

Fourier-Space Approach to Physics-Informed Magnetization Reconstruction from Nitrogen-Vacancy Measurements

Alexander Setescak^{1,2,*}, Florian Bruckner¹, Dieter Suess¹, Young-Gwan Choi^{3,4}, Hayden Binger³, Lotte Boer³, Chenhui Zhang⁵, Hyunsoo Yang⁵, Claire Donnelly³, Uri Vool³, and Claas Abert¹

¹Faculty of Physics, University of Vienna, 1090 Vienna, Austria

²University of Vienna, Vienna Doctoral School in Physics, 1090 Vienna, Austria

³Max Planck Institute for the Chemical Physics of Solids, 01187 Dresden, Germany

⁴Department of Physics, University of Ulsan, Ulsan 44610, Republic of Korea

⁵Department of Electrical and Computer Engineering, National University of Singapore, Singapore 117583, Singapore



(Received 1 April 2026; revised 28 July 2026; accepted 8 September 2026; published 14 October 2026)

Reconstructing magnetization textures from nitrogen-vacancy (NV) magnetometry stray-field measurements is a challenging, fundamentally ill-posed inverse problem, further complicated by the unknown effective distance between sensor and magnetic material. Here we show that incorporating a micromagnetic energy functional directly into the inversion filters out unphysical, high-energy configurations, while Fourier-space upward continuation of the stray field allows us to simultaneously fit the distance. Applied to measurements of the van der Waals ferromagnet $\text{Fe}_{3-x}\text{GaTe}_2$, it recovers an effective distance estimate of approximately 81 nm and low-energy configurations that reproduce the observed field. More broadly, embedding physics directly into the reconstruction turns ill-posed magnetic inverse problems into transparent, interpretable reconstructions, with applicability well beyond NV magnetometry.

DOI: 10.1103/nc5g-3yw9

Introduction— Many recent developments in stray-field measurement techniques have enabled the exploration of nanoscale magnetic textures. Among these, nitrogen-vacancy (NV) magnetometry [Fig. 1(a)] has emerged as a powerful tool: It uses nitrogen-vacancy centers in diamond as atomic-scale quantum sensors. The fluorescence (light emission) from these sensors changes based on their spin state, allowing for the highly sensitive and quantitative measurement of magnetic fields [1,2]. However, inferring the underlying magnetization from measured stray fields constitutes an ill-posed inverse problem [1,3,4]: infinitely many magnetization configurations can produce the same stray field [Fig. 1(b)]. This ambiguity is a hallmark of magnetic inverse problems, also posing significant challenges in the reconstruction of current densities from magnetic field maps [5]. Approaches to solving this inverse problem have evolved across different methodologies. Early work by Hug *et al.* [6] established quantitative magnetic force microscopy (qMFM) by determining an instrument calibration function to analytically deconvolve the stray field data. This approach effectively reconstructs the magnetic surface

charge equivalent, providing a quantitative map of the field sources at the sample surface [7]. In contrast, Bruckner *et al.* [4] formulated the reconstruction as a variational optimization task. Their approach utilizes the adjoint method and Tikhonov regularization to reconstruct the 3D magnetization by minimizing the field residual via gradient-based optimization. Yao *et al.* [8] additionally incorporate topological charge constraints to successfully reconstruct different magnetic quasiparticles. Simultaneously, the integration of deep learning has further advanced the field. Physically informed neural networks have proven effective in reconstructing magnetization configurations [9,10]. The network's architecture naturally enforces smoothness, while the field based forward loss enables the recovery of the magnetization. However, when applied to fully three-dimensional spin textures, successful reconstruction remