIC estimator, and first-order sensitivity analysis are established. The added value lies in their joint, condition-qualified use to identify failure mechanisms, not in a new metric, network, or perturbation theorem.

II. Related Work

A. Classical and Dynamic-Range Phase Recovery

Temporal unwrapping integrates wrapped increments and is exact only when all adjacent noiseless phase differences remain within the principal interval [3]. State-space filtering adds a dynamical prior [4], while reliability-sorted two-dimensional (2-D) unwrapping can route around inconsistent sensor–time gradients [5]. DAS-specific work has extended dynamic range through differential–unwrapping–integration [6], 2-D transport-of-intensity processing [7], and multi-frequency acquisition [8]. These techniques alter continuity, processing, or acquisition assumptions. Here they motivate separate frequency and amplitude stresses; the acquisition model itself remains fixed.

B. Learning Representations for Wrapped Phase

Learning-based unwrapping has used continuous regression, integer fringe-order classification, global spatial dependencies, and physics-guided unrolling: PhaseNet and PhaseNet 2.0 learn 2-D representations [9], [10], and later work adds transformer dependency modeling [11], deep unrolling [12], and multi-scale phase-aware classification [13]. These are methodological context, not direct baselines here; the implemented instances use recurrent and U-Net-style backbones [14], [15], so factor-control statements are restricted to that family and transfer to attention-based backbones is open.

C. Measured DAS Array Processing and the Deployment Gap

Finite gauge length, strain transfer, and the surrounding medium shape DAS response. Field and modeling studies have documented gauge-dependent spatial filtering, imperfect amplitude transfer, and medium-dependent response of helically wound cables [16], [17]. These mechanisms are absent from the point-channel model used here.

Measured submarine and near-surface studies nevertheless establish the relevance of downstream array processing [18]–[21]. Their deployment conditions cannot provide the continuous phase truth and paired nominal/actual manifolds required by the present diagnosis. The simulator is therefore used for internal mechanism separation, not as a calibrated model of those field systems.

D. Nearest Neighbors and the Present Scope

First-order analysis of MUSIC under sensor gain, phase, and manifold errors is classical [22], [23], yielding expectation-level error expressions; the signed tangent diagnostic used here belongs to that family, and no new perturbation theorem is claimed. The differences are operational: branch information is retained through to the task, the decomposition is checked per record against paired MUSIC outcomes, and the experiments delimit where the first-order description holds, degrades, or fails, including a cancellation that expectation-level analysis conceals.

Two works are the closest neighbors. Sun *et al.* unwrap DAS phase and carry it to time-difference-of-arrival localization with a field test [7], validating one method end-to-end on measured data; here several locked instances are compared under shared truth, noise, and seeds, and the layers that disagree are reported. Chen *et al.* resolve wrapped phase-difference ambiguity inside the estimator [24]; here the estimator is fixed and the object of study is how upstream branch decisions propagate into it.

Across the reviewed evaluation-level neighbors, wrap ambiguity, measured end-to-end performance, and manifold sensitivity are addressed separately. The distinguishing scope here is their condition-locked conjunction: retained branch information is carried into a target-frequency array task, signed recovery and nominal increments are checked against paired MUSIC outcomes, and both cancellation and breakdown of the local approximation are retained as results. This is a bounded operational difference, not a priority claim.

III. Signal Model and Evaluation Framework

A. Idealized Helical Phase-Array Observation and Normalized MUSIC

The ith of 60 point channels lies on an analytical helix of radius $R = 0.25$ m and pitch $p = 0.02$ m, with 0.1 m discrete arc-length spacing. Its tangent lies approximately $0.73°$ from the circumferential direction, and the 59 intervals span approximately 3.76 turns and 0.075 m axially. The spacing is computational rather than a modeled gauge length; reported helically wound cables use winding angles farther from the circumferential limit [25], [26], so this geometry is treated as a small-pitch diagnostic fixture. Let $\mathbf{r}_i$ and $\mathbf{t}_i$ denote its position and unit tangent. For azimuth θ, elevation φ, frequency f, sound speed c, and propagation direction $\mathbf{u}(\theta, \varphi)$, the nominal steering response is

$$a_i(\theta, \varphi, f) = (\mathbf{t}_i^\top \mathbf{u})^2 \exp\left(-j2\pi f \frac{\mathbf{r}_i^\top \mathbf{u}}{c}\right). \tag{2}$$

The factor $(\mathbf{t}_i^\top \mathbf{u})^2$ is a pointwise longitudinal projection for one far-field plane wave with homogeneous propagation and perfect strain transfer. The fixture omits gauge-length averaging, cable scattering, cable–fiber coupling, and measured transfer behavior, all of which can alter DAS directivity [16], [17], [25], [27]. The resulting $\phi_i(t)$ is a normalized phase-array surrogate rather than a calibrated differential optical phase, pressure, or strain measurement:

$$\phi_i(t) = A|a_i| \sin(2\pi f t + \angle a_i), \tag{3}$$

$$z_i(t) = \exp[j\phi_i(t)] + n_i(t), \quad n_i(t) \sim \mathcal{CN}(0, \sigma^2), \tag{4}$$

where A is a reference phase-modulation amplitude in radians, $\sigma^2 = 1/\mathrm{SNR}$ defines unit-phasor I/Q-observation SNR rather than acoustic SNR, and the sampling rate is $f_s = 2$ kHz. Write the noisy principal-phase observation as $\tilde{\phi}_{w,i} = \angle z_i$; unlike $\phi_{w,i}^{(0)}$, it does not obey an exact integer decomposition. The scale $[2\,\mathrm{SNR}]^{-1/2}$ provides a local reference for qualitative comparison with I/Q demodulation analyses [28]; no optical power budget is implied. These controlled abstractions provide exact branch truth and paired nominal/actual manifold counterfactuals, enabling the signed record-level diagnosis used below.

