

Supplementary Information for: A unified rate-dependent anisotropic hysteresis model for grain-oriented, non-oriented, and directional electrical steels

Mario Bendra ^{1,*} and Florian Poltschak ¹

¹JKU Hoerbiger Research Institute for Smart Actuators, Altenberger Strasse 66, Science Park 1, 4040 Linz, Austria

*mario.bendra@jku.at

ABSTRACT

This document is the Supplementary Information for the manuscript referenced above. It contains tables of fitted parameters for every material studied, per-loop residuals, the moment-closure verification against the Ramesh–Jiles–Bi 1997 anisotropic-hysteresis benchmark, the permeability overlay for the Li and Luo amplitude sweep, the discrete Boltzmann convergence study, and the parameter-identifiability map. All numerical results are reproducible from the public `hysterpy` code repository cited in the main text.

This Supplementary Information accompanies the main manuscript and contains: a model-naming legend (**Note S1**) clarifying which model is used where, fitted parameter sets for every material (**Tables S1–S4**), per-loop residuals and identifiability summaries (**Tables S5–S6**), the per-ring cross-model comparison, unification ablation, and cross-validation summary (**Tables S7–S9**), supporting figures (**Figures S1–S13**: S^2 geometry, unification ablation, the 1 Hz fit, cross-model overlays, the cost-landscape multimodality study, the cross-model RMSE summary, the multistart-basin scatter, the 1D Maxwell skin-depth verification, the Ramesh–Jiles–Bi consistency check, the third-harmonic excitation demonstration, the HEFMAG external-dataset blind cross-validation, and the LLG-direction infrastructure validation), three further explanatory notes (**Notes S2–S4**) on the MagHyst[®] measurement protocol, the lat–lon grid convergence, and the rationale for fitting the 1 Hz quasi-static loop separately, a cross-validation and cost-landscape-identifiability study (**Note S5**), and a 1D Maxwell-coupled derivation of the cubic-anisotropy sigmoid corner frequency (**Note S6, Fig. S9**). All numerical results are reproducible from the public code repository (`hysterpy`, MIT license) cited in the main text Data Availability statement.

Supplementary Note S1: Model naming and where each model is used

The paper compares the present *vector Fokker–Planck* (*vector FP*) model against six published baselines. The naming convention used throughout, and the `hysterpy` class behind each, is listed below. The main result of the paper relies solely on the first entry. The rest appear as cross-model baselines in main-text Table 2, main-text Table 3, and Supplementary Table S7.

- **vector FP** S^2 (`hysterpy.DTHMS2Model`): the present model. Fokker–Planck equation on the unit sphere S^2 , Chang–Cooper well-balanced flux, Zeeman + uniaxial K_u + cubic $K_1(f)$ free energy, polycrystalline orientation distribution. Used for: main-text Fig. 3, main-text Table 1, the “This work” column of main-text Table 2.
- **scalar FP** S^1 (`hysterpy.DTHMModel`): the strict $S^2 \rightarrow S^1$ reduction of the vector FP model: same Chang–Cooper flux, anisotropy reduced to a uniaxial scalar, no cubic invariant and used for: the rate-independent isotropic reduction on the seven NGO rings (Tables S3, S7 and Fig. S7).
- **DTHM-Simplified** (`hysterpy.DTHMSimplifiedModel`): Jiles–Atherton-family ODE solver with separate M_{irr} tracking. Used as a baseline only.
- **J–A** (`hysterpy.JilesAthertonModel`): classical Jiles–Atherton ODE¹.
- **Chwastek** (`hysterpy.ChwastekModel`): modified J–A with mixing-rule β replacing c^2 .
- **DDD** (`hysterpy.DDDModel`): D’Aloia–Di Francesco–De Santis elastic-collision model³.

- **Preisach** (`hysterpy.PreisachModel`): classical Preisach with Gaussian hysteron density⁴.
- **P-I** (`hysterpy.PrandtlIshlinskiiModel`): Prandtl–Ishlinskii play-operator pile⁵.

Throughout this paper, “this work” denotes the vector Fokker–Planck S^2 model with the parameter set of main-text Table 1. The “isotropic reduction” on the NGO rings is the same model with $K_u = K_1 = 0$ enforced, which is mathematically identical to scalar FP on S^1 .

Supplementary Note S2: MagHyst[®] measurement protocol

The three non-oriented ring-core measurements (X10CrAlSi7, X20Cr13, S235JR) reported in the main text were acquired in-house on the MagHyst[®] Modular system of Ilmenauer Mechatronik GmbH⁶. The closed magnetic circuit consists of a toroidal ring sample with primary and secondary windings. The primary is current-driven by a software-defined reference $i_{\text{prim}}(t)$ following a triangular waveform of period $T = 1/f$, the secondary picks up the induced voltage $u_{\text{ind}}(t) = -d\Psi/dt$, and the flux linkage is recovered by numerical integration of the open-circuit secondary voltage. Sample geometry parameters (mean radius r_{Fe} , iron cross-section A_{Fe} , primary turns N_{prim} , secondary turns N_{sec}) translate (i_{prim}, Ψ) to material-frame $(H, B) = (N_{\text{prim}} i_{\text{prim}} / (2\pi r_{\text{Fe}}), \Psi / (N_{\text{sec}} A_{\text{Fe}}))$. The MagHyst[®] device exports the measured loop as an XML record containing the full $(t, i_{\text{prim}}(t), u_{\text{ind}}(t), B(t), H(t))$ time series at a sampling rate of typically 100 kS/s. The measurement frequency is the inverse of the four-cycle MeasureTime field in the XML. We use the third loop of each four-cycle record for fitting, discarding the first two as warmup. This is consistent with the $n_{\text{warmup}} = 2$ convention used in all model simulations⁷.

Supplementary Note S3: Lat–lon grid convergence study

The Chang–Cooper well-balanced flux discretization is exact in the sense that the discrete Boltzmann distribution is the exact kernel of the discrete Fokker–Planck operator, regardless of grid resolution. Convergence to the continuum limit is therefore a question of grid quadrature accuracy, not of stability or operator bias. The convergence of the discrete $\langle m_x \rangle = M/M_s$ at $H/a = 1$ to the continuum Langevin function $\mathcal{L}(1) = 0.3130$ as a function of N_θ at fixed $N_\phi = 2N_\theta$ is summarised numerically below:

$$N_\theta \in \{16, 24, 32, 48, 64\} \Rightarrow \frac{|\langle m_x \rangle_{\text{disc}} - \mathcal{L}(1)|}{\mathcal{L}(1)} \in \{1.6 \times 10^{-3}, 5 \times 10^{-4}, 3 \times 10^{-4}, 1 \times 10^{-4}, 5 \times 10^{-5}\}. \quad (\text{S1})$$

The reported main-text S^2 results use $(N_\theta, N_\phi) = (32, 64)$, which gives 0.03 % accuracy in the Langevin recovery and is sufficient for all materials in our validation set ($H/a < 5$). Sharper-distribution problems (e.g. ferromagnets near the Curie point or low-field limit of permanent magnets) may require higher N_θ .

Supplementary Note S4: Quasi-static 1 Hz loop fit

The 1 Hz Baghel measurement is fitted separately from the 50 to 200 Hz dynamic loops. The reason is geometric. At fixed $\tau_R D = 1$, the canonical Boltzmann limit is reached when $\omega \tau_R \ll 1$ (period much longer than relaxation). For our fitted τ_R values ($\tau_R \approx 0.02$ s in dynamic fits), $\omega \tau_R \approx 0.13$ at 1 Hz versus $\omega \tau_R \approx 6.3$ at 50 Hz: the 1 Hz loop is essentially anhysteretic. The same $(M_s, a, \alpha, K_u, K_1)$ shared parameters that fit the dynamic loops are found by an independent fit of the 1 Hz loop with $\tau_R \rightarrow 0^+$ in the model. This split is consistent with established practice in the dynamic Jiles–Atherton literature^{7,8} and with the Kalmykov–Coffey analysis of the orientational Fokker–Planck relaxation⁹.

Supplementary Note S5: Generalisation across frequency splits, $K_1(f)$ functional form, and cost-landscape identifiability

The main result of the main text (Methods § Joint fit, Results Fig. 3) shares a base (M_s, a, α, K_u, D) plus the three sigmoid constants $(K_1^{\text{low}}, f_c, n)$ across all eight RD+TD dynamic loops at 50 to 200 Hz. We test three orthogonal robustness questions on the same dataset and the same shared-base architecture: (a) does the fit generalize across frequency splits? (b) is the sigmoid form for $K_1(f)$ defensible against alternative monotone parametrisations? (c) does the 38-dimensional cost surface admit unique identification, or are there competing local minima?

Cross-frequency generalisation. Two frequency-split runs probe extrapolation in both directions. A *forward-extrapolation* fit jointly fits the 50 Hz RD+TD pair only and predicts the higher frequencies via the $K_1(f)$ sigmoid, the $\tau_R(f) = \tau_R^{50}(50/f)$ scaling, and $H_{\text{pin}}^{\text{TD}}(f) = H_{\text{pin}}^{50}(f/50)$ of the main text. Mean RMSE at 50 Hz is 0.043 T. A *reverse-extrapolation* fit holds the shared base fixed and refits only the four $K_1(f)$ sigmoid constants $(K_1^{\text{low}}, f_c, n, \tau_{\text{const}})$ on the 100 to 200 Hz loops, then predicts 50 Hz. Mean fit cost at the held-in frequencies is 0.072 T. Both splits reproduce the held-out frequencies at residuals comparable to the main joint fit (Table S9), confirming that the $K_1(f)$ sigmoid is a transferable scaling rather than a frequency-by-frequency refit in disguise.

Alternative $K_1(f)$ functional forms. The same reverse-extrapolation split was re-run with the sigmoid replaced by a monotonically *increasing* Hill function $K_1(f) = K_1^{\max} f^n / (f_0^n + f^n)$. The Hill form lands at fit cost 0.092 T versus the sigmoid's 0.072 T (a 28 % degradation under identical budget). This selects the monotonically decreasing sigmoid form as the data-preferred functional shape, consistent with the skin-depth-mediated cubic-anisotropy washout interpretation of the main text (Methods § Frequency-dependent effective cubic anisotropy).

Leave-one-out cross-validation. A 4-fold leave-one-out cross-validation of the sigmoid form refits $(K_1^{\text{low}}, f_c, n, \tau_{\text{const}})$ on three of the four frequencies and tests on the held-out one. Per-fold fit costs range 0.062 to 0.077 T. The sigmoid constants themselves span $K_1^{\text{low}} \in 1098 \text{ to } 5655 \text{ J/m}^3$ and $f_c \in 37 \text{ to } 182 \text{ Hz}$ across the four folds, reflecting the well-known three-parameter sigmoid degeneracy: the physically observable quantity is the curve $K_1(f)$, not the three parameters that index it. Evaluated at the measurement frequencies, the leave-one-out $K_1(f)$ predictions all sit within the main-fit trajectory $\pm 30\%$, supporting the sigmoid as a phenomenological surrogate for the cubic-anisotropy washout rather than as a derived intrinsic constant.

Cost-landscape identifiability. To quantify practical identifiability, we ran a 20-seed cold-start multistart on the full 38-dimensional joint cost surface (uniform-random initial population over the full parameter bounds, popsize = $3 \cdot 38 = 114$, maxiter = 40, total wall 16.94 h on 24 cores). The resulting 20 best-of-seed parameter sets cluster into three basins (Fig. S8): (i) *basin A* (this work, $N = 14$ of 20): $K_1^{\text{low}} \in 955 \text{ to } 2685 \text{ J/m}^3$, $f_c \in 38 \text{ to } 215 \text{ Hz}$. The within-basin coefficient of variation $\sigma(\cdot)/\langle\cdot\rangle$ is 0.010 on M_s , 0.036 to 0.043 on $\tau_R(f)$, 0.041 on ω_{Goss} , and 0.136 on the physically observable $K_1(50 \text{ Hz})$, while the individual sigmoid constants $(K_1^{\text{low}}, f_c, n)$ span CoV 0.34 to 0.40 along the $K_1(f)$ curve (the same parametric degeneracy as in the cross-validation study above). (ii) *Basin B* ($N = 2$): elevated cubic anisotropy at $K_1^{\text{low}} \sim 1.2 \times 10^4 \text{ J/m}^3$ (within $4\times$ of the Fe-Si single-crystal value $4.8 \times 10^4 \text{ J/m}^3$) with sharp sigmoid falloff $n \approx 3.8$ that collapses $K_1(200 \text{ Hz})$ to within 10 J/m^3 . (iii) *Basin B'* ($N = 4$): near-constant $K_1 \in 150 \text{ to } 850 \text{ J/m}^3$ across the measurement band, with corner frequency f_c above 400 Hz so the sigmoid never turns off.

The three basins are degenerate to within $\Delta\text{cost} \approx 0.03$ at the multistart budget (best-of-seed costs 0.096, 0.107, 0.102 for A, B, B' respectively), but they are physically distinguishable. Basin A produces the textbook HGO Goss-textured RD+TD coupling reported in the main text. Basin B implies single-crystal-magnitude cubic anisotropy that would require $\sim 4\times$ higher saturation fields to verify than the Baghel dataset spans. Basin B' has K_1 effectively switched off, contradicting the Goss-texture cubic structure. Restricting the parameter search to the basin-A neighborhood is therefore a physically motivated identifiability constraint, not a fitting artifact. The Bingham-spread anisotropic-texture variant of basin B (basin C of Supplementary Fig. S6) returns mean RMSE 1.1 T and is rejected on the goodness-of-fit basis there. All multistart outputs and the basin analysis are reproducible from the public companion repository cited in the main-text Code Availability statement.

Supplementary Note S6: 1D Maxwell-coupled analysis of the cubic-anisotropy sigmoid corner frequency

The main text introduces a phenomenological frequency-dependent effective cubic anisotropy

$$K_1(f) = \frac{K_1^{\text{low}}}{1 + (f/f_c)^n}, \quad K_1^{\text{low}} = 1619 \text{ J/m}^3, \quad f_c = 99.0 \text{ Hz}, \quad n = 2.81, \quad (\text{S2})$$

to close the rolling-direction residual at 150 Hz and 200 Hz. Methods § Frequency-dependent effective cubic anisotropy gives the closed-form linear skin-depth consistency check: at the published HGO operating-knee permeability $\mu_r \approx 1.35 \times 10^4$, the lamination ratio $d/\delta = 1$ lands at $f \approx 99 \text{ Hz}$, matching the fitted corner. This Note reports a direct 1D Maxwell-coupled numerical simulation that confirms the corner-frequency match, quantifies the depth gap between the fitted sigmoid and linear skin theory, and identifies the operative mechanism.

Setup. We solve the 1D coupled magnetic-diffusion problem $\sigma \mu_0 \partial_t(H + M) = \partial_z^2 H$ (Methods § Frequency-dependent effective cubic anisotropy) inside one Hi-B 27M-OH lamination of thickness $d = 0.30 \text{ mm}$ and bulk conductivity $\sigma = 2.1 \times 10^6 \text{ S/m}$. The cross-section is discretized into $N_z = 17$ uniformly spaced slices. Each slice carries an independent rotation-only Fokker–Planck constitutive state on S^2 (lat–lon grid 16×32) using the shared parameters $(M_s, a, \alpha, K_u, D, \tau_R)$ of this work and the *same* cubic anisotropy $K_1 = K_1^{\text{low}}$ for every slice and every frequency. The driving field $H_{\text{lab}}(t) = 250 \text{ A/m} \sin(2\pi ft)$ is applied as Dirichlet boundary condition at $z = \pm d/2$. Time integration uses operator-splitting: per step the per-slice Fokker–Planck operator updates $M^{n+1}(z)$ from $H^n(z)$, then a backward-Euler tridiagonal solve advances $H^{n+1}(z)$ given the new $M^{n+1}(z)$ and the boundary value. We take 160 time steps per period and report the third period after two warm-up cycles.

Spatial field profile. Supplementary Fig. S9(a) shows the steady-state $H_{\text{internal}}(z)/H_{\text{lab,peak}}$ at the surface-field peak, for two representative frequencies. At 50 Hz the interior field is nearly uniform across the lamination ($\leq 1\%$ depression at the center). At 200 Hz the centre-line value drops to $\sim 0.84 H_{\text{lab,peak}}$, the signature of incipient skin shielding at $d/\delta \approx 1.5$.

Corner-frequency match. The physically operational quantity is the bulk-averaged attenuation $A_{\text{num}}(f) = \max_t |\langle H(z, t) \rangle_z| / \max_t |H_{\text{lab}}(t)|$. Across the six measured frequencies the numerical $A_{\text{num}}(f)$ matches the analytical linear-theory envelope $|\tanh(\gamma d/2)/(\gamma d/2)|$ with $\gamma = (1 + j)/\delta$ to better than 1 %, validating the numerical solver against textbook Maxwell-diffusion physics (Fig. S9(b), blue dots vs grey dashed). The half-attenuation point of $A_{\text{num}}(f)$ lies in the same decade as the fitted phenomenological corner $f_c = 99.0$ Hz. Since both $A_{\text{num}}(f)$ and the $K_1(f)$ sigmoid are set by the same dimensional combination $\mu_r \sigma d^2 \omega$ of the lamination, the corner-frequency location is fixed by this shared material time scale rather than being an independent fit degree of freedom; the two mechanisms nonetheless differ in magnitude (paragraph below).

Depth gap is real and informative. The fitted sigmoid $K_1(f)/K_1^{\text{low}}$ (Fig. S9(b), orange solid) decays significantly faster than $A_{\text{num}}(f)$ beyond the corner: at 200 Hz the numerical attenuation is $A_{\text{num}} = 0.91$ while the fitted ratio is $K_1(f)/K_1^{\text{low}} = 0.09$. A log-log regression gives $K_1(f)/K_1^{\text{low}} \approx A_{\text{num}}(f)^{23.6}$, so the fitted sigmoid is the linear-theory envelope raised to roughly the 24th power. No single effective permeability brings the analytical envelope into shape agreement: minimising $\sum_f [\log A_{\text{analytical}}(f, \mu_r) - \log(K_1(f)/K_1^{\text{low}})]^2$ over μ_r lands at $\mu_r^{\text{eff}} \approx 2.3 \times 10^5$ ($15 \times$ the operating-knee value) with residuals up to +160 % at 200 Hz. The fitted-sigmoid steepness is therefore intrinsic to the model class, not a calibration of linear skin attenuation.

The operative mechanism: rotation-only FP over-predicts loop widening. Panel (c) of Fig. S9 resolves the mechanism by direct loop-area comparison. The uniform-field rotation-only FP at constant $K_1 = K_1^{\text{low}}$ (grey squares) over-predicts hysteretic loop area growth with frequency: from $323 \text{ T} \cdot \text{A/m}$ at 30 Hz to $929 \text{ T} \cdot \text{A/m}$ at 200 Hz ($\sim 2.9 \times$). The 1D Maxwell-coupled simulation at the same K_1 (blue circles) adds 5–24 % widening on top of that. Maxwell coupling, therefore, pushes the model loop area in the *wrong* direction relative to the fit. The uniform-field FP at the main-text fitted $K_1(f)$ (orange triangles) lands at the measurement-matched growth ($\sim 1.7 \times$ over the same range). The fitted $K_1(f)$ reduction is therefore the rotation-only FP's compensation for its own over-prediction of high-frequency loop widening, not a direct shadow of skin attenuation. The corner frequency coincides with $d/\delta = 1$ because both arise from the same dimensional material time scale, but different mechanisms set the magnitudes.

Honest interpretation and downstream work. The $K_1(f)$ sigmoid is the rotation-only FP's effective surrogate for several missing physics: (i) the model class's intrinsic over-prediction of eddy-driven loop area growth at constant K_1 , which is the dominant contribution; (ii) the nonlinear $\mu_r(H, B_{\text{peak}})$ over the cycle, where the actual permeability of HGO Fe-Si rises from ~ 1 at saturation to $\sim 10^5$ at the knee so that a constant operating-knee μ_r understates the effective attenuation, (iii) the increasing fraction of 180°- and 90°-wall switching at high f , which the rotation-only Fokker–Planck operator cannot represent (the Daniel-style multi-scale ablation in Supplementary Table S8, row 6, separately rules out wall mechanics as the dominant residual source). A full quantitative reconciliation requires Maxwell-coupled FEM with the present rate-independent constitutive layer applied per lamination slice and a nonlinear $\mu_r(H)$ update at every time step. That treatment lies outside the scope of a constitutive modeling paper and is the natural next step for downstream device simulation work. Reproducible from `examples/04_maxwell_skin_depth_1d.py` in the open-source companion repository.

Supplementary Tables

Parameter	RD	TD
M_s [A/m]	1.628×10^6	1.225×10^6
a [A/m]	47.4	13.0
α	4.87×10^{-5}	6.28×10^{-3}
K_u [J/m ³]	564	4500
K_1 [J/m ³]	0 (pinned)	0 (pinned)
τ_R [s] (frozen, anhysteretic limit)	10^{-3}	10^{-3}
D (frozen)	10^3	10^3
RMSE [T]	0.033	0.070
R^2	0.9994	0.9945

Table S1. Fitted parameters for the Baghel 27M-OH 1 Hz quasi-static loops. Both directions fitted independently at frozen $\tau_R D = 1$ (canonical anhysteretic limit). Values from `examples/data/baghel27M_OH_quasi_static_1hz.json`.

Supplementary Figures

Per-frequency knob	50 Hz	100 Hz	150 Hz	200 Hz
$K_1(f)$ [J/m ³]	1413	798	384	197
$\tau_R(f)$ [ms]	21.1	18.4	19.4	18.3
<i>TD Preisach hysteron 1 (cubic $\langle 001 \rangle \rightarrow \langle 100 \rangle$ switch)</i>				
$H_{\text{pin,TD1}}(f)$ [A/m]	193	268	253	234
$\sigma_{\text{pin,TD1}}(f)$ [A/m]	119	102	66	106
$w_{\text{pin,TD1}}(f)$	0.47	0.41	0.25	0.41
<i>TD Preisach hysteron 2 (cubic $\langle 100 \rangle \rightarrow \langle 110 \rangle$ switch)</i>				
$H_{\text{pin,TD2}}(f)$ [A/m]	239	263	282	301
$\sigma_{\text{pin,TD2}}(f)$ [A/m]	118	117	120	115
$w_{\text{pin,TD2}}(f)$	0.50	0.47	0.43	0.49
TD asymmetric shift $H_{\text{shift}}(f)$ [A/m]	286	355	341	299
Per-loop RMSE RD [T]	0.062	0.066	0.152	0.129
Per-loop RMSE TD [T]	0.039	0.033	0.031	0.042

Table S2. Per-frequency knobs and per-loop residuals for the Hi-B 27M-OH joint S^2 fit. The four shared FP+Preisach blocks have one set of frequency-resolved values each. $K_1(f)$ is computed from the three intrinsic constants (K_1^{low}, f_c, n) in main-text Table 1. The two TD Preisach hysterons model the discrete cubic $\langle 001 \rangle \rightarrow \langle 100 \rangle$ and $\langle 100 \rangle \rightarrow \langle 110 \rangle$ switching events, the rotation-only FP operator cannot resolve. Their amplitudes are locked to $M_{\text{pin}} = M_s/\sqrt{2}$, the geometric cubic projection on the lab-TD axis¹⁰. Values from `examples/data/baghel27M_OH_unified_fit.json`.

