Maxwell's Perspective

Observations, learned kernels and the analytic functions we infer

J. R. Landers · November 2026

Starting question. If two observers meet the same magnetic world through different experiments and different mathematical tools, will they arrive at the same analytic function? Our first two notes [1,2] approached this question operationally: learn a predictor from observations, translate it into mathematics, construct an explicit function, and test what that function predicts. A new hidden-wire experiment lets us place the recovered expression beside a known theoretical solution. Their close agreement is a useful result. The way that agreement was obtained is the subject of this note.

1. A perspective before a formula. Imagine receiving a table of positions, currents and compass-like readings. The field's formula is absent. You can interpolate the table, draw curves through it, propose relationships, and ask what another observation would show. Every step introduces a representation. Eventually a short equation may do what the table could not: describe a whole family of observations at once.

This is the sense in which we borrow "Maxwell's perspective." In his work on physical lines of force, Maxwell compared mechanical hypotheses with observed magnetic and electrical phenomena [3]. His later field theory joined experimentally motivated relations and connected electromagnetic disturbances with light [4]. Our experiment reproduces one inference problem, not that historical achievement. Maxwell's intermediates included field pictures, analogies and mathematical constraints; ours include a kernel predictor and a finite expression search. Both routes connect observations to mathematical consequences.

In modern vacuum notation, one relevant field equation is

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}. \tag{1}$$

Here $\mathbf{E}$ and $\mathbf{B}$ are electric and magnetic fields, $\mathbf{J}$ is current density, and μ_0, ϵ_0 are vacuum constants. Under steady conditions, an ideal infinite straight wire and cylindrical symmetry give

$$\oint \mathbf{B} \cdot d\boldsymbol{\ell} = \mu_0 I, \qquad B_\varphi(r) = \frac{\mu_0 I}{2\pi r}. \tag{2}$$

I is the enclosed signed current, r distance from the wire, and φ the direction around it. This familiar solution follows from the steady Ampère law together with geometry. It is not the full Maxwell system.

2. Put the functions beside each other. For positions x, y in metres, current I in amperes and fields in microtesla, our simulator uses the coefficient 0.2. The first fixed 1%-noise run recovers:

Hidden ideal-wire field	Kernel-derived analytic field
$\mathbf{B}_* = \dfrac{0.2I}{x^2+y^2}\begin{pmatrix}-y\\ x\\ 0\end{pmatrix}$	$\mathbf{B}_\theta = \dfrac{I}{1.003060x^2+y^2}\begin{pmatrix}-0.198992y\\ 0.199601x\\ 0\end{pmatrix}.$

(3)

The same algebraic structure has emerged, with small coefficient differences. The learner was given neither expression. The displayed decimals are rounded; all scores use the saved full-precision function. Agreement here concerns one particular magnetic response, within a supplied algebraic search class. The simulated world obeys the wire law by construction; recovering it tests our inference process without independently validating electromagnetism.

3. Three experiments, one operational arc. The shared route is

$$\text{data} \longrightarrow \text{kernel} \longrightarrow \text{mathematical structure} \longrightarrow \text{analytic function} \longrightarrow \text{test}. \tag{4}$$

For an input $\mathbf{u}$, the common predictor has the form

$$\widehat{\mathbf{B}}_K(\mathbf{u}) = \overline{\mathbf{B}} + \sum_n k_\ell(\mathbf{u}, \mathbf{u}_n)\boldsymbol{\alpha}_n, \qquad k_\ell = \exp\left[-\frac{\|S^{-1}(\mathbf{u}-\mathbf{u}_n)\|^2}{2\ell^2}\right]. \tag{5}$$

The coefficients come from regularized fitting. Validation selects the smoothing length ℓ and penalty. In the wire study S contains training input standard deviations; the earlier spatial studies use their stated coordinate scales. The similarity kernel supplies smoothness, not a magnetic field law.

	Note 1: hidden profile	Note 2: measured room	This note: hidden wire
Observations	Dimensionless $B(x), J(x)$ samples	Planar positions and three magnetic channels	Cartesian positions, signed current and three magnetic channels
Bridge	Sparse differential equation	Supplied magnetic dipole response	Sparse polynomial/rational search
Construct	Solve the selected ODE	Assemble equivalent sources	Evaluate the selected rational function
Judge	Hidden analytic profile	Held-out measured readings	Hidden field at unseen conditions

In note 1, a field-only search selects $B' = p - qB^2$. Separation of variables, on the branch fixed by a learned reference value, yields

$$B_{\text{sym}}(x) = A\tanh\left(\frac{x-x_c}{L}\right), \qquad A = \sqrt{p/q}, \quad L = (pq)^{-1/2}. \tag{6}$$

The hidden profile has $(A, L, x_c) = (1.7, 0.9, 0)$. The noisy recovery gives $(1.699495, 0.899139, -0.000644)$, with 0.045% field error. The equation grammar and integration rule are supplied; the selected equation, scales and center are inferred. The companion source recovers a sech^2 shape.

In note 2, public Aalto recordings [5] supply real vector measurements, but no analytic room-field oracle. A theoretical representation replaces the ODE:

$$\mathbf{B}_D(\mathbf{r}) = \mathbf{b}_0 + \sum_j \left(\frac{3\mathbf{R}_j\mathbf{R}_j^{\mathsf{T}}}{R_j^5} - \frac{\mathbb{I}_3}{R_j^3}\right)\mathbf{q}_j, \qquad \mathbf{R}_j = \mathbf{r} - \mathbf{s}_j. \tag{7}$$

The source positions $\mathbf{s}_j$ belong to an assumed subsurface grid; effective moments $\mathbf{q}_j$ and background $\mathbf{b}_0$ are fitted to kernel outputs. The source-free magnetic structure is supplied. Stable reference-sensor RMSE falls from 8.090 to 5.225; a direct-data dipole control gives 5.312, in recording units. Equivalent sources do not identify the room's actual objects, and planar observations do not fix a unique 3D field.

For the wire, we supply general monomials $\boldsymbol{\phi}_d = \{x^a y^b I^c : a+b+c \le d\}$ instead of dipoles. The candidate is

$$B_{\theta,j}(\mathbf{u}) = \frac{\boldsymbol{\phi}_d(\mathbf{u})^{\mathsf{T}}\mathbf{p}_j}{\boldsymbol{\phi}_d(\mathbf{u})^{\mathsf{T}}\mathbf{q}}, \qquad t_{n,j}Q(\mathbf{u}_n) - P_j(\mathbf{u}_n) \approx 0, \quad \mathbf{t}_n = \widehat{\mathbf{B}}_K(\mathbf{u}_n). \tag{8}$$

Sparse homogeneous fitting and field-value refinement construct the candidate; validation observations select it. Polynomials are also allowed. As in note 2, targets are teacher values at actual training inputs, not extra observations. The extracted expression is a new model, not an exact symbolic rewrite of the Gaussian sum.

4. What happened when the world was hidden? We fit 480 training observations, validate on 180, and score 800 independent points per test. The main domain has radii 0.25–2 m, all angles and $|I| = 0.5$–4 A. Main noise is independent Gaussian noise per component, with standard deviation 1% of pooled clean training component RMS; its scale is withheld from selection. Three fixed seeds are run for each of five conditions: no noise, 1% noise, 5% noise, 80 training points, and a restricted angular wedge.

The learner sees x, y, I and B_x, B_y, B_z, with no radial features, derivatives or permeability. Validation selects among Gaussian kernels, then among general polynomials through degree 4 and rational expressions through degree 2. The sparsest expression within 5% of the best validation MSE is retained, with a fixed numerical floor. The identical expression search is independently repeated on raw observations. All choices finish before test labels are opened.

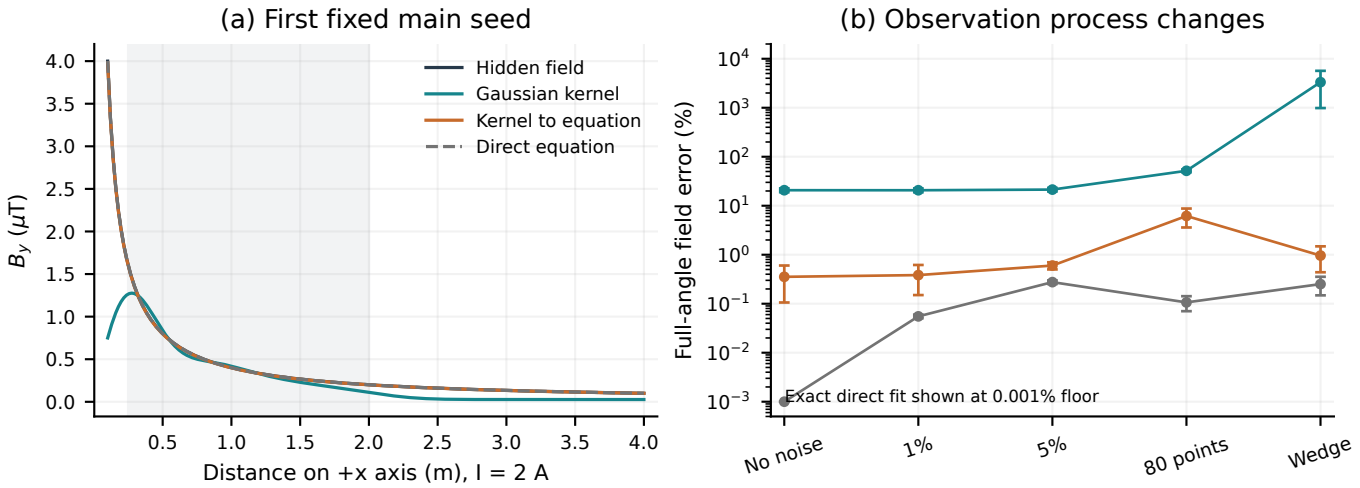


Figure 1. (a) A fixed cut at $y = 0$, $I = 2$ A, for the first main seed; shading marks the observed radius range. (b) Full-angle errors under five observation conditions; points and bars are means and standard deviations across three seeds. "Wedge" training/validation sees only angles $\pm\pi/6$. The noiseless direct fit is at numerical precision and displayed at the stated floor. All selected expressions, including failures of accurate coefficient recovery, are retained in the experiment archive.

The main-case table reports $100\|\mathbf{B}_{\text{model}} - \mathbf{B}_*\|_2/\|\mathbf{B}_*\|_2$, averaged over all three seeds:

Unseen conditions	Kernel	Kernel → equation	Direct equation
Same domain	20.439	0.383	0.056
Smaller radii	71.538	0.381	0.056
Larger radii	97.839	0.379	0.055
Larger currents	97.366	0.380	0.055
Moved wire; no update	86.991	89.869	89.999

The extracted fields recover the wire's monomial support in all 15 runs. Current proportionality and inverse-distance scaling follow exactly from that selected support; circular symmetry is approximate because coefficients differ. Errors remain small beyond the observed distances and currents. Yet direct equation fitting is better: 0.056% versus 0.383% in the main domain. The kernel stage is not shown necessary. Its main fits select the smallest searched length, so its 20.439% error is a result for this search, not an established optimum for kernel learning. Even without measurement noise, the finite kernel-to-equation route retains 0.353% mean field error; exact samples do not remove its approximation choices.

Sparse sampling increases the distilled error to 6.23%. Wedge-only observations give 0.158% within the wedge and 0.961% over all angles. A nonmagnetic polynomial control recovers a polynomial, rather than a wire shape. Moving the wire without supplying its new position defeats all frozen models. These outcomes measure what the process retains, what its grammar can express, and which changing conditions its inputs omit.

5. The experiment helps define the function. A measurement is already a relationship between a world and an apparatus. Write an experiment's observations schematically as

$$\mathbf{b}_i = H_e[F](\mathbf{u}_i) + \boldsymbol{\epsilon}_i, \qquad f_e^*(\mathbf{u}) = \mathbb{E}[\mathbf{b} \mid \mathbf{u}, e]. \tag{9}$$

F denotes a candidate physical state or field history; H_e includes the channels, calibration, orientation and averaging of experiment e. For squared prediction loss, the population target is the conditional mean f_e^*. A finite learner approximates it through its chosen representation. Changing the available inputs or the observation operator can change that target, even when the physical world is unchanged.

The wire gives a sharp example. Keep both signs of current but remove I from the learner's inputs. If its conditional mean is zero, then

$$\mathbb{E}[\mathbf{B}_* \mid x, y] = \frac{0.2\,\mathbb{E}[I \mid x, y]}{x^2 + y^2}(-y, x, 0) = \mathbf{0}. \tag{10}$$

Zero is then the optimal mean predictor for those incomplete inputs. Individual readings still have nonzero fields and irreducible error for that predictor. The zero function and the current-conditioned wire function answer different statistical questions. This is a theoretical consequence of the setup, not an additional fitted experiment. Similarly, spatial averages, missing orientation or magnetic history can require a different effective description.

6. Would the two routes converge? They can converge in predictive content if they target the same calibrated quantity, include the relevant conditions, sample a sufficiently informative domain, and use representations and estimators capable of consistent recovery. They need not converge to identical notation. For two experimental routes, observational equivalence over a tested domain means

$$F_1 \sim F_2 \quad \Longleftrightarrow \quad H_e[F_1](\mathbf{u}) = H_e[F_2](\mathbf{u}) \quad \text{for every tested } (e, \mathbf{u}). \tag{11}$$

This identifies a unique field only when those observations distinguish the candidate fields. Finite noisy samples generally leave alternatives; the planar sensor study particularly leaves unobserved 3D structure. Differential equations, source expansions and integral kernels may also express the same field while retaining different parameters. Our small wire experiment demonstrates close recovery within its grammar, not a convergence theorem.

Neither an analytic expression nor an ML predictor becomes "more real" through its form. Both earn empirical scope by surviving tests. Maxwell's equations carry claims across many configurations; one recovered wire function addresses a narrow slice. Two functions may be equally adequate on that slice and disagree elsewhere. Their disagreement gives the next experiment: vary the source geometry, resolve a missing direction, or observe time dependence and induction. A useful theory connects those experiments with fewer new assumptions.

Takeaway. The three studies make an empirical route to explicit mathematics inspectable. The function we recover depends jointly on observations, accessible variables and mathematical choices. Our wire expressions nearly meet at a familiar law; our sensor reconstruction supplies a physical family; our first note solves a selected equation. The next question is which recovered relationships remain stable when the experiment changes. That stability can support a shared theory without requiring a privileged formula or a unique hidden source arrangement.

References and reproducibility.

[1] J. R. Landers, *From samples to an equation* (2026). [2] J. R. Landers, *From a learned kernel to a magnetic field* (2026).

[3] J. C. Maxwell, *On physical lines of force* (1861–62), opening discussion.

[4] J. C. Maxwell, *A dynamical theory of the electromagnetic field* (1865), doi:10.1098/rstl.1865.0008; Parts I and VI.

[5] A. Solin et al., *Modeling and interpolation of the ambient magnetic field by Gaussian processes* (2018), doi:10.1109/TRO.2018.2830326.

Protocol, all fits, sensor-stage audit and tests: `experiments/hidden_wire/`. Build all notes: `python scripts/build_papers.py`.