

Observability and Identifiability in Multi-Pose Acoustic Free-Field Reconstruction

A mathematical framework for rotating-DUT room measurements, independent source information, and qualified hybrid estimation

Nico Germanos · PHYSICAL LAB

Research monograph · 17 September 2026

Contents

Abstract	5
1. Scope, evidence, and the reconstruction question	5
1.1 The problem before the algorithm	5
1.2 Three different levels of claim	6
1.3 What the paper deliberately does not claim	6
2. Physical acquisition, coordinates, and conventions	7
2.1 Six configurations, not 432 stationary microphones	7
2.2 Acoustic and Fourier conventions	8
2.3 A reference frame is part of a phase target	8
2.4 Notation and operator domains	8
3. Linear target identifiability and quotient spaces	9
3.1 The basic theorem	9
3.2 Quotient-space interpretation	10
3.3 Why an excellent fit can coexist with an arbitrarily wrong target	10
4. Stability, near-nullspaces, and nonlinear targets	10
4.1 The smallest data-to-target amplification	10
4.2 Exact uniqueness is not stable continuation	11
4.3 Nonlinear targets and local Jacobians	12
5. Regularization and transformed measurements	12
5.1 A unique penalized solution is not unique information	12
5.2 Algebraic transformations preserve or discard information	13
5.3 Finite-difference bias is a separate issue	13
6. New observations, new nuisance parameters, and effective information	13
6.1 The correct restriction to a missing subspace	13
6.2 Nuisance-projected target theorem	14
6.3 A sensor that appears informative but is not	14
6.4 Profiling nuisance in a likelihood	14
7. Spherical waves and the stationary six-ring control	15
7.1 A local representation, not an automatic component label	15
7.2 Azimuthal block decomposition	15

7.3 Reproduced nullspace witnesses	17
8. Rotations, shared source states, and axisymmetric sectors	17
8.1 Rotation as a representation on angular functions	17
8.2 What yaw cannot do	18
8.3 Diversity and parameter sharing are different mechanisms	18
8.4 The rotating-room coupling	19
9. Off-axis translation and multipole mixing	19
9.1 Motion of the source origin	19
9.2 Addition theorem for a translated source	20
9.3 Far-field translation and phase reference	21
10. Source-room gauges and target-dependent invariances	21
10.1 Overlapping measurement subspaces	21
10.2 Two explicit room-model examples	22
10.3 Calibration gauges	22
10.4 Normalized targets are nonlinear and have different gauges	22
11. Statistical information, priors, and target uncertainty	23
11.1 Complex observations and real parameters	23
11.2 A target bound is not a coefficient covariance with zeros inserted	23
11.3 Data information and prior information must be reported separately	24
11.4 Correlated uncertainty	24
12. Robustness to geometry and model discrepancy	24
12.1 Operator error acts through the unknown source	24
12.2 Singular-value perturbation is informative but not sufficient	25
12.3 Source and model inadequacy	25
12.4 What calibration can and cannot rescue	25
13. Local source-state measurements and spatial sampling	25
13.1 A measurement close to a source is not source truth	25
13.2 Source complexity before sensor count	25
13.3 Symmetry can defeat repeated axial sampling	26
13.4 Co-rotation and common-mode contamination	27
14. Independent velocity, surface data, and equivalent sources	28
14.1 Pressure and a true radial derivative	28
14.2 Boundary representation and LDV	29
14.3 Equivalent-source methods in the same framework	29
15. Coupled radiators and electrical information	29
15.1 A low-frequency system is not necessarily one source	29
15.2 What terminal voltage and current constrain	30
15.3 Source-class validation	31
16. Optical geometry, registration, and room-plane information	31
16.1 A geometric sensor measures a different parameter block	31
16.2 Range and plane derivatives	31
16.3 Why horizontal coverage can be insufficient	31
16.4 Two-tier metrology	32
17. Geometry-constrained reflections and low-frequency room modes	32
17.1 Image-source construction	32

17.2 A phase budget is not a complete inversion tolerance	33
17.3 Room modes and dimension sensitivity	34
17.4 Transfer, diffraction, and scene changes	34
18. Reference-source identification and its calibrated subspace	35
18.1 What a known reference actually identifies	35
18.2 Repetition versus source-space coverage	35
18.3 Transferability and derivative estimation	36
18.4 Empirical Green functions and reciprocity	37
19. Operational free field versus same-drive anechoic response	37
19.1 Two measurands that must not be conflated	37
19.2 A source-loading bound	37
19.3 Product implications	38
20. Cross-frequency structure and physics-informed factorization	38
20.1 Shared structure can stabilize a valid model	38
20.2 A smooth exact gauge	39
20.3 Tensor organization is not automatic physical separation	39
21. Time gating: exact error decomposition and qualification	40
21.1 The identity behind every gate	40
21.2 Why a fixed crossover rule is inadequate	40
21.3 Identifying the direct response is part of the inverse problem	41
21.4 Phase-consistent targets	41
22. Bias-aware hybrid estimation	41
22.1 Correlated estimates of one complex target	41
22.2 A minimax bound for unknown bounded bias	41
22.3 Vector fields and joint fusion	42
22.4 A decision map rather than a mandatory crossover	42
23. Target-aware sensor and experiment design	43
23.1 Optimize useful information, not raw signal level	43
23.2 A monotonicity result with explicit assumptions	43
23.3 Fixed layouts across frequency	43
23.4 Pose count and source diversity	44
24. Numerical evidence and publication rechecks	44
24.1 What was actually rerun	44
24.2 Matched architecture comparison	45
24.3 Existing real-data evidence	46
24.4 Reading numerical figures correctly	46
25. A testable hybrid architecture and the remaining experiment	46
25.1 The minimal framework supported by the mathematics	46
25.2 The next physical question	47
25.3 Independent forward validation	47
25.4 Open questions stated as decisions	47
26. Conclusions	47
Appendix A. Rotation conventions and an independent construction	48
A.1 Orthonormal harmonics	48
A.2 Generator construction	49

A.3 Perturbations of pose	49
Appendix B. Green-function derivatives and independent velocity	49
B.1 Gradient and Hessian without nested numerical differentiation	49
B.2 Inverting the radial two-family system	50
B.3 Complete boundary information versus sparse points	50
Appendix C. Spherical translation, convergence, and truncation	51
C.1 Radial Green-function jump	51
C.2 Truncation and practical checks	51
C.3 Center choice as model gauge or physical parameter	51
Appendix D. Efficient nuisance elimination and target metrics	51
D.1 Why whitening precedes Euclidean projection	51
D.2 Joint old/new correlation	52
D.3 Nullspace leakage and target ambiguity	52
D.4 Reconstructible target factors	52
Appendix E. A fully worked two-channel source–room example	52
E.1 Exact recovery with a known differential signature	52
E.2 Calibration sensitivity	53
E.3 A restricted model can still be useful	53
Appendix F. Nonlinear constraints, fibers, and local geometry	53
F.1 Exact recovery lives on observation fibers	53
F.2 Constraints change the relevant tangent directions	54
F.3 Remainders behind a Jacobian uncertainty estimate	54
Appendix G. Correlated measurement blocks and calibration floors	55
G.1 Orthogonalizing new measurements without double counting	55
G.2 Profiling and marginalizing are different operations	55
G.3 A common error does not average away	56
Appendix H. Generalized target spectra and bounded-error inference	56
H.1 The generalized Rayleigh quotient	56
H.2 A finite-dimensional minimax obstruction	56
H.3 Combining data, operator, and target uncertainty	57
Appendix I. Reference-source stencils, derivative noise, and transfer range	57
I.1 What a nearby source stencil can identify	57
I.2 Bias and variance compete as spacing changes	57
I.3 Predicting a finite DUT neighborhood	58
Appendix J. Vector fusion, conditional coverage, and gate bounds	58
J.1 Minimum covariance among unbiased linear fusion rules	58
J.2 Vector bias is not automatically removed by covariance optimality	59
J.3 Conditional probability coverage is a model statement	59
J.4 A deterministic gate bound and the information it needs	59
Appendix K. Reproducibility, claim status, and reading guide	60
K.1 Provenance and numerical reconstruction	60
K.2 Formal-result index	60
K.3 Evidence-status summary	60
K.4 Suggested use of the monograph	61

Abstract

Recovering a loudspeaker’s free-field response from measurements in a room is not merely a problem of fitting a sufficiently flexible sound-field model. Different source and room states can generate indistinguishable observations while producing different desired outputs. This monograph develops a target-specific mathematical framework for that distinction and applies it to the PHYSICAL LAB acquisition: six sequential room-fixed receiver configurations and 72 known orientations of a rotating device under test. The treatment separates exact identifiability, numerical conditioning, statistical uncertainty, model discrepancy, and physical validation.

The central finite-dimensional theorem characterizes target recovery by the inclusion $\ker A \subseteq \ker B$, with equivalent quotient-space and factorization formulations. Its consequences are developed through explicit impossibility examples, nuisance-projected augmentation, spherical-wave representations, azimuthal block decomposition, Wigner rotation matrices, and translated-source expansions. A structural argument proves a rank upper bound of 30 for the declared 32-coefficient stationary six-ring control at degree three; independent numerical assembly attains that bound. The two invisible directions alter outgoing-field targets. Controlled source rotation and eccentricity can improve finite-model information, but neither guarantees stable source–room separation.

Further results characterize the information retained by pressure-derived gradients, the calibrated subspace of a reference-source room experiment, the distinction between normalized magnitude and complex directivity, and the effect of shared calibration errors. Independent radial velocity is related to a nonvanishing spherical-wave Wronskian, with explicit sampling and conditioning qualifications. Optical registration and acoustic boundary estimation enter separate measurement blocks. A source-loading bound distinguishes direct-field continuation of the source operating in the room from the counterfactual anechoic response at unchanged electrical drive. Finally, an exact gating-error identity and a bias-aware fusion rule support a frequency-dependent hybrid estimator without assuming a universal crossover frequency.

The numerical evidence is drawn from the supplied Phase-I, Phase-II, extended, and architecture-closure research archives. Selected computations are replayed and additional deterministic mathematical checks are supplied. Existing real data support angular prediction of total room pressure over their inspected development band; they do not certify free-field component truth. The resulting architecture is therefore a conditional, testable measurement framework, not a physically validated general room-removal instrument.

Keywords: inverse acoustics; target identifiability; sound-field separation; spherical harmonics; rotating source; nullspace; near-field sensing; geometric registration; reference-source calibration; hybrid estimation.

1. Scope, evidence, and the reconstruction question

1.1 The problem before the algorithm

A frequency response measured in a room is the response of a coupled experiment: an electrically driven device, an acoustic environment, a measurement geometry, and a calibrated acquisition chain. A change in any of these can change the observations. It is consequently insufficient to call an inverse result “free field” without specifying which part of the experiment is held fixed in that counterfactual. This manuscript asks which target is determined by the available measurements, which additional observations can determine missing target directions, and how much uncertainty remains when those observations are imperfect.

The established numerical programme [2–5] provides unusually direct evidence for this question. A stationary six-ring model can reproduce measured traces while leaving two outgoing/regular combinations unobserved. Subsequent work correctly reinterprets the actual experiment as a rotating source in a fixed room and adds controlled source-state, geometric, electrical, and reference-source observations. The present contribution is a unified mathematical exposition of that evidence, with proofs of the underlying finite-dimensional claims and explicit boundaries on their physical interpretation. The elementary linear-algebra theorems are not presented as newly discovered general mathematics. Their role is to make the project’s inference chain auditable.

The principal objects are a forward observation map F , a physical state q , and a requested target $\mathcal{T}$. In a linear model, these reduce to $y = Aq$ and $t = Bq$. In a nonlinear model, such as one that estimates source coefficients and reflection parameters jointly, the correct notation remains $y = F(q)$. A Jacobian can describe local sensitivity but must not silently replace the global nonlinear identification problem.

$$y = F(q) + \varepsilon + \delta_{\text{mod}}, \quad t = \mathcal{T}(q). \quad (1)$$

Here ε is acquisition error under a stated uncertainty model and δ_{mod} is discrepancy between the physical experiment and the chosen representation. Their distinction is essential. Random repeatability does not measure systematic model inadequacy; a visually smooth result does not identify the source and room separately.

1.2 Three different levels of claim

A mathematical statement applies to an explicitly declared class of operators or physical models. A numerical result evaluates a finite discretization under specified coordinates, source support, normalization, noise, and nuisance assumptions. A physical claim concerns the actual manufactured device and room. The first two can motivate the third but cannot replace its independent validation.

This monograph uses the supplied acquisition contract [1], the Phase-I scientific record [2], Phase-II source-state studies [3], extended hybrid studies [4], and architecture-closure source [5]. The earlier 34-page publication [6] supplies editorial context, not an independent replication. New publication checks [7] independently assemble selected operators and test identities. Numerical values are traced to saved matrices or JSON/CSV records rather than inferred from earlier chat summaries. Historical test totals are not accumulated into a fictitious count of physical experiments.

No new low-frequency local-microphone, optical room-scan, reference-source, or independent component-truth measurements are introduced. The existing development cache used in the closure replay contains $6 \times 72 \times 9$ complex observations at native frequencies from approximately 500 to 1500 Hz. Assertions about the actual source class below that band would therefore exceed the supplied data.

1.3 What the paper deliberately does not claim

The numerical programme does not establish one universal local microphone count, a universal near-field standoff, a universal gating crossover, or a finalized bill of materials. It does not establish that one measured reference-source transfer function characterizes an arbitrary room operator. New MATLAB and COMSOL execution remain unperformed in this publication run. Physical source–room and same-drive anechoic certification remain withheld. These qualifications are not a retreat from the mathematics: they identify exactly which conclusions follow from it and which need an experiment.

The frozen production baseline remains V3.63. The publication checks operate in separate research copies; they neither create V3.64 nor inspect or spend reserved blind-validation seeds. Older per-

centage estimates of project maturity are not used as scientific evidence.

2. Physical acquisition, coordinates, and conventions

2.1 Six configurations, not 432 stationary microphones

The frozen contract records six sequential receiver configurations. During each angular scan the microphone is fixed in the room, while the DUT rotates in nominal 5-degree steps. Seventy-three traces per scan include a seam repeat; 72 orientations are unique. Thus 438 stored traces correspond to 432 unique receiver–pose combinations, not 432 simultaneously available positions in a stationary room field. The nominal radius/height pairs are (1, 0), (1.5, 0), (1.5, −0.25), (1, −0.25), (1, −0.5), and (1.5, −0.5) metres, with height referenced to the declared DUT datum [1].

These coordinates are manual engineering values. They are not laser-calibrated positions. The contract’s 10–20 mm placement ranges and approximately one-degree angular range are sensitivity assumptions, not measured Gaussian standard deviations. Scanwise position and orientation errors are correlated across all poses of a scan. The constant settings of an amplifier do not, by themselves, establish a traceable complex terminal-voltage reference.




Figure 1: Room-fixed receiver configurations and source motion. Sequential receiver configurations. Nominal coordinates are not metrological truth. The offset is illustrative, not a frequency-independent measured center. [1,7]

Let $\mathbf{x}_s$ be a room-fixed microphone position, $\mathbf{t}$ the turntable datum, $R_n \in SO(3)$ the active DUT

rotation, and $\mathbf{d}$ a source-expansion origin relative to the DUT datum. The same physical point expressed in the rotating source frame is

$$\mathbf{r}_{s,n} = R_n^\top(\mathbf{x}_s - \mathbf{t}) - \mathbf{d}. \quad (2)$$

Conversely, a point with local coordinate ℓ_i on a co-rotating fixture occupies

$$\mathbf{x}_{i,n}^{\text{loc}} = \mathbf{t} + R_n \ell_i. \quad (3)$$

The room does not rotate under either coordinate change. Its boundary geometry can remain stationary while its excitation changes with DUT direction and acoustic-center motion. Replacing the room field by one invariant external field in the rotating frame is therefore an additional model assumption, not a geometrical identity.

2.2 Acoustic and Fourier conventions

We use the time dependence $\text{Re}\{p(\mathbf{x}, \omega)e^{i\omega t}\}$, Fourier transform $H(\omega) = \int h(t)e^{-i\omega t} dt$, $k = \omega/c$, and outgoing Green function

$$G_0(\mathbf{x}, \mathbf{s}; k) = \frac{e^{-ik|\mathbf{x}-\mathbf{s}|}}{4\pi|\mathbf{x}-\mathbf{s}|}, \quad (\nabla^2 + k^2)G_0 = -\delta_{\mathbf{s}}. \quad (4)$$

The outgoing spherical radial function is consequently $h_l^{(2)} = j_l - iy_l$. Using $h_l^{(1)}$ while retaining the same time convention would reverse the radiation phase and is not a harmless typographical change. Linear momentum balance gives

$$i\omega\rho_0\mathbf{u} = -\nabla p. \quad (5)$$

Unless otherwise stated, coordinates are in metres, frequency in hertz, pressure in pascals, and velocity in metres per second. A monopole of volume velocity Q has $p = i\omega\rho_0 Q G_0$. Some research matrices absorb the prefactor into a source coefficient; those coefficients must not be relabelled as volume velocity without restoring the units.

Spherical coordinates use polar angle $\theta \in [0, \pi]$ measured from $+z$ and azimuth φ . DUT yaw is denoted α , not θ . Hermitian transpose is $*$; ordinary transpose is $^\top$. The nullspace basis is N ; the number of angular samples is N_α . This avoids using one symbol for two unrelated quantities.

2.3 A reference frame is part of a phase target

For magnitude-only patterns, a coordinate-origin change may be unobservable or irrelevant. A complex far-field pattern has an origin-dependent directional phase. A global translation that is merely a change of coordinates requires consistent transformation of both observations and target, not necessarily an extra sensor. Conversely, a target defined relative to an external mechanical datum requires that datum to be measured. This distinction prevents an arbitrary coordinate convention from being mistaken for a physically missing degree of freedom.

The effective acoustic center also needs care. A fitted phase center may depend on frequency, radiator, and observation direction; it need not be a unique physical point on the diaphragm. A distributed source model with fixed mechanical coordinates may be more meaningful than allowing an unconstrained center to move independently at every frequency.

2.4 Notation and operator domains

Symbol	Meaning and qualification
$p, \mathbf{u}, v_n$	Complex pressure, particle velocity, and source-surface normal velocity
q, a, b, η	Joint state, source coefficients, room coefficients, and nuisance parameters
A, B, C, D	Linear observation, linear target, added-observation, and nuisance operators
$F, \mathcal{T}$	Nonlinear observation and target maps
N	A basis of a declared measurement nullspace; metric and threshold supplied
$R_n, \mathcal{D}^l$	Physical pose rotation and its degree- l harmonic representation
$\mathbf{x}_s, \mathbf{t}, \mathbf{d}$	Room-fixed receiver, turntable datum, and body-frame source-origin offset
θ, φ, α	Polar direction, azimuth direction, and DUT yaw angle
ϑ	Real statistical parameter vector, distinct from polar angle
G_0, G	Outgoing free-space Green kernel and a discretized room-transfer operator
Σ, Λ	Measurement covariance and separately declared prior precision
S, H, B_S	State metric, whitened observation matrix, and metric-adjusted target
P_D	Orthogonal projector after nuisance-block whitening
w, β_L, β_H	Convex fusion weight and nonnegative bias-magnitude bounds
k, ω, c, ρ_0	Wavenumber, angular frequency, sound speed, and mean fluid density

Bold geometric vectors and explicit frame definitions are retained when phases are compared. Matrix column orders and any coefficient rescaling are part of the reproducibility data, not implicit conventions.

3. Linear target identifiability and quotient spaces

3.1 The basic theorem

Let $X = \mathbb{C}^n$, $Y = \mathbb{C}^m$, and $Z = \mathbb{C}^p$ with their specified norms. The observation and target are linear maps $A : X \rightarrow Y$ and $B : X \rightarrow Z$. We say the target is identifiable if every pair of states generating the same noiseless observation generates the same target. Full coefficient identifiability is the special case $B = I$.

Theorem 1 — Equivalent characterizations of linear target identifiability. The following are equivalent: (i) $Aq_1 = Aq_2$ implies $Bq_1 = Bq_2$ for all states; (ii) $\ker A \subseteq \ker B$; (iii) there is a linear map K on $\text{range } A$ satisfying $B = KA$; (iv) $B = BA^\dagger A$; and (v) $\text{range } B^* \subseteq \text{range } A^*$.

$$\boxed{\ker A \subseteq \ker B \iff B = KA \iff B = BA^\dagger A.} \quad (6)$$

Proof. If targets are unique and $v \in \ker A$, then $A(q + v) = Aq$ for every q . Hence $B(q + v) = Bq$ and $Bv = 0$. Conversely, if the kernel inclusion holds, equality of two observations implies $q_1 - q_2 \in \ker A \subseteq \ker B$, which gives equality of their targets.

Define $K(Aq) = Bq$. This definition is independent of the representative q precisely because differ-

ences between representatives lie in $\ker A$. It is linear on $\text{range } A$, proving the factorization. Since $A^\dagger A$ is the orthogonal projector onto $(\ker A)^\perp$, the kernel inclusion says that B annihilates $I - A^\dagger A$, proving (iv). Taking orthogonal complements in finite dimensions gives the row-space inclusion (v). Conversely, that row-space inclusion again implies that every vector orthogonal to the rows of A is annihilated by B . $\square$

The theorem concerns every state in the declared linear class. A source positivity condition, passivity restriction, norm bound, known support, or nonlinear physical constraint changes that class and may change the conclusion. Such restrictions must be stated rather than silently introduced by the reconstruction algorithm.

3.2 Quotient-space interpretation

Define $q_1 \sim q_2$ if $Aq_1 = Aq_2$. Each observation identifies the equivalence class $[q] = q + \ker A$. The quotient $X/\ker A$ is isomorphic to $\text{range } A$ through $[q] \mapsto Aq$. The target descends to this quotient if $[q] \mapsto Bq$ is independent of representative, exactly the kernel condition in Theorem 1.

$$X \xrightarrow{\pi} X/\ker A \xrightarrow{\tilde{A}} \text{range } A, \quad B = \tilde{B}\pi. \quad (7)$$

This formulation separates an unidentifiable physical description from an identifiable practical output. A harmless gauge can remain in the coefficients while cancelling in the target. Conversely, a visually plausible coefficient vector selected by a solver need not represent the correct equivalence-class member for a target that changes along the nullspace.

3.3 Why an excellent fit can coexist with an arbitrarily wrong target

Corollary 1 — A two-state impossibility bound. If $Av = 0$ and $Bv \neq 0$, then for any estimator $\hat{t}(y)$ and any real a , at least one of the states q and $q + av$ has target error at least $|a|\|Bv\|/2$ at their common observation.

$$\max\{\|\hat{t}(Aq) - Bq\|, \|\hat{t}(A(q + av)) - B(q + av)\|\} \geq \frac{|a|}{2}\|Bv\|. \quad (8)$$

Proof. The distance between the two targets is $|a|\|Bv\|$. By the triangle inequality it cannot exceed the sum of their distances from the same estimated target. One of those distances must therefore be at least half the target separation. $\square$

Without a bound on admissible a , no uniformly finite data-only worst-case error exists. This is not a statement about a particular optimizer failing. It says that the observations do not contain enough information to select the correct target. A regularizer can restrict a , but then the guarantee depends on that restriction.

Physical interpretation. An outgoing/source contribution and an external/room contribution can compensate exactly at the microphones. Their sum fits perfectly while the predicted direct field away from those microphones changes. Testing held-out total pressure is helpful, but a component target must still be tested against its own null directions.

4. Stability, near-nullspaces, and nonlinear targets

4.1 The smallest data-to-target amplification

Use a declared parameter metric $q = S\xi$, noise covariance $\Sigma \succ 0$, and whitened operator $H = \Sigma^{-1/2}AS$. Let $B_S = BS$. Parameter metrics make the scaling used in singular-value comparisons explicit; noise whitening reflects actual measurement reliability. These roles are distinct.

$$y_w = \Sigma^{-1/2}y, \quad H = \Sigma^{-1/2}AS, \quad t = B_S\xi. \quad (9)$$

Theorem 2 — Sharp finite-dimensional target stability. If $\ker H \subseteq \ker B_S$, then the minimal possible operator norm among linear maps K with $KH = B_S$ is

$$\gamma_B = \sup_{H\xi \neq 0} \frac{\|B_S\xi\|}{\|H\xi\|} = \|B_S H^\dagger\|_2. \quad (10)$$

Proof. Theorem 1 defines a unique map on $\text{range } H$. Its norm is the supremum in the middle expression. Extend it by zero on the orthogonal complement of that range. The resulting map is $B_S H^\dagger$, and adding an arbitrary action on the orthogonal complement cannot reduce the norm of the already prescribed restriction. $\square$

For $H = U_r \text{diag}(\sigma_1, \dots, \sigma_r) V_r^*$, the target-noise map is

$$K_B = B_S V_r \text{diag}(\sigma_1^{-1}, \dots, \sigma_r^{-1}) U_r^*. \quad (11)$$

A small singular value is damaging only to the extent that its right-singular direction changes the target. The coefficient condition number alone therefore need not rank architectures correctly. Conversely, setting small singular values to zero does not establish that their target effect was negligible; that is an additional decision requiring a source bound or prior.




Figure 2: Full rank does not imply stable target recovery. Illustrative finite-dimensional family, not an experimental SFS error curve. [7]

4.2 Exact uniqueness is not stable continuation

Consider $H_\epsilon = \text{diag}(1, \epsilon)$ and the target equal to the second coefficient. Every $\epsilon > 0$ gives full rank, but $\gamma_B = 1/\epsilon$. Arbitrarily small noise can create a large target error when ϵ is small. This is the simplest model of the distinction between an eccentricity-induced rank increase and practically useful information.

In infinite dimensions the separation is sharper. On ℓ^2 , let $A(q_1, q_2, \dots) = (q_1, q_2/2, q_3/3, \dots)$ and $B = I$. Both kernels are zero, so the target is exactly unique. Yet $\|Ae_n\| = 1/n$ while $\|Be_n\| = 1$.

No bounded inverse-to-target map exists. Thus the linear kernel theorem remains an algebraic uniqueness statement, but stability in a Hilbert-space model needs a uniform inequality $\|Bq\| \leq \gamma\|Aq\|$ or an appropriate regularity restriction. A finite modal truncation is one such modelling choice, not a proof that arbitrary physical fields are recoverable.

The numerical SVD is also subject to finite precision and scaling. A threshold such as $10^{-10}\sigma_1$ is an operational classification, not a universal physical boundary. LAPACK's error-analysis framework motivates treating small singular values and poorly separated subspaces cautiously [10]. The project additionally evaluates target responses and uncertainty assumptions rather than interpreting a thresholded rank in isolation.

4.3 Nonlinear targets and local Jacobians

Normalized directivity is not a linear target operator. For a general observation map F and target $\mathcal{T}$, exact target identifiability means

$$F(q_1) = F(q_2) \implies \mathcal{T}(q_1) = \mathcal{T}(q_2). \quad (12)$$

At a differentiable operating point q_0 , a necessary condition for a differentiable local factorization is $\ker J_F(q_0) \subseteq \ker J_{\mathcal{T}}(q_0)$. This tests infinitesimal invisible directions. It does not automatically prove the global implication in (12).

For example, $F(x) = x^2$ has a nonzero derivative away from zero but cannot distinguish x from $-x$ globally. Conversely, $F(x) = x^3$ is globally injective although its derivative vanishes at zero; its inverse is poorly conditioned there. Under constant-rank assumptions, annihilation of the measurement-fiber tangents at every point of a connected local fiber implies that the target is constant on that fiber. One Jacobian at one nominal state establishes less.

Physical interpretation. A full-rank source/room Jacobian supports a local sensitivity calculation near its assumed parameters. It does not rule out a second remote solution, establish source-class adequacy, or turn normalized complex directivity into a linear quantity. All 41- and 58-real-parameter architecture comparisons later in this manuscript retain this local qualification.

5. Regularization and transformed measurements

5.1 A unique penalized solution is not unique information

Consider the quadratic estimator

$$\hat{q}_\lambda = \arg \min_q \{\|Aq - y\|_2^2 + \lambda\|Lq\|_2^2\}, \quad \lambda > 0. \quad (13)$$

Proposition 1 — Penalized uniqueness. The objective in (13) has a unique minimizer for every y if and only if $\ker A \cap \ker L = \{0\}$. Along a vector in $\ker A$, all additional discrimination is provided by the penalty.

Proof. The Hermitian Hessian is proportional to $A^*A + \lambda L^*L$. For any v ,

$$v^*(A^*A + \lambda L^*L)v = \|Av\|^2 + \lambda\|Lv\|^2. \quad (14)$$

It is positive for every nonzero v exactly when the intersection of the two kernels is trivial. If $Av = 0$, the residual is unchanged under $q \mapsto q + av$, so only the penalty distinguishes those states. $\square$

This does not make regularization undesirable. It is indispensable in many ill-conditioned acoustic inverse problems. The distinction concerns the language of the conclusion: a prior-conditioned

estimate is not a measurement-only identification. Smoothness, minimum norm, sparse support, and a restricted source manifold can each be useful, provided their physical adequacy and target bias are evaluated separately.

5.2 Algebraic transformations preserve or discard information

Suppose $y = Aq + \varepsilon$ with covariance $\Sigma \succ 0$, and form $z = Ty$. A finite-difference gradient, pressure-derived velocity, or redundant augmentation belongs to this class. Its covariance is $T\Sigma T^*$, which can be singular when T adds redundant rows.

Proposition 2 — Information under a deterministic linear transform. Put $H = \Sigma^{-1/2}A$ and $M = T\Sigma^{1/2}$. Then

$$A^*T^*(T\Sigma T^*)^\dagger TA = H^*P_{\text{range } M^*}H \preceq H^*H. \quad (15)$$

Equality holds whenever the columns of H lie in $\text{range } M^*$; in particular, it holds for an injective T .

Proof. A singular-value decomposition of M gives $M^*(MM^*)^\dagger M = P_{\text{range } M^*}$. Substitute $TA = MH$. The difference from H^*H is $H^*(I - P)H$, a positive semidefinite matrix. If T is injective, M has full column rank and $P = I$. $\square$

Thus retaining both original pressures and their differences cannot add measurement information when the full covariance is used. Discarding the original pressures can lose information, but cannot create it. Treating the differences as independent new channels replaces the correct singular correlated covariance by an unjustified diagonal matrix and overstates precision.

5.3 Finite-difference bias is a separate issue

For symmetrically placed microphones and a smooth pressure field,

$$\frac{p(x + d/2) - p(x - d/2)}{d} = p'(x) + \frac{d^2}{24}p'''(x) + O(d^4). \quad (16)$$

For a plane wave this derivative estimate carries a sinc response; for a curved near field, higher derivatives depend on source distance and modal order as well as frequency. Short baselines reduce truncation bias but suppress the differential signal and amplify calibration error relative to it. Fitting raw pressures at their actual coordinates avoids imposing the finite-difference approximation, but does not remove the underlying information or noise limitations. An independently measured velocity is different: it is a new observation, even though an exact continuous pressure field would determine velocity through (5).

6. New observations, new nuisance parameters, and effective information

6.1 The correct restriction to a missing subspace

Let the columns of N span $\ker A$. If an additional observation has no new unknowns, the remaining old-state ambiguity is $q \mapsto q + N\xi$ with $CN\xi = 0$. To recover all old coefficients requires CN to be injective. To recover only the target requires less:

$$\ker(CN) \subseteq \ker(BN). \quad (17)$$

A target-changing-nullspace dimension is not automatically a minimum number of physical micro-

phones. A single microphone at several registered, independent states can contribute several rows, while many colocated or symmetry-equivalent microphones may contribute almost the same row direction.

6.2 Nuisance-projected target theorem

A realistic additional sensor can introduce local room contamination, gain, phase, or reference-session parameters. Write

$$y = Aq, \quad z = Cq + D\eta. \quad (18)$$

Theorem 3 — Additional observations with unrestricted nuisance. After whitening the new block, let $C_w = \Sigma_z^{-1/2}C$, $D_w = \Sigma_z^{-1/2}D$, and $P_D = I - D_w D_w^\dagger$. The augmented system determines the old target Bq if and only if

$$\boxed{\ker(P_D C_w N) \subseteq \ker(BN)}. \quad (19)$$

Proof. Any old-state difference compatible with y has the form $N\xi$. It remains compatible with z if some nuisance difference $\delta\eta$ satisfies $C_w N\xi + D_w \delta\eta = 0$. Such a nuisance exists exactly when $C_w N\xi$ lies in range D_w , equivalently $P_D C_w N\xi = 0$. The target is unique exactly when every such ξ also satisfies $BN\xi = 0$. $\square$

Whitening is irrelevant to exact noiseless rank when invertible, but matters to practical information and projection geometry. If errors of old and new measurements are correlated, the combined likelihood must be whitened or conditioned jointly; separate block whitening does not model those correlations.

6.3 A sensor that appears informative but is not

Take $y = a + b$ and $z = a + c$. The original null direction is proportional to $(1, -1)$, and the new sensor directly sees a . Nevertheless, the augmented map on (a, b, c) has the invisible direction $(1, -1, -1)$. Every increase in source amplitude can be compensated in both room nuisance terms. The correct projected operator vanishes because the new nuisance spans the entire one-dimensional observation.

This example is directly relevant to a proposed near-field microphone assigned an arbitrary independent reflection contribution, or a reference-source session assigned an arbitrary gain. It does not show that such sensors are useless. It shows that their nuisance structure must be constrained by calibration, shared physics, or additional observations before a net information gain can be claimed.

6.4 Profiling nuisance in a likelihood

For whitened data $y_w = J_q \delta q + J_\eta \delta \eta + e$, minimizing over $\delta \eta$ gives the profiled objective $\|P_\eta(y_w - J_q \delta q)\|^2$, where $P_\eta = I - J_\eta J_\eta^\dagger$. The effective information is

$$\mathcal{J}_{q|\eta} = J_q^* P_\eta J_q = \mathcal{J}_{qq} - \mathcal{J}_{q\eta} \mathcal{J}_{\eta\eta}^\dagger \mathcal{J}_{\eta q}. \quad (20)$$

This is the Schur-complement form under the appropriate range conditions. It makes a practical design principle explicit: the useful part of a new sensor is not its total signal level, but the part of its source sensitivity that cannot be explained by nuisance variations. Strong near-field pressure can still be poorly discriminating if all source changes look like calibration or local reflection changes.

7. Spherical waves and the stationary six-ring control

7.1 A local representation, not an automatic component label

In a source-free spherical annulus, a Helmholtz field can be expanded into separated radial and angular functions. With the convention of Section 2, the finite model used for the stationary control is

$$p(r, \theta, \varphi) = \sum_{l=0}^L \sum_{m=-l}^l [a_{lm} h_l^{(2)}(kr) + b_{lm} j_l(kr)] Y_{lm}(\theta, \varphi). \quad (21)$$

There are $(L+1)^2$ coefficients in each family, hence $2(L+1)^2$ in the joint model. The orthonormal spherical harmonics follow the usual complex convention [8]. The functions j_l are regular at the origin; $h_l^{(2)}$ have an outgoing radiation character and are singular at the origin. An external field can be represented by regular functions inside a ball without its external singularities. Neither this convergence domain nor a truncation order can be ignored.

In particular, “regular” does not mean “purely incoming”: $j_l = (h_l^{(1)} + h_l^{(2)})/2$. A finite outgoing/regular split is a useful local parameterization, but attaching the physical labels DUT and room requires a justified source domain, external domain, and forward model. A finite regular expansion valid near the DUT need not be valid throughout an arbitrary room or at a continued virtual receiver beyond its convergence region.

At $L = 3$, each family has 16 coefficients. For six radius/height traces sampled at 72 azimuths, evaluating (21) gives the 432×32 matrix used in the mandatory control [2–5]. Its column scaling changes a numerical conditioning metric but not the exact kernel dimension.

7.2 Azimuthal block decomposition

For the stationary ring model, each trace has constant spherical radius r_s and polar angle θ_s , while azimuth varies. Since $Y_{lm}(\theta, \varphi)$ has angular factor $e^{im\varphi}$, the discrete Fourier transform over uniformly sampled azimuth isolates orders m when $N_\alpha > 2L$. An orientation convention can reverse the sign of the Fourier index but not the ranks.

Proposition 3 — Rank decomposition of fully sampled axisymmetric traces. For S complete coaxial pressure traces and an angular sampling count exceeding $2L$, the stationary mixed-wave operator is unitarily row-equivalent, up to zero Fourier rows, to blocks C_m of dimension $S \times 2(L - |m| + 1)$. Consequently,

$$\text{rank } A = \sum_{m=-L}^L \text{rank } C_m, \quad \text{rank } C_m \leq \min\{S, 2(L - |m| + 1)\}. \quad (22)$$

Proof. Reorder columns by azimuthal order. Apply the unitary discrete Fourier transform separately to each trace. Orthogonality of the sampled characters annihilates every order other than the matching m , leaving the radial/polar coefficients of that order at the S traces. Because row transformations are unitary and column permutations are invertible, ranks are preserved and add over the resulting blocks. $\square$

For $S = 6$ and $L = 3$, the coefficient counts from $m = -3$ to 3 are 2, 4, 6, 8, 6, 4, 2. The $m = 0$ block therefore has at least two null directions. Even infinitely dense azimuth sampling on the same six traces cannot add missing radial or height constraints within that sector.

Corollary 2 — Structural upper bound for the degree-three control. The declared six-trace mixed model has $\text{rank } A \leq 30$ and $\dim \ker A \geq 2$.

Proof. Sum the seven block bounds in (22): $2 + 4 + 6 + 6 + 6 + 4 + 2 = 30$. The model has 32 columns. $\square$

This is an analytic upper bound, not a proof that every six-ring geometry at every frequency attains rank 30. Special locations, frequencies, or degeneracies can reduce the rank further. The lower bound needed for exact equality is supported here by the numerical block ranks for the actual declared test coordinates and frequency.




Figure 3: Stationary six-ring singular spectrum. 1000 Hz, $L=3$, 432×32 complex operator, column-norm scaling; threshold $1e-10$. [7]




Figure 4: Azimuthal block dimensions. Dimension bound is analytic; achieved block ranks are numerical for this particular geometry/frequency. [7]

7.3 Reproduced nullspace witnesses

The publication independently assembles the 1000 Hz matrix from the nominal coordinates and harmonic definition. It matches the saved closure operator, attains numerical rank 30 at relative threshold 10^{-10} , and has two target-changing directions. The normalized Frobenius null residual is approximately 2.26×10^{-16} ; the earlier 4.17×10^{-16} result belongs to another saved numerical basis/normalization. Reproducing a null subspace does not require reproducing roundoff-level digits of a particular basis [7].

The computed null vectors have negligible support outside $m = 0$. Their outgoing parts evaluated away from the traces are nonzero. Figure 5 displays a meridional target at 2 m; each witness is normalized for visualization, not assigned a physical measured amplitude. The joint outgoing-plus-regular pressure perturbation is invisible on the original rows even though its outgoing continuation is visible.




Figure 5: An invisible data perturbation with a visible outgoing target. Each arbitrary null-basis vector is normalized separately for display. The joint field change vanishes at the measurement rows, not at all spatial points. [7]

Physical interpretation. More observations of the same angular character cannot recover a radial degree of freedom that the six traces do not distinguish. This is a controlled negative result for a finite stationary model. It is not the final verdict on the actual rotating-DUT acquisition, whose room excitation changes with pose.

8. Rotations, shared source states, and axisymmetric sectors

8.1 Rotation as a representation on angular functions

For $R \in SO(3)$, define the action $U(R)f(\hat{\mathbf{x}}) = f(R^{-1}\hat{\mathbf{x}})$. Surface measure on the unit sphere is rotation-invariant, so $U(R)$ preserves the L^2 inner product. Rotations also commute with the spherical Laplacian. The subspace of degree- l harmonics is therefore invariant, and its matrix representation is

$$\mathcal{D}_{mm'}^l(R) = \int_{\mathbb{S}^2} \overline{Y_{lm}(\Omega)} Y_{lm'}(R^{-1}\Omega) d\Omega. \quad (23)$$

Proposition 4 — Harmonic rotation and the invariant yaw sector. For the preceding convention, $\mathcal{D}^l(R)$ is unitary and satisfies $\mathcal{D}^l(R_1 R_2) = \mathcal{D}^l(R_1) \mathcal{D}^l(R_2)$. For a rotation $R_z(\alpha)$,

$$\mathcal{D}_{mm'}^l(R_z(\alpha)) = \delta_{mm'} e^{-im\alpha}. \quad (24)$$

Proof. Unitarity follows by expanding two rotated harmonics in an orthonormal basis and using invariance of the sphere integral. The composition identity follows from $U(R_1)U(R_2)f = f((R_1 R_2)^{-1}\Omega)$. A yaw rotation replaces φ by $\varphi - \alpha$, so the factor $e^{im\varphi}$ becomes $e^{-im\alpha} e^{im\varphi}$. $\square$

Thus a shared source coefficient vector transforms as $a_l(R) = \mathcal{D}^l(R)a_l$. For general rotations, Appendix A gives a generator construction and an independent evaluation check. These conventions matter: passive/active rotation or transposed indexing can otherwise produce a code that appears plausible while representing the wrong physical pose.




Figure 6: Rotational characters and the invariant $m=0$ sector. Real parts shown; a complex character also contains its imaginary component. [7]

8.2 What yaw cannot do

At $m = 0$, the yaw multiplier is exactly one. A centered source’s axisymmetric component is unchanged by orientation diversity alone. For an outgoing-only model, there are $L + 1$ such coefficients. At $L = 6$ they already exceed six independent receiver traces, giving at least one axisymmetric coefficient ambiguity in the centered idealized geometry. The consequence for a particular target still depends on whether that target changes along the missing direction.

The same statement must not be extended to all tangential derivatives. An axisymmetric function has zero azimuthal derivative $\partial/\partial\varphi$, but may have a nonzero polar derivative $\partial/\partial\theta$. Vertical or polar diversity can therefore observe some $m = 0$ structure that purely azimuthal sampling cannot. This is why a radial/vertical/non-collinear sensor comparison is more informative than a generic label such as “tangential sensing.”

8.3 Diversity and parameter sharing are different mechanisms

If a separate coefficient vector a_n is assigned to each pose, increasing pose count increases both observations and unknowns. Imposing one shared a reduces the admissible state dimension. Rotating

that shared state changes the rows of the measurement operator. Both can improve identifiability, but only the latter is new geometric or orientation diversity at fixed unknown count.

A fair ablation therefore compares equal-dimensional models when attributing a rank increase to motion. A linked direct/image-source model can have fewer unknowns than an unrestricted outgoing-plus-external split. An improvement produced by replacing 32 independent coefficients with 16 linked coefficients must not be credited entirely to the rotation. The archived Phase-II analyses explicitly separate these effects [3].

8.4 The rotating-room coupling

A useful schematic model is $y_{s,n} = [G_{0,s,n}(g) + G_{r,s,n}(g, \beta)]a + \varepsilon$, where g denotes geometry and β boundary or reflection parameters. The source state a is shared, but reflected propagation is pose-dependent. If both a and β are unknown, products such as $\beta_j G_{j,s,n} a$ make the model bilinear. Treating the entire problem as linear requires either fixing one block, enlarging the state with consistency constraints, or using a local Jacobian. The monograph never interprets the rank of that local Jacobian as a global uniqueness theorem.

9. Off-axis translation and multipole mixing

9.1 Motion of the source origin

If the source-expansion origin is offset by $\mathbf{d}$ from the rotation axis, its room position is $\mathbf{s}_n = \mathbf{t} + R_n \mathbf{d}$. A room-fixed receiver then experiences both a changing source orientation and a changing propagation vector. Even a locally axisymmetric source can generate pose-dependent path length through this motion.

For a scalar source and receiver displacement vector $\mathbf{r} = \mathbf{x} - \mathbf{s}$, the source-position sensitivity is

$$\nabla_{\mathbf{s}} G_0 = (ik + r^{-1}) G_0 \hat{\mathbf{r}}. \quad (25)$$

The first term is phase sensitivity and the second amplitude sensitivity. At large kr , phase dominates; close to a source, the inverse-distance term remains important. A small center displacement can therefore lift a finite-model degeneracy, but the resulting weak signature may be smaller than position or phase-calibration uncertainty. The derivative also shows why an unmeasured acoustic-center shift cannot automatically be treated as a known source of beneficial diversity.




Figure 7: Eccentricity can lift a finite-model near-null direction. L=3 outgoing source plus six receiver-specific static nuisances; 22 columns. Not the 32-column stationary control. [3]

The plotted offset sweep uses a 22-column outgoing-source/static-nuisance model, not the 32-column stationary control. It illustrates the mechanism without merging incompatible parameter counts. In a separate matched 32-parameter direct/image study [3], eccentricity at 1 kHz increases numerical rank from 30 to 32 while retaining a column-scaled condition number of approximately 1.27×10^5 . A formal lifting of the nullspace is therefore not enough to establish useful recovery under realistic error.

9.2 Addition theorem for a translated source

Proposition 5 — Exterior expansion of a translated monopole. For $r = |\mathbf{x}| > |\mathbf{d}| = d$, real $k > 0$, and the stated outgoing convention,

$$G_0(\mathbf{x}, \mathbf{d}; k) = -ik \sum_{l=0}^{\infty} \sum_{m=-l}^l j_l(kd) h_l^{(2)}(kr) Y_{lm}(\hat{\mathbf{x}}) \overline{Y_{lm}(\hat{\mathbf{d}})}. \quad (26)$$

Derivation. Expand the Green function in spherical harmonics about the observation origin. Each radial coefficient solves the separated Helmholtz equation, is regular on the smaller-radius side of the source and outgoing on the larger-radius side. Continuity and the Green-function derivative jump determine the product $j_l(kr_<)h_l^{(2)}(kr_>)$ and its factor $-ik$. The harmonic addition theorem completes the angular dependence [8,9]. At $d = 0$, only $l = 0$ remains; using $h_0^{(2)}(kr) = ie^{-ikr}/(kr)$ verifies the normalization in (4). Appendix B details the radial argument and the scope of finite truncation.

A translated monopole excites infinitely many orders about the old origin. Likewise, translating a higher multipole mixes orders. A finite-degree source space is not generally closed under translation, so a truncated translation matrix has approximation error even with perfect arithmetic. For a general outgoing basis function, a constructive coefficient definition is

$$T_{LM,lm}(\mathbf{d}; k) = \frac{\int_{\mathbb{S}^2} \overline{Y_{LM}(\Omega)} h_l^{(2)}(k|r_0\Omega - \mathbf{d}|) Y_{lm}(r_0\widehat{\Omega} - \mathbf{d}) d\Omega}{h_L^{(2)}(kr_0)}. \quad (27)$$

Here r_0 must enclose the translated singularity and lie in a valid exterior region. For the full outgoing expansion, the coefficient is independent of admissible r_0 by uniqueness of the separated exterior radial representation. A finite approximation must be checked for order convergence. The actual publication check evaluates (26) at four exterior directions for $L = 0$ through 16; its error decreases to near floating-point precision in that chosen case.




Figure 8: Convergence of the translated monopole expansion. 400 Hz, offset (0.12,0.03,-0.02) m, four exterior directions at radius 2 m. Pointwise finite test, not a universal truncation bound. [7]

9.3 Far-field translation and phase reference

For a physical translation of the source by $\mathbf{d}$ without changing its orientation, $|r\Omega - \mathbf{d}| = r - \Omega \cdot \mathbf{d} + O(r^{-1})$. Thus the far-field amplitude is multiplied by $e^{ik\Omega \cdot \mathbf{d}}$. A passive shift of the coordinate origin by δ , in contrast, multiplies the reported far-field amplitude by $e^{-ik\Omega \cdot \delta}$. Confusing these operations can turn a legitimate change of reference into an apparent acoustic phase error.

Physical interpretation. Eccentric motion can be an additional controlled excitation of spatial information, but only if its geometry is known and the translated source class remains adequate. The present evidence supports using it explicitly in the operator, not relying on an approximate center estimate as a substitute for metrology.

10. Source–room gauges and target-dependent invariances

10.1 Overlapping measurement subspaces

For a linear split $y = A_s a + A_r b$, let $\mathcal{S} = \text{range } A_s$ and $\mathcal{R} = \text{range } A_r$. A nonzero signature in their intersection can be attributed to either block. This ambiguity is separate from a nullspace internal to one block.

Proposition 6 — Dimension of shared-signature ambiguity. The joint rank satisfies

$$\text{rank}[A_s \ A_r] = \text{rank } A_s + \text{rank } A_r - \dim(\mathcal{S} \cap \mathcal{R}). \quad (28)$$

If both blocks are injective, each independent signature in the intersection induces one independent source–room exchange direction.

Proof. The column space of the concatenated matrix is $\mathcal{S} + \mathcal{R}$. The dimension formula for the sum of two subspaces gives (28). When the blocks are injective, every shared signature has unique preimages within each block, giving the joint null vector $(\delta a, -\delta b)$. If blocks are not injective, their internal kernels create further ambiguity. $\square$

This result explains why enlarging a room residual can improve training fit while reducing component identifiability. The numerical residual measures distance to the joint column space, whereas decomposition depends on the overlap between the component column spaces.

10.2 Two explicit room-model examples

For a centered rotating source, an $m = 0$ contribution is constant in yaw. If each receiver has an unrestricted static complex room value, any such source signature at the six receivers can be absorbed into those six nuisance values. More yaw samples repeat the same exchange freedom.

A Fourier room residual is more flexible:

$$r_s(\alpha) = \sum_{j=-K}^K c_{s,j} e^{ij\alpha}. \quad (29)$$

When its angular characters overlap those of the centered source, the corresponding signatures can be exchanged. At $K \geq L$, a receiver-specific residual can span every yaw character of a centered degree- L source. This is a negative control for unconstrained component fitting, not a recommendation to forbid all residual models. A restricted residual, measured boundary information, or shared reflection physics may retain useful distinctions.

10.3 Calibration gauges

An unknown common complex scale can be traded against source coefficients. Additional independent channels can break that gauge only if their reference relationship is itself established. Treating a source monitor as perfectly calibrated or taking a reference speaker’s voltage-to-volume-velocity ratio as exact is additional information, not a free consequence of adding a sensor.

The 58-parameter closure comparison retains reference-session nuisance variables even in architectures without reference observations. Some of those unobserved nuisance directions are harmless for the DUT target. Hence comparing only full parameter ranks across architectures can be misleading. The data-only target-null directions and prior-conditioned target covariance are the meaningful reported quantities [5].

10.4 Normalized targets are nonlinear and have different gauges

Let the complex directional field be $p_j = b_j q$ and choose a nonzero reference $p_r = b_r q$. Define complex normalized directivity $d_j = p_j/p_r$ and magnitude directivity $m_j = |p_j|/|p_r|$.

$$\delta d_j = \frac{b_j - d_j b_r}{p_r} \delta q, \quad \delta \log m_j = \text{Re} \left(\frac{\delta p_j}{p_j} - \frac{\delta p_r}{p_r} \right). \quad (30)$$

Proposition 7 — Common-scale and origin invariances. Both normalized targets are invariant under a common nonzero complex multiplier $p \mapsto \gamma p$. Under a passive origin shift δ , magnitude

directivity is unchanged, while complex directivity transforms as

$$d'_j = e^{-ik(\Omega_j - \Omega_r)\delta} d_j. \quad (31)$$

Proof. The common multiplier cancels in the complex ratio and its modulus cancels in the magnitude ratio. The origin-shift factor follows from the far-field phase transformation derived in Section 9, applied separately to direction j and the reference. Its modulus is one but its relative phase is generally nonzero. $\square$

A claim that magnitude normalization removes an ambiguity while complex normalization does not must therefore identify the actual gauge—such as directional datum phase—not attribute it to an arbitrary common phase that would cancel in both. Near zeros of p_r or p_j , the derivative formulas become ill-conditioned and need explicit masking or a different target norm. No amount of normalization guarantees a stable phase at a null.

11. Statistical information, priors, and target uncertainty

11.1 Complex observations and real parameters

For a proper complex Gaussian error with $\mathbb{E}[\varepsilon\varepsilon^*] = \Sigma_c$, a real parameter vector ϑ and complex Jacobian $J_c = \partial F / \partial \vartheta$, the real Fisher information is

$$\mathcal{J}_\vartheta = 2 \operatorname{Re}(J_c^* \Sigma_c^{-1} J_c). \quad (32)$$

This factor of two is necessary under the stated complex-RMS convention. The equivalent real observation stacks real and imaginary parts, with covariance

$$\Sigma_R = \frac{1}{2} \begin{bmatrix} \operatorname{Re} \Sigma_c & -\operatorname{Im} \Sigma_c \\ \operatorname{Im} \Sigma_c & \operatorname{Re} \Sigma_c \end{bmatrix}, \quad J_R = \begin{bmatrix} \operatorname{Re} J_c \\ \operatorname{Im} J_c \end{bmatrix}. \quad (33)$$

Then $J_R^T \Sigma_R^{-1} J_R$ equals (32). If parameters themselves are unconstrained complex coefficients, they must also be real-stacked, or one can work directly with the complex covariance of a linear estimator. Mixing these conventions leads to incorrect uncertainty comparisons between pressure, velocity, and geometry channels.

11.2 A target bound is not a coefficient covariance with zeros inserted

For an identifiable linear target and whitened unit-covariance noise, the covariance of the minimum-norm data-to-target map is $K_B K_B^*$. Its trace gives expected squared Euclidean target error. A corresponding local Cramér–Rao bound uses the target Jacobian and nuisance-profiled information under the required regularity and unbiasedness assumptions.

If the data leave a target-changing nullspace and no prior or bound controls it, a pseudoinverse of the singular information matrix does not represent finite uncertainty in that direction. Its zero eigenvalues in the discarded nullspace mean “not represented by this inverse,” not “known exactly.” The closure implementation returns an unbounded/unsupported target status instead of reporting a falsely finite full-target error [5,7].

For a linear complex observation with known covariance, one useful relative summary is

$$\epsilon_{\text{RMS}} = \frac{\sqrt{\operatorname{tr}(K_B \Sigma K_B^*)}}{\|t_0\|_2}. \quad (34)$$

The denominator t_0 and target grid are part of the definition. This is not a pointwise percentage error, a dB error, or a 95% coverage interval. Ratios become fragile near target zeros and should be accompanied by absolute or energy-normalized errors.

11.3 Data information and prior information must be reported separately

A Gaussian prior with precision Λ gives local posterior precision

$$\mathcal{J}_{\text{post}} = \mathcal{J}_R^T \Sigma_R^{-1} \mathcal{J}_R + \Lambda. \quad (35)$$

A full-rank posterior can arise because the prior supplies information absent from the data. Its covariance is conditional on the prior, the linearization, and the model class. It is not a measured reconstruction error and is not a data-only proof of identifiability. If the combined prior and likelihood retain a target-changing nullspace, even the posterior target uncertainty is not a proper finite covariance without further specification.

This distinction is decisive in the architecture studies. The quoted 1–2% target values for some registered local-sensor architectures are prior-conditioned analytic uncertainties. They must not be compared as though they were the experimentally measured accuracy of alternative commercial hardware systems.

11.4 Correlated uncertainty

A shared phase reference, common source drift, room temperature, turntable axis, or scanwise placement error couples many observation rows. The covariance propagation law includes cross terms; treating every repeated pose as independent would overstate the information [11]. For an error model $e_i = e_c + e_{i,\text{ind}}$ with independent zero-mean components,

$$\text{Var}(\bar{e}) = \sigma_c^2 + \frac{\sigma_{\text{ind}}^2}{N_\alpha}. \quad (36)$$

Only the independent part decreases with the number of poses. Covariance is not, however, a universal substitute for model discrepancy. An unknown coherent reflection or a wrong source model may produce a persistent bias outside the stochastic model used in (36).

12. Robustness to geometry and model discrepancy

12.1 Operator error acts through the unknown source

Let the nominal whitened model be $y_w = H\xi + e$, the actual forward operator $H + \Delta H$, and the actual target operator $B_S + \Delta B$. For a nominally exact target map $K_B H = B_S$,

$$\hat{t} - t_{\text{true}} = K_B e + K_B \Delta H \xi + K_B \delta_{\text{mod}} - \Delta B \xi. \quad (37)$$

If $\|\xi\| \leq M$, a first-order norm bound is

$$\|\hat{t} - t_{\text{true}}\| \leq \|K_B\|(\|e\| + M\|\Delta H\| + \|\delta_{\text{mod}}\|) + M\|\Delta B\|. \quad (38)$$

The state bound M is not optional. A geometrical perturbation multiplies the unknown source; it is not generally equivalent to a fixed additive noise independent of that source. When geometry is explicitly estimated as nuisance, its Jacobian and covariance must be propagated jointly rather than counted twice as both a state uncertainty and separate white noise.

12.2 Singular-value perturbation is informative but not sufficient

For a matrix perturbation E , $|\sigma_i(H + E) - \sigma_i(H)| \leq \|E\|_2$. Thus if the smallest useful singular value is comparable to the plausible perturbation norm, its apparent nonzero value should not be interpreted as robust information. But a worst-case arbitrary matrix perturbation can be much more pessimistic than physically structured coordinate errors. Both the normalization and the allowed perturbation family matter.

A local coordinate-error model has $\Delta H \approx \sum_j H_{,g_j} \delta g_j$. Interval boxes, correlated Gaussian parameter models, and a physically measured repeatability distribution lead to different uncertainty statements. Exhaustively checking corners of a linearized box can bound a convex norm of that linear response; it is not a proof for the nonlinear geometry map far from the linearization point.

12.3 Source and model inadequacy

Write the physical field as the contribution of retained source coefficients plus an unresolved component. Increasing the source class can reduce approximation bias but weaken identifiability. Reducing it can stabilize the matrix while excluding a real diaphragm mode or port contribution. Therefore source-order selection must depend on prediction at withheld spatial and pose observations, not just training residual or matrix condition.

The earlier Phase-II stress tests are retained as a warning [3]: large manual receiver-coordinate errors can destroy an otherwise successful exact-geometry near-field-assisted reconstruction. The newer restricted prior-conditioned comparisons do not erase those failures. They describe different assumptions and must not be used to claim general robustness of the existing manual setup.

12.4 What calibration can and cannot rescue

A measured calibration may reduce uncertainty in gain, phase, group delay, sensor positions, and reference strength. It does not make an inadequate radiation basis exact or transform an unmodelled obstacle into a known reflection coefficient. The proper question is whether the remaining, measured calibration and discrepancy budget lies inside a target gate. “Calibration can correct it” is not a substitute for specifying which parameter is calibrated, with which independent observation, and how its residual enters (37).

13. Local source-state measurements and spatial sampling

13.1 A measurement close to a source is not source truth

A local pressure vector is represented as $z_{\text{loc}} = C_{\text{loc}} a + D_{\text{loc}} \eta + \varepsilon_{\text{loc}} + \delta_{\text{loc}}$. The first term is its sensitivity to the retained source state. Other radiators, coherent room contributions, uncertain microphone transfer functions, and model error can remain. Source proximity may improve the ratio of desired signal to room contamination, but it does not by itself provide a calibrated boundary condition or a unique mapping to diaphragm volume velocity.

Under a validated scalar radiator model, one complex local pressure can constrain one complex source signature strongly. Under a richer source model, one point observes only one linear combination of several spatial modes. The relevant criterion is the nuisance-profiled restriction in Theorem 3, not the label “near field.” A fixed fixture is attractive because it controls relative geometry, not because it makes every local sample reflection-free.

13.2 Source complexity before sensor count

At 63 Hz, the extended campaign compares declared source classes under its stated noise conventions [4]. Two scalar sources with a restricted external field have about 0.43% analytic relative target

RMS uncertainty with no added local observations. A five-source-coefficient woofer-mode/port class reaches about 1.27% with two added local observations. In a broader 22-coefficient joint class, two, four, and six added observations give approximately 56.44%, 19.70%, and 11.58%, respectively. These are total added observations, not observations per radiator.




Figure 9: Source complexity controls required local information. 63 Hz, exact stated source/room classes and noise assumptions. Different source classes are not alternative fits to measured real-DUT data. [4]

The comparison is not evidence that the real DUT has any one of these classes. It shows why a universal sensor-count answer cannot be derived without source-model adequacy. A compact model can be both informative and wrong; a rich model can be physically plausible but unobservable. The practical task is to find the smallest model that predicts withheld observations within a measured uncertainty/discrepancy allowance.

13.3 Symmetry can defeat repeated axial sampling

Proposition 8 — Axis invisibility of odd disk modes. Consider a circular radiating surface in a rigid infinite baffle. Let the kernel from each surface point to an axial pressure observer depend only on its radius. A prescribed velocity mode odd under a reflection of one transverse coordinate contributes zero pressure at every point on the center axis.

Proof. Pair each surface point (x, y) with $(-x, y)$. The axial propagation kernel and surface measure are identical for the pair, while the odd velocity changes sign. Their contributions cancel in the radiation integral. Moving the observation along the same axis changes the common kernel but not this cancellation. $\square$

The disk study [3] makes this concrete. Three pressure points at different center-axis distances observe only two of four prescribed modes; five distributed points observe all four in the tested class. Four distributed ideal surface-velocity samples also observe all four, while one central velocity sample observes only one. These comparisons establish finite modal visibility, not equal-cost or equal-noise superiority of LDV over pressure.




Figure 10: Spatial symmetry and source-mode observability. Prescribed four-mode baffled disk, 80 mm radius, 1 kHz. Ideal LDV is not compared at an equal physical noise/cost budget. [3]

A single spatial measurement vector at a single frequency is one snapshot. Repeating a scalar multiple of the same drive generates a rank-one snapshot matrix regardless of how many diaphragm modes contributed to that pattern. Its SVD cannot count independent physical source modes without additional modelling or independent source states. Registered held-out points test a proposed basis; separately controlled excitations can test independent source directions. Frequency stacking requires a frequency-dependent model rather than silently treating different frequencies as repeated draws of one fixed matrix.

13.4 Co-rotation and common-mode contamination

Co-rotating local sensors have a fixed relative direct-field geometry. A useful additive decomposition is $z_i(\alpha) = s_i + r_i(\alpha) + \varepsilon_i(\alpha)$. Averaging over uniformly sampled yaw retains the $m = 0$ component of every term, including the room:

$$\bar{z}_i = s_i + \bar{r}_i + \bar{\varepsilon}_i. \quad (39)$$

A floor reflection can be invariant in yaw in an idealized geometry and therefore survive every average. The executed three-amplitude control observes only two independent amplitudes with one local sensor over 72 poses [4]. More poses may improve the independent-noise term while leaving a source/constant-reflection exchange unresolved.




Figure 11: Pose averaging cannot remove a common error floor. Assumed correlated uncertainty model; not measured calibration repeatability. [4]

Physical interpretation. Co-rotation addresses the user’s practical concern about tedious and poorly repeatable manual scanning. It is a strong candidate for a registered workflow. Its legitimate benefits are geometry control, common timing, and useful repeated observations; its mean is not automatically the direct source signal.

14. Independent velocity, surface data, and equivalent sources

14.1 Pressure and a true radial derivative

Suppose spherical harmonic coefficients of both pressure and radial derivative are known on a complete sphere in a source-free annulus. For each (l, m) the two radial families satisfy

$$\begin{bmatrix} p_{lm} \\ \partial_r p_{lm} \end{bmatrix} = \begin{bmatrix} h_l^{(2)}(kr) & j_l(kr) \\ kh_l^{(2)'}(kr) & kj_l'(kr) \end{bmatrix} \begin{bmatrix} a_{lm} \\ b_{lm} \end{bmatrix}. \quad (40)$$

Proposition 9 — Nonvanishing radial Cauchy determinant. For real $k, r > 0$, the determinant of the matrix in (40) equals $i/(kr^2)$ and is nonzero.

Proof. The spherical Bessel Wronskian is $j_l y_l' - j_l' y_l = 1/z^2$, with $z = kr$ [9]. Substituting $h_l^{(2)} = j_l - iy_l$ gives $h_l^{(2)} j_l' - j_l h_l^{(2)'} = i/z^2$. Multiplication of the derivative row by k gives $ik/z^2 = i/(kr^2)$. $\square$

This establishes per-mode uniqueness for the stated complete-sphere or sufficiently sampled finite-mode Cauchy data. It does not establish that one local p–u sensor separates arbitrary room fields. The angular coefficients must themselves be observable, both channels need a common reference and noise model, and the finite radial matrix can still be poorly conditioned. Appendix C examines these qualifications quantitatively.

True particle velocity relates to the derivative through (5) and is a genuinely independent finite observation when separately measured. This is the distinction exploited in pressure–velocity separation methods [13,14]. Velocity computed from the same retained pressure samples has only the

information characterized by Proposition 2.

14.2 Boundary representation and LDV

Let Ω be a bounded source-free acoustic region, with outward normal n . Green's second identity and (4) give, for $\mathbf{x} \in \Omega$,

$$p(\mathbf{x}) = \int_{\partial\Omega} [G_0(\mathbf{x}, \mathbf{y}) \partial_n p(\mathbf{y}) - p(\mathbf{y}) \partial_n G_0(\mathbf{x}, \mathbf{y})] dS_{\mathbf{y}}. \quad (41)$$

For an exterior radiator problem, the inner boundary normal of the acoustic domain points toward the radiator; using the outward radiator normal instead reverses the corresponding signs. This convention must be settled before comparing a boundary-element implementation with a source-surface velocity measurement.

For an ideal rigid infinite baffle and normal velocity into its radiating half-space, the Rayleigh representation is

$$p(\mathbf{x}) = \frac{i\omega\rho_0}{2\pi} \int_S v_n(\mathbf{y}) \frac{e^{-ik|\mathbf{x}-\mathbf{y}|}}{|\mathbf{x}-\mathbf{y}|} dS_{\mathbf{y}}. \quad (42)$$

A dense calibrated surface-velocity measurement can therefore constrain the radiating state more directly than a local pressure surrogate under that model. Sparse LDV samples still require a surface basis and a sampling matrix. An unobserved odd mode at a central laser point is no less unobserved because the instrument is optical. Normal-direction registration, phase reference, surface access, acoustic loading, and the validity of the boundary representation remain part of the uncertainty budget. COMSOL's baffled-piston model is a useful future canonical forward check, not an execution result of this monograph [18].

14.3 Equivalent-source methods in the same framework

An equivalent-source model places source coefficients q_j at prescribed interior locations $\mathbf{s}_j$ and predicts

$$p(\mathbf{x}_i) = \sum_{j=1}^J G_0(\mathbf{x}_i, \mathbf{s}_j) q_j, \quad y = A_{\text{ESM}} q. \quad (43)$$

A continued pressure or far-field target uses a different evaluation matrix B_{ESM} . Theorem 1 then applies directly: a fit at the measurement positions does not guarantee the continuation if B_{ESM} changes on $\ker A_{\text{ESM}}$. Source placement, regularization, and support restrictions control conditioning and model bias. Increasing the number of equivalent sources can improve interpolation while leaving continuation dependent on priors. Classical near-field holography provides the broader context for such spatial continuation problems [19].

15. Coupled radiators and electrical information

15.1 A low-frequency system is not necessarily one source

A bass-reflex loudspeaker has at least an active diaphragm contribution and a port contribution in a low-order model. A passive-radiator system has two mechanically coupled radiating states. A multiwoofer system may need common and differential source coordinates. The lowest electrical branch is therefore not automatically the only acoustic degree of freedom relevant below the gating band.

For a two-radiator model, a target can be written $t = b_w Q_w + b_p Q_p$. The relative error from omitting the secondary radiator is

$$\epsilon_{\text{omit}} = \frac{\|b_p Q_p\|}{\|b_w Q_w + b_p Q_p\|}. \quad (44)$$

This can be large when the secondary radiator dominates or when the denominator is reduced by cancellation. The archived synthetic vented example gives approximately 95.8% error near its tested 50.84 Hz tuning-region point [4]. That frequency and error are not values measured on the user’s actual loudspeaker.




Figure 12: Secondary-radiator omission in coupled low-frequency models. Synthetic coupled systems. Large relative errors can also reflect destructive interference in the denominator. [4]

Whether a port needs its own sensor depends on whether its state is already constrained adequately by a coupled electromechanical model and the available measurements. “Every port requires a microphone” and “the woofer alone is always enough” are both unjustified general rules. A practical first fixture can provide a second channel for the port or passive radiator when that is the dominant missing source state.

15.2 What terminal voltage and current constrain

In a linear lumped moving-coil approximation, with all quantities referred to a consistent coordinate,

$$V = Z_e I + (Bl)v, \quad F = (Bl)I, \quad (45)$$

and the mechanical/acoustic load relates F to v . If Z_e and Bl are known, $v = (V - Z_e I)/(Bl)$ estimates a coil-related velocity. The subtraction can be ill-conditioned when the motional contribution is small relative to the electrical impedance term. Parameter uncertainty, heating, frequency dependence, and nonlinearities all matter.

The coil velocity is not a complete spatial velocity field of a breaking-up diaphragm. Nor does one terminal impedance uniquely determine arbitrary port flow, cabinet radiation, or an unobserved differential mode of multiple radiators. Electrical measurements can nevertheless stabilize an appropriate low-order state model and supply a drive reference. Their information value must be

tested in the same joint source/nuisance parameterization as local pressure rather than treated as a free exact source truth.

15.3 Source-class validation

The next physical decision is whether a compact low-frequency source signature predicts withheld registered observations. A one-amplitude model, a woofer-plus-port model, and a small modal model should be compared with the same held-out targets and measured calibration budget. AIC or BIC can assist when their likelihood and effective parameter assumptions are met, but correlated deterministic development data do not automatically meet those assumptions. Predictive residual structure and independently registered observations are often more informative than a single model-selection score.

16. Optical geometry, registration, and room-plane information

16.1 A geometric sensor measures a different parameter block

Let g denote geometry and β acoustic boundary parameters. A room operator is $G(g, \beta)$, while optical observations have the form $z_{\text{opt}} = L(g) + \varepsilon_{\text{opt}}$. Optical rows can remove geometric ambiguities and reduce operator uncertainty. They do not directly measure impedance, complex reflection coefficients, absorption, or diffusion.

The matched architecture studies deliberately compare this geometric information with acoustic and calibration information [4,5]. They do not infer the value of a scanner from its point density alone. A dense cloud with poor visibility or an uncertain transform to the acoustic frame can leave the target-relevant geometry weakly constrained.

16.2 Range and plane derivatives

For a plane $n^\top x = h$ with unit normal n , and a laser ray $x = o + tu$, the ideal range is

$$t = \frac{h - n^\top o}{n^\top u}, \quad n^\top u \neq 0. \quad (46)$$

Its first-order perturbation obeys

$$\delta t = \frac{\delta h - \delta n^\top o - n^\top \delta o - t(\delta n^\top u + n^\top \delta u)}{n^\top u}. \quad (47)$$

Grazing incidence makes the denominator small. Plane-normal errors, scan-angle errors, scanner-origin errors, and range noise are therefore not interchangeable. The unit-normal constraint leaves two local orientation degrees of freedom per plane, plus one offset. Registration to the DUT/receiver frame adds its own rigid-transform uncertainty.

16.3 Why horizontal coverage can be insufficient

A single horizontal scan samples an intersection line of each visible wall with the scan plane. It need not identify all out-of-plane wall tilts, and it generally does not observe the floor and ceiling. Three horizontal heights can add wall constraints without sampling the horizontal boundaries themselves.

In the declared six-plane, 18-geometric-parameter synthetic room, the rank calculation gives 8 observed parameters for one horizontal scan, 12 for three heights, and 18 for elevation-diverse or spherical coverage [4]. Known face labels are used for this observability calculation. It is not an experimental proof that arbitrary real-scene segmentation succeeds.




Figure 13: Room-plane geometry requires adequate angular coverage. Six unconstrained planes, 18 parameters. Known face labels used for rank analysis; not an arbitrary-scene segmentation result. [4]

This supports elevation diversity rather than a blanket demand for an expensive scanner. The correct specification is target-dependent: which surfaces, normals, distances, and transforms must be constrained to the required acoustic uncertainty? Some local variables may need tighter registration than large room boundaries.

16.4 Two-tier metrology

A broad room reference scan can generate candidate surfaces and paths, while separate local metrology establishes the turntable axis, receiver coordinates, source fixture, and acoustic datum. The resulting error blocks should be propagated jointly. A cylindrical or other known reference artifact can assist frame registration, but an axisymmetric artifact alone cannot determine every rotation without suitable features or extra measurements. Two range measurements are likewise not a general six-degree-of-freedom pose solution; their sufficiency depends on the known geometry and remaining unknowns.

The paper does not select a particular laser or LiDAR product. The numerical studies justify investigating a coverage-first room scan and tighter local registration; they do not establish a universal millimetre requirement or a final low-cost hardware acceptance.

17. Geometry-constrained reflections and low-frequency room modes

17.1 Image-source construction

For a plane $n^\top x = h$, reflection of source point $\mathbf{s}$ produces the image

$$\mathbf{s}' = \mathbf{s} + 2(h - n^\top \mathbf{s})\mathbf{n}, \quad L = |\mathbf{x} - \mathbf{s}'|, \quad \tau = L/c. \quad (48)$$

A candidate frequency-domain path has factor $\beta(f)e^{-ikL}/(4\pi L)$, with visibility and finite-surface geometry checked separately. For ideal infinite rigid or pressure-release planar boundaries, image

constructions can be exact wave solutions rather than high-frequency ray approximations [15]. It is therefore incorrect to reject every image-source method solely because frequency is low. In contrast, a small list of scalar specular coefficients is not generally an exact model for finite, impedance, diffuse, or resonant boundaries.

Scanned geometry can constrain path lengths and candidate directions while acoustic observations estimate boundary response. A model such as

$$y = \left[G_0(g) + \sum_j \beta_j(f) G_j(g) \right] a + r_{\text{res}} + \varepsilon \quad (49)$$

remains bilinear in unknown source state and reflection coefficients. A richer residual can still absorb source signatures, as Section 10 demonstrates. Echo-based geometry studies motivate careful path association [16]; a detected IR event is not automatically a uniquely labelled wall reflection.

17.2 A phase budget is not a complete inversion tolerance

Proposition 10 — Conservative single-plane offset phase bound. With fixed normal and a small offset perturbation δh , the image-source displacement has norm $2|\delta h|$. Hence the path-length perturbation obeys $|\delta L| \leq 2|\delta h|$, giving

$$|\delta\phi| = k|\delta L| \leq 2k|\delta h|, \quad |\delta h| \leq \frac{c\phi_{\text{max}}}{4\pi f}. \quad (50)$$

Proof. The image point in (48) changes by $2\delta h \mathbf{n}$. The reverse triangle inequality bounds the change in its distance to the receiver by that displacement. Multiplication by k gives the phase bound. $\square$

For $c = 343$ m/s and $\phi_{\text{max}} = 10$ degrees, the bound is about 47.6 mm at 100 Hz, 4.76 mm at 1 kHz, and 0.476 mm at 10 kHz. These are kinematic single-path bounds. In an inverse problem, unknown reflection phase can absorb some changes, while cancellation, resonances, and weak target directions can amplify others.




Figure 14: Conservative single-reflection plane-offset phase bound. Single-path kinematic bound. Not a global inverse-problem or room-modal tolerance. [7]

17.3 Room modes and dimension sensitivity

For the ideal rectangular rigid room, modal frequencies are

$$f_{n_x n_y n_z} = \frac{c}{2} \sqrt{\frac{n_x^2}{L_x^2} + \frac{n_y^2}{L_y^2} + \frac{n_z^2}{L_z^2}}. \quad (51)$$

Proposition 11 — First-order dimension sensitivity. For a nonzero mode,

$$\frac{\partial f_n}{\partial L_x} = -\frac{c^2 n_x^2}{4 f_n L_x^3}. \quad (52)$$

Proof. Differentiate the squared expression $f_n^2 = (c^2/4) \sum_i n_i^2/L_i^2$, holding c and the other dimensions fixed. $\square$

A damped modal contribution is proportional to $(\omega_n^2 - \omega^2 + i2\zeta_n \omega_n \omega)^{-1}$ under the stated time convention. Near a lightly damped resonance, even a modest shift of ω_n relative to the linewidth can produce a large complex field change. Long wavelength therefore does not automatically imply insensitivity to room dimensions.




Figure 15: Low-frequency modal sensitivity to room dimensions. Smoothed 0.20 m source, damped Neumann room model. Not point-source or actual-room truth. [4]

The archived converged smoothed-source example shows approximately 17.7% field change near 31.18 Hz for a 20 mm dimension perturbation [4]. Its source is an explicitly smoothed 0.20 m Neumann heat-kernel source with damping, not an arbitrarily truncated point-source singularity or the actual DUT. Modal-order convergence and the source model are part of that result.

17.4 Transfer, diffraction, and scene changes

A scanned empty room and the room with a large DUT, supports, local fixtures, and an optical scanner in place need not have the same transfer operator. Source substitution can change scattering as well as source directivity. Device-aware image-source research explicitly addresses such transfer effects [17]. The optical apparatus should therefore be removed or retained consistently, with the scene difference either modelled or bounded. Registration cannot correct an unacknowledged change

of physical boundary conditions.

18. Reference-source identification and its calibrated subspace

18.1 What a known reference actually identifies

Let $G \in \mathbb{C}^{m \times n}$ be an unknown linear room-transfer operator on a declared source-coordinate space. Columns of $Q \in \mathbb{C}^{n \times r}$ are known calibration source states. The measurements are

$$Y_{\text{ref}} = GQ. \quad (53)$$

Theorem 4 — Reference-source subspace identifiability. For an otherwise unrestricted G , the action Ga on a specified source state a is uniquely determined by (53) if and only if $a \in \text{range } Q$.

Proof. If $a = Qc$, then $Ga = Y_{\text{ref}}c$. Although c need not be unique when Q has a kernel, any difference satisfies $Q\delta c = 0$, hence $Y_{\text{ref}}\delta c = GQ\delta c = 0$; the target is unique.

If $a \notin \text{range } Q$, there is a vector $v \in \ker Q^*$ with $v^*a \neq 0$. Choose any nonzero receiver vector u and set $\Delta G = uv^*$. Then $\Delta GQ = 0$ but $\Delta Ga = u(v^*a) \neq 0$. Thus G and $G + \Delta G$ agree on the calibration and disagree on the DUT action. $\square$

This is an exact theorem for unrestricted finite matrices. It does not construct two physically realizable rooms with arbitrary prescribed Green operators, nor does it rule out stronger conclusions under known Helmholtz, reciprocity, geometry, and boundary parameter constraints. A low-dimensional parametric room may be identified by fewer reference excitations if its own joint Jacobian is observable. That is a different theorem with additional assumptions.

18.2 Repetition versus source-space coverage

A reference monopole at one location samples one source direction. Repeating it 72 times improves independent-noise averaging, not the rank of the excited source space. Multiple positions or controlled source patterns can add directions, but the source strengths must actually be known. A free-field voltage calibration of the reference source does not automatically establish its in-room volume velocity or remove session gain uncertainty.

For a local first-order representation,

$$G(\mathbf{s} + \delta\mathbf{s}) = G(\mathbf{s}) + \nabla_{\mathbf{s}}G(\mathbf{s}) \cdot \delta\mathbf{s} + O(\|\delta\mathbf{s}\|^2). \quad (54)$$

The feature matrix has columns $[1, \delta x, \delta y, \delta z]$. Repeating the center has rank one, three vertical points rank two, a horizontal layout at most three, and suitable tetrahedral or axial-three-dimensional positions rank four. These are feature ranks of a Taylor model, not a universal requirement for four physical devices.




Figure 16: Reference-source spatial feature coverage. First-order Taylor basis $[1, dx, dy, dz]$; not a universal minimum number of hardware sources. [5]

18.3 Transferability and derivative estimation

Small reference stencils make derivative estimation sensitive to source-placement and complex calibration errors; large stencils increase Taylor truncation error. The closure study compares that trade-off using an exact two-plane synthetic transfer model and 16 deterministic stress scenarios per configuration [5]. For a seven-point first-order calibration predicting a held neighborhood of radius 125 mm, the post-hoc best tested median errors are about 1.47% at 63 Hz, 6.35% at 250 Hz, and 67.1% at 1 kHz. These are model-specific grid results, not universal best standoffs or statistical coverage guarantees.




Figure 17: Reference-source transferability versus spatial stencil. Seven-point first-order calibration; 125 mm held neighborhood; 16 deterministic stress scenarios, not statistical confidence coverage. [5]

A scalar correction formed from $G_0(\mathbf{x}, \mathbf{s}_0)/G_{\text{room}}(\mathbf{x}, \mathbf{s}_0)$ is exact only for its matched scalar-source case, away from problematic denominator zeros. Applying it to a displaced or distributed source does not generally yield the correct direct field. Theorem 4 explains the source-space reason; the Taylor experiments show a finite-model transfer consequence.

18.4 Empirical Green functions and reciprocity

Reciprocity can relate $G(\mathbf{x}, \mathbf{s})$ and $G(\mathbf{s}, \mathbf{x})$ in an appropriate linear reciprocal medium with compatible boundary conditions. It does not turn one source position into a complete Green function over all possible DUT source distributions. Spatial interpolation, reduced-order wave models, or geometric boundary priors can extend the calibrated neighborhood, but each adds structure that must survive mismatch tests.

Physical interpretation. A reference source is potentially useful as a targeted room-calibration measurement. It is not a universal room-removal curve. Its value must be evaluated jointly with source-space coverage, known in-room strength, scene substitution, and the uncertainty of transferring its operator to the actual rotating DUT.

19. Operational free field versus same-drive anechoic response

19.1 Two measurands that must not be conflated

Let a_{room} denote the source state during the room measurement, and a_{free} the state produced by the same electrical command in a free-field environment. Direct propagation of the measured operational state gives $B_0 a_{\text{room}}$. The same-drive anechoic target is $B_0 a_{\text{free}}$. They coincide only when the source state is unchanged, or when their difference is known to be insignificant for the requested target.

$$t_{\text{oper}} = B_0 a_{\text{room}}, \quad t_{\text{anechoic}} = B_0 a_{\text{free}}. \quad (55)$$

Acoustic loading can modify cone or port motion; substituting a reference source can also change the scattered field. Correctly removing the reflected propagation component does not, by itself, prove the stronger same-drive counterfactual.

19.2 A source-loading bound

Suppose a source-coordinate impedance model gives $Z_0 a_{\text{free}} = F$ and $(Z_0 + \Delta Z) a_{\text{room}} = F$, with all generalized forces and impedances referred to consistent coordinates. Assume Z_0 is invertible and define $E = Z_0^{-1} \Delta Z$.

Proposition 12 — Bounded source-loading effect. If $\|E\| \leq r < 1$ in a submultiplicative induced norm, then

$$a_{\text{room}} = (I + E)^{-1} a_{\text{free}}, \quad \frac{\|a_{\text{room}} - a_{\text{free}}\|}{\|a_{\text{free}}\|} \leq \frac{r}{1 - r}. \quad (56)$$

Proof. The Neumann series gives $(I + E)^{-1} = \sum_{j=0}^{\infty} (-E)^j$ and norm at most $(1 - r)^{-1}$. Subtracting a_{free} yields $-(I + E)^{-1} E a_{\text{free}}$, with the stated bound. For a scalar negative-real perturbation $E = -r$, equality is attained. $\square$




Figure 18: Operational source state and same-drive anechoic state. Requires a common source coordinate and invertible reference impedance. This is not a measured correction for the real DUT. [7]

This is a sufficient model-based bound. It does not measure r for the actual loudspeaker. A terminal-drive model may first need electromechanical elimination before it takes the common-force form assumed above. Source-state monitoring measures a_{room} ; an independent loading model or validation must connect it to a_{free} .

19.3 Product implications

The operational direct-field target can already be useful for radiation diagnostics, directivity comparison under declared conditions, and simulation correlation. The stronger anechoic-at-same-drive claim needs additional qualification. A future instrument should state which target it reports rather than use “free field” as an undifferentiated label. This target distinction is one of the most consequential outcomes of the architecture-closure analysis [5].

20. Cross-frequency structure and physics-informed factorization

20.1 Shared structure can stabilize a valid model

A source state may be represented by a low-order frequency basis, $a(f_k) = \sum_j V_{kj} \theta_j$, rather than estimated independently at every bin. Resonance models, causal rational functions, or smooth bases can reduce variance if the true response lies sufficiently close to their span. The gain is conditional on this restriction. It does not constitute new measured information at an otherwise invisible state direction.

For a correctly specified linear subspace and white noise, orthogonal projection reduces expected noise energy in proportion to the retained dimension. But the discarded component of a richer truth becomes bias. At a physical resonance or crossover, excessive smoothing can remove exactly the feature that the reconstruction is intended to characterize.

20.2 A smooth exact gauge

Proposition 13 — Shared frequency bases can preserve component ambiguity. If source and room amplitudes enter as $y = Va + Vb$ for the same full or partial basis V , then $(a, b) \mapsto (a + c, b - c)$ leaves all frequency observations unchanged for every c .

Proof. $V(a + c) + V(b - c) = Va + Vb$. Thus $[V \ V](c, -c)^\top = 0$ independently of the rank or smoothness of V . $\square$

The same algebra holds if the common basis is built from causal impulse responses. Causality by itself therefore does not universally resolve component exchange. Stronger physical constraints may help, but their restrictions and residual gauges must be derived for the particular model.

The closure study retains this negative control and a model-mismatch test [5]. Its analytic white-noise reference has an RMS reduction from about 2.48% to 0.393% under projection, whereas a separate deterministic richer-truth test gives about 25.9% median error after an inadequate coupled model versus 2.12% without that restrictive coupling. These are different experiments and statistical interpretations; they must not be merged into one performance claim.




Figure 19: Cross-frequency stabilization versus model bias. Deterministic development scenarios. Do not interpret these medians as white-noise statistical confidence coverage. [5]

20.3 Tensor organization is not automatic physical separation

The natural data organization is $Y_{s,n,k}$ for receiver, pose, and frequency. Known rotations and repeated source support can be exploited without flattening away that structure. A physics-informed tensor or reduced-order model can improve efficiency, enable shared parameters, and reveal residual structure.

However, if source and room admit the same separable factor, transferring it from one component to the other leaves the tensor sum unchanged. Algebraic uniqueness of a chosen factorization, when available, also does not automatically identify its factors as direct source and room. The physical map from factors to the target must be validated. Thus tensor methods are permitted representations within the observability framework, not an alternative to it.

21. Time gating: exact error decomposition and qualification

21.1 The identity behind every gate

Let the impulse response be $h = h_d + h_r$, where h_d is the chosen direct-source target response and h_r is the remainder. Let $w(t)$ be the analysis window. The gated estimate is $H_g = \mathcal{F}\{wh\}$.

Proposition 14 — Exact gating-error identity. For signals and windows for which the transforms exist,

$$H_g - H_d = \mathcal{F}\{(w - 1)h_d\} + \mathcal{F}\{wh_r\}. \quad (57)$$

Proof. Expand $wh = w(h_d + h_r)$, subtract h_d , and apply linearity of the Fourier transform. $\square$

The first term removes genuine source response. The second retains room response. Excluding an identified reflection therefore does not establish that the complete direct response is retained. Conversely, keeping a long decay may preserve source ringing and also include reflections. Both terms must be qualified for the requested complex target.

21.2 Why a fixed crossover rule is inadequate

The characteristic spectral smoothing scale of a finite window is related to $1/T_g$, but that does not establish a universal lowest valid response frequency. If a known short direct pulse is entirely contained in the window and no room contribution is retained, its Fourier transform can be accurate even at frequencies with periods much longer than the pulse. A ringing direct response can instead be biased at relatively high frequencies when its tail is removed.

The synthetic gating studies use identical 3 ms direct-to-reflection separation and a 2.55 ms usable gate with different prescribed source responses [4]. Their passing higher-frequency grid regions begin at about 50 Hz, 312 Hz, and 2.63 kHz for a short pulse, a low-frequency decay, and long high-frequency ringing. These are properties of those known test signals and gates. They are not a measurement-resolution guarantee, a real-DUT crossover, or a general instruction to use millisecond windows for arbitrary low-frequency loudspeaker responses.




Figure 20: The same gate need not have the same useful frequency band. 3 ms direct/reflection gap, 2.55 ms rectangular gate; synthetic known direct responses. No real transition frequency is inferred. [4]

21.3 Identifying the direct response is part of the inverse problem

The next IR feature after an initial pulse may be a room reflection, another radiator, diffraction, crossover ringing, or an extended source response. Geometric delay candidates help narrow the alternatives but do not label them uniquely without additional evidence. Room scans, reference-source measurements, pose dependence, and known source dynamics can contribute to that association.

The supplied real development data do not independently establish the complete direct waveform or the first reflected event. Their real transition field therefore remains unspecified. A deployable gate qualification must estimate or bound both terms in (57) using information available in a real measurement—not an oracle that reads the known synthetic direct response.

21.4 Phase-consistent targets

A gate and an inverse reconstruction must refer to the same source origin, target position/direction, complex electrical reference, and time origin before fusion. Arbitrarily shifting or minimum-phase-reconstructing one estimate changes the target. A magnitude-only splice may be useful for a narrower application, but it is not equivalent to recovery of a complex polar field. The distinction should be reflected in both reported confidence and output metadata.

22. Bias-aware hybrid estimation

22.1 Correlated estimates of one complex target

Let two estimates of the same complex scalar target be $\hat{p}_L = p + b_L + e_L$ and $\hat{p}_H = p + b_H + e_H$. The labels suggest a source-assisted and gated estimator but do not impose a fixed frequency split. Assume zero-mean stochastic errors with variances v_L, v_H and cross-covariance $c = \mathbb{E}[e_L \bar{e}_H]$. For a real convex weight w ,

$$\hat{p} = w\hat{p}_L + (1 - w)\hat{p}_H, \quad 0 \leq w \leq 1. \quad (58)$$

When biases vanish, the mean-squared error is $w^2 v_L + (1 - w)^2 v_H + 2w(1 - w) \operatorname{Re} c$. Shared raw data, timing, or geometry generally make c nonzero. Treating the two estimates as independent can therefore overstate the benefit of combining them.

22.2 A minimax bound for unknown bounded bias

Suppose only $|b_L| \leq \beta_L$ and $|b_H| \leq \beta_H$ are known. Their worst combined bias magnitude for a convex weight is $w\beta_L + (1 - w)\beta_H$, attained by aligned complex biases.

Proposition 15 — Scalar convex bias-aware weight. With valid covariance and nonnegative bias bounds, the worst-bias quadratic risk is

$$R(w) = w^2 v_L + (1 - w)^2 v_H + 2w(1 - w) \operatorname{Re} c + [w\beta_L + (1 - w)\beta_H]^2. \quad (59)$$

If $D = v_L + v_H - 2 \operatorname{Re} c + (\beta_L - \beta_H)^2 > 0$, its minimizer on $[0, 1]$ is

$$w^* = \operatorname{clip}_{[0,1]} \left(\frac{v_H - \operatorname{Re} c + \beta_H(\beta_H - \beta_L)}{v_L + v_H - 2 \operatorname{Re} c + (\beta_L - \beta_H)^2} \right). \quad (60)$$

Proof. Expand (59) as a constant plus a linear term and Dw^2 . Positive semidefiniteness of the covariance ensures $v_L + v_H - 2 \operatorname{Re} c \geq 0$, so $D \geq 0$. Setting the derivative to zero yields the interior solution; convexity gives endpoint clipping. If $D = 0$, the two stochastic errors have zero difference

variance and the bias bounds coincide; the risk is constant in w under the stated valid covariance. $\square$

If correlation is unknown, Cauchy–Schwarz gives $\text{Re } c \leq \sqrt{v_L v_H}$. The conservative choice uses this upper bound for a convex combination. If the bias is known as a complex value, one can first correct it, propagating the uncertainty in that correction. Unknown unbounded bias does not support a finite risk guarantee merely because the observed variance is small.




Figure 21: Bias-aware hybrid weights under different correlation assumptions. Prescribed bias/covariance functions, not an inferred real-data reliability estimator. [5]

22.3 Vector fields and joint fusion

For two unbiased estimates of a vector target, stack them as $z = Et + e$, with $E = [I \ I]^T$ and joint covariance $\Sigma_z \succ 0$. The generalized least-squares estimator is

$$\hat{t} = (E^* \Sigma_z^{-1} E)^{-1} E^* \Sigma_z^{-1} z. \quad (61)$$

Its derivation minimizes the whitened residual and uses the target’s common embedding in both observations. Singular covariance requires a consistent subspace treatment rather than a numerical inverse of redundant data. Bias bounds and model incompatibility also require more than the unbiased formula.

