Magnetic Fields from Measurements

J. R. Landers

November 2026

Abstract

How does a physical field learn from measurements, and what does smoothing add? From 500 indoor magnetic readings, a Gaussian map achieves held-out vector RMSE 8.09, compared with 5.22 for dipoles fitted to that map and 5.31 for dipoles fitted directly to readings, in reference recording units. Random forests and boosted stumps score 10.56 and 14.70. The physical form accounts for most of the observed gain. At fixed settings, the dipole field is exactly a weighted sum of the original readings; those weights expose substantial sensitivity changes hidden by similar prediction errors. In a simulated wire experiment, direct equation fitting outperforms the kernel-mediated route in all five regimes. The results separate the supplied representation, its empirical coefficients, and the contribution of intermediate smoothing.

1 From an experiment to a field

Switch on a current beside a compass and the needle turns. Maxwell's theory connected such observations through field equations [1]. Here we ask the inverse question: what do measurements teach a field model whose equations are already supplied? A robot's magnetic readings can fit a Gaussian map without a field law, or a sum of dipole responses that obeys magnetostatics by construction. Smoothing with the Gaussian before fitting dipoles raises two questions: how does the resulting field use each reading, and how much does the smoother contribute beyond the physical form?

We rewrite a background plus 100 fitted dipoles as an exact linear combination of 500 readings, then compare it with dipoles fitted directly to those readings under identical settings. Tree baselines broaden the predictive comparison. A simulated wire asks whether a learner can also select the field's formula from a candidate library; Appendix A checks recovery of a scalar differential law. In these tests, the physical form supplies most of the gain, while the measurement weights reveal what smoothing changes.

2 What the robot records

The Aalto robot mapped roughly $6\,\mathrm{m} \times 6\,\mathrm{m}$, recording time t in seconds, planar position $\mathbf{x}(t) = (x(t), y(t))$ in metres, and calibrated, orientation-corrected magnetic vectors $\mathbf{B} = (B_x, B_y, B_z)^\mathsf{T}$, with B_z vertical. Solin et al. used a scalar-potential Gaussian process [2]; later studies addressed sequential estimation [3] and large-scale local maps [4]. We predict all components from position alone; time only orders predictions along the route. Superscript $^\mathsf{T}$ denotes transpose.

2.1 A Gaussian map from the readings

At each of n planar training positions $\mathbf{x}_i$, the robot records a three-component vector $\mathbf{b}_i$. A Gaussian kernel of length scale $\ell > 0$ assigns any two positions a similarity $k_\ell(\mathbf{x}, \mathbf{x}') = \exp(-\|\mathbf{x} - \mathbf{x}'\|^2/(2\ell^2))$, where $\|\cdot\|$ is the Euclidean norm. At a query position $\mathbf{x}_*$, the kernel map adds one weighted bump per

training position to the sample mean $\overline{\mathbf{B}}$:

$$\widehat{\mathbf{B}}_K(\mathbf{x}_*) = \overline{\mathbf{B}} + \sum_{i=1}^{n} k_\ell(\mathbf{x}_*, \mathbf{x}_i)\boldsymbol{\alpha}_i. \tag{1}$$

Here $\boldsymbol{\alpha}_i$ sets the magnitude and direction of the bump centered at $\mathbf{x}_i$.

Because kernel ridge regression is linear in the measured vectors, every prediction in Eq. (1) is a possibly signed weighted sum of the raw readings. To make those weights explicit, stack the training readings into the $n \times 3$ matrix Y with rows $\mathbf{b}_i^\mathsf{T}$, let I_d denote the d-dimensional identity, and define the averaging row $m^\mathsf{T} = \mathbf{1}^\mathsf{T}/n$ and centering projector $P = \mathbf{1}m^\mathsf{T}$, where $\mathbf{1}$ is the n-vector of ones. Every row of PY equals the mean field $\overline{\mathbf{B}}^\mathsf{T} = m^\mathsf{T}Y$, leaving centered readings $Z = (I_n - P)Y$. Writing K for the $n \times n$ kernel matrix with entries $K_{ai} = k_\ell(\mathbf{x}_a, \mathbf{x}_i)$ and k_*^T for the query row with entries $k_\ell(\mathbf{x}_*, \mathbf{x}_i)$, kernel ridge regression with effective penalty $\rho > 0$ (including numerical jitter) minimizes $\|K\alpha - Z\|_F^2 + \rho\,\mathrm{tr}(\alpha^\mathsf{T}K\alpha)$, where $\|\cdot\|_F$ is the Frobenius norm and tr is the trace. Since K is positive semidefinite, solving $(K + \rho I_n)\alpha = Z$ yields the coefficient matrix $\alpha = (K + \rho I_n)^{-1}Z$ whose ith row is $\boldsymbol{\alpha}_i^\mathsf{T}$. Substituting α into $\widehat{\mathbf{B}}_K(\mathbf{x}_*)^\mathsf{T} = \overline{\mathbf{B}}^\mathsf{T} + k_*^\mathsf{T}\alpha$ gives both the query prediction and the matrix T of smoothed values at the training positions:

$$\widehat{\mathbf{B}}_K(\mathbf{x}_*)^\mathsf{T} = g_*^\mathsf{T}Y, \qquad g_*^\mathsf{T} = m^\mathsf{T} + k_*^\mathsf{T}(K + \rho I_n)^{-1}(I_n - P), \tag{2}$$

$$T = HY, \qquad H = P + K(K + \rho I_n)^{-1}(I_n - P). \tag{3}$$

Notice that g_*^T assigns reading i the same scalar weight g_{*i} across all three field components, so components never mix. Row a of the smoother matrix H gives those weights at training position $\mathbf{x}_a$. The resulting smoothed vectors in T, one at each measured position, serve as the targets for our two-stage dipole fit.

3 Giving the map a textbook form

In a static region of air with no electric currents and uniform permeability, Maxwell's magnetic equations require $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = 0$, where the derivatives are spatial. The local field is therefore the negative gradient of a harmonic scalar potential. Equivalent-source methods exploit this structure to build magnetic maps from scattered measurements, sometimes placing sources at multiple depths to capture different spatial scales [5, 6]. Here we place M equivalent dipoles on a plane below the observation surface $z = 0$. Dipole j sits at $\mathbf{s}_j$ with effective moment $\mathbf{q}_j$, which absorbs the factor $\mu_0/(4\pi)$ in SI units (μ_0 is the vacuum permeability). At observation point $\mathbf{r}$, write $\mathbf{R} = \mathbf{r} - \mathbf{s}_j$ and $R = \|\mathbf{R}\| > 0$. Using $\nabla R^{-3} = -3\mathbf{R}/R^5$, the negative gradient of the dipole potential gives

$$-\nabla\left(\frac{\mathbf{q}_j^\mathsf{T}\mathbf{R}}{R^3}\right) = \underbrace{\left(\frac{3\mathbf{R}\mathbf{R}^\mathsf{T}}{R^5} - \frac{I_3}{R^3}\right)}_{D(\mathbf{r},\mathbf{s}_j)}\mathbf{q}_j.$$

Combining M fitted moments $\widehat{\mathbf{q}}_j$ with a fitted uniform background $\widehat{\mathbf{b}}_0$ produces the textbook superposition

$$\widehat{\mathbf{B}}_D(\mathbf{r}) = \widehat{\mathbf{b}}_0 + \sum_{j=1}^{M} D(\mathbf{r}, \mathbf{s}_j)\widehat{\mathbf{q}}_j. \tag{4}$$

Because the dipole potential equals $-\mathbf{q}_j^\mathsf{T}\nabla(1/R)$ (with ∇ in $\mathbf{r}$) and $1/R$ is harmonic for $R > 0$, the potential is harmonic everywhere above the source plane; its negative gradient therefore has zero divergence and zero curl by construction.

Equations (1) and (4) share the same additive structure, a background plus a sum of responses, but their building blocks behave very differently. A Gaussian response is a scalar bump centered at a

measurement, whereas a dipole response is a 3×3 matrix centered at an assumed subsurface source. These local, uncoupled bumps and long-range, component-coupled physical responses give individual readings different reach across the room.

3.1 Fitting the moments and background

At each training point $\mathbf{r}_a = (\mathbf{x}_a, 0)$, assemble the block row $D_a = [D(\mathbf{r}_a, \mathbf{s}_1) \ \cdots \ D(\mathbf{r}_a, \mathbf{s}_M)]$, and define D_* analogously at a query point $\mathbf{r}_* = (\mathbf{x}_*, 0)$. Stack the D_a into D, the moments into $\mathbf{q}$, and the fitting targets $\mathbf{t}_a$ into $\mathbf{t}$. Let $C = (I_n - P) \otimes I_3$ subtract the component-wise mean from stacked vectors (with $\otimes$ the Kronecker product), and let S be the diagonal matrix of column norms of CD, floored at 10^{-12} to equalize column scaling. For background $\mathbf{b}_0$ and ridge penalty $\lambda > 0$, the fit solves

$$\min_{\mathbf{b}_0, \mathbf{q}} \sum_a \|\mathbf{t}_a - \mathbf{b}_0 - D_a\mathbf{q}\|^2 + \lambda\|S\mathbf{q}\|^2. \tag{5}$$

For any fixed $\mathbf{q}$, setting the gradient with respect to $\mathbf{b}_0$ to zero yields $\mathbf{b}_0 = \bar{\mathbf{t}} - \bar{D}\mathbf{q}$, where bars denote averages over the training positions. Writing $A = CDS^{-1}$ and $\mathbf{w} = S\mathbf{q}$ reduces Eq. (5) to the standard ridge problem $\min_{\mathbf{w}} \|C\mathbf{t} - A\mathbf{w}\|^2 + \lambda\|\mathbf{w}\|^2$, whose normal equations $(A^\mathsf{T}A + \lambda I_{3M})\mathbf{w} = A^\mathsf{T}C\mathbf{t}$ define a linear map L from targets to moments:

$$\widehat{\mathbf{q}} = L\mathbf{t}, \qquad L = S^{-1}(A^\mathsf{T}A + \lambda I_{3M})^{-1}A^\mathsf{T}C, \qquad \widehat{\mathbf{b}}_0 = \bar{\mathbf{t}} - \bar{D}\widehat{\mathbf{q}}. \tag{6}$$

In code, the ridge solve is evaluated by singular value decomposition. The direct fit sets $\mathbf{t}_a = \mathbf{b}_a$ (the raw readings), whereas the two-stage fit sets $\mathbf{t}_a$ to the Gaussian predictions from Eq. (3) at the same training positions.

3.2 Testing the maps on robot routes

We train on 500 readings from Invensense trials 1–2 and select hyperparameters on trial 3. To test both interpolation and robustness, we evaluate on two held-out sets: trials 4–5 in the same stable room, and trials 6–9 recorded after metal tool shelves were moved into the space. A rectangular gap $[1.5, 2.5]\,\mathrm{m} \times [-2, -1]\,\mathrm{m}$ (with a 0.15-m buffer) is excluded from training and validation so that held-out routes cross an unobserved patch. To test cross-sensor transfer, Nexus phone and Trivisio recordings are placed on the Invensense reference scale using sensor-wise affine alignments fitted on trial 1; all field values are reported in reference recording units. Because this project previously examined both test sets, we treat the indoor comparison as exploratory.

Validation on trial 3 selects a Gaussian length scale of 0.35 m and $\rho = 0.1 + 10^{-10}$ (ridge plus jitter). For the dipoles, we search square grids spanning the training footprint with grid sides $4, 7, 10$, depths $0.25, 0.5, 1$ m, and $\lambda \in \{10^{-5}, 10^{-3}, 0.1, 1\}$. Two-stage validation error selects a 10×10 grid at depth 1 m with $\lambda = 10^{-3}$. We hold that grid and penalty fixed for the direct fit so that only the fitting targets change. A constant baseline predicting the mean training vector everywhere provides a floor for comparison.

Tree baselines use identical arrays and calibration. A 150-tree forest predicts the components jointly; boosted depth-one trees fit them separately. Eighteen forest settings vary depth, leaf size, and feature fraction; nine stump settings vary stages and learning rate. Mean trial-3 RMSE across seeds 11, 23, and 37 selects unrestricted forest depth, leaf size one, and feature fraction 0.7, and 200 boosting stages at rate 0.03. Test scores average per-seed errors; complete grids are archived with the results.

Prediction error. Across N test samples, vector RMSE is $(N^{-1}\sum_{i=1}^{N}\|\widehat{\mathbf{B}}_i - \mathbf{b}_i\|^2)^{1/2}$, where $\widehat{\mathbf{B}}_i$ is the prediction and $\mathbf{b}_i$ the measurement at sample i. On the stable Invensense trials, passing from the Gaussian map to the two-stage dipole map cuts error from 8.09 to 5.22, a 35% reduction from the same 500 readings (Table 1). Yet the direct fit achieves 5.31, within 0.09 units of the two-stage fit, showing

Table 1: Vector RMSE against held-out measurements in reference recording units. Scores pool trials within each environment; tree entries average three seeds. Two-stage and direct columns denote the dipole fits.