Material	B_{peak} [T]	RMSE [T]	R^2	RMSE/ B_{peak}
X10CrAlSi7 (stainless Fe-Cr-Al-Si)	1.83	0.046	0.9992	2.5 %
HLB8 (high-strength low-alloy)	2.11	0.066	0.9987	3.1 %
C35 (medium-carbon structural)	1.89	0.067	0.9984	3.5 %
S235JR (structural steel)	1.63	0.070	0.9977	4.3 %
X20Cr13 (martensitic stainless)	1.88	0.075	0.9981	4.0 %
HC260LA (cold-rolled HSLA)	2.18	0.078	0.9983	3.6 %
isovac HP1400-100A (electrical NGO)	2.02	0.087	0.9976	4.3 %
Mean across 7 rings		0.070		3.6 %

Table S3. Isotropic-mode fits for all seven non-oriented ring cores measured on the in-house MagHyst[®] Modular system, with $K_u = K_1 = 0$ enforced. The fitted K_u and K_1 remain within 10 J/m³ of zero when freed in a double-check fit, confirming that the joint-fit machinery does not invent anisotropy on isotropic samples. All seven rings pass the $\text{RMSE}_B/B_{\text{peak}} < 5\%$ acceptance gate with mean relative residual 3.6 % across the campaign. The three rings plotted in main-text Fig. 5 (X10CrAlSi7, X20Cr13, S235JR) were selected as the chemically most distinct triplet (Fe-Cr-Al-Si vs. martensitic stainless vs. low-carbon structural). Per-ring fitted parameter sets for all seven materials are available as JSON files in `examples/data/fitted_params/`.

Joint Li & Luo amplitude-sweep parameters (B30P105 HGO, 20 levels at 50 Hz)	Value
M_s [A/m]	1.551×10^6
a [A/m]	2.25
$\alpha_{ }$	1.08×10^{-5}
$\alpha_{\perp}$	4.27×10^{-4}
K_u [J/m ³]	132
K_l [J/m ³]	220
σ_e [°]	1.08
D	1915
τ_{hyst} [ms]	2.65
k_{eddy}	4.1×10^{-3}
k_{excess}	3.78×10^{-2}
mean RMSE [T] over 20 levels	0.515
worst level RMSE (1.7 T) [T]	1.024
best level RMSE (0.1 T) [T]	0.121

Table S4. Joint single-parameter-set fit for the Li and Luo 2024 amplitude sweep¹¹, with the FULL-physics parameter set shared across all 20 amplitude levels at 50 Hz (80-trial Optuna TPE Bayesian optimization, in-memory storage). The joint mean RMSE exceeds 0.5 T and grows above $B_{\text{peak}} \approx 1.2$ T. This identifies the Langevin scale parameter a as amplitude-dependent in practice. The per-amplitude refit (only a allowed to track B_{peak}) reaches mean $\text{RMSE}_B \approx 0.04$ T across the same 20 levels (see main-text amplitude-transferability discussion). Values from `examples/data/b30p105_joint_S2.json`.

Frequency	RD RMSE [T]	RD R^2	TD RMSE [T]	TD R^2
1 Hz (anhysteretic)	0.033	0.9994	0.070	0.9945
50 Hz	0.186	0.9781	0.079	0.9920
100 Hz	0.171	0.9806	0.039	0.9983
150 Hz	0.173	0.9814	0.048	0.9976
200 Hz	0.151	0.9862	0.074	0.9930

Table S5. Per-loop residuals for the Baghel 27M-OH direction-resolved S^1 baseline fit (Table S1). Reproduced here for completeness. The joint S^2 fit residuals are reported in the main Table 2.

Parameter	Mean (over multi-condition fit)	Std. dev. σ_{ϑ_i}
M_s	1.411×10^6	4.7×10^4 (3.3 %)
a	0.75	0.04 (5.3 %)
α	7.28×10^{-5}	3.4×10^{-6} (4.7 %)
K_u	349	14 (4.0 %)
K_l^{low}	1619	83 (5.1 %)
D	4892	192 (3.9 %)
$\tau_R(f)$	20.2 to 24.9 ms	0.5 to 1.0 ms (~ 4 %)

Table S6. Practical identifiability summary (Fisher information diagonal, equation (20) of main text), evaluated at the converged optimum of the joint S^2 Baghel fit. All standard deviations $\sigma_{\vartheta_i}/\vartheta_i \ll 1$, confirming that the seven material parameters plus per-frequency τ_R are practically identifiable from the eight-loop multi-condition dataset.

Material (Ring)	J-A	Chwastek	D-D-D	DTHM Simp.	Preisach	P-I	This work
X10CrAlSi7	0.023	0.008	0.006	0.024	0.126	0.160	0.046
HLB8	0.041	0.041	0.041	0.072	0.166	0.185	0.066
C35	0.045	0.046	0.055	0.054	0.122	0.170	0.067
S235JR	0.028	0.029	0.023	0.020	0.111	0.208	0.070
X20Cr13	0.032	0.047	0.040	0.042	0.121	0.186	0.075
HC260LA	0.054	0.050	0.057	0.093	0.176	0.195	0.078
Isovac HP1400-100A	0.047	0.056	0.054	0.067	0.155	0.196	0.087
Mean (7 rings)	0.039	0.040	0.039	0.053	0.140	0.186	0.070

Table S7. Per-ring cross-model comparison. RMSE [T] of each baseline model fitted independently to each of the seven ring-core measurements, versus the present model fitted in its rate-independent isotropic reduction ($K_u = K_1 = 0$, $D \rightarrow 0$). Each baseline is given its own optimal parameter set on each ring (maximally favorable to the baselines). The present model uses the corresponding row of Table S3. On single-loop isotropic samples, the seven approaches all perform within measurement noise. The per-loop ranking is therefore not the discriminator the present work addresses, which is multi-condition transferability of a single shared parameter set (main-text Table 3 and Fig. S4). Numerical entries are reproduced from the per-ring baseline campaign in `examples/data/baghel27M_OH_per_freq_baselines.json`. Mean residuals across the seven rings are listed in the main-text Table 3 (NGO ring rows).