Recovered real phase, rather than $\exp(j\hat{\phi})$, enters the task chain so that integer-cycle errors are not erased a second time; all methods share the same segmentation, windowing, and target-bin extraction (Section V). With noise projector $\mathbf{P}_n$, MUSIC uses the scale-normalized denominator

$$D_N(\theta) = \frac{\mathbf{a}^H(\theta)\mathbf{P}_n\mathbf{a}(\theta)}{\mathbf{a}^H(\theta)\mathbf{a}(\theta)}, \qquad P_N(\theta) = D_N^{-1}(\theta), \tag{5}$$

which is invariant to nonzero scalar rescaling of the steering vector [29]. The grid minimum is refined by a local quadratic fit to D_N, not its reciprocal. All reported task results concern a single source and a one-dimensional azimuth scan with elevation fixed to the simulated value; they are not joint azimuth–elevation or multisource results.

B. Itoh Margin and Conditional Local-Noise Reference

With $A_{\mathrm{eff},i} = A|a_i|$, the exact finite-record maximum of the noiseless adjacent increment (supplement) yields the phase-independent sufficient condition

$$2A_{\mathrm{eff}}^{\max}\left|\sin\left(\frac{\pi f}{f_s}\right)\right| < \pi, \qquad A_{\mathrm{eff}}^{\max} = \max_i A|a_i|, \tag{6}$$

which for $f \ll f_s$ reduces to $fA_{\mathrm{eff}} < f_s/2$. It explains why frequency and amplitude jointly control temporal continuity, while remaining only a local condition; spatial gradients and noise can still violate recovery assumptions.

For one unit-amplitude complex IQ sample in circular Gaussian noise, the high-SNR phase bound is [30]

$$\mathrm{var}(\hat{\phi}) \geq \frac{1}{2\mathrm{SNR}}. \tag{7}$$

We use $[2\mathrm{SNR}]^{-1/2}$ only as a conditional pointwise observation- noise reference after correct branch selection. It is not a field-level bound for biased recovery, correlated samples, model mismatch, or wrap-count accuracy. In particular, no single-snapshot CRB is claimed for the probability of selecting the correct integer branch.

C. Complementary Phase and Branch Metrics

Let $\mathcal{W}(x) = ((x+\pi) \bmod 2\pi) - \pi$. From truth ϕ, noisy observation $\tilde{\phi}_w$, and estimate $\hat{\phi}$, define the noise-relative nearest- integer branch labels

$$k_{\mathrm{true}} = \mathrm{round}\,\frac{\phi - \tilde{\phi}_w}{2\pi}, \qquad k_{\mathrm{pred}} = \mathrm{round}\,\frac{\hat{\phi} - \tilde{\phi}_w}{2\pi}. \tag{8}$$

Here "true" denotes the reference label from simulated continuous truth and the same noisy input, not $k^{(0)}$. Absolute RMSE preserves the branch; global alignment permits one field-wide integer offset,

$$\begin{aligned} m^* &= \arg\min_{m\in\mathbb{Z}} \|\hat{\phi} + 2\pi m - \phi\|_F^2, \\ e_{\mathrm{GA}} &= (NT)^{-1/2}\|\hat{\phi} + 2\pi m^* - \phi\|_F, \end{aligned} \tag{9}$$

which still penalizes spatially or temporally varying branch errors. The samplewise branch-label mismatch density is

$$\rho_{\mathrm{br}} = \frac{1}{NT}\sum_{i,t}\mathbb{I}[k_{\mathrm{pred},i,t} \neq k_{\mathrm{true},i,t}]. \tag{10}$$

This density counts mismatched samples. It is not an event rate and does not distinguish one persistent cycle slip from several isolated branch errors. Raw/aligned branch-error energy gives the squared-error fraction on branch-wrong samples before/after m^*; snapshot correlation gives target-frequency spatial consistency. The supplementary material specifies both ratios, aligned labels, zero-denominator rules, and the correlation formula, and Supplementary Table SIII reports their archived values for representative E7 conditions, identifying the rows where each layer changes the diagnosis. No metric substitutes for another.

D. Signed Recovery and Manifold-Error Diagnostics

Let $\mathbf{X}^{\mathrm{true}}$ and $\hat{\mathbf{X}}$ be true and recovered target-frequency snapshots, and let $\mathbf{P}_{\mathrm{nom}} = \mathbf{a}\mathbf{a}^H/(\mathbf{a}^H\mathbf{a})$. The nominal-subspace residual splits exactly as

$$\mathbf{E}_{\mathrm{rec},\perp} = (\mathbf{I} - \mathbf{P}_{\mathrm{nom}})(\hat{\mathbf{X}} - \mathbf{X}^{\mathrm{true}}), \tag{11}$$

$$\mathbf{E}_{\mathrm{nom}} = (\mathbf{I} - \mathbf{P}_{\mathrm{nom}})\mathbf{X}^{\mathrm{true}}, \tag{12}$$

$$\mathbf{E}_{\mathrm{tot}} = (\mathbf{I} - \mathbf{P}_{\mathrm{nom}})\hat{\mathbf{X}} = \mathbf{E}_{\mathrm{rec},\perp} + \mathbf{E}_{\mathrm{nom}}. \tag{13}$$

If $\mathbf{P}_{\mathrm{act}}$ uses the actual perturbed manifold, then $\mathbf{E}_{\mathrm{nom}} = \mathbf{E}_{\mathrm{mis}} + \mathbf{E}_{\mathrm{oracle}}$ with $\mathbf{E}_{\mathrm{mis}} = (\mathbf{P}_{\mathrm{act}} - \mathbf{P}_{\mathrm{nom}})\mathbf{X}^{\mathrm{true}}$ and $\mathbf{E}_{\mathrm{oracle}} = (\mathbf{I} - \mathbf{P}_{\mathrm{act}})\mathbf{X}^{\mathrm{true}}$. Every squared-norm identity retains its cross term, e.g.,

$$\|\mathbf{E}_{\mathrm{tot}}\|_F^2 = \|\mathbf{E}_{\mathrm{rec},\perp}\|_F^2 + \|\mathbf{E}_{\mathrm{nom}}\|_F^2 + 2\,\mathrm{Re}\langle\mathbf{E}_{\mathrm{rec},\perp}, \mathbf{E}_{\mathrm{nom}}\rangle_F. \tag{14}$$

For local angular diagnostics, the nominal steering derivative is evaluated by a five-point central difference with a 0.1° step and expressed per radian. For nominal steering $\mathbf{a}$, define $\boldsymbol{\alpha} = \mathbf{a}^H\mathbf{X}^{\mathrm{true}}/(\mathbf{a}^H\mathbf{a})$, $\mathbf{g} = (\mathbf{I} - \mathbf{P}_{\mathrm{nom}})\partial\mathbf{a}/\partial\theta$, and $\mathbf{B} = \mathbf{g}\boldsymbol{\alpha}$. The signed first-order bearing prediction for any residual $\mathbf{E}$ is

$$\widehat{\Delta\theta}_{\mathrm{lin}} = \frac{\mathrm{Re}\langle\mathbf{B}, \mathbf{E}\rangle_F}{\|\mathbf{B}\|_F^2}. \tag{15}$$

This is a diagnostic approximation, not a causal identification device. E7 compares its recovery, nominal, and total predictions with corresponding signed MUSIC errors. A prediction is undefined when $\|\mathbf{g}\|_2^2$ or $\|\mathbf{B}\|_F^2$ is at or below the fixed 10^{-12} numerical floor. This absolute guard is tied to the archived steering and unnormalized DFT scaling; it is not a physical tolerance or a scale-invariant degeneracy criterion. No E7 record triggered the guard.

IV. Representative Implemented Methods

The locked instances are as follows. M0 is a wrapped-identity metric sanity check, M1 integrates wrapped temporal increments, M2 is a fixed generic random-walk Kalman diagnostic, and M3 applies reliability-sorted 2-D unwrapping. M4 is a two-layer, 64-hidden-unit per-sensor LSTM; continuous-2D is a group-normalized U-Net-style sensor–time field with continuous output; and M5 shares that backbone but predicts a bounded wrap count. M2 is not an optimized sinusoidal tracker, and M5 is one tested instance rather than the paper's contribution.

The learned 2-D inputs use $[\cos\tilde{\phi}_w, \sin\tilde{\phi}_w]$ with fixed gradient and guide channels, in 256-sample windows at stride 128. Reflection padding mirrors local edge content independently at each sensor-axis end; it does not make sensors 0 and 59 adjacent, but it imposes a mirror-boundary prior that is not a physical array model, so edge-specific transfer remains a limitation. A record shorter than 256 samples is reflection-padded to one model window; no future samples are used. Architecture, losses, training support, checkpoint selection, and secondary controls are in the supplement.

All learned variants use AdamW, initialization seeds 17, 29, and 43, and validation-selected checkpoints, and no full-test result selects an epoch, seed, loss, or threshold. Training data, initialization count, validation selection, and test pairing are aligned, but capacity and optimizer are not: M4 has 50,753 parameters against about 2.67 million, a 52.6-fold gap. No M4-to-2-D difference is therefore attributed uniquely to topology, representation, or capacity; F1 supports only within-backbone comparisons, and a capacity-matched cross-family study was not added post hoc. Zero initialization SD means three checkpoints produced the same reported value, not that one run was made.

V. Experimental Protocol

A. Locked Suite and Evidence Identity

The simulator uses 60 elements, $f_s = 2$ kHz, and the nominal geometry in Section III. Training conditions span 0–40 dB, 50–800 Hz, 0.5–10 rad, azimuths 0–90°, and elevations $-30°$–30°. V0 checks metric behavior; V1 validates the conditional IQ noise reference; E1–E3 vary SNR, frequency, and amplitude; E4 and F1 examine downstream error and implemented factors. These frozen results are primary evidence. Legacy M4 trajectories and the old E5 timing experiment are supporting or exploratory only and are not used for the scientific conclusions. Supplementary Table SII gives the independent unit and archived replication count of every suite; the reported V0–F1 transitions are Monte Carlo means at grid points fixed in the archived full-run configuration, not estimates of a continuous transition boundary.

B. Mismatch Profiles and D1-G Realization Resampling

D0 is matched. D1-P fixes $+1.5\%$ sound-speed error and sensor-response log-gain/phase standard deviations of 0.03 and 0.02 rad. D1-G combines 2% radius error, 5% pitch error, 1-mm position jitter, and $0.5°$ tangent jitter. D2-J combines larger geometry and sound-speed changes with calibration, correlated-noise, drift, fading, 5% sensor dropout, and IQ non-idealities. The supplementary profile table gives every magnitude, distribution, correlation law, seed, and fixed/random status. E7/E8 use one archived fixed realization per profile; Monte Carlo replication varies noise, not arrays. Separately, 50 D1-G seeds keep the radius/pitch offsets fixed and resample position/tangent jitter. These are sensitivity conditions, not measured hardware tolerances. Recovery and MUSIC retain the nominal model while the perturbed model generates truth.

C. E7 Error-Propagation Matrix

E7 contains 26 deduplicated conditions: a matched grid (D0, $A \in \{3, 6, 8, 10\}$ rad, 20 dB, 200 Hz, 30°) and a profile grid (D0, D1-G, D1-P, D2-J; $A \in \{3, 10\}$ rad; azimuths $0°, 30°, 60°$). Ten method instances (M3 plus three checkpoints each for M4, continuous-2D, and M5) share truth, noise, and 300 seeds per condition, giving 78,000 records. Regression, paired bootstrap intervals (2,000 resamples of the 300 paired record indices), and sign agreement are computed only within a fixed condition and checkpoint; checkpoints are never pooled as a bootstrap population. The nominal component is deterministic inside a fixed profile/condition, so its within-condition slope, R^2, and interval are undefined and it is reported as a point comparison rather than assigned artificial precision.

D. E8 Condition- and Record-Length-Qualified Outcome Matrix

E8 fixes 20 dB, 200 Hz, 30°, amplitudes 3 and 10 rad, and record lengths $N \in \{100, 500, 2000, 10000\}$. D1-G and D2-J are additionally shown at $N = 2000$. Twelve method instances share 300 accuracy seeds per condition, giving 43,200 accuracy records. Primary outcomes are phase-failure rate, DOA MAE, bias, P95 absolute error, and DOA-failure rate; RMSE is supporting. The MUSIC extractor uses 20 nonoverlapping Hann-windowed segments, zero-pads each segment to at least 64 DFT points, and reads the nearest target-frequency bin; with the 256-sample model window and padding rule, these jointly define the $N < 256$ evaluation.

The primary failure rules are $\rho_{\mathrm{br}} > 10^{-3}$ (phase) and absolute DOA error $> 1°$ (task). These classification rules were fixed for the archived full run to make outcomes comparable across this study's conditions; they are not hardware acceptance criteria, certification limits, or safety thresholds, and no deployed system is claimed to require them. Sensitivity is checked across a 3×3 threshold grid whose checkpoint-level empirical rates are tabulated in the supplementary material; all main results use the primary pair, and the other eight combinations are sensitivity checks rather than retuned decision rules.

Supplementary Table SI summarizes the uncertainty scope. In particular, 0/300 has a 95% Wilson upper limit of 1.26%, and profile/checkpoint variability is not absorbed into the fixed-condition noise intervals.

E. Reproducibility and Quality Gates

Full E7 and E8 outputs match their run-manifest SHA-256 values. The E7 audit finds 78,000 unique condition–method–seed keys, 26 conditions, no geometric degeneracy, no undefined linear prediction, and maximum signed MUSIC closure 7.11×10^{-15} degrees. The largest vector identity residual is 1.82×10^{-12} in norm. E8 contains 43,200 unique accuracy keys. Frozen and overlay configuration hashes, code and checkpoint hashes, and output hashes are recorded in the independent research package.

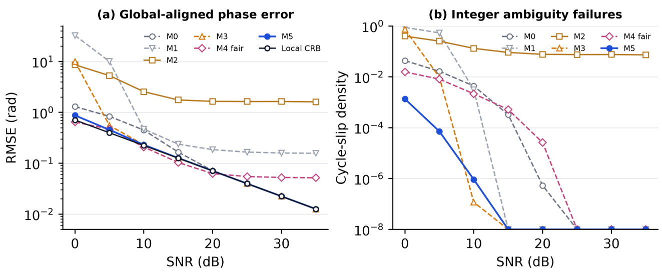

Fig. 1. Two-level SNR evaluation at $A = 3$ rad, 200 Hz, and 30°. (a) Globally aligned phase RMSE with the conditional local IQ noise reference. (b) Samplewise branch-label mismatch density (labeled "cycle-slip density" in the archived plot). The panels are complementary; agreement with the local reference is interpreted only after branch errors vanish.

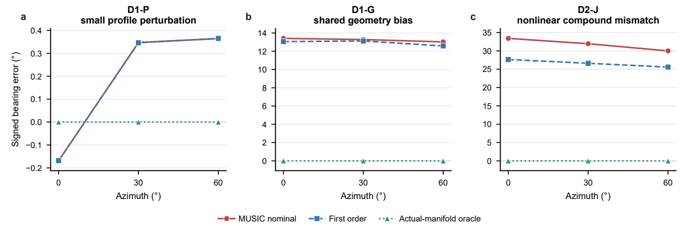

Fig. 2. Profile-level nominal-manifold diagnosis. Nominal MUSIC bias, first-order prediction, and the actual-manifold oracle are shown separately for the locked propagation, geometry, and compound profiles. The near-zero oracle localizes the reported bias within the same simulator; it is not field validation or a single-parameter tolerance.

VI. Results I: Evaluation Validity and Observed Transitions

A. Circular Error and the Local Reference Are Conditional

V0 supplies a direct counterexample. At 30 dB and 8 rad, wrapped identity M0 has circular RMSE 0.0224 rad but absolute RMSE 3.77 rad and branch-mismatch density 0.36. Thus a locally small circular discrepancy coexists with widespread wrong branches. V1 separately verifies the observation model: empirical single-IQ- sample phase standard deviations at 20, 30, and 40 dB differ by less than 0.5% from $[2\mathrm{SNR}]^{-1/2}$. These two results jointly show why the local reference is useful but insufficient.

At 30–40 dB, $A = 3$ rad, M0, M3, and M5 all approach the local-noise reference when the wrap count is trivial, so an M5/reference ratio near 1.0001 is not distinctive: any branch-correct estimator approaches the same conditional scale here. Fig. 1 places local error and branch-mismatch density side by side rather than treating either as a complete score.

B. Frequency, Amplitude, and Implemented-Factor Results

The increment relation behind (6) predicts joint frequency–amplitude stress. E2 observes that at $A = 3$ rad, increasing frequency from 500 to 800 Hz changes M3 absolute RMSE from 0.409 to 7.459 rad, whereas M5 remains near 0.071 rad. E3 observes a separate amplitude transition: at 200 Hz, M3 is 0.071 rad at 5 rad, increases to 0.922 rad at 6 rad, and reaches 7.687 rad at 10 rad. Supplementary Fig. S2 visualizes these tested-grid transitions; neither the 5–6 rad interval nor the sampled frequencies define a general phase-recovery boundary.