Environment	Instrument	Constant	Gaussian	Forest	Stumps	Two-stage	Direct
Stable	Invensense	16.48	8.09	10.56	14.70	5.22	5.31
	Nexus phone	16.32	10.34	11.77	14.72	8.83	8.85
	Trivisio	16.42	8.03	10.20	14.62	5.56	5.71
Changed	Invensense	17.62	8.64	10.82	16.15	7.33	7.36
	Nexus phone	27.53	24.16	25.29	27.01	23.81	23.79
	Trivisio	17.24	8.69	10.58	16.01	7.43	7.50

that the dipole form itself accounts for almost the entire improvement. The same pattern holds across the Nexus phone and Trivisio sensors, as well as after the metal shelves are introduced, where errors rise across all models and remain large on the phone recordings. Neither tree baseline reaches the Gaussian or dipole scores on any instrument in either environment. On stable reference trials, forest seed standard deviation is 0.03 units and stump predictions are identical across seeds. This extends the predictive comparison beyond one smooth statistical baseline; it does not establish an optimum over tree methods or a universal advantage for physical constraints.

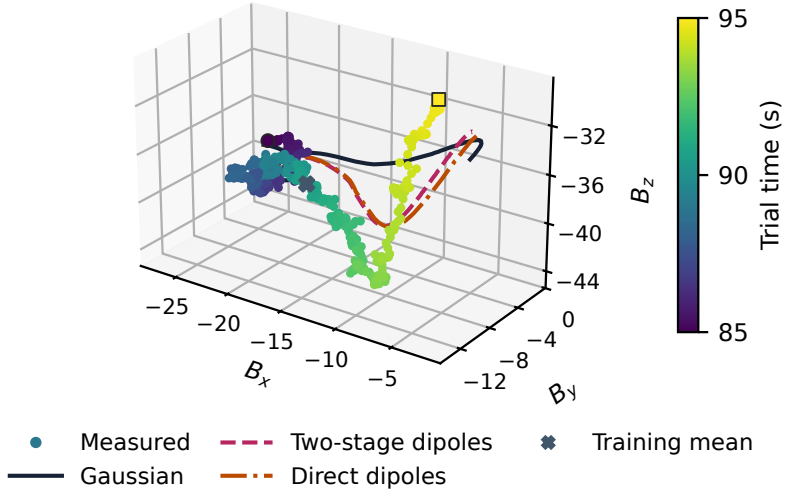


Figure 1: Magnetic vectors from 501 samples at 85–95 s, near the midpoint of held-out Invensense trial 4. Color shows measurement time, from the start circle to the end square. The lines evaluate the fitted models at the same recorded positions. Axes show field components in reference recording units.

Figure 1 shows ten seconds of the held-out route in field space. All three maps underestimate the measured B_z minimum of -44.3: the Gaussian reaches -41.4, and the dipole maps -41.5 to -41.6. The dipole curves nearly overlap, but their measurement weights reveal what the smoothing changed.

4 One field, two sums

Equation (4) builds the dipole field from 100 assumed sources. Once the measurement positions, source geometry, and penalties are fixed, however, the fitted moments and background are linear functions of the targets. That means the exact same predicted field can be rewritten as a weighted sum over the 500 original readings.

Partition L in Eq. (6) into $3M \times 3$ block columns L_a so that $L\mathbf{t} = \sum_a L_a \mathbf{t}_a$. Substituting $\widehat{\mathbf{q}}$ and $\widehat{\mathbf{b}}_0$

into Eq. (4) expresses the dipole prediction directly in terms of its targets:

$$\widehat{\mathbf{B}}_D(\mathbf{r}_*) = \bar{\mathbf{t}} + (D_* - \bar{D})L\mathbf{t} = \sum_a U_a(\mathbf{r}_*)\mathbf{t}_a, \qquad U_a(\mathbf{r}_*) = \frac{I_3}{n} + (D_* - \bar{D})L_a. \tag{7}$$

Each $U_a(\mathbf{r}_*)$ is a 3×3 matrix giving the direct fit's weight on raw reading a. In the two-stage fit, each target is itself a Gaussian average $\mathbf{t}_a = \sum_i H_{ai}\mathbf{b}_i$; substituting this into Eq. (7) yields the two-stage measurement weights:

$$\widehat{\mathbf{B}}_{\text{two}}(\mathbf{r}_*) = \sum_i W_i(\mathbf{r}_*)\mathbf{b}_i, \qquad W_i(\mathbf{r}_*) = \sum_a H_{ai}U_a(\mathbf{r}_*). \tag{8}$$

Equating the source sum in Eq. (4) with the measurement sum in Eq. (8) gives the central identity

$$\widehat{\mathbf{B}}_{\text{two}}(\mathbf{r}_*) = \widehat{\mathbf{b}}_0 + \sum_{j=1}^{M} D(\mathbf{r}_*, \mathbf{s}_j)\widehat{\mathbf{q}}_j = \sum_{i=1}^{n} W_i(\mathbf{r}_*)\mathbf{b}_i. \tag{9}$$

The left-hand side tracks the superposition of physical dipole responses; the right-hand side tracks each piece of empirical data into that exact same field.

This identity makes the effect of smoothing transparent. In the two-stage fit, reading i feeds target a through the scalar H_{ai}, and U_a carries that target to the query, shifting reading i's weight by $W_i - U_i = \sum_a (H_{ai} - \delta_{ai})U_a$ (where $\delta_{ai} = 1$ for $a = i$ and zero otherwise). At fixed positions and settings, perturbing reading i by Δ shifts the direct prediction by $U_i\Delta$ and the two-stage prediction by $W_i\Delta$. Unlike the scalar weights g_{*i} of the Gaussian map, the 3×3 matrices U_i and W_i have off-diagonal entries that carry perturbations from one measured component into another predicted component. Meanwhile, both fits preserve any constant vector offset because centering removes it from the moment fit and the background restores it, giving $\sum_i U_i = \sum_i W_i = I_3$.

Geomagnetic inversions use analogous linear weights to assess spatial resolution and variance [7]. In machine learning, self-distillation refits a model to its own predictions, which can both suppress noise and weaken signal [8, 9]. Our two-stage procedure distills once while changing the model family from kernel to dipoles, and Eq. (8) quantifies the exact resulting shift in every reading's weight.

4.1 The effect of one reading

Using coordinates alone, we select the trial-4 position nearest the gap's center: $(1.990, -1.316)$ m, 0.482 m from the nearest training reading. Figure 2 compares Gaussian weights g_{*i} with dipole vertical weights $[W_i]_{zz}$. Gaussian influence concentrates near the query; the dipole weights spread across the room.

The two maps diverge most sharply at the ringed reading $(2.109, -0.849)$ m, just outside the gap. Raising that single reading's B_z by one unit pushes predicted B_z *up* by $+0.579$ in the Gaussian map, but pushes it *down* by -0.025 in the direct dipole fit and -0.085 in the two-stage dipole fit. Although a Gaussian kernel value is positive, its fitted regression weights can have either sign. The reversal here is a property of the complete regularized fits, not a consequence of kernel positivity. The dipole basis enforces $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = 0$ above the source plane while coupling components and distributing influence across positions. Here, the direct dipole fit already flips the sign, and pre-smoothing makes the response more than three times as negative, showing that nearly identical route RMSEs can coexist with a substantial change in how a specific reading is used. Across all three models, the analytical weights match the archived query prediction and full refit perturbations in all three components to within 10^{-9} recording units.

The global reach of the physical model also appears in the weight distribution: readings more than 1 m from the query carry 33.0% of the absolute vertical weight in the Gaussian fit, but 67.2% in the two-stage dipole fit. While both signed sums equal one, the sum of absolute vertical weights nearly doubles from 3.562 to 6.680, reflecting stronger cancellation between positive and negative weights. Cross-component coupling is likewise substantial, with the largest absolute weight from a horizontal reading to predicted B_z reaching 0.469 in the two-stage fit.

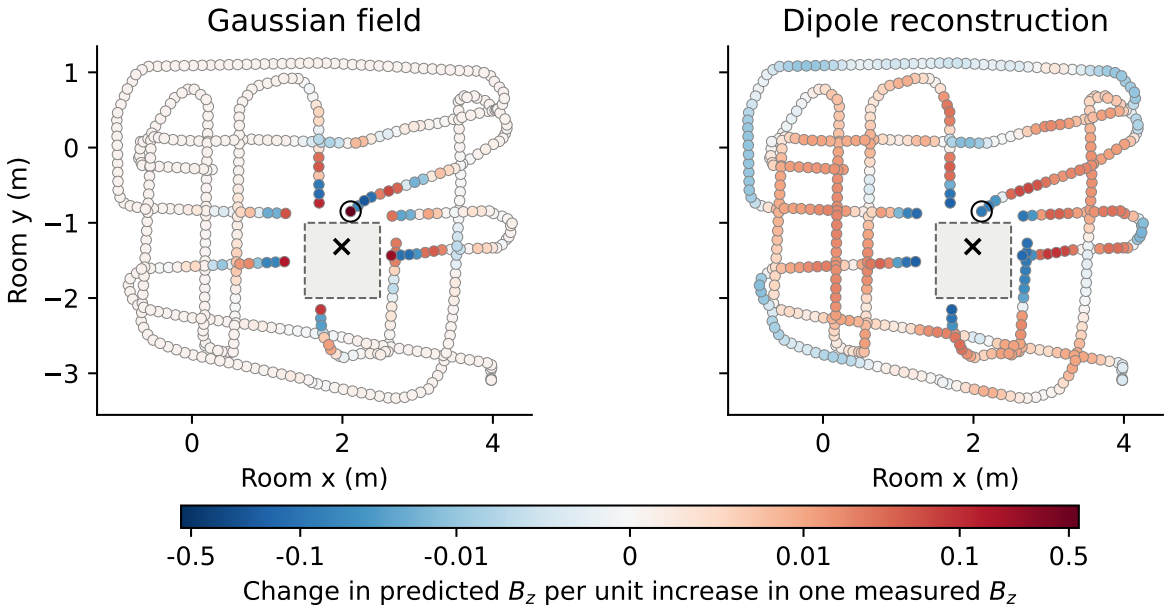


Figure 2: Effect of increasing one measured B_z by one unit on predicted B_z at the cross. Dots mark training positions, the dashed box marks the spatial gap, and a ring marks the reading with the largest absolute difference in vertical weights between models at the query. The panels share a signed color scale, linear within ± 0.01 and logarithmic beyond. Larger absolute weights are drawn on top.

5 Hiding the law in a wire experiment

The room fit supplies dipole responses. A simulated wire instead asks the learner to choose a formula from a general library, with an exact field for scoring.

For a steady signed current I along the z axis of an ideal infinite wire, Ampère's law gives $\oint \mathbf{B}\cdot d\boldsymbol{\ell} = \mu_0 I$, where $d\boldsymbol{\ell}$ is the oriented line element. By cylindrical symmetry, the field is tangent to circles of radius $r = \sqrt{x^2+y^2}$ with signed tangential magnitude $B_\varphi = \mu_0 I/(2\pi r)$ and unit tangent $(-y, x, 0)^\mathsf{T}/r$. In microtesla, with x, y in metres and I in amperes, the simulator evaluates

$$\mathbf{B}_*(x, y, I) = 0.2I\frac{(-y, x, 0)^\mathsf{T}}{x^2+y^2}. \tag{10}$$

The learner observes sampled inputs $\mathbf{u} = (x, y, I)$ paired with noisy field readings; for kernel fitting, the three inputs are standardized on the training set.

Both the two-stage and direct routes search the same candidate library: polynomial fields through degree 4 and rational fields of degree 1 or 2. Each rational candidate takes the form $p_c(\mathbf{u})/q(\mathbf{u})$ for components $c = 1, 2, 3$, sharing a polynomial denominator q with separate polynomial numerators p_c. Given target t_{ic} at input $\mathbf{u}_i$, we initialize by solving $t_{ic}q(\mathbf{u}_i) - p_c(\mathbf{u}_i) \approx 0$ via homogeneous least squares, prune small coefficients, fix the denominator scale, and refit to the field values by nonlinear least squares. As before, the two-stage route fits to Gaussian kernel predictions whereas the direct route fits to raw readings. Each route selects the sparsest candidate within 5% of its lowest validation mean squared error (plus a numerical tolerance), and both are scored against the exact field on independent test points.

Gaussian process smoothing is often used to denoise data before symbolic regression [10], while input ranges [11] and complexity–accuracy trade-offs [12] strongly affect which expressions are recovered. Holding the search space and selection rule fixed isolates the effect of pre-smoothing.

The main case uses 480 training and 180 validation readings across radii 0.25–2 m, all angles, and currents $|I| = 0.5$–4 A, corrupted by independent Gaussian noise at 1% of the pooled RMS of the clean

Table 2: Hidden-wire error on independent inputs in each training domain. For each seed, vector RMSE is divided by the root mean square of the true field magnitudes. Entries are percentages averaged over three seeds; the noiseless direct error rounds to zero.

Case	Gaussian	Two-stage expression	Direct expression
Noiseless	20.42	0.35	0.00
Main, 1% noise	20.44	0.38	0.06
Higher noise, 5%	21.18	0.60	0.28
Sparse, 1% noise	51.10	6.23	0.10
Limited angles, 1%	3.62	0.16	0.04

training components. The sparse regime cuts training to 80 readings, and limited angles restrict azimuth to $\pm\pi/6$ radians; every test domain uses 800 fresh points.

All 15 two-stage fits (five cases across three seeds) select the exact same monomials. Here is the generating field beside the reconstruction from the first seed in the main case, with fitted coefficients rounded for display:

$$\begin{array}{ll} \text{Generating field} & \text{Field recovered through the kernel} \\ \dfrac{0.2I}{x^2+y^2}\begin{pmatrix}-y\\ x\\ 0\end{pmatrix} & \dfrac{I}{1.003060x^2+y^2}\begin{pmatrix}-0.198992y\\ 0.199601x\\ 0\end{pmatrix}. \end{array}$$

Normalizing the y^2 coefficient to one puts every recovered expression in the form $I(-ay, bx, 0)^{\mathsf{T}}/(cx^2+y^2)$. Two physical symmetries follow exactly from that form: doubling the current doubles the field, and doubling both coordinates at fixed current halves it. Circular symmetry, which also requires $a = b$ and $c = 1$, holds approximately.

The main-case rational fit reduces Gaussian error from 20.44% to 0.38%, retaining 0.38% error at larger radii and unseen currents. Direct fitting is better in every regime: 0.06% in the main case and 0.10% with sparse data, versus 6.23% for the two-stage route. The Gaussian struggles with rapid near-wire variation, reaching 51.10% error with 80 readings. Since the library contains the true rational family, direct fitting avoids inheriting that approximation error.

Missing inputs impose a different limit. If signed currents are symmetric at each position and magnitude but current is omitted, the position-conditioned mean field is zero. Under squared error, zero is then the best position-only prediction. Even a suitable equation cannot restore unobserved source information.

6 What the readings determine

The dipole form accounts for most of the room-data gain over the Gaussian; direct equation fitting wins every wire regime. Tree baselines also trail dipoles, although their model classes and tuning budgets differ.

The indoor experiment also makes visible how empirical data enter a textbook field. Collecting Eqs. (1), (3), and (9) places the three models side by side. Superscripts d and s mark coefficients fitted to raw and smoothed targets, and we write $k_{*i} = k_\ell(\mathbf{x}_*, \mathbf{x}_i)$ and $D_{*j} = D(\mathbf{r}_*, \mathbf{s}_j)$:

	Sum of fitted responses		Sum over training readings
Gaussian	$\overline{\mathbf{B}} + \sum_i k_{*i}\boldsymbol{\alpha}_i$	$=$	$\sum_i g_{*i}\mathbf{b}_i$
Direct dipoles	$\widehat{\mathbf{b}}_0^d + \sum_j D_{*j}\widehat{\mathbf{q}}_j^d$	$=$	$\sum_i U_i(\mathbf{r}_*)\mathbf{b}_i$
Two-stage dipoles	$\widehat{\mathbf{b}}_0^s + \sum_j D_{*j}\widehat{\mathbf{q}}_j^s$	$=$	$\sum_i W_i(\mathbf{r}_*)\mathbf{b}_i.$

Each dipole row expresses one field through 100 assumed sources or 500 actual readings. Theory supplies the responses; observations determine their fitted coefficients. The right column gives exact sensitivity to each reading. Physical validity, equivalence of formulas, and predictive accuracy remain separate claims. A companion wave study [14] illustrates this: phase stumps and a phase forest both satisfy Maxwell exactly, but the forest fits the pulse better. Its propagation law is supplied, as magnetostatics is here.

Scope and limitations. These conclusions describe the specific estimators and domains tested here. The indoor comparison covers one room and previously examined trials, and trial 3 selected the shared dipole settings for the two-stage fit; a comparison of fully tuned procedures would tune the direct fit separately, and a scalar-potential Gaussian process offers another physics-based baseline [2]. Moreover, both grid size and depth reached the upper limits of our search, planar measurements leave three-dimensional sources underdetermined, and route-based scores measure accuracy along robot paths rather than uniformly across the room [13].

Data and code. The Aalto recordings are pinned to commit `b45da3b`. The project repository contains the companion notes, code, fixed protocols, saved predictions, and fitted dipole coefficients. The additional baselines and complete validation grids are retained in the tree-comparison directory.

A A scalar check with a known differential law

The dimensionless scalar check uses $B_*(x) = 1.7\tanh(x/0.9)$ and $J_*(x) = 0.8\,\mathrm{sech}^2(x/0.9)$, with $\mathrm{sech}\,u = 1/\cosh u$. Gaussian models fit 80 samples per channel with 0.1% amplitude noise. Sparse regression of the fitted derivative selects $B' = aJ$, recovering $a_* = 85/36$ to 0.022% relative error. Integration from b_0 at $x_0 = -3$ gives $b_0 + a\int_{x_0}^{x} \widehat{J}(u)\,du$.

Restricting the search to powers of B alone selects $B' = p - qB^2$ with $p, q > 0$. Separating variables on $|B| < A$ gives the closed-form solution

$$B(x) = A\tanh((x - x_c)/L), \qquad J(x) = (p/a)\,\mathrm{sech}^2((x - x_c)/L), \qquad A = \sqrt{p/q}, \quad L = 1/\sqrt{pq},$$

The fitted field at the median training position fixes x_c. Noisy-test errors are 0.045% for the solved field, 0.085% for the Gaussian, and 0.229% for the integral reconstruction. Only the field-only equation also recovers the generating function's analytic shape.

References

[1] J. C. Maxwell. A dynamical theory of the electromagnetic field. *Philosophical Transactions of the Royal Society of London*, 155:459–512, 1865. doi:10.1098/rstl.1865.0008.

[2] A. Solin, M. Kok, N. Wahlström, T. B. Schön, and S. Särkkä. Modeling and interpolation of the ambient magnetic field by Gaussian processes. *IEEE Transactions on Robotics*, 34(4):1112–1127, 2018. doi:10.1109/TRO.2018.2830326.

[3] P. E. Chang, P. Verma, S. T. John, A. Solin, and M. E. Khan. Memory-based dual Gaussian processes for sequential learning. *Proceedings of the 40th International Conference on Machine Learning*, PMLR 202:4035–4054, 2023. Proceedings.

[4] I. Abdul-Raouf, V. Gay-Bellile, C. Joly, S. Bourgeois, and A. Paljic. Large scale mapping of indoor magnetic field by local and sparse Gaussian processes. *Proceedings of the 8th Conference on Robot Learning*, PMLR 270:2104–2119, 2025. Proceedings.

[5] D. A. Emilia. Equivalent sources used as an analytic base for processing total magnetic field profiles. *Geophysics*, 38(2):339–348, 1973. doi:10.1190/1.1440344.

[6] I. Uppal and R. Holme. Dual-layer gradient-boosted equivalent sources for magnetic data. *Geophysical Journal International*, 243(3):ggaf359, 2025. doi:10.1093/gji/ggaf359.

[7] M. D. Hammer and C. C. Finlay. Local averages of the core–mantle boundary magnetic field from satellite observations. *Geophysical Journal International*, 216(3):1901–1918, 2019. doi:10.1093/gji/ggy515.

[8] H. Mobahi, M. Farajtabar, and P. L. Bartlett. Self-distillation amplifies regularization in Hilbert space. *Advances in Neural Information Processing Systems*, 33, 2020. Proceedings.

[9] M. Wu, A. Y. Yang, and Q. Sun. Why self-distillation helps and hurts: Denoising vs. signal forgetting. *Proceedings of the 43rd International Conference on Machine Learning*, PMLR 306:137334–137367, 2026. Proceedings.

[10] J. Hsin, S. Agarwal, A. Thorpe, L. Sentis, and D. Fridovich-Keil. Symbolic regression on sparse and noisy data with Gaussian processes. arXiv:2309.11076, 2024.

[11] Y. Matsubara, N. Chiba, R. Igarashi, and Y. Ushiku. Rethinking symbolic regression datasets and benchmarks for scientific discovery. *Journal of Data-centric Machine Learning Research*, 2024. Article.

[12] K. S. Fong and M. Motani. Pareto-optimal fronts for benchmarking symbolic regression algorithms. *Proceedings of the 42nd International Conference on Machine Learning*, PMLR 267:17392–17410, 2025. Proceedings.

[13] A. M. J.-C. Wadoux, G. B. M. Heuvelink, S. de Bruin, and D. J. Brus. Spatial cross-validation is not the right way to evaluate map accuracy. *Ecological Modelling*, 457:109692, 2021. doi:10.1016/j.ecolmodel.2021.109692.

[14] J. R. Landers. Decision stumps as exact Maxwell fronts. Companion research note, 2026. Companion paper.