From a learned kernel to a magnetic field

Real-data reconstruction with equivalent dipoles

J. R. Landers · November 2026

Starting question. Can a kernel learned from real magnetic measurements be converted into an explicit magnetic-field function whose theoretical shapes predict unseen measurements? The earlier note [1] verified a workflow on a hidden synthetic field: samples became a Gaussian kernel, then a differential equation, then its solved analytic function. Here we test a different bridge: project a measured vector-field kernel onto a family supplied by magnetostatic theory, and judge the result against held-out measurements. There is no known analytic ground-truth curve for the room.

1. Measurements become a vector-field kernel. The public Aalto recordings [2] provide tracked planar positions and three calibrated, orientation-corrected magnetic channels from a moving robot. Let $\mathbf{x}_i = (x_i, y_i)$ be position in metres and $\mathbf{b}_i$ its measured field vector. Each channel uses the same Gaussian similarity kernel:

$$\widehat{\mathbf{B}}_K(\mathbf{x}) = \overline{\mathbf{B}} + \sum_i \boldsymbol{\alpha}_i k_\ell(\mathbf{x}, \mathbf{x}_i), \qquad k_\ell(\mathbf{x}, \mathbf{x}_i) = e^{-\|\mathbf{x}-\mathbf{x}_i\|^2/(2\ell^2)}. \tag{1}$$

The coefficient matrix is $(K + \rho I)^{-1}(Y - \mathbf{1}\overline{\mathbf{B}}^{\mathsf{T}})$, where Y stacks the measured vectors. Validation selects ℓ and ρ. All field values retain the reference instrument's recording units; independent SI calibration is unavailable. We use the full vector, retaining directional information that a scalar magnitude profile discards.

2. Theory supplies a magnetic response kernel. In static, source-free air with constant permeability, assume $\nabla \cdot \mathbf{B} = 0$ and $\nabla \times \mathbf{B} = 0$. Locally $\mathbf{B} = -\nabla\Phi$ with $\nabla^2\Phi = 0$. An explicit harmonic potential is

$$\Phi(\mathbf{r}) = -\mathbf{b}_0 \cdot \mathbf{r} + \sum_{j=1}^{M} \frac{\mathbf{q}_j \cdot \mathbf{R}_j}{R_j^3}, \qquad \mathbf{R}_j = \mathbf{r} - \mathbf{s}_j, \quad R_j = \|\mathbf{R}_j\|. \tag{2}$$

Here $\mathbf{s}_j$ is an equivalent source position, $\mathbf{q}_j$ its effective dipole moment, and $\mathbf{b}_0$ a uniform background. Differentiating gives [3]

$$\boxed{\mathbf{B}_D(\mathbf{r}) = \mathbf{b}_0 + \sum_{j=1}^{M} \underbrace{\left(\frac{3\mathbf{R}_j\mathbf{R}_j^{\mathsf{T}}}{R_j^5} - \frac{I_3}{R_j^3} \right)}_{D(\mathbf{r}, \mathbf{s}_j)} \mathbf{q}_j.} \tag{3}$$

With SI calibration, $\mathbf{q}_j = \mu_0 \mathbf{m}_j/(4\pi)$ for magnetic moment $\mathbf{m}_j$; here this scale is absorbed into the fitted coefficients. Every term satisfies the two assumed field equations away from its source. The Gaussian k_ℓ measures similarity; the tensor D maps source moments to magnetic response. They are different kernels connected by the fit below.

For one vertical dipole at $(a, 0, -h)$, evaluated on $(x, 0, 0)$,

$$B_z(x) = b_z + q\frac{2h^2 - (x-a)^2}{[(x-a)^2 + h^2]^{5/2}}, \qquad B_x(x) = b_x + 3qh\frac{x-a}{[(x-a)^2 + h^2]^{5/2}}. \tag{4}$$

Superposition produces peaks, troughs and sign changes without requiring one monotone tanh curve. These shapes are supplied physics; their weights are inferred.