At the matched hard point ($A = 10$ rad, 20 dB, 200 Hz), M0 and M3 retain circular RMSE 0.0709 rad while their absolute RMSEs are 4.338 and 7.687 rad and their branch-mismatch densities are 0.413 and 0.595. Their local modulo-2π appearance hides branch failure. Shared-data M4 (labeled "M4 fair" in archived plots) has 3.151 ± 0.296 rad across three initialization means, while continuous-2D and M5 are 0.0667 and 0.0736 rad in the implemented factor controls. The largest matched pointwise phase-RMSE gain therefore belongs to the tested 2-D continuous instance, not uniquely to wrap-count output; because the 1-D and 2-D controls differ in capacity and topology, this is a single-family factor observation, not an architecture-independent ablation. Complete factor, class, and boundary metrics are in the supplement.

VII. Results II: Failure Modes and Phase-to-DOA Propagation

A. Paired Phase and DOA Outcomes

Under matched E4-H conditions at 20 dB, 200 Hz, and 30°, M3 changes from 0.0709 rad phase RMSE and 0.00957° DOA RMSE at $A = 3$ rad to 7.687 rad and 2.089° at $A = 10$ rad. Shared-data M4 at 10 rad has 3.151 rad phase RMSE and 1.903° DOA RMSE. The M5 instance has 0.0736 rad and 0.00324°. Both outcomes can be driven by SNR, amplitude, or mismatch; their association alone does not establish isolated phase-to-DOA causality. The 90° scan endpoint, where hard branch failure can yield either zero error or an endpoint flip, is reported as an estimator-boundary special case rather than excluded.

The mismatch results show a qualitatively different limitation. With accurate phase recovery, D1-P introduces signed nominal MUSIC biases from $-0.169°$ to $0.364°$. The fixed D1-G profile produces 13.024°–13.407° across tested azimuths. In 50 independently resampled position/tangent-jitter profiles with fixed radius/pitch offsets, its three-azimuth median |bias| of 13.279° lies at the 40th empirical percentile (resampled median 13.290°; Fig. S1). That spread therefore bounds jitter-realization variability only: the +2% radius and +5% pitch offsets associated with the bias remain a single design point, and the check does not establish how the bias would move under other systematic geometry errors. D2-J produces 30.0°–33.41° under one fixed realization. The actual-manifold oracle remains near zero. Thus calibration can dominate despite accurate branch recovery; the oracle is internal validation, not external evidence.

B. Condition-Qualified First-Order Evidence

Figs. 2 and 3 compare first-order predictions and signed MUSIC increments without pooling conditions. Supplementary Table SVII stratifies the 150 D0, D1-P, and D1-G condition–checkpoint groups with no record-level exceedance of the fixed 10^{-3} branch-mismatch threshold; their pooled median total MAE of 1.19×10^{-3} degrees is dominated by

TABLE I

Failure regimes identified by the locked offline diagnostic.

Regime	Diagnostic signature and evidence	Development implication
Hidden branch	Low circular; high absolute/mismatch (V0; matched hard D0)	Verify branch integrity before local fit
Recovery-limited	Phase and bearing fail together (matched hard M1/M3)	Improve recovery, not the manifold
Manifold-limited	Phase clean; shared bias (D1-G; oracle near zero)	Audit geometry/manifold calibration
Cancellation	Opposite signed components (D1-G M4 seed 17)	Inspect components, not total error alone
Nonlinear stress	First-order sign/magnitude fail (D2-J)	Do not apply a linear correction
Phase/task split	Failure labels disagree (E7/E8 census)	Report both diagnostic layers

D0/D1-P and must retain its condition labels. These strata are result-defined and descriptive, not proof of zero branch-error events. At D0, $A = 3$ rad, 30°, and M3, the recovery slope is 1.00005 (paired-bootstrap 95% CI $[0.99991, 1.00019]$) and prediction MAE is 9.47×10^{-6} degrees ($[8.63 \times 10^{-6}, 1.03 \times 10^{-5}]$).

For the deterministic nominal component, D1-P linear predictions differ from MUSIC by only 0.0010°–0.0014°. Under D1-G, the linear prediction is 12.582°–13.122°, within 0.157°–0.442° of the MUSIC bias; this range is an across-azimuth spread, not a confidence interval. The first-order term therefore preserves the direction and magnitude of the D1-G bias while the actual-manifold oracle remains near zero. This supports a nominal-manifold-dominated interpretation for the compound D1-G condition; it does not identify radius, pitch, position, or tangent error separately.

E7 also reveals why a small total error need not mean robustness. At D1-G, $A = 10$ rad, and 30°, the common nominal MUSIC error is +13.279°. M4 seed 17 adds a recovery increment of approximately −12.584°, leaving only +0.694° total error. E8 shows that this cancellation is not stable across the three locked M4 checkpoints: the family MAE is 4.410 ± 5.853°. A lower error for this checkpoint is therefore an interaction between two large opposite-signed errors, not calibration robustness.

The approximation fails more severely in D2-J. Across 60 D2-J condition–checkpoint groups, median recovery-prediction MAE is 1.93°, recovery sign agreement is only 0.048, and total-prediction MAE is 4.40°. The median total R^2 is approximately 0.40 over the 50 groups with defined R^2; the remaining ten groups have undefined R^2 and are reported as such rather than assigned zeros. For M3 at D2-J, $A = 10$ rad, and 30°, the recovery MUSIC increment is −94.51° while the linear prediction is −32.21°; the total MUSIC and linear errors are −62.55° and −5.59°. D2-J is therefore an unresolved/nonlinear regime.

VIII. Results III: Condition-Specific Task Outcomes

Fig. 4 reports condition-specific DOA MAE, phase-failure rate, and DOA-failure rate. It tests whether the evaluation layers agree and whether the method instances separate into distinct behavioral regimes rather than a single ranking.