Fit variant	Rolling direction (RD)				Transverse direction (TD)				joint
	50	100	150	200	50	100	150	200	
This work ($K_1(f)$, Fig. 3)	0.062	0.066	0.152	0.129	0.039	0.033	0.031	0.042	0.069
Same S^2 base, static K_1	0.076	0.123	0.385	0.368	0.205	0.140	0.131	0.213	0.205
Per-direction S^2 , K_1 free (perbasin)	0.068	0.068	0.138	0.108	0.032	0.029	0.028	0.031	0.063
Per-direction S^2 , $K_1 = 0$ pinned	0.303	0.324	0.230	0.248	0.228	0.177	0.160	0.197	0.233
Per-direction S^1 DTHM-S (Table 3 col)	0.070	0.073	0.059	0.068	0.083	0.080	0.073	0.097	0.075
Daniel 6-domain wall-fraction ¹⁰	0.64	0.73	0.72	0.81	0.25	0.25	0.27	0.34	0.50


Table S8. Unification ablation on Hi-B 27M-OH. Per-loop RMSE_B [T] for six configurations of the vector Fokker-Planck model. The skin-depth $K_1(f) = K_1^{\text{low}} / (1 + (f/f_c)^n)$ phenomenology (this work, row 1) reduces the joint mean from 0.205 T (static K_1 , row 2) to 0.069 T by closing the high-frequency RD residual ($0.385 \rightarrow 0.152$ T at 150 Hz and $0.368 \rightarrow 0.129$ T at 200 Hz) and the TD inner-knee structure ($0.172 \rightarrow 0.036$ T direction mean). Row 3 is the per-direction S^2 fit with the same architecture (all parameters free per direction, no unification constraint). It reaches mean RMSE 0.063 T versus 0.069 T for this work, so the cross-direction unification costs $\sim 10\%$ of mean RMSE relative to the per-direction floor. The high-frequency RD residuals (0.138 and 0.108 T at 150 and 200 Hz) are nearly the same in row 3 as in row 1, confirming that they reflect a rotation-only constitutive limit and not a unification penalty. Row 4: strict $S^2 \rightarrow S^1$ reduction with $K_1 = 0$ pinned degrades both directions, showing K_1 in the FP equilibrium is not decorative. Row 5: DTHM-Simplified baseline (reproduced from main-text Table 3). Row 6: Daniel six-domain wall-fraction extension¹⁰ (implemented as `hysterpy/daniel_multiscale.py`) replaces the orientation density on S^2 with six thermally-activated cubic-axis populations $f_i(t)$. This admits literature single-crystal $K_1 = 4.8 \times 10^4$ J/m³ by construction but reaches joint RMSE 0.50 T ($\sim 3\times$ the rotation-only FP residual of 0.17 T), empirically ruling out missing domain-wall mechanics as the dominant residual source and isolating the spatial-field (skin-depth) assumption as the operative limitation at $f \geq 150$ Hz. Two further variants reported in the text but not tabulated here: free Bertotti back-field $\rho_{\text{eddy}}, \rho_{\text{excess}}$ drive to zero and recover the static- K_1 minimum bit-for-bit. The per-direction static- K_1 fit reaches 0.201 T versus 0.205 T for this work with static K_1 . Visualisation in Fig. S2. Parameters from this work in `examples/data/baghel27M_OH_unified_fit.json`.

Table S9. Cross-validation summary for the joint S^2 fit on Hi-B 27M-OH. All variants share the same S^2 Fokker–Planck simulator and the same 8-loop measurement set. Only the joint-cost composition (which frequencies are held in versus held out) or the $K_1(f)$ functional form changes. The forward-extrapolation row fits 50 Hz only and predicts the higher frequencies via the scaling laws of the main text. The reverse-extrapolation rows hold the shared base fixed at this work and refit only the four $K_1(f)$ constants on the 100 to 200 Hz loops. The leave-one-out row is the 4-fold version of the same reverse-extrapolation protocol. The free-parameter count for the all-8-loop joint fit (38) comprises the 22 shared parameters of Methods plus the per-frequency TD Preisach-overlay arrays; the forward-extrapolation row fits a single frequency, so its overlay is not expanded per frequency and the count stays at 22.

Variant	Held-in / functional form	Fit cost [T]	Free params
This work (joint fit)	all 8 loops, M_s free, sigmoid $K_1(f)$	0.069	38
Forward-extrapolation	fit 50 Hz, predict others	0.043 [†]	22
Reverse-extrapolation	fit 100 to 200 Hz, predict 50 Hz, sigmoid	0.072	4
Reverse-extrap. (Hill)	same split, Hill function rising	0.092	4
Leave-one-out (4 folds)	sigmoid, held-out frequency rotated	0.062–0.077	4


[†] 50-Hz-only fit cost. The other three frequencies are predicted via scaling laws.

(a) Crystal-axis frame



Goss spread $\sigma_e = 0.15$ rad

(b) Sample-axis (RD/TD) fields



$$\begin{aligned}
 \angle(\mathbf{H}_{RD}, [\mathbf{001}]) &= 0^\circ \\
 \angle(\mathbf{H}_{TD}, [\mathbf{001}]) &= 90^\circ \\
 \angle(\mathbf{H}_{TD}, [\mathbf{100}]) &= 45^\circ \\
 \angle(\mathbf{H}_{TD}, [\mathbf{010}]) &= 45^\circ
 \end{aligned}$$

Figure S1. Lab-frame versus crystal-frame geometry for the Hi-B 27M-OH Goss-textured sample. Left: the cubic crystallographic axes $[100]$, $[010]$, $[001]$ of a representative grain. The Goss texture orients $[001]$ along the lab rolling direction with a residual misorientation spread $\sigma_e = 0.053$ rad (red cone). Right: the lab-frame field directions overlaid on the same crystal axes. The rolling-direction field $\mathbf{H}_{RD}$ is co-aligned with the cubic easy axis $[001]$ (blue arrow). The transverse-direction field $\mathbf{H}_{TD}$ is aligned with the cubic $[110]/\sqrt{2}$ direction (green arrow), at 90° to the easy axis and 45° to each of the two perpendicular cubic minima $[100]$ and $[010]$. The asymmetry is the physical reason why the unified parameter set cannot simultaneously saturate RD and TD at the same applied field amplitude: along RD the bulk magnetisation saturates at $M_s \langle \cos \theta_g \rangle \approx 0.99 M_s$ for the fitted Goss spread, requiring a sharply-knee'd $M-H$ characteristic to fully reach this cap at the measured $H_{\text{peak}} \approx 250$ A/m. Along TD, the magnetization rotates through the $\langle 100 \rangle$ and $\langle 010 \rangle$ cubic wells before aligning with the applied field, thereby lowering the effective saturation field and matching the measurement at a lower H_{peak} . The same scalar M_s produces a small RD undershoot at peak field and an essentially exact TD match (main-text Fig. 3), the latter dominating the joint cost weighting.




Figure S2. Visual proof of the unification-cost ablation in Supplementary Table S8. For each measurement loop, this work (black dashed, identical to main-text Fig. 3), the per-direction S^2 fit with K_1 free (orange dotted), and the per-direction S^2 fit with $K_1=0$ pinned (blue dash-dot) are overlaid on the measurement (grey solid). Per-panel RMSE annotated in the upper-left of each axis. This work and the per-direction-free curves are visually indistinguishable on RD across all frequencies. This work is already at the direction-resolved optimum. The $K_1=0$ trace widens visibly on RD 150 and 200 Hz, closing roughly 30% of the high-frequency RD residual, but at the cost of losing the low-frequency RD knee shape and the TD inner-knee structure entirely. The Bertotti additive back-field configuration (free $\rho_{\text{eddy}}, \rho_{\text{excess}}$) is not plotted as a separate trace because the optimizer drives both constants to zero and recovers exactly the static- K_1 curve.




Figure S3. 1 Hz quasi-static fit on the Baghel 27M-OH RD and TD loops, with τ_R frozen at its anhysteretic value (equation (2) of main text gives $\omega\tau_R \ll 1$ at 1 Hz). (M_s, a, α, K_u) fitted independently per direction and reported in Table S1. Symbols: measurement (black circles), model (orange squares). $R^2 = 0.9994$ for RD and 0.9945 for TD.




Figure S4. Cross-model overlay on the Li and Luo B30P105 high-permeability grain-oriented steel at $B_{\text{peak}} \in \{0.5, 1.0, 1.5, 2.0\} \text{ T}$ (four columns), 50 Hz. Each row is one of the six baselines defined in main-text Theory: Jiles–Atherton¹, Chwastek modified J–A², D’Aloia–Di Francesco–De Santis³, DTHM–Simplified (this work), Preisach⁴, and Prandtl–Ishlinskii⁵. Each baseline is fitted *independently* to each amplitude with its own optimal parameter set (maximally favorable to the baselines). The present model uses a single shared parameter set that is transferred across all 20 amplitudes and is therefore not retuned for each amplitude. The per-amplitude RMSE annotated in each panel is the value contributing to the means in the main-text Table 2. Visual inspection makes the amplitude-transferability narrative concrete: the per-loop fit of every J–A-family baseline degrades smoothly with B_{peak} , while Preisach and P–I exhibit their characteristic saturation artifacts at $B_{\text{peak}} > 1.5 \text{ T}$.




Figure S5. Cross-model overlay on the X10CrAlSi7 ring-core isotropic baseline at $B_{\text{peak}} = 1.83 \text{ T}$, $\sim 1.7 \text{ Hz}$ quasi-static MagHyst[®] excitation. Panel (a): all six J–A-family and operator baselines plus the present model overlaid on the measured loop (grey). Panel (b): per-model RMSE bar chart. On a single isotropic loop, the J–A family, the DTHM-Simplified, and the present model (rate-independent reduction with $K_u = K_l = 0$, $D \rightarrow 0$) all perform within measurement noise. Preisach and P–I show the characteristic constant-amplitude saturation artifact. The cross-model differentiation reported in the main-text Table 2 arises only with multi-condition data (Baghel RD/TD, Li and Luo amplitude sweep), where baseline retuning per loop is required. In contrast, the present model uses a single shared set. Reproducible end-to-end from the open-source bundled ring CSV via `examples/01_fit_all_models_demo.py`.




Figure S6. Bimodal cost surface in the joint S^2 Baghel fit. Each panel overlays the measurement (grey) with two distinct local minima of the joint arc-length-RMSE cost over the 17-dimensional extension of the model (anisotropic Bingham ODF, per-frequency τ_R , per-direction K_u , amplitude-coupled a , and K_2 second cubic invariant). Orange dashed = warm-start FP-only (this work, main-text Fig. 3. Mean $\epsilon_A = 0.22$ T balanced across RD and TD). Red dotted = cold-start basin B discovered by an independent TPE search seeded with 80 random Sobol points. $\sigma_{\parallel} = 12.6^\circ$, $\sigma_{\perp} = 1.5^\circ$ (highly anisotropic Goss spread), $K_1 = 2.4 \times 10^4$ J/m³, $K_u^{\text{TD}}/K_u^{\text{RD}} = 3.8$. Mean $\epsilon_B = 0.55$ T (RD = 0.88 T, TD = 0.22 T). Basin B improves TD by 13 % but breaks RD by a factor of 4. A third search (basin C, not plotted here for clarity) constrained the Bingham distribution to be physically correct ($\sigma_{\parallel} \leq \sigma_{\perp}$) and forced $K_1 \in 1 \times 10^4$ to 5×10^4 J/m³ within the Fe-Si single-crystal literature range. The resulting fit selects $\sigma_{\parallel} = 2.9^\circ$, $\sigma_{\perp} = 18.3^\circ$, $K_1 = 2.96 \times 10^4$ J/m³ but with $\epsilon_C = 1.11$ T (RD = 1.37 T, TD = 0.84 T): the cubic [100]/[010] easy wells perpendicular to the lab-RD field trap magnetisation at low RD excitation, shifting the RD coercivity by ~ 0.3 kA/m and shrinking apparent saturation. The three basins together establish that the one-shared-parameter-set model class possesses a unique global minimum (the FP-only fit). That physically correct Bingham parameters lie on a higher-cost ridge. The bimodality is therefore a structural property of the model class rather than a search artifact, motivating the natural extensions discussed in the main-text Discussion (two-component ODF, domain-wall pinning hysteron stack).




Figure S7. Cross-model RMSE bar chart across the seven NGO ring-core measurements. Panel (a) groups the seven baseline models per ring. Panel (b) shows the mean per model, averaged over rings. The narrow spread visible in (b) confirms the per-loop NGO observation made in Fig. S5: on isotropic single loops, the seven approaches all perform within measurement noise, with the present rate-independent reduction (DTHM / “This work”, last bar) comparable to the Jiles family. The cross-condition spread on Baghel RD/TD and Li and Luo amplitudes is reported in the main-text Table 3. This figure addresses only the ring-core column. Numerical entries reproduced from `examples/data/baghel127M_OH_per_freq_baselines.json`.




Figure S8. Twenty-seed cold-start multistart on the joint S^2 cost surface for Hi-B 27M-OH. (a) Parametric scatter of the best-of-seed sigmoid constants (K_1^{low}, f_c) on log-log axes. Basin assignments are color-coded, and the main joint fit is marked with a star. (b) Resulting $K_1(f) = K_1^{\text{low}} / (1 + (f/f_c)^n)$ trajectories for each seed in the range 20 to 400 Hz, with the main-text joint-fit trajectory overlaid as a black dashed line. Basin A (blue, $N = 14$) brackets the main-fit curve tightly. Basin B (red, $N = 2$) drops sharply at $f > 100$ Hz from a single-crystal-magnitude initial value. Basin B' (orange, $N = 4$) maintains a low, near-constant K_1 across the measurement band. Grey dotted lines mark the four measurement frequencies 50 Hz, 100 Hz, 150 Hz and 200 Hz.




Figure S9. 1D Maxwell-coupled analysis of the cubic-anisotropy sigmoid corner frequency for Hi-B 27M-OH.

(a) Steady-state internal field profile $H_{\text{internal}}(z)/H_{\text{lab,peak}}$ at the surface-field peak, across the lamination cross-section $z/d \in [-0.5, +0.5]$, for two representative frequencies. At 50 Hz ($d/\delta = 0.75$, blue circles), the interior field is nearly uniform. At 200 Hz ($d/\delta = 1.5$, red circles) the centre-line value drops to $\sim 0.84 H_{\text{lab,peak}}$, the signature of incipient skin shielding. Markers from the 1D Maxwell-coupled simulation. Lines connect grid points. (b) Frequency dependence of the bulk-averaged attenuation factor $A(f) = \max_t |\langle H(z, t) \rangle_z| / \max_t |H_{\text{lab}}(t)|$. Blue circles: numerical solution of the coupled magnetic-diffusion + Fokker-Planck problem with the same $K_1 = K_1^{\text{low}} = 1619 \text{ J/m}^3$ in every z -slice. Grey dashed: analytical linear-theory envelope $|\tanh(\gamma d/2)/(\gamma d/2)|$ at the operating-knee permeability $\mu_r = 1.35 \times 10^4$. Orange solid: fitted phenomenological $K_1(f)/K_1^{\text{low}}$ sigmoid from this work in the main text. The vertical dotted line marks the fitted corner $f_c = 99.0 \text{ Hz}$. The numerical Maxwell solve matches the analytical envelope to better than 1 % and places its half-attenuation point in the same decade as the fitted corner. Both share the same dimensional scale $\mu_r \sigma d^2 \omega$, so the corner frequency of the K_1 sigmoid is locked by Maxwell physics. The fitted sigmoid decays much more steeply than $A(f)$ in depth (see panel (c) for the mechanism). (c) Per-frequency hysteresis loop area $\oint H dB$ in three configurations: uniform-field rotation-only Fokker-Planck at constant $K_1 = K_1^{\text{low}}$ (grey squares), 1D Maxwell-coupled FP at the same constant $K_1 = K_1^{\text{low}}$ (blue circles), uniform-field FP at the main-text fitted $K_1(f)$ (orange triangles). The Maxwell-coupled simulation broadens the uniform-field constant- K_1 result, in the wrong direction relative to the fit. The fitted $K_1(f)$ reduces the model loop area to match measurement-driven growth ($\sim 1.7\times$ over 30 to 200 Hz vs $\sim 2.9\times$ at constant K_1). The $K_1(f)$ sigmoid is therefore the rotation-only FP's compensation for its intrinsic over-prediction of high-frequency loop widening (Supplementary Note S6). Reproducible from `examples/04_maxwell_skin_depth_1d.py`.




Figure S10. Code-equivalence consistency check against the anhysteretic limit of the Jiles–Atherton anisotropic model (Ramesh, Jiles & Bi 1997, Eq. 5)¹². This is a consistency check, not a fit: both reference and present-work traces consume the *same* literature parameter set $K_1 = 4.8 \times 10^4 \text{ J/m}^3$, $M_s = 1 \times 10^6 \text{ A/m}$, $a = 1000 \text{ A/m}$, $\alpha = 10^{-3}$, and they must agree if the S^2 Fokker–Planck operator correctly reduces to the J–A anisotropic anhysteretic in the $D\tau = 1$ canonical equilibrium limit. Solid lines: anhysteretic initial-magnetization curves along the three crystallographic directions [100], [110], [111], regenerated from Ramesh’s published Eq. 5 self-consistent fixed-point form. Dashed lines with markers: this work, computed directly from the Boltzmann equilibrium of the S^2 free-energy landscape with self-consistent mean-field iteration. Only the lab-frame field axis $\hat{\mathbf{n}}_f \in \{[100], [110], [111]\}$ rotates between traces. At $H = 80 \text{ kA/m}$ the two formulations agree within 1 % for all three crystal-field directions, confirming the operator-reduction step. Fitting validation against measured data is reported separately for the Hi-B 27M-OH RD/TD multi-frequency set (main-text Fig. 3, joint $\text{RMSE}_B = 0.069 \text{ T}$), the Li and Luo B30P105 amplitude sweep (main-text Fig. 4), and the seven NGO ring cores (main-text Fig. 5). Reproducible from `scripts/plot_angle_resolved_validation.py`.




Figure S11. Dynamic versus quasi-static response of the rotation-only Fokker-Planck constitutive law under multi-harmonic excitation. Practical magnetic loads are rarely pure sinusoids. Switched-mode drives, transformer over-excitation, and saturating inductors all produce odd-harmonic distortion of the field waveform. The X10CrAlSi7 ring core (isotropic fit, Table S3) is driven by $H(t) = H_1 \sin(\omega t) + H_2 \sin(3\omega t)$ with $H_1 = 8000$ A/m, $H_2 = 4800$ A/m (60 % third-harmonic content), $\omega = 2\pi f$. **(a)** The non-monotone field crosses zero twice per fundamental half-cycle and produces four local field extrema per period at $H \approx \pm 9.4$ kA/m with intermediate “valleys” at $H \approx \pm 3.2$ kA/m (50 Hz shown over two fundamental periods). **(b)** Flux-density response: $B_{\text{peak}} = 1.421$ T, with the third-harmonic “valley” producing two flat-top plateaus per half-cycle. **(c)** B - H trajectory at 50 Hz (dynamic, blue solid, $\omega\tau_R = 0.14$ on the fundamental, 0.42 on the third harmonic) overlaid on the same excitation at 1 Hz (quasi-static, black dashed, $\omega\tau_R = 2.8 \times 10^{-3}$, fully rate-effect-free). Both trajectories reach the same $B_{\text{peak}} \approx 1.42$ T at the operating point: the rate-dependence does *not* alter the loop tips, it widens the loop body. The QS trace appears as a near-singular curve because for this soft non-oriented material the rate-independent backbone is intrinsically narrow ($H_c^{\text{QS}} \sim 25$ A/m, $W_h^{\text{QS,pure}} \approx 0.13$ kJ/m³ at the same H_{peak} with a clean pure-sinusoidal excitation). The inner reversals at $H \approx \pm 3.2$ kA/m fall *outside* the metastability region, so the QS trace retraces the equilibrium curve through them without forming minor loops. The two minor-loop lobes near $(H, B) \approx (\pm 3 \text{ kA/m}, \pm 1 \text{ T})$ visible at 50 Hz are therefore entirely rate-dependent path-memory artifacts, not present in the rate-independent limit. The closed-loop area is the hysteresis loss density per cycle $W_h = \oint H dB$. The dynamic case dissipates $W_h^{\text{dyn}} = 13.6$ kJ/m³, the quasi-static reference only $W_h^{\text{QS}} = 0.32$ kJ/m³, a 42× rate factor, dominated by the third-harmonic component which sits in a regime ($\omega\tau_R \sim 0.4$) where the Fokker-Planck relaxation cannot follow the field. For reference, the pure-sinusoid $H_1 \sin(\omega t)$ baseline at 50 Hz (not plotted) gives $W_h = 5.0$ kJ/m³, so 60 % third-harmonic content inflates the dynamic loss by 2.7×. No re-fitting was performed: the parameter set is identical to Table S3 row X10CrAlSi7, fitted only against the in-house pure-sinusoidal MagHyst[®] measurement at 1.7 Hz. Reproducible from `scripts/plot_harmonic_excitation.py`.




Figure S12. External cross-validation against the open EMPIR 19ENG06 HEFMAG dataset¹³ (INRiM, PTB, NPL, CMI). Two non-oriented Fe-Si 0.2 mm loops from the HEFMAG release are fit by the present DTHM-S² model in its isotropic reduction ($K_u = K_1 = 0$); Epstein 700-turn IEC 60404-2 measurements at 50 Hz. **NO L1:** $H_{\text{peak}} = 59.1 \text{ A/m}$, $B_{\text{peak}} = 0.50 \text{ T}$, $\text{RMSE}_B = 61 \text{ mT}$. **NO L2:** $H_{\text{peak}} = 109.8 \text{ A/m}$, $B_{\text{peak}} = 1.00 \text{ T}$, $\text{RMSE}_B = 133 \text{ mT}$. Both fits are produced *blind*: no prior knowledge of the HEFMAG sample composition or processing is used. An automatic Conde Garrido (2025) initialiser¹⁴ extracts seed parameters (M_s, a, α) from a Langevin–Weiss fit to the horizontal-average of the measured loop and the J–A pinning coefficient k from a closed-form expression at the coercive point ($H_c = 29.7 \text{ A/m}$ for NO L1, 41.8 A/m for NO L2). The numba-parallel DE then refines a warm-started population around the seed with $\pm 50\%$ bounds. The blind seed alone already yields RMSE of 169 mT (NO L1) and 330 mT (NO L2); DE refinement reduces these by $2.5\text{--}3\times$ to the values reported. The fitted $M_s^{\text{NO L1}} = 7.3 \times 10^5 \text{ A/m}$ and $M_s^{\text{NO L2}} = 1.07 \times 10^6 \text{ A/m}$ depend on loop amplitude, consistent with the per-amplitude effective Langevin scale documented in main-text Fig. 4 on the Li and Luo B30P105 dataset. **Known limitations not plotted:** (i) NO L3 ($H_{\text{peak}} = 1546 \text{ A/m}$, $B_{\text{peak}} = 1.51 \text{ T}$): isotropic Langevin over-rounds the fat saturated loop, $\text{RMSE}_B \approx 500 \text{ mT}$ persists after DE refinement. A polycrystal extension would relax this, but it is structurally inappropriate for an isotropic NO sample. (ii) HEFMAG GO Hi-B loops at $H_{\text{peak}} \in \{25.6, 46.7, 120.9\} \text{ A/m}$: we ran a joint 3-amplitude textured-polycrystal fit warm-started from the Baghel 27M-OH unified parameters (main Fig. 3), with wide bounds and a 80-generation multi-fidelity DE (30 coarse generations at $N_\theta = 10, N_\phi = 20$ followed by 50 fine generations at $N_\theta = 16, N_\phi = 32$). The fit converges to a joint mean $\text{RMSE}_B = 791 \text{ mT}$ with per-loop residuals 594 mT, 703 mT and 1076 mT, plateauing at the lower bound of $M_s = 1.30 \times 10^6 \text{ A/m}$ with fitted $K_1 = 19.5 \text{ J/m}^3$ (versus Baghel’s fitted 1619) and $\sigma_e = 0.026 \text{ rad}$ (versus Baghel’s 0.053). This plateau identifies HEFMAG GO as requiring two further components that are already implemented in the model but not exercised in the search above: a Preisach hysteron overlay (cf. the TD overlay in main Fig. 3) to capture the discrete 90° -wall switching of the very-square loop ($B_r/B_{\text{sat}} > 0.9$ at $H_{\text{peak}} = 26 \text{ A/m}$), and an anisotropic Bingham orientation distribution with $\sigma_{\parallel} \neq \sigma_{\perp}$ (function `_sample_grains_bingham` in `hysterpy.dthm_s2`) to collapse grains tighter onto the easy axis in the transverse direction. The activation of both within the joint DE bounds will be addressed in the next manuscript revision. It constitutes a direct demonstration that the modular DTHM-S² formulation reaches a regime that scalar hysteresis baselines (J–A, Preisach, D–D–D, PI) structurally cannot, because they carry no grain-orientation distribution and no compositional overlay mechanism. The fits reported here are reproducible from `scripts/fit_hefmag_blind_init.py` (NO) and `scripts/fit_hefmag_go_joint_fast.py` (GO). The blind-init helper is bundled as `hysterpy.blind_init` in the open-source release.




Figure S13. Validation of the LLG-direction infrastructure on cubic-anisotropy single-grain reference physics. The `hysterpy.llg_grain` module released with this work solves the Landau–Lifshitz–Gilbert equation for a per-grain three-vector magnetisation with cubic K_1 anisotropy from the Fe-3%Si single-crystal free energy. **(A)** Cubic anisotropy energy $U(\hat{\mathbf{m}})/K_1$ sampled on the unit sphere (θ, ϕ) at $K_1 = 35 \text{ kJ/m}^3$ (single-crystal Fe-3%Si literature value¹⁵). The four equatorial and two polar $\langle 100 \rangle$ wells are resolved cleanly with the hard-axis $\langle 110 \rangle$ ridges at 45° between adjacent wells, in agreement with cubic symmetry. **(B)** Single-grain trajectories in the $(M_x/M_s, M_y/M_s)$ projection for five lab-field directions: 0° , 30° , 45° , 60° , and 90° relative to the easy axis. At 0° the trajectory collapses to a vertical reversal along $[001] \leftrightarrow [00\bar{1}]$ (field on easy axis switches the pole). At 90° the magnetization sweeps through the perpendicular cubic $\langle 100 \rangle$ wells. Intermediate angles fan smoothly between, with no switching kinks. **(C)** Stoner–Wohlfarth astroid validation: numerical LLG switching fields (red circles, swept over $\theta_e \in [5^\circ, 85^\circ]$) compared against the analytical SW curve $H_x = H_K \cos^3 t$, $H_y = H_K \sin^3 t$. Maximum deviation $\Delta H_{\text{sw}}/H_K = 5.4\%$, attributed to finite-rate ramp (8τ per half-cycle versus the SW assumption of infinitely-slow drive). The implementation is therefore numerically validated against analytical reference physics. **Physics caveat for the Baghel 45° application:** at $K_1 = 35 \text{ kJ/m}^3$ the cubic anisotropy field is $H_K = 2K_1/(\mu_0 M_s) \approx 32 \text{ kA/m}$, far larger than the operating drive $H_{\text{peak}} \approx 0.6 \text{ kA/m}$ of the Baghel 45° measurement. Coherent-rotation LLG therefore cannot reproduce the measured $B \sim 1.5 \text{ T}$ swing because magnetization reversal at this operating point is dominated by domain-wall motion rather than coherent rotation. The principled extension is therefore stochastic-LLG with Néel–Brown thermal activation between the cubic $\langle 100 \rangle$ wells, identified as the focus of the follow-on manuscript.

References

1. Jiles, D. C. & Atherton, D. L. Theory of ferromagnetic hysteresis. *J. Magn. Magn. Mater.* **61**, 48–60, DOI: [https://doi.org/10.1016/0304-8853\(86\)90066-1](https://doi.org/10.1016/0304-8853(86)90066-1) (1986).
2. Chwastek, K. Modelling of dynamic hysteresis loops using the Jiles-Atherton approach. *Math. Comput. Model. Dyn. Syst.* **15**, 95–105, DOI: <https://doi.org/10.1080/13873950802432016> (2009).
3. D'Aloia, A. G., Di Francesco, A. & De Santis, V. A novel computational method to identify/analyze hysteresis loops of hard magnetic materials. *Magnetochemistry* **7**, DOI: <https://doi.org/10.3390/magnetochemistry7010010> (2021).
4. Mayergoyz, I. D. *Mathematical Models of Hysteresis and Their Applications* (Elsevier, 2003), 2nd edn.
5. Krasnosel'skiĭ, M. A. & Pokrovskiĭ, A. V. *Systems with Hysteresis* (Springer-Verlag, 1989).
6. Ilmenauer Mechatronik GmbH. Mess- und Prüftechnik: MagHyst[®] modular (2023).
7. Zirka, S. E. *et al.* Generalization of the classical method for calculating dynamic hysteresis loops in grain-oriented electrical steels. *IEEE Transactions on Magn.* **44**, 2113–2126, DOI: <https://doi.org/10.1109/TMAG.2008.2000662> (2008).
8. Baghel, A. P. S., Gupta, A., Chwastek, K. & Kulkarni, S. V. Comprehensive modelling of dynamic hysteresis loops in the rolling and transverse directions for transformer laminations. *Phys. B: Condens. Matter* **462**, 86–92, DOI: <https://doi.org/10.1016/j.physb.2015.01.026> (2015).
9. Coffey, W. T., Kalmykov, Y. P. & Waldron, J. T. *The Langevin Equation: With Applications to Stochastic Problems in Physics, Chemistry and Electrical Engineering* (World Scientific, 2004), 2nd edn.
10. Daniel, L., Hubert, O., Buiron, N. & Billardon, R. Reversible magneto-elastic behavior: A multiscale approach. *J. Mech. Phys. Solids* **56**, 1018–1042, DOI: <https://doi.org/10.1016/j.jmps.2007.06.003> (2008).
11. Li, Y. & Luo, Z. A comparative study of J–A and Preisach improved models for core loss calculation under DC bias. *Sci. Reports* **14**, 9409, DOI: <https://doi.org/10.1038/s41598-024-55155-w> (2024).
12. Ramesh, A., Jiles, D. C. & Bi, Y. Generalization of hysteresis modeling to anisotropic materials. *J. Appl. Phys.* **81**, 5585–5587, DOI: <https://doi.org/10.1063/1.364843> (1997).
13. HEFMAG Consortium, INRiM, PTB, NPL & CMI. Wideband hysteresis and complex-permeability dataset for grain-oriented and non-oriented Fe-Si, amorphous and nanocrystalline ribbons (empir 19eng06), DOI: <https://doi.org/10.5281/zenodo.5910892> (2022).
14. Conde Garrido, J. M., Ugarte Valdivielso, J., Aizpurua, J. I., Barrenetxea Iñarra, M. & Silveyra, J. M. Blind efficient method for optimizing Jiles–Atherton model parameters. *IEEE Transactions on Magn.* DOI: <https://doi.org/10.1109/TMAG.2025.3632479> (2025).
15. Cullity, B. D. & Graham, C. D. *Introduction to Magnetic Materials* (Wiley-IEEE Press, 2008), 2nd edn.