3. The learned kernel determines equivalent moments. Set the observation plane to $z = 0$ and place a regular source grid at $z = -h$. This is an assumed representation, not a measured floor depth. At the n training positions, stack the kernel predictions into $\mathbf{t} \in \mathbb{R}^{3n}$ and the response blocks into $D \in \mathbb{R}^{3n\times 3M}$. Let C subtract the per-component mean across positions, and let diagonal S contain the column norms of CD. The implemented regularized projection is

$$A = CDS^{-1}, \qquad \widehat{\mathbf{w}} = \arg\min_{\mathbf{w}}\{\|A\mathbf{w} - C\mathbf{t}\|_2^2 + \lambda\|\mathbf{w}\|_2^2\},$$
$$\widehat{\mathbf{q}} = S^{-1}(A^\mathsf{T}A + \lambda I)^{-1}A^\mathsf{T}C\mathbf{t}, \qquad \widehat{\mathbf{b}}_0 = \mathrm{mean}_i\left[\widehat{\mathbf{B}}_K(\mathbf{x}_i) - (D\widehat{\mathbf{q}})_i\right]. \tag{5}$$

An SVD evaluates this expression. Substitution into Eq. (3) produces an explicit three-dimensional magnetic field, observed and tested on its planar restriction. For fixed geometry and penalties this is a linear transformation of the learned kernel's outputs. The background is unpenalized; column scaling specifies the regularization convention. No measured currents or derivatives are available, and no synthetic physical observations enter fitting.

4. Real-data experiment. Use Invensense trials 1–2 for fitting, trial 3 for selection, trials 4–5 for stable-environment evaluation, and trials 6–9 for changed-environment stress. Training uses 500 positions, capped at 250 per trial. Exclude $[1.5, 2.5] \times [-2, -1]$ m and a 0.15 m buffer from training and validation. The original kernel selects $\ell = 0.35$ m and $\rho = 0.1$. Other instruments use the existing trial-1-only affine alignment to the reference frame and scale. Evaluation uses the same measured observations for all models.

Fit moments only to kernel predictions at actual training positions. These targets are derived model outputs, not additional observations. Search grid sides $4, 7, 10$, depths $0.25, 0.5, 1$ m, and penalties $10^{-5}, 10^{-3}, 0.1, 1$. Choose minimum measured source-validation vector RMSE among the 36 candidates; freeze all choices before scoring the test labels. Selection gives 100 dipoles at depth 1 m with $\lambda = 0.001$: 303 fitted coefficients including the background. Validation error falls from 8.63 to 6.18.

5. Predictive results. Report vector error $\sqrt{N^{-1}\sum_i \|\mathbf{b}_i - \widehat{\mathbf{B}}_i\|_2^2}$, in reference recording units. Stable trials 4–5 are pooled by instrument:

Instrument	Constant	Kernel	Kernel → dipoles	Direct dipoles
Invensense	16.48	8.09	5.22	5.31
Nexus phone	16.32	10.34	8.83	8.85
Trivisio	16.42	8.03	5.56	5.71

The reconstructed field reduces source-instrument error by 35.4%. Its three component correlations with measured test values are 0.938, 0.908, 0.958. A matched control fits moments directly to training measurements at the same selected geometry and penalty. Its similar error shows that the supplied physical representation accounts for much of the gain; the kernel stage has not been shown necessary. The control is not separately optimized.

With the fitted fields frozen, changed-environment trials 6–9 give:

Instrument	Constant	Kernel	Kernel → dipoles	Direct dipoles
Invensense	17.62	8.64	7.33	7.36
Nexus phone	27.53	24.16	23.81	23.79
Trivisio	17.24	8.69	7.43	7.50

The large phone errors remain visible. Neither predictor establishes transfer across arbitrary environmental changes. These acquisitions were examined in earlier experiments: this is exploratory test reuse, not a fresh confirmatory study.