For an angular field, cross-direction covariance can matter. Applying a scalar blend independently at every angle may ignore shared source and geometry errors and produce discontinuous patterns. A practical implementation can begin with a conservative scalar or bandwise weight but should preserve shared reference and validity masks.

22.4 A decision map rather than a mandatory crossover

Let $\mathcal{R}_m(f)$ be a qualified target-risk bound for method m . Selection can be expressed as $m^*(f) = \arg \min_{m \in \mathcal{V}(f)} \mathcal{R}_m(f)$, where $\mathcal{V}(f)$ contains only methods whose assumptions are supported at that frequency. If that set is empty or every risk exceeds the declared target gate, the output is unsupported. This is different from forcing a continuous-looking curve through every frequency.

The closure hybrid implementation tests the algebra and invalid-input handling with specified covariance and bias functions [5,7]. It does not infer those functions from the actual DUT data. Establishing a deployable uncertainty estimator is therefore a remaining scientific task, not a hidden completed feature.

23. Target-aware sensor and experiment design

23.1 Optimize useful information, not raw signal level

A design variable ξ can include receiver positions, local sensor standoff, pose subset, reference-source positions, and scan elevation. It changes both source sensitivity and nuisance sensitivity. Design criteria should therefore use nuisance-profiled, noise-whitened information on the target-relevant state rather than unscaled singular values of a mixed-unit matrix.

D-optimality maximizes a log determinant, E-optimality increases the smallest relevant eigenvalue, and A-optimality reduces a trace of covariance. These are useful frameworks [12], but their coefficient metrics may differ from the actual acoustic objective. A target-weighted covariance or worst target-noise gain is often more directly interpretable for directivity or virtual-receiver reconstruction.

23.2 A monotonicity result with explicit assumptions

Proposition 16 — Marginal value of a calibrated independent row. Let $P \succ 0$ be a parameter covariance under fixed nuisance and prior assumptions. Adding an independent scalar observation with sensitivity row c , noise variance σ^2 , and no new unknowns gives

$$P' = P - Pc^*(\sigma^2 + cPc^*)^{-1}cP. \quad (62)$$

For a linear target B , the covariance reduction is positive semidefinite and its trace equals

$$\text{tr}\{B(P - P')B^*\} = \frac{\|BPc^*\|_2^2}{\sigma^2 + cPc^*}. \quad (63)$$

Proof. Apply the rank-one matrix inversion identity to $P^{-1} + c^*c/\sigma^2$. The subtracted matrix is an outer product divided by a positive scalar. Multiplication by B preserves positive semidefiniteness and taking its trace yields (63). $\square$

This monotonic improvement does not apply unchanged when a sensor introduces new unknown calibration or contamination parameters, or when the covariance is wrongly assumed independent. In those cases the joint likelihood and nuisance projection must be recomputed. It also does not guarantee that a greedy finite-candidate search finds the global best nonlinear hardware geometry.

23.3 Fixed layouts across frequency

A hardware layout must remain physically realizable across the operating band. Optimizing a different set of microphone positions independently at every frequency would provide an unattainable upper benchmark, not a product design. The closure study therefore evaluates fixed subsets of 18 candidate locations over eight frequencies from 40 to 1000 Hz [5]. Its greedy choices are provisional finite-grid designs under one restricted source class and assumed calibration.




Figure 22: One fixed sensor layout evaluated across frequency. Greedy finite candidate search with four complex source coefficients and assumed calibrated geometry. Sensor positions remain provisional. [5]

The worst conditional RMS over that grid decreases to approximately 1.78%, 1.53%, 1.40%, and 1.36% for one through four local observations. This diminishing return motivates a small experimental fixture. It does not establish that the actual source is adequately represented or that one or two permanent sensors are sufficient.

23.4 Pose count and source diversity

More poses can improve averaging and distinguish additional angular characters until the relevant subspace is saturated. More source states can improve reference-operator coverage only if they are genuinely independent and known. Repeatedly scaling one drive does not increase excitation rank in a linear regime. Separate driver excitation may help characterize a source basis, but changing electrical termination or crossover conditions can also change the operating system. The validation must retain that distinction.

An adaptive design can choose the next registered observation expected to reduce target uncertainty most. The choice remains a design prediction until checked with withheld data. It should not consume a reserved blind-validation set to tune the architecture.

24. Numerical evidence and publication rechecks

24.1 What was actually rerun

The FULL bundle contains the four primary research reports, handoffs and contracts, numerical workbooks, the earlier paper, and the full nested research archives. The 42 file entries in its supplied full SHA-256 manifest were checked. Decisive primary code, saved arrays, result tables, and mathematical specifications were inspected; this is not a claim that every line of every historical source archive was independently audited.

A separate copy of the closure source was used to replay six selected stages: the stationary controls, reference-source study, frequency/fusion study, real-data replay, regression suite, and saved-result verifier. All six returned exit code zero. The regression suite reports 36 passed methods and the

saved-result verifier 218 passing checks. The latter verifies supplied saved architecture/design results; this selected replay does not claim to rerun every original optimization stage.

The additional deterministic publication suite independently assembles the stationary operator and tests Wigner rotations, spherical translation, the radial Wronskian, transformed information, nuisance projection, real/complex Fisher conventions, reference-source counterexamples, fusion minimization, gating, origin invariance, covariance monotonicity, and source loading. It contains 47 passing checks with no failures [7]. These are algebraic and numerical publication checks, not 47 physical experiments and not proof of every statement in an unrestricted acoustic model.

24.2 Matched architecture comparison

The closure comparison keeps 58 real parameters in a common state representation, including source, reflection, geometry/calibration, and reference-session nuisance variables. It reports data-only ranks and target ambiguity separately from prior-conditioned uncertainty. At 1 kHz, the acoustic baseline has about 10.75% conditional target RMS; room-plane data alone about 10.64%; local registration about 2.26%; registered woofer/port observations about 1.47%; and combinations with one or three reference positions about 1.20% and 1.13%, respectively [5].




Figure 23: A matched 58-real-parameter architecture comparison. Local Jacobian, assumed priors and exact model. Prior-conditioned uncertainties are not measured reconstruction errors. [5]

Those values cannot be read as universal sensor rankings. The model already assumes source support and a finite reflection parameterization. Exact local likelihoods, shared calibration nuisances, and explicit priors are part of the comparison. The earlier 41-parameter result with two local observations gives approximately 1.78% under a different comparison; it should not be substituted for the 58-parameter registered woofer/port value.

The main inference is narrower and useful: under a compact source/room representation, improving local registration and obtaining a little calibrated source information can provide more target benefit than refining room offsets alone. Reference calibration can add conditional information, but its marginal benefit depends on source-space transfer and nuisance calibration.

24.3 Existing real-data evidence

The replay uses the existing $6 \times 72 \times 9$ development cache, not newly acquired signals. Over approximately 500–1500 Hz, the 54 receiver/frequency cases give about 1.47% median and 2.33% 90th-percentile error for held-angle prediction of total pressure [3,5]. The held angle is a prediction exercise within previously used development data, not a newly reserved blind test.




Figure 24: Existing real-data angular prediction. Previously used $6 \times 72 \times 9$ development cache; not new measurements, low-frequency validation, or component truth. [5]

This demonstrates useful angular structure at the observed receiver configurations. It does not establish direct/reflected decomposition, predict arbitrary new receiver locations, validate local low-frequency sensing, or fix the real gating transition. Earlier held-receiver tests were much less accurate and remain part of the primary evidence [3]. Their failures are not narrated as a long development diary here, but their consequence is retained: successful angular interpolation is not free-field component certification.

24.4 Reading numerical figures correctly

Every numerical figure in this monograph includes its model scope and exact data source in the accompanying figure-data package. Analytic uncertainty, deterministic stress quantiles, synthetic reconstruction error, and real measured prediction error are different objects. Prior-conditioned uncertainty is not a confidence interval unless its probabilistic assumptions and calibration are established. A finite sample maximum is not a worst-case theorem. A low-rank Jacobian is not a global nonlinear proof. These distinctions are part of the result, not optional footnotes.

25. A testable hybrid architecture and the remaining experiment

25.1 The minimal framework supported by the mathematics

The strongest current direction is a registered, target-aware hybrid. The existing six-configuration multi-pose data remain the main spatial acquisition. A small fixed or co-rotating local fixture can add low-frequency source information where the source model and target require it. Room geometry constrains surfaces and candidate propagation paths; tighter local registration establishes sensitive source/receiver transforms. Electrical or acoustic reference data are added only where they

constrain a specific remaining target direction. Gating is used where both source truncation and retained-room error are qualified.

This is not a prescription that every system must contain LiDAR, reference source, p-u, LDV, and multiple MEMS. Each extra modality must justify its nuisance-profiled marginal information. The simplest system that meets a measured target-risk budget is preferable to a theoretically richer system with uncalibrated extra degrees of freedom.

25.2 The next physical question

A small registered low-frequency source-signature experiment is more discriminating than another broad enumeration of hardware concepts. It should offer one and two local pressure channels, a retained far/reference channel, and a common complex electrical/time reference. For a system with a significant port or passive radiator, the second local channel tests that secondary state; otherwise it tests a non-collinear source location. Additional registered challenge observations can be withheld from fitting without being declared permanent product sensors.

The acceptance question is whether a low-dimensional source class predicts those withheld observations within measured geometry, calibration, noise, and discrepancy allowances. Failure directs the next measurement toward a specific missing mode or nuisance. Success qualifies that source class and experiment, not every loudspeaker. Manual freehand repositioning is not the proposed normal workflow.

A simple elevation-diverse room scan can be a separate experiment tied to the same datum. It should test geometry and reflection candidates, not identify wall impedance optically. A reference source is a further controlled branch if uncertainty remains in the relevant room subspace. Source substitution, loading, and reference strength require their own transfer tests.

25.3 Independent forward validation

MATLAB can independently verify the new operators, transforms, and conditioning conventions. COMSOL can provide controlled forward truth for canonical sources, simple reflection geometries, source loading, and fixture scattering. Neither has been newly executed for this monograph. An independent forward solver should use a mesh/domain/PML/tolerance audit and a source model sufficiently different from the inverse model to detect inverse crimes. The most important physical qualification ultimately needs independently known direct or free-field target data.

25.4 Open questions stated as decisions

The real low-frequency source class and permanent sensor count remain undetermined. The real inverse-to-gate transition remains unspecified. A general causal multi-frequency boundary inverse, global nonlinear robustness beyond the tested local models, and source-state transfer to same-drive anechoic operation remain open. Each is a concrete research question with a measurable consequence, not an excuse to declare the entire method impossible or complete.

The publication's purpose is to make those decisions sharply testable. It supports targeted mathematical closure and physical validation rather than replacing the existing framework with another unconstrained fit.

26. Conclusions

A pressure fit and a free-field target are different mathematical objects. For finite linear models, exact target identifiability is equivalent to the target being invariant on every measurement-null direction. The quotient-space and factorization formulations make this statement constructive; the

sharp data-to-target norm adds the stability requirement that rank alone lacks. For nonlinear and normalized targets, exact fiber invariance and local Jacobian analysis must be distinguished.

The stationary six-ring control contains an analytically unavoidable axisymmetric dimension deficit at degree three. Independent assembly attains rank 30 in the 32-column declared case, and the remaining two directions change outgoing targets. Correctly modelling the real rotating-DUT experiment changes the available information, but centered yaw leaves $m = 0$ unchanged and flexible room models can reintroduce source–room exchange. Off-axis motion provides potentially useful spatial modulation at the price of geometry sensitivity and translation-model requirements.

Independent local pressure, velocity, surface motion, geometry, and reference-source observations can constrain different missing directions. Their contribution must be evaluated after accounting for new nuisance variables and measurement covariance. A reference source identifies an unrestricted room operator only on its excited source subspace; source loading separates operational direct-field continuation from same-drive anechoic prediction. Neither a smooth frequency model nor repeated co-rotating samples automatically removes a shared coherent ambiguity.

The gating identity exposes both lost direct response and retained room response. A correlated, bias-aware fusion rule can combine qualified estimates of the same target, but its covariance and bias inputs must themselves be supported. The emerging architecture is consequently frequency-dependent and conditional: source assistance and geometric constraints where needed, qualified gating where justified, and an explicit unsupported status where the information does not meet the target requirement.

The mathematical framework is substantially more explicit than a residual-driven reconstruction procedure. The numerical rechecks support its declared finite-model statements. The next decisive evidence is physical: a registered small-source-signature experiment, measured uncertainty, and independent target qualification. V3.63 remains frozen, no final hardware is selected, and no real physical component or same-drive anechoic certification is asserted.

Appendix A. Rotation conventions and an independent construction

A.1 Orthonormal harmonics

The complex harmonics are orthonormal on the unit sphere:

$$\int_{\mathbb{S}^2} \overline{Y_{lm}(\Omega)} Y_{l'm'}(\Omega) d\Omega = \delta_{ll'} \delta_{mm'}. \quad (64)$$

A convention including the Condon–Shortley phase satisfies $Y_{l,-m} = (-1)^m \overline{Y_{lm}}$. Software interfaces differ in angle order and normalization, so a library name alone is insufficient documentation. The independent assembly uses polar θ and azimuth φ and records the (l, m) column order explicitly [7,8].

For the active rotation action used in Section 8,

$$Y_{lm'}(R^{-1}\Omega) = \sum_{m=-l}^l Y_{lm}(\Omega) \mathcal{D}_{mm'}^l(R). \quad (65)$$

This equation fixes the coefficient orientation. With row-vector evaluation $Y_l(\Omega)$, it reads $Y_l(R^{-1}\Omega) = Y_l(\Omega) \mathcal{D}^l(R)$. Confusing the evaluation transform with the coefficient transform can insert a transpose or conjugation error that is not visible in power-only plots.

A.2 Generator construction

Order the basis by $m = -l, \dots, l$. Define J_z diagonal with entries m , and the raising matrix by

$$(J_+)_{m+1,m} = \sqrt{l(l+1) - m(m+1)}, \quad J_- = J_+^*, \quad J_y = \frac{J_+ - J_-}{2i}. \quad (66)$$

For Euler angles with the active composition $R = R_z(\alpha)R_y(\beta)R_z(\gamma)$,

$$\mathcal{D}^l(R) = e^{-i\alpha J_z} e^{-i\beta J_y} e^{-i\gamma J_z}. \quad (67)$$

The generators are Hermitian in this convention, so each factor is unitary. The z factors immediately recover the characters in (24). This construction avoids inserting a long special-function formula without checking its convention.

The independent numerical test uses four nonsymmetric unit directions and a nontrivial three-angle rotation. It checks both $D^*D = I$ and the evaluation identity for degrees zero through four. Passing unitarity alone would be insufficient: an incorrect sign or inverse rotation can still yield a unitary matrix. The independent spatial evaluation is what resolves that ambiguity.

A.3 Perturbations of pose

For a small yaw error $\delta\alpha$, a single angular character changes by

$$e^{-im(\alpha+\delta\alpha)} - e^{-im\alpha} = -ime^{-im\alpha}\delta\alpha + O(\delta\alpha^2). \quad (68)$$

Thus high- $|m|$ content is more sensitive to angle error. The $m = 0$ component has no first-order or exact yaw modulation, but can still vary through off-axis translation or a changing room coupling. A pose error therefore affects both orientation and, for an eccentric source, the physical source location. These contributions must be differentiated together when propagating an actual turntable error.

Appendix B. Green-function derivatives and independent velocity

B.1 Gradient and Hessian without nested numerical differentiation

Put $r = |\mathbf{x} - \mathbf{s}|$, $e = (\mathbf{x} - \mathbf{s})/r$, and $g(r) = e^{-ikr}/(4\pi r)$. Direct differentiation gives

$$g' = (-ik - r^{-1})g, \quad g'' = (-k^2 + 2ik/r + 2/r^2)g. \quad (69)$$

The receiver-coordinate gradient and Hessian are

$$\nabla_{\mathbf{x}} G_0 = g'e, \quad \nabla_{\mathbf{x}}^2 G_0 = (g'' - g'/r)ee^T + (g'/r)I. \quad (70)$$

The source-coordinate gradient has the opposite sign; its Hessian is the same radial Hessian because two source derivatives remove that sign. These identities provide a direct implementation of pressure/velocity sensitivities and geometry Jacobians. Finite differences remain useful as an independent check with step-size studies, not as the only definition of a nested derivative.

This mattered in the extended programme: nested numerical differentiation initially made an ideal velocity channel appear to remove a vertical-datum gauge. Analytic derivatives restored the correct local rank [4]. The erroneous intermediate result is not counted as information gain. A sensitivity implementation must preserve known invariances, not merely return finite numbers.

B.2 Inverting the radial two-family system

Writing $h = h_l^{(2)}(kr)$, $j = j_l(kr)$ and $d_r = \partial_r p_{lm}$, Proposition 9 gives the explicit inverse

$$\begin{bmatrix} a_{lm} \\ b_{lm} \end{bmatrix} = \frac{1}{ik/(kr)^2} \begin{bmatrix} kj_l' & -j_l \\ -kh_l^{(2)'} & h_l^{(2)} \end{bmatrix} \begin{bmatrix} p_{lm} \\ d_r \end{bmatrix}. \quad (71)$$

The radial functions and derivatives in this expression are evaluated at kr . The nonzero determinant ensures algebraic invertibility, but the inverse entries can be large relative to sensor noise. Pressure and radial derivative also have different units. Dividing the derivative row by k supplies one dimensionless algebraic comparison; physical performance needs the actual velocity-channel covariance and the factor $-i\omega\rho_0$ from (5).




Figure 25: Independent radial derivative: nonzero determinant, varying conditioning. Rows are pressure and derivative divided by k . Column normalization is a declared algebraic metric, not a physical sensor-noise model. [7]

The plotted radial matrix condition is therefore not a prediction of a commercial p-u probe's accuracy. It illustrates how a formal Wronskian result and a practical conditioning problem coexist. The independent publication test evaluates the Wronskian for degrees zero through six at several positive arguments. It is a numerical check of a cited analytic identity, not its proof.

B.3 Complete boundary information versus sparse points

The boundary representation (41) involves traces over a surface. A finite sensor set is equivalent only after specifying a finite model or a sampling theorem. Inferring complete boundary data from a few pressure/velocity points is itself an inverse problem. Ideal full-surface Cauchy information can separate radial families per harmonic, while a local probe may leave angular modes or global source/room nuisance unresolved.

Similarly, an LDV supplies a measured line-of-sight component of surface velocity. Converting it to normal velocity requires geometry and directional information. A single-component optical observation is not automatically the complete boundary normal velocity needed by the acoustic model.

Appendix C. Spherical translation, convergence, and truncation

C.1 Radial Green-function jump

Expand $G_0(\mathbf{x}, \mathbf{d})$ in $Y_{lm}(\hat{\mathbf{x}})$. Away from $r = d$, each radial coefficient obeys

$$\frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{du_l}{dr} \right) + \left(k^2 - \frac{l(l+1)}{r^2} \right) u_l = 0. \quad (72)$$

Regularity at zero selects $j_l(kr)$ on the inner side. Outgoing radiation selects $h_l^{(2)}(kr)$ on the outer side. Integrating the distributional equation across $r = d$ fixes the derivative jump. The Wronskian $j_l h_l^{(2)'} - j_l' h_l^{(2)} = -i/(kd)^2$ then gives the coefficient $-ik$ multiplying $j_l(kr_<)h_l^{(2)}(kr_>)$. This yields (26) after angular projection.

The expansion domain is not optional. An observation inside the translated source radius requires the other $r_</r_>$ ordering; a general extended source requires an enclosing radius appropriate to its entire support. A local regular-wave room expansion likewise has a convergence domain determined by external singularities and boundaries. Continuing outside that domain can invalidate a model even if its finite matrix is well-conditioned.

C.2 Truncation and practical checks

For $d/r < 1$, increasing the expansion order improves the represented translated monopole in the tested exterior region. The order needed depends on kd , observation distance, target tolerance, and near-singular configurations. A small kd suggests that low orders may dominate, but is not a universal error certificate for a distributed radiator or an inverse continuation.

The independent check uses $f = 400$ Hz, $\mathbf{d} = (0.12, 0.03, -0.02)$ m and four exterior observations at radius 2 m. Order 16 gives a relative error below the predeclared 10^{-11} check tolerance [7]. Figure 8 records every truncation order; no smooth extrapolation is substituted for the computed values.

A robust implementation can compare orders L , $L + \Delta L$, and $L + 2\Delta L$ for the target of interest, rather than choose one order solely from a visual field plot. The comparison should be performed on the continued target as well as the measured trace. An inverse operator can amplify a truncation discrepancy that is small on the acquisition surface.

C.3 Center choice as model gauge or physical parameter

Changing the expansion origin and transforming all coefficients exactly is a reparameterization of the same field. Estimating a frequency-dependent center while also allowing an unrestricted coefficient expansion can therefore introduce redundant degrees of freedom. In a truncated model, the equivalence is only approximate and center optimization may trade against omitted modes.

Consequently, a fitted “acoustic center” can absorb source-model discrepancy. Its apparent frequency variation is not always a literal moving radiator. A mechanical reference plus a declared modal basis is preferable when it supports the intended target, while any effective phase-center fit should carry a frequency, direction, and model qualification.

Appendix D. Efficient nuisance elimination and target metrics

D.1 Why whitening precedes Euclidean projection

For $z = Cq + D\eta + e$ with covariance $\Sigma \succ 0$, the weighted residual is $\|\Sigma^{-1/2}(z - Cq - D\eta)\|^2$. A projector orthogonal in the unweighted Euclidean coordinates need not remove the least-squares nuisance correctly. Whiten first or use the equivalent weighted projector. In whitened coordinates,

$$\hat{\eta}(q) = D_w^\dagger(z_w - C_w q), \quad \min_{\eta} \|z_w - C_w q - D_w \eta\|^2 = \|P_D(z_w - C_w q)\|^2. \quad (73)$$

This profiles deterministic nuisance rather than assigning it a probabilistic prior. If nuisance has a measured or assumed prior covariance, augment the residual with the corresponding prior rows and identify their origin separately. A bounded unknown discrepancy is not automatically equivalent to such a Gaussian prior.

D.2 Joint old/new correlation

For stacked errors with covariance blocks Σ_{yy} , Σ_{yz} and Σ_{zz} , the new conditional innovation has covariance

$$\Sigma_{z|y} = \Sigma_{zz} - \Sigma_{zy} \Sigma_{yy}^{-1} \Sigma_{yz}. \quad (74)$$

The innovation also changes the state sensitivity: $z - \Sigma_{zy} \Sigma_{yy}^{-1} y$ has operator $C - \Sigma_{zy} \Sigma_{yy}^{-1} A$, along with its nuisance block. This is one way to avoid treating two estimates derived from the same microphone record as independent. Singular blocks require a support-consistent generalized formulation; blindly inverting a nearly singular covariance can itself create spurious information.

D.3 Nullspace leakage and target ambiguity

A useful normalized numerical diagnostic is $\|HN\|_F / (\|H\|_2 \|N\|_F)$, accompanied by $\|B_S N\|$ and the rank or singular spectrum of $B_S N$. The first quantity checks the null witness. The second describes its target relevance. Basis changes within the nullspace alter individual vectors and plots but not the null subspace or the existence of target-changing directions.

An exact target-nullspace dimension can be expressed as $\text{rank}(B_S N)$ for a full basis of $\ker H$. A near-nullspace depends on a declared threshold or uncertainty scale. Reporting its basis without the parameter metric, whitening, ordering, and threshold is insufficient for a later sensor comparison.

D.4 Reconstructible target factors

When the target is identifiable, $B_S H^\dagger$ is a directly useful object. It predicts noise amplification without requiring all coefficients to be unique. A generalized singular-value analysis of (B_S, H) can characterize the same ratio in (10), but zero-denominator target-changing directions must be retained as infinite sensitivity rather than numerically discarded.

For a practical discrete design, one can separately store the full data-only null basis, a near-null basis at the declared threshold, and a target-weighted covariance under explicit priors. These are complementary objects. Combining them into a single undocumented “confidence score” obscures the difference between missing information, weak information, and prior-supported estimates.

Appendix E. A fully worked two-channel source–room example

E.1 Exact recovery with a known differential signature

Consider a single-frequency model with source amplitude a , room amplitude b , primary pressure $y = a + b$, and a calibrated local observation $z = \alpha a + \beta b$. The two-row matrix and determinant are

$$\begin{bmatrix} y \\ z \end{bmatrix} = \begin{bmatrix} 1 & 1 \\ \alpha & \beta \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix}, \quad \det A = \beta - \alpha. \quad (75)$$

If $\alpha \neq \beta$, the source is

$$a = \frac{\beta y - z}{\beta - \alpha}, \quad b = \frac{z - \alpha y}{\beta - \alpha}. \quad (76)$$

If $\alpha = \beta$, the new channel is merely a multiple of the old mixture and adds no source–room discrimination. High local signal level alone is therefore not the correct design objective; the difference between source and room sensitivities matters.

With independent measurement errors of variances σ_y^2 and σ_z^2 , the source variance is

$$\text{Var}(\hat{a} - a) = \frac{|\beta|^2 \sigma_y^2 + \sigma_z^2}{|\beta - \alpha|^2}. \quad (77)$$

For correlated errors $c_{yz} = \mathbb{E}[e_y \bar{e}_z]$, the numerator becomes $|\beta|^2 \sigma_y^2 + \sigma_z^2 - 2 \text{Re}(\beta c_{yz})$. Thus a matrix can be invertible and still amplify errors strongly when α and β are close. This is the smallest possible example of target-relevant conditioning.

E.2 Calibration sensitivity

Differentiating (76) at fixed measured y, z gives

$$\frac{\partial a}{\partial \alpha} = \frac{a}{\beta - \alpha}, \quad \frac{\partial a}{\partial \beta} = \frac{b}{\beta - \alpha}. \quad (78)$$

These derivatives refer to changing the assumed calibration coefficients while holding the observations fixed. If the true forward coefficients change instead while the inverse remains fixed, the corresponding perturbation enters with the appropriate opposite forward-model sign. Either interpretation exposes the same denominator. A small residual can therefore coexist with a calibration-driven source error when the differential signature is weak.

If α and β are both treated as unrestricted unknowns, two complex observations no longer determine the state. Their calibration or physical constraint is part of the information. Adding an arbitrary local contamination c to z restores the example of Section 6 and can remove the apparent benefit entirely.

E.3 A restricted model can still be useful

This toy system is not asserted to be the actual loudspeaker/room model. It illustrates the logical structure of a useful restricted method: validate the source and room signatures, measure their calibration uncertainty, verify that their difference is resolvable, and evaluate a target-specific error bound. A second local sensor is justified when it separates an actual remaining state direction or constrains a nuisance—not merely because “two is better than one.”

For a larger source class, replace a, b by vectors and the scalar denominator by the relevant singular spectrum of the nuisance-projected operator. The same reasoning then underlies the formal framework, local sensor tests, and architecture comparisons in the main text.

Appendix F. Nonlinear constraints, fibers, and local geometry

F.1 Exact recovery lives on observation fibers

The linear kernel criterion is unusually simple because every fiber of a linear observation map is an affine translate of the same vector space. In a nonlinear model this is no longer true. Let $\mathcal{M}$

be a smooth admissible-state manifold, $F : \mathcal{M} \rightarrow \mathbb{R}^m$ the observation map, and $\mathcal{T} : \mathcal{M} \rightarrow \mathbb{R}^p$ the requested target. Complex data can be real-stacked before this discussion. The exact criterion is

$$F(q_1) = F(q_2) \implies \mathcal{T}(q_1) = \mathcal{T}(q_2), \quad q_1, q_2 \in \mathcal{M}. \quad (79)$$

This condition is global over the declared admissible set. Its meaning is unchanged by reparameterization: all physical states consistent with the same observation must give the same target. A full-rank numerical Jacobian at one fitted state is neither a proof of this statement nor a test of distant competing solutions. Conversely, a vanishing derivative at an isolated singular point need not mean that the nonlinear map is many-to-one. The functions $x \mapsto x^2$ and $x \mapsto x^3$ illustrate these two different cautions.

Proposition 17 — Constancy on connected regular fibers. Suppose F has constant rank on a neighborhood $U \subset \mathcal{M}$, its fibers in U are connected, and $\mathcal{T}$ is continuously differentiable. Then $\mathcal{T}$ is constant on every such fiber if its differential annihilates the fiber tangent space at every point in U :

$$\ker dF_q \subseteq \ker d\mathcal{T}_q, \quad q \in U. \quad (80)$$

Proof. The constant-rank theorem makes each fiber an embedded submanifold with tangent space $\ker dF_q$. A connected manifold is path connected by piecewise smooth paths. For a path γ in one fiber, differentiation gives $(\mathcal{T} \circ \gamma)' = d\mathcal{T}_\gamma \gamma' = 0$. Integrating on each path segment shows that the target has the same value at its endpoints. The converse holds along all tangent paths if the target is constant on the fiber. $\square$

The quantifier “at every point” and the connectedness assumption matter. For $F(x) = x^2$, each nonzero point has a nonsingular scalar derivative, but the positive and negative branches are separate members of a global two-point fiber. This proposition is therefore a local regular-fiber result, not an unrestricted uniqueness theorem for the nonlinear acoustic room inverse problem.

F.2 Constraints change the relevant tangent directions

Suppose admissible states are locally described by $q = \psi(\zeta)$, with tangent matrix $R = d\psi_{\zeta_0}$. Linearized observations and target use the same tangent restriction:

$$\delta y = J_F R \delta \zeta, \quad \delta t = J_{\mathcal{T}} R \delta \zeta. \quad (81)$$

A known source support, calibrated reflection law, or restricted diaphragm basis can reduce the tangent dimension. That reduction can make a previously ambiguous target identifiable inside the restricted model. It must not be reported as newly acquired data. If the physical source leaves that manifold, its discrepancy enters the error budget even when the reduced Jacobian is perfectly conditioned.

For normalized complex directivity, division by a reference pressure also makes the target locally singular near a reference null. The correct response is not to hide the singularity with an arbitrary numerical floor. It is to choose a different qualified reference or report that this normalization is unreliable over that band. A small absolute pressure error can correspond to a very large ratio error.

F.3 Remainders behind a Jacobian uncertainty estimate

A local covariance result assumes that the nonlinear remainder is negligible over the uncertainty region. If J_F is Lipschitz with constant L_F along the segment from q_0 to $q_0 + \delta q$, integration of the

derivative gives

$$\|F(q_0 + \delta q) - F(q_0) - J_F(q_0)\delta q\| \leq \frac{L_F}{2} \|\delta q\|^2. \quad (82)$$

The same construction applies to the target. These constants are not established by one Jacobian evaluation. They must be bounded analytically, sampled through a validated nonlinear model, or treated as unresolved. In this project, the large manual-geometry perturbations of the earlier campaign cannot be declared harmless merely because a tightly registered local linear model has small covariance. The perturbation scale and the range over which the linearization is accurate are part of the scientific claim.

Appendix G. Correlated measurement blocks and calibration floors

G.1 Orthogonalizing new measurements without double counting

Let old observations be $y = Aq + e_y$ and new observations be $z = Cq + D\eta + e_z$. Assume their zero-mean errors have known positive-definite block covariance, with $\Sigma_{yz} = \Sigma_{zy}^*$. Define the innovation

$$z_{\text{in}} = z - \Sigma_{zy}\Sigma_{yy}^{-1}y = (C - \Sigma_{zy}\Sigma_{yy}^{-1}A)q + D\eta + e_{\text{in}}. \quad (83)$$

The innovation covariance and old/new cross covariance are

$$\Sigma_{\text{in}} = \Sigma_{zz} - \Sigma_{zy}\Sigma_{yy}^{-1}\Sigma_{yz}, \quad \text{Cov}(e_{\text{in}}, e_y) = 0. \quad (84)$$

Proposition 18 — Correlated augmentation has an equivalent innovation experiment.

The observation pair (y, z) and the pair (y, z_{in}) contain the same data because their block-triangular transform is invertible. Under the stated covariance, the latter admits separate whitening of the old block and the innovation block. On an old null direction $N\xi$, the innovation source operator restricts to $CN\xi$ because $AN = 0$.

Proof. The inverse map is $z = z_{\text{in}} + \Sigma_{zy}\Sigma_{yy}^{-1}y$. Expanding the covariance gives the Schur complement and zero cross covariance. Substitution of $AN = 0$ proves the final assertion. The target test then profiles D after innovation whitening, rather than treating the old and new channels as independent before the transform. $\square$

This result separates exact geometric coupling from statistical gain. Correlation does not change CN on an exact old nullspace, but it changes the effective noise and therefore the practical value of each newly observed direction. Treating two strongly correlated estimates as independent can make the reported covariance much too small without altering the apparent matrix rank.

G.2 Profiling and marginalizing are different operations

An unrestricted nuisance is removed by a projector or a likelihood minimization. A random nuisance with known covariance is instead marginalized into the observation covariance. For η independent of measurement noise and with covariance Σ_η ,

$$\Sigma_{z,\text{marg}} = \Sigma_e + D\Sigma_\eta D^*. \quad (85)$$

The first treatment states that the nuisance can take any value; the second supplies a probabilistic constraint. In the limit of very large nuisance variance, information in its range tends to be lost. With finite variance, a posterior estimate can remain finite even though the data alone have an ambiguity. Both models can be useful, but they answer different questions. The architecture ledgers

must state which was used for channel calibration, local room contamination, and reference-session gains.

G.3 A common error does not average away

Let a repeated scalar observation have independent variance σ_e^2 and a shared error with variance σ_c^2 . For an arithmetic mean of N repetitions,

$$\text{Var}(\bar{z}) = \frac{\sigma_e^2}{N} + \sigma_c^2. \quad (86)$$

This follows by summing all N^2 covariance terms: only N independent-noise terms survive, whereas the shared component appears in every pair. A deterministic common bias adds another error term rather than a variance that decreases with N . Thus a co-rotating sensor can improve repeatability and expose pose-dependent contamination while still retaining a floor from shared calibration or a pose-invariant reflected contribution. More poses cannot replace a missing calibration anchor.

Appendix H. Generalized target spectra and bounded-error inference

H.1 The generalized Rayleigh quotient

For a whitened matrix H and target B_S , assume that B_S annihilates $\ker H$. Restrict to the complement of that kernel. Squared target amplification is a generalized Rayleigh quotient:

$$\gamma^2 = \sup_{v: Hv \neq 0} \frac{v^* B_S^* B_S v}{v^* H^* H v}. \quad (87)$$

Stationary directions satisfy the generalized eigenvalue equation

$$B_S^* B_S v = \lambda H^* H v. \quad (88)$$

On the identifiable quotient, the largest eigenvalue is γ^2 . Equivalently, write $H = U_r \Sigma_r V_r^*$ and compute the largest singular value of $B_S V_r \Sigma_r^{-1}$. This is usually more stable than forming normal equations, which square the condition number. A generalized SVD gives another representation of the same comparison between measurement and target norms.

The distinction from $\kappa(H)$ is important. A nearly invisible parameter direction is not necessarily harmful if the target is also nearly insensitive to it. Conversely, a modest-looking coefficient residual can be unacceptable for a target that amplifies precisely that direction. The coefficient metric, whitening covariance, and target norm must therefore be attached to every plotted spectrum.

H.2 A finite-dimensional minimax obstruction

Let $\mathcal{Q}$ be an admissible state set, which may include explicit physical bounds. Define its target modulus under observation tolerance δ by

$$\omega_{\mathcal{Q}}(\delta) = \sup_{q_1, q_2 \in \mathcal{Q}: \|A(q_1 - q_2)\| \leq \delta} \|B(q_1 - q_2)\|. \quad (89)$$

Proposition 19 — Bounded-noise lower bound. If the observation error can have any value of norm at most ε , no estimator can have a uniform worst-case target error below $\omega_{\mathcal{Q}}(2\varepsilon)/2$.

Proof. Choose two admissible states whose observations differ by at most 2ε . Their midpoint observation is consistent with either state under an error of norm at most ε . The estimator must return the same value for that observation. By the triangle inequality, its error for at least one state is at least half the distance between the two targets. Taking the supremum, or a sequence approaching it when not attained, gives the bound. $\square$

For an unrestricted linear state space with a target-changing null direction, this modulus is infinite even at zero observation tolerance. A bounded source-energy class can make it finite, but then the bound derives partly from the physical restriction. This provides a precise way to distinguish a practical restricted Normal Model from a claim of universal separation.

H.3 Combining data, operator, and target uncertainty

Let a nominal estimator satisfy $KH = B_S$, while the true whitened forward matrix and target matrix are $H + \Delta H$ and $B_S + \Delta B$. For $y_w = (H + \Delta H)\xi + d + e$, its target error is exactly

$$Ky_w - (B_S + \Delta B)\xi = K(e + d) + (K\Delta H - \Delta B)\xi. \quad (90)$$

Here d is a model discrepancy and e measurement error. If $\|\xi\| \leq M$, the induced-norm bound is

$$\|\delta t\| \leq \|K\|(\|e\| + \|d\| + M\|\Delta H\|) + M\|\Delta B\|. \quad (91)$$

The factor M is essential. A fixed coordinate uncertainty does not imply a fixed output error independently of the strength and spatial complexity of the source. Near-field gradients, strong cancellation between sources, or a target normalization near a null can enlarge the effective bound. Conversely, a carefully restricted low-order source with stable reference geometry can have a much smaller uncertainty budget. This is why a quoted sub-millimetre coordinate requirement cannot be transferred uncritically between all source classes and targets.

Appendix I. Reference-source stencils, derivative noise, and transfer range

I.1 What a nearby source stencil can identify

For a fixed receiver and frequency, write the room transfer as a function of source position $g(\mathbf{s})$. A first-order local representation around $\mathbf{s}_0$ uses four coefficients: one value and three derivatives. For displacements $\delta\mathbf{s}_j$, the calibration design has rows

$$X_j = [1, \delta s_{j,x}, \delta s_{j,y}, \delta s_{j,z}], \quad y_j = X_j\beta + r_j + e_j. \quad (92)$$

The rank of X determines which of these four local features can be estimated. Repeating the origin changes noise averaging but not its rank. A vertical stencil measures only a value and vertical derivative; a horizontal stencil misses a derivative normal to its plane. Four affinely independent points span the first-order features. These statements concern this Taylor model, not the number of source states needed to identify an arbitrary spatial Green operator.

I.2 Bias and variance compete as spacing changes

For a scalar directional function with three bounded derivatives, central differencing gives

$$\hat{g}'(0) = \frac{g(h) + e_+ - g(-h) - e_-}{2h}, \quad |\text{bias } \hat{g}'| \leq \frac{M_3 h^2}{6}. \quad (93)$$

The derivative is taken along one specified direction. If the two errors are zero mean and independent with equal complex variance σ^2 , their contribution to derivative variance is $\sigma^2/(2h^2)$. Smaller spacing reduces truncation bias but magnifies pressure/reference noise. Source-position error, unequal source strengths, and common gain errors add other terms and generally change the optimum.

Proposition 20 — Conditional central-stencil optimum. Under only the preceding bounded-third-derivative and independent equal-noise assumptions, a conservative mean-square derivative error is bounded by

$$R(h) = \frac{M_3^2 h^4}{36} + \frac{\sigma^2}{2h^2}, \quad h_* = \left(\frac{3\sigma}{M_3} \right)^{1/3}. \quad (94)$$

Proof. The squared bias bound plus the variance gives R . For $M_3, \sigma > 0$, setting its derivative to zero yields $h^6 = 9\sigma^2/M_3^2$. The risk diverges as h tends to zero or infinity within this algebraic model, and its positive stationary point is the unique minimum. The bound must only be used within the neighborhood where M_3 is valid. If the feasible neighborhood ends earlier, the constrained optimum lies at its boundary. $\square$

This is not a recommended physical spacing. The derivative bound, measurement covariance, source-position errors, and admissible local expansion region have not been measured for the actual room. The result explains the competing mechanisms behind the stored reference-transfer sweep instead of treating its best grid point as a universal optimum.

I.3 Predicting a finite DUT neighborhood

For a first-order Green expansion, a Hessian bound M_2 gives the pointwise transfer remainder

$$|g(\mathbf{s}_0 + \delta\mathbf{s}) - g(\mathbf{s}_0) - \nabla g(\mathbf{s}_0) \cdot \delta\mathbf{s}| \leq \frac{M_2}{2} \|\delta\mathbf{s}\|^2. \quad (95)$$

Increasing calibration stencil coverage does not by itself make the first-order model accurate over an extended woofer/port system. The target neighborhood and its remainder must be checked. A richer but physically constrained room model, more source positions, or a higher-order local expansion may improve transfer, but each introduces additional parameters or assumptions. The reference-source subspace theorem and the target-modulus bound remain useful checks on such extensions.

Appendix J. Vector fusion, conditional coverage, and gate bounds

J.1 Minimum covariance among unbiased linear fusion rules

Let two aligned estimates be stacked into $x = Et + e$, where $E = [I \ I]^\top$ and the joint error covariance Σ_x is positive definite. The generalized least-squares fusion matrix is

$$K_* = (E^* \Sigma_x^{-1} E)^{-1} E^* \Sigma_x^{-1}, \quad K_* E = I. \quad (96)$$

Proposition 21 — Covariance-optimal vector fusion. For any other unbiased linear rule K satisfying $KE = I$, put $\Delta = K - K_*$. Its covariance exceeds that of K_* in the positive-semidefinite order:

$$K \Sigma_x K^* = K_* \Sigma_x K_*^* + \Delta \Sigma_x \Delta^* \succeq K_* \Sigma_x K_*^*. \quad (97)$$

Proof. Since $\Delta E = 0$ and $\Sigma_x K_*^* = E(E^* \Sigma_x^{-1} E)^{-1}$, both cross terms in the covariance expansion vanish. The remaining quadratic term is positive semidefinite. $\square$

This proof does not require independent errors between the estimates. It does require a valid joint covariance and a common target. Two reconstructions from the same pressure data are generally correlated. If the covariance is singular because one estimate is an exact transform of the other, fusion must be formulated on the supported observation subspace rather than blindly inverting the matrix or assuming another independent measurement.

J.2 Vector bias is not automatically removed by covariance optimality

Write a general two-estimate fusion as $\hat{t} = W\hat{t}_L + (I - W)\hat{t}_H$. For Euclidean bias bounds $\|b_L\| \leq \beta_L$, $\|b_H\| \leq \beta_H$, a conservative bound is

$$\|Wb_L + (I - W)b_H\| \leq \|W\|\beta_L + \|I - W\|\beta_H. \quad (98)$$

Unlike the scalar convex case, this upper bound need not be attained by one pair of bias vectors because the maximizing output directions of the two matrices may differ. A vector minimax design must respect that distinction. The manuscript's scalar formula is exact for its specified real convex weight and complex disk-shaped bias sets; it should not be advertised as the unrestricted matrix minimax solution.

J.3 Conditional probability coverage is a model statement

For an unbiased real-stacked Gaussian target error with positive-definite covariance P_t and dimension d , the standard ellipsoid obeys

$$\Pr\left((\hat{t} - t)^\top P_t^{-1}(\hat{t} - t) \leq \chi_{d,1-\alpha}^2\right) = 1 - \alpha. \quad (99)$$

This identity is conditional on the distributional model, covariance, estimator, and zero-bias assumptions. A deterministic trigonometric stress ensemble does not establish Gaussian coverage. A misspecified source basis or an unmodelled reflected component can shift the estimator outside this ellipsoid even when repeatability is excellent. Known bias regions may enlarge a confidence set by a set sum; unknown bias cannot be converted into a finite confidence guarantee merely by plotting a narrow posterior band.

J.4 A deterministic gate bound and the information it needs

The Fourier triangle inequality yields a uniform absolute-frequency bound on the two gating-error contributions:

$$\sup_{\omega} |H_{\text{gate}} - H_{\text{direct}}| \leq \|(w - 1)h_{\text{direct}}\|_{L^1} + \|wh_{\text{room}}\|_{L^1}. \quad (100)$$

The inequality is elementary but operationally revealing. Both terms require knowledge or bounds on the components one is trying to separate. Agreement between two gates is not enough to measure them: both windows can omit the same direct ringing or retain the same early reflection. Similarly, agreement between a gated and an inverse estimate does not rule out shared bias. A deployable hybrid system needs independent qualification, a validated physical model, or conservative external bounds for those missing quantities. The current numerical examples demonstrate the mathematics of fusion, not a completed real-data error sensor.

Appendix K. Reproducibility, claim status, and reading guide

K.1 Provenance and numerical reconstruction

The supplied FULL-bundle manifest is verified at the archive-file level. The original nested archives remain unmodified. Selected closure scripts run in a separate copied directory, and the independent publication script has no hardware interface or pseudorandom seed generation. All 25 figures have exact CSV backing data and a registry recording source file, model scope, and interpretation. The companion workbooks expose those numeric arrays and a claim ledger.

Publication rechecks include the explicit stationary matrix, coefficient metric, null basis, singular values, Fourier-block ranks, Wigner evaluation residuals, and translation convergence. A normalized null residual at floating-point scale demonstrates numerical consistency of that finite witness. It does not replace an analytic proof for arbitrary frequency, geometry, or physical model. The analytic rank bound and its numerical attainment are deliberately separate claims.

The source package contains the manuscript source, figure generator, deterministic mathematical checks, selected replay driver and logs, and compact numerical witness files. Large original research archives are referenced by their supplied names and hashes rather than redundantly copied into the publication ZIP. Reproducing the selected historical scripts requires the original FULL bundle; reproducing the new mathematical checks uses the included compact control witness and source.

K.2 Formal-result index

Theorem 1 establishes linear target identifiability and factorization. Theorem 2 gives the sharp finite-dimensional target-noise gain. Theorem 3 gives target recovery after augmentation with unrestricted nuisance. Theorem 4 states the calibrated subspace of an unrestricted reference-source room operator. Corollaries 1 and 2 provide the two-state impossibility bound and the stationary degree-three rank upper bound.

Propositions 1–21 address penalized uniqueness, transformed data, Fourier blocks, harmonic rotation, translated monopoles, source–room intersections, normalized target invariances, odd-mode invisibility, radial Cauchy data, plane-offset phase sensitivity, room-mode dimension sensitivity, source loading, shared frequency gauges, gating, hybrid weighting, and marginal target information. Propositions 17–21 give regular-fiber constancy, correlated innovation equivalence, a bounded-noise lower bound, a conditional reference-stencil optimum, and covariance-optimal vector fusion. These are finite-model or convention-specific statements; their hypotheses are part of each result.

K.3 Evidence-status summary

Result or claim	Status in this manuscript	Boundary of the claim
Linear target kernel criterion	Proved	Linear maps and declared admissible state class
Sharp finite-dimensional target gain	Proved	Stated norms and observable target
Nuisance-projected augmentation	Proved	Unrestricted linear nuisance; covariance handled consistently
Stationary six-ring rank at most 30	Proved	Complete coaxial traces, $L = 3$, adequate azimuth sampling
Stationary rank equals 30 in the control	Numerically reverified	Supplied nominal geometry, 1 kHz and declared threshold
Two target-changing control directions	Numerically reverified	Saved outgoing continuation target
Centered yaw leaves $m = 0$ unmodulated	Proved	Stated active rotation and harmonic convention

Result or claim	Status in this manuscript	Boundary of the claim
Eccentricity improves practical recovery universally	Not established	Finite rank lifting does not imply robustness
Reference source identifies arbitrary room action	Disproved in unrestricted matrix class	Structural room priors can change the class
Local sensor count for the actual DUT	Open	No new real LF source-class data
Real inverse/gating transition	Unsupported	No independent direct waveform/reflection labels
Bias-aware scalar fusion formula	Proved and numerically checked	Valid covariance, stated bias bounds and common target
Historical real angular prediction	Replayed development evidence	Total pressure, approximately 500–1500 Hz
New independent MATLAB/COMSOL execution	Not run	Specifications and historical evidence are not new execution
Physical source-room / same-drive anechoic truth	Not certified	Independent component truth and loading qualification absent

K.4 Suggested use of the monograph

For the fundamental identifiability issue, read Sections 3–7 and the counterexamples. For multi-pose physics, read Sections 8–10 and Appendices A–C. For sensor architecture, read Sections 6 and 11–19. For the hybrid estimator, read Sections 20–23 and Appendix E. For the actual evidence boundary and next experiment, read Sections 24–26 and Appendix K. Appendices F–J deepen nonlinear fiber recovery, correlated measurements, bounded-error inference, reference-source stencils, and vector fusion. The notation and theorem index supplies equation numbers and symbol definitions for readers who wish to follow a proof independently of the engineering narrative.

References and source inventory

- [1] PHYSICAL LAB. *Frozen Acquisition Contract V2*. 14 September 2026. Supplied file `PHYSICAL_LAB_SFS_ACQUISITION_CONTRACT_V2_FROZEN_2026-09-14.json`. Primary project protocol; nominal geometry, sequential receiver configurations, unique poses, and declared uncertainty. No public URL asserted.
- [2] PHYSICAL LAB. *Sound-Field Separation: Scientific Due Diligence and Research Reconstruction*. Phase-I report and source evidence, September 2026. Supplied `PHYSICAL_LAB_SFS_PHASE1_SCIENTIFIC_REPORT.pdf` and historical primary archive. Stationary outgoing/regular control and target-changing nullspace. Primary project source, not an independent external publication.
- [3] PHYSICAL LAB. *Multi-Pose Observability and Near-Field Source-State Reconstruction*. Phase-II report and research-source archive, September 2026. Supplied `PHYSICAL_LAB_SFS_PHASE2_FINAL_REPORT.pdf` and `PHYSICAL_LAB_SFS_PHASE2_RESEARCH_SOURCE.zip`. Numerical JSON/NPZ records are identified per figure in the companion registry.
- [4] PHYSICAL LAB. *Extended Free-Field, Source-State, Room-Geometry and Hybrid Reconstruction Study*. September 2026. Supplied `PHYSICAL_LAB_SFS_EXTENDED_FREEFIELD_REPORT.pdf` and `PHYSICAL_LAB_SFS_EXTENDED_RESEARCH_SOURCE.zip`. Primary numerical evidence for the source-class, optical, co-rotating, modal-room, and gating studies.
- [5] PHYSICAL LAB. *Architecture Closure Research Campaign*. September 2026. Supplied `PHYSICAL_LAB_SFS_ARCHITECTURE_CLOSURE_REPORT.pdf`, research package, handoff, `docs/PARAMETERS_58.json`, `docs/CONFIG.json`, and associated matrices/results. Main source for the 58-real-parameter comparison, reference-source study, fixed sensor design, and existing real-data replay.

- [6] PHYSICAL LAB. *Hybrid Free-Field Reconstruction Paper*. Earlier 34-page manuscript, September 2026. Supplied PHYSICAL_LAB_SFS_HYBRID_FREEFIELD_RECONSTRUCTION_PAPER.pdf, editable DOCX, and publication figure data. Editorial baseline; not counted as independent validation of [2–5].
- [7] PHYSICAL LAB. *Mathematical Monograph Publication Rechecks*. 17 September 2026. Included monograph_checks.py, supplement_checks.py, COMBINED_MATHEMATICAL_AUDIT.json, INDEPENDENT_STATIONARY.npz, figure registry, replay logs, and claim ledger. Forty-seven deterministic mathematical checks; six selected closure replay stages. No physical acquisition, MATLAB, or COMSOL execution.
- [8] NIST. *Digital Library of Mathematical Functions*, §14.30, Spherical and Spheroidal Harmonics. Official mathematical definitions, orthogonality, symmetry, and addition relations. Accessed 17 September 2026. <https://dlmf.nist.gov/14.30>
- [9] NIST. *Digital Library of Mathematical Functions*, §10.50, Wronskians and Cross-Products; §10.60, Sums. Official spherical-Bessel identities and addition relations. Accessed 17 September 2026. <https://dlmf.nist.gov/10.50> and <https://dlmf.nist.gov/10.60>
- [10] Anderson, E., et al. *LAPACK Users' Guide*, Third Edition. SIAM, 1999. Online chapter: Error Bounds and Condition Numbers for the Singular Value Decomposition. Primary numerical-library documentation. Accessed 17 September 2026. <https://www.netlib.org/lapack/lug/node96.html>
- [11] Joint Committee for Guides in Metrology. *Evaluation of Measurement Data—Guide to the Expression of Uncertainty in Measurement*. JCGM 100:2008. Official BIPM edition, particularly covariance and the law of propagation of uncertainty. https://www.bipm.org/documents/20126/2071204/JCGM_100_2008_E.pdf
- [12] Joshi, S., and Boyd, S. “Sensor Selection via Convex Optimization.” *IEEE Transactions on Signal Processing* 57(2), 451–462 (2009). Primary author publication page and abstract inspected. https://web.stanford.edu/~boyd/papers/sensor_selection.html
- [13] Ma, F., Zhang, W., and Abhayapala, T. D. “Real-time separation of non-stationary sound fields on spheres.” arXiv:1901.05122 (2019). Primary author abstract and metadata inspected; pressure/radial-velocity sound-field formulation. <https://arxiv.org/abs/1901.05122>
- [14] Fernandez-Grande, E., Jacobsen, F., and Leclère, Q. “Sound field separation with sound pressure and particle velocity measurements.” *Journal of the Acoustical Society of America* 132(6), 3818–3825 (2012). DOI: 10.1121/1.4763988. Primary institutional abstract inspected; no claim to reproduce the complete article’s implementation. <https://doi.org/10.1121/1.4763988>
- [15] Allen, J. B., and Berkley, D. A. “Image method for efficiently simulating small-room acoustics.” *Journal of the Acoustical Society of America* 65(4), 943–950 (1979). Original article PDF inspected. DOI: 10.1121/1.382599. https://www.umiacs.umd.edu/~ramani/cmssc828d_audio/AllenBerkley79.pdf
- [16] Dokmanić, I., Parhizkar, R., Walther, A., Lu, Y. M., and Vetterli, M. “Acoustic echoes reveal room shape.” *Proceedings of the National Academy of Sciences* 110(30), 12186–12191 (2013). Primary article; room-echo geometry context, not validation of the present acquisition. <https://pmc.ncbi.nlm.nih.gov/articles/PMC3725100/>
- [17] Xu, Z., Herzog, A., Loder Meyer, A., Habets, E. A. P., and Prinn, A. G. “Simulating room transfer functions between transducers mounted on audio devices using a modified image source method.” arXiv:2309.03486 (2023). Primary author abstract and metadata inspected; transducer geometry/diffraction and transfer-model context. <https://arxiv.org/abs/2309.03486>
- [18] COMSOL. *Baffled Piston Radiation*, Application Gallery model 10137. Official canonical forward-model description. Accessed 17 September 2026; no COMSOL computation performed

for this manuscript. <https://www.comsol.com/model/baffled-piston-radiation-10137>

[19] Maynard, J. D., Williams, E. G., and Lee, Y. “Nearfield acoustic holography: I. Theory of generalized holography and the development of NAH.” *Journal of the Acoustical Society of America* 78(4), 1395–1413 (1985). DOI: 10.1121/1.392911. Primary bibliographic/abstract record checked; complete article not re-audited here. <https://doi.org/10.1121/1.392911>