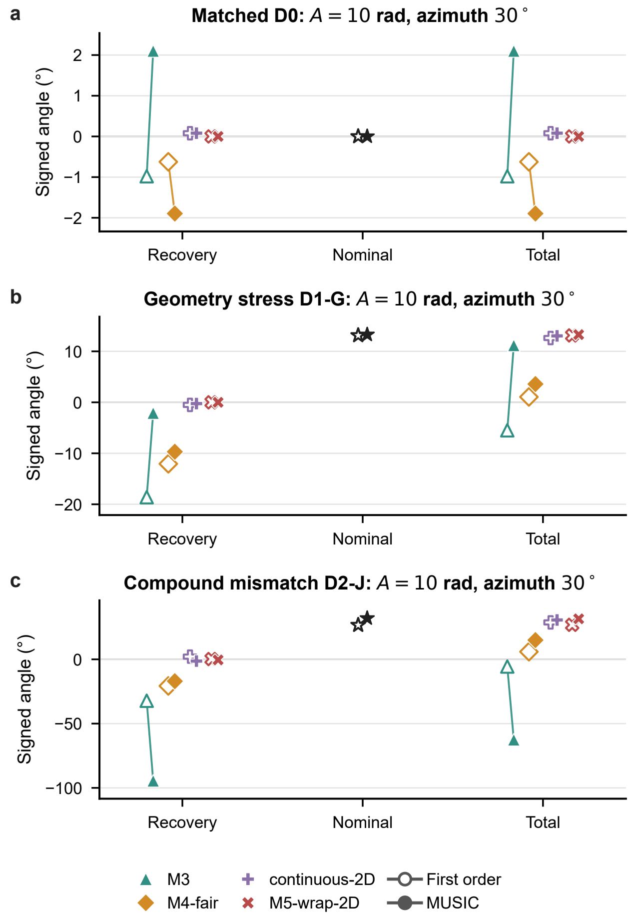

Fig. 3. Compact E7 signed-error decomposition for D0, fixed D1-G, and D2-J at $A = 10$ rad and 30°. D0 is matched; D1-G combines geometry perturbations; D2-J adds calibration, noise, dropout, fading, drift, and IQ stress. Open/filled markers are first-order/MUSIC outcomes; no pooled regression is shown.

A descriptive census of locked condition–instance cells confirms both directions of disagreement. In E7, 53 of 260 cells are phase-clean/task-fail (45 D1-G and eight D2-J), whereas three D0 cells are phase-fail/task-success; the corresponding E8 counts are 17 of 144 (16 D1-G and one D2-J) and two D0 cells. Each listed cell is unanimous over its 300 records. These design-cell counts are not independent population samples or occurrence probabilities; the complete list and provenance are in the supplement.

A. Matched Easy, Hard, and Locked-Pipeline Short-Record Conditions

In the easy matched condition (D0, $A = 3$ rad, $N = 10000$), M1, M3, and all three M5 checkpoints have DOA MAE near 0.0036° with zero observed 1° failures. Thus several recovery routes give nearly identical reported MAE and the same empirical failure label here; no formal equivalence margin or global method ordering is implied.

At D0, $A = 10$ rad, and $N = 10000$, that easy-condition similarity collapses into three distinct regimes; Supplementary Table SIV adds the per-instance P95 and failure rates.

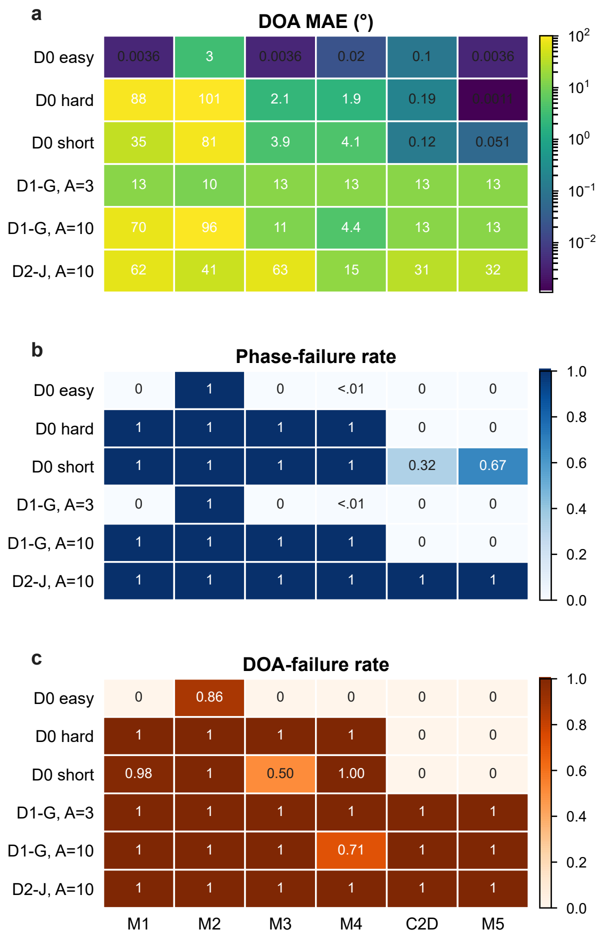

Fig. 4. E8 condition-specific DOA MAE and phase/task failure matrices under the primary thresholds. Entries are fixed methods or three-checkpoint means; every cell is defined and none imputed. No overall ranking is implied.

M1 undergoes algorithmic overflow across the temporal-wrap boundary, producing a grossly wrong bearing with 88.242° MAE—a categorical unwrap-overflow regime rather than a graded accuracy change. M3 and shared-data M4 occupy an intermediate overflow regime with 2.077° and $1.852 \pm 0.789^\circ$ MAE and unit failure rate. Continuous-2D and M5 remain within the non-overflow regime at $0.185 \pm 0.028^\circ$ and $0.0011 \pm 0.0002^\circ$ with zero observed failures. These are condition-specific behavioral regimes, not a monotonic ranking or evidence of general superiority.