6. Do the shapes agree?

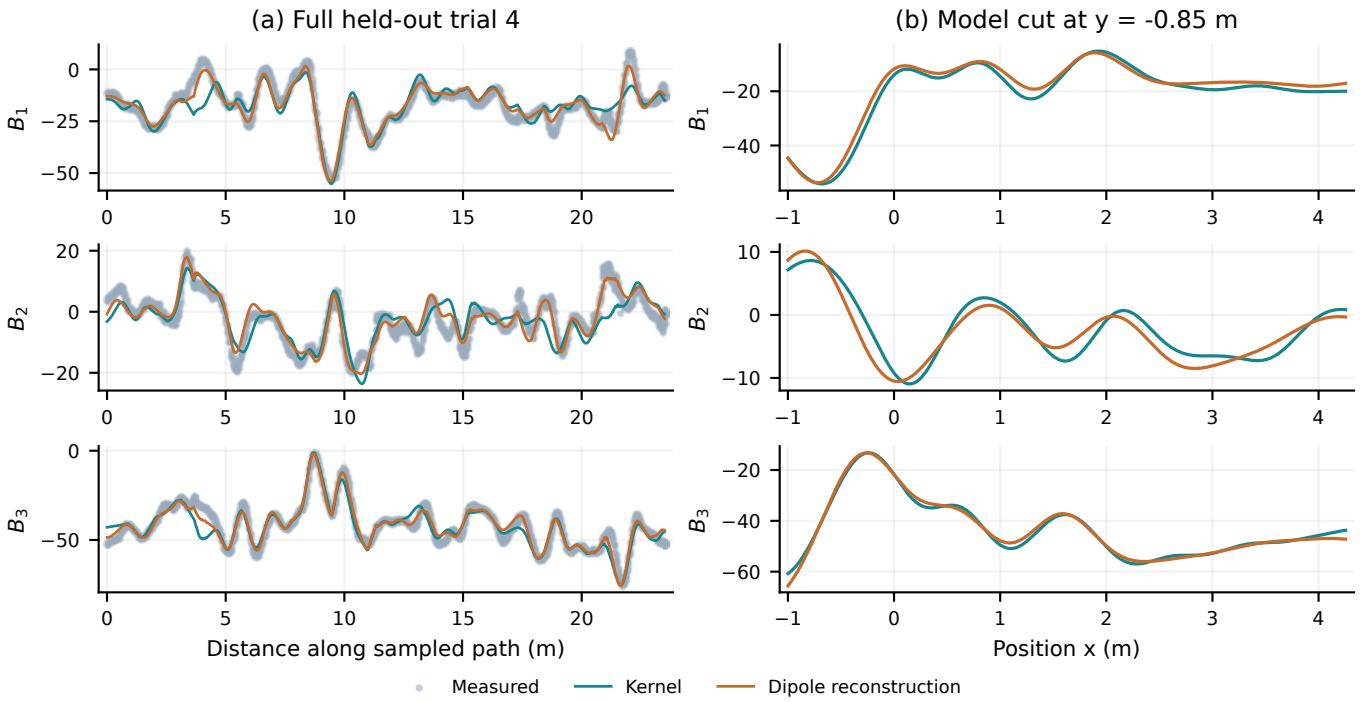


Figure 1. All three field components, in reference recording units. (a) The full first held-out acquisition, Invensense trial 4: measured values, learned Gaussian kernel and equivalent dipole reconstruction. Distance is accumulated between retained path positions. (b) A fixed model cut at $y = -0.85$ m, the median training coordinate. The two reconstructed shapes agree broadly; no independent measured ground-truth curve is available exactly on this cut. Tables pool both stable test trials rather than selecting the displayed acquisition.

7. What has been derived? Equations (2)–(5) give a complete route from the real-data kernel to a theoretical magnetic-field function. The reconstructed field obeys the supplied source-free magnetostatic equations away from the equivalent sources, and predicts held-out values better than the original kernel in this experiment. Three-dimensional curl and divergence identities were checked numerically; saved coefficients reproduce the reported predictions and errors.

This differs from the earlier synthetic recovery [1]. There, a restricted equation search recovered the hidden function family and was scored against an analytic oracle. Here, the physical family is supplied and its coefficients are inferred. The evidence is agreement with measured values; satisfying the imposed field equations does not independently discover them.

Planar measurements do not determine a unique three-dimensional continuation or the actual source arrangement. The fitted dipoles are equivalent sources, not identified room objects. Depth and grid size both reach the largest searched values, so neither is an established physical estimate or an established search optimum. Source-depth sensitivity, stability across training trials and independent public data with three-dimensional positions are the next tests of robustness.

References and reproducibility.

[1] J. R. Landers, *From samples to an equation: A magnetic-like kernel benchmark* (2026). Companion note in `papers/01-kernel-to-equation/`.

[2] A. Solin et al., *Modeling and interpolation of the ambient magnetic field by Gaussian processes*, IEEE Trans. Robotics (2018), doi:10.1109/TRO.2018.2830326. Public recordings; pinned archive, upstream calibration and orientation correction retained.

[3] NASA CCMC, *Potential Form of the Dipole Field Equations*.

Build all notes: `python scripts/build_papers.py`. Protocol, coefficients, provenance and scores: `experiments/magnetic_mapping/dipole_field/`.