At $N = 100$, the real record duration is $100/f_s = 0.05$ s—a deliberately extreme boundary condition. The 2-D methods reflection-pad those 100 observed samples to 256 inputs. Continuous-2D and M5 have $0.122 \pm 0.036^\circ$ and $0.051 \pm 0.008^\circ$ MAE with zero observed DOA failures. Nevertheless, two of the three M5 checkpoints exceed the 10^{-3} branch-mismatch threshold; their mean mismatch density is approximately 0.00167 and branch-wrong samples contribute about 93% of squared absolute phase error. This is an explicit metric non-equivalence: a phase-failure label can coexist with a task-success label.

The extractor uses 20 nonoverlapping five-sample Hann segments (half of a 200-Hz cycle), zero-padded to 64 points; 200 Hz is therefore read at the 187.5-Hz bin. Zero padding interpolates the DFT but adds no physical resolution, so this is a short-window projection, not a resolved spectrum. Because the recovered phase field is real valued, that bin also receives the conjugate component: for the five-sample Hann window the response at -200 Hz is attenuated by only about 0.67 relative to the main term, so an $N = 100$ snapshot can carry a significant conjugate steering contribution. The ratio is specific to this record length; it falls to about 0.003 at $N = 500$ and is negligible at $N \geq 2000$, where the segment places 200 Hz on an exact bin. The $N = 100$ outcome is therefore an extractor-specific locked-pipeline observation: it is not evidence of a pure record-length effect, of extractor-independent phase/task divergence, of general short-record robustness, or of real-time or deployment feasibility. Exact-frequency or analytic-signal extraction and segment-count sensitivity remain untested.

B. Mismatch Is Not an Accuracy Ranking Opportunity

At D1-G, $A = 3$ rad, and $N = 2000$, all principal instances are near 13° DOA MAE even though most have zero phase-failure rate. This agrees with E7's shared nominal bias. At $A = 10$ rad, continuous-2D and M5 remain near 13.03° and 13.28°; M4's lower family mean is explained by the unstable cancellation above. Under D2-J, M1, M3, shared-data M4, continuous-2D, and M5 have 100% failure at the 1° threshold for both 3 and 10 rad; M2 has 299/300 failures at 3 rad and 300/300 at 10 rad. At 10 rad, shared-data M4, continuous-2D, and M5 have $14.885 \pm 4.441^\circ$, $30.607 \pm 0.577^\circ$, and $31.515 \pm 0.749^\circ$ MAE. These differences lie inside a common failure region and do not constitute a useful winner ranking. The supplementary threshold table shows that the short-record phase label changes with the branch-mismatch threshold, whereas the D1-G and D2-J task-failure labels remain unchanged from 0.5° to 2° in the reported M5 checks.

IX. Discussion and Limitations

A. Practical Significance

The evaluation defines a staged, truth-dependent offline diagnostic. It first checks branch labels and target-frequency spatial consistency, then separates recovery and nominal-manifold bearing components before interpreting total task error. Under the fixed D1-G stress, an approximately 13° nominal bias persists across 50 jitter resamples. This value is an illustrative simulator checkpoint, not a cable-calibration tolerance. Within that checkpoint, improving recovery alone cannot remove the nominal bias.

B. Scope of Interpretation

Within the archived simulator, separating branch, recovery, manifold, and total errors identifies failure modes that an aggregate score conceals. Absolute RMSE and branch-mismatch density expose most matched-grid branch failures. Signed D1-G components reveal cancellation, and the phase/task labels disagree in both directions. The cancellation is also consistent with an under-capacity or out-of-regime checkpoint, given its $\pm 5.853^\circ$ spread. The available evidence does not distinguish these explanations or establish that the taxonomy is invariant to another array geometry, extractor, or DOA estimator.

C. Limitations

Evidence comes from one helix, one normalized-MUSIC task, fixed compound profiles, and no measured records. The model omits gauge averaging, transfer, coupling, multipath, and measured noise. Intervals cover fixed-profile noise; E8 continuous outcomes and the full threshold grid lack corresponding interval/census analyses. Comparators are not capacity-matched, and the $N = 100$ result is extractor-specific. Under mismatch, the tangent at the generating azimuth also differs from a sequential expansion at the nominal-MUSIC solution. These constraints limit the protocol to offline simulator diagnosis.

X. Conclusion

In the evaluated simulator, circular error can conceal a wrong branch, nominal-manifold error can remain after accurate recovery, and signed components can cancel. The first-order diagnostic agrees with MUSIC locally but not under D2-J, while E7/E8 contain phase/task disagreements in both directions. These findings are specific to the tested fixture, profiles, extractor, thresholds, checkpoints, and known-elevation MUSIC task.

Data and Code Availability

Code, configurations, checkpoints, record-level results, and regeneration scripts are archived with SHA-256 manifests and reproduction commands at `https://doi.org/10.5281/zenodo.21424624`.

Acknowledgment

During the preparation of this work, the authors used AI-assisted tools to support language editing, organization of manuscript drafts, and consistency checks across manuscriptand supplementary files. The authors reviewed, edited, and verified all AI-assisted output,including numerical claims and references, and take full respon -sibility for the content of the publication.

References

[1] J. Wei, W. Gong, J. Xing, and H. Xu, "Distributed acoustic sensing technology in marine geosciences," *Intell. Mar. Technol. Syst.*, vol. 2, no. 1, p. 26, 2024.

[2] A. Xenaki, P. Gerstoft, E. Williams, and S. Abadi, "Overview of distributed acoustic sensing: Theory and ocean applications," *J. Acoust. Soc. Am.*, vol. 158, no. 1, pp. 801–825, 2025.

[3] K. Itoh, "Analysis of the phase unwrapping algorithm," *Appl. Opt.*, vol. 21, no. 14, p. 2470, 1982.

[4] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.

[5] M. Arevallilo Herráez, D. R. Burton, M. J. Lalor, and M. A. Gdeisat, "Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path," *Appl. Opt.*, vol. 41, no. 35, pp. 7437–7444, 2002.

[6] C. Fan, H. Li, T. He, S. Zhang, B. Yan, Z. Yan, and Q. Sun, "Large dynamic range optical fiber distributed acoustic sensing (DAS) with differential-unwrapping-integral algorithm," *J. Lightwave Technol.*, vol. 39, no. 22, pp. 7274–7280, 2021.

[7] J. Sun, Y. Wang, J. Zhang, Y. Liang, G. Zhang, A. Wan, S. Zhang, Z. Ye, Y. Zhou, Q. Jing, Y. Rao, H. Wang, and Z. Wang, "2-d phase unwrapping in DAS based on transport-of-intensity-equation: Principle, algorithm and field test," *J. Lightwave Technol.*, vol. 42, no. 18, pp. 6490–6500, 2024.

[8] J. Ji, J. Chen, Z. Xiao, Z. Li, and Z. He, "Strain measurement range extension in distributed acoustic sensor based on multi-frequency phase unwrapping," *J. Lightwave Technol.*, vol. 44, no. 1, pp. 362–371, 2026.

[9] G. E. Spoorthi, S. Gorthi, and R. K. S. S. Gorthi, "PhaseNet: A deep convolutional neural network for two-dimensional phase unwrapping," *IEEE Signal Process. Lett.*, vol. 26, no. 1, pp. 54–58, 2019.

[10] G. E. Spoorthi, R. S. S. Gorthi, and S. Gorthi, "PhaseNet 2.0: Phase unwrapping of noisy data based on deep learning approach," *IEEE Trans. Image Process.*, vol. 29, pp. 4862–4872, 2020.

[11] Y. Quan, X. Yao, Z. Chen, and H. Ji, "Phase unwrapping via fully exploiting global and local spatial dependencies," *Opt. Laser Technol.*, vol. 181, p. 111872, 2025.

[12] Z. Chen, Y. Quan, and H. Ji, "Unsupervised deep unrolling networks for phase unwrapping," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2024, pp. 25 182–25 192.

[13] S. Shu, T. Zhang, J. Niu, P. Hu, Y. Zhang, and B. Guan, "MSPANet: A multi-scale phase-aware network for spatial phase unwrapping via fringe order classification," *Opt. Express*, vol. 34, no. 6, pp. 10 919–10 942, 2026.

[14] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, 1997.

[15] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.*, 2015, pp. 234–241.

[16] S. Hendi, M. Gorjian, G. Bellefleur, C. D. Hawkes, and D. White, "Investigation of the effects of surrounding media on the distributed acoustic sensing of a helically wound fibre-optic cable with application to the new afton deposit, british columbia," *Solid Earth*, vol. 14, no. 1, pp. 89–99, 2023.

[17] Y. Duan, S. Du, T. Chen, C. Guo, S. Wu, and L. Liang, "Quantitative analysis of phase response enhancement in distributed acoustic sensing systems using helical fiber winding technology," *Sensors*, vol. 25, no. 23, p. 7289, 2025.

[18] H. Matsumoto, E. Araki, T. Kimura, G. Fujie, K. Shiraishi, T. Tonegawa, K. Obana, R. Arai, Y. Kaiho, Y. Nakamura, T. Yokobiki, S. Kodaira, N. Takahashi, R. Ellwood, V. Yartsev, and M. Karrenbach, "Detection of hydroacoustic signals on a fiber-optic submarine cable," *Sci. Rep.*, vol. 11, p. 2797, 2021.

[19] A. S. Douglass, S. H. Abadi, and B. P. Lipovsky, "Distributed acoustic sensing for detecting near surface hydroacoustic signals," *JASA Express Lett.*, vol. 3, no. 6, p. 066005, 2023.

[20] J. Rychen, P. Paitz, P. Edme, K. Smolinski, J. Brackenhoff, and A. Fichtner, "Test experiments with distributed acoustic sensing and hydrophone arrays for locating underwater sound sources," *arXiv preprint arXiv:2306.04276*, 2023.

[21] M. P. A. van den Ende and J.-P. Ampuero, "Evaluating seismic beamforming capabilities of distributed acoustic sensing arrays," *Solid Earth*, vol. 12, pp. 915–934, 2021.

[22] H. Srinath and V. U. Reddy, "Analysis of MUSIC algorithm with sensor gain and phase perturbations," *Signal Process.*, vol. 23, no. 3, pp. 245–256, 1991.

[23] B. Friedlander, "A sensitivity analysis of the MUSIC algorithm," *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 38, no. 10, pp. 1740–1751, 1990.

[24] H. Chen, T. Ballal, and T. Y. Al-Naffouri, "DOA estimation with non-uniform linear arrays: A phase-difference projection approach," *IEEE Wireless Commun. Lett.*, vol. 10, no. 11, pp. 2435–2439, 2021.

[25] B. N. Kuvshinov, "Interaction of helically wound fibre-optic cables with plane seismic waves," *Geophys. Prospect.*, vol. 64, pp. 671–688, 2016.

[26] L. Cong, J. Lu, J. Duan, and R. Zhang, "Directional response and field deployment of distributed acoustic sensing cables in near-horizontal coal-mine boreholes," *Front. Built Environ.*, vol. 12, p. 1869882, 2026.

[27] S. P. Näsholm, K. Iranpour, A. Wuestefeld, B. D. E. Dando, A. F. Baird, and V. Oye, "Array signal processing on distributed acoustic sensing data: Directivity effects in slowness space," *J. Geophys. Res. Solid Earth*, vol. 127, p. e2021JB023587, 2022.

[28] X. Lu, M. A. Soto, P. J. Thomas, and E. Kolltveit, "Evaluating phase errors in phase-sensitive optical time-domain reflectometry based on I/Q demodulation," *J. Lightw. Technol.*, vol. 38, no. 15, pp. 4133–4141, 2020.

[29] R. O. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Trans. Antennas Propag.*, vol. 34, no. 3, pp. 276–280, 1986.

[30] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Englewood Cliffs, NJ, USA: Prentice Hall, 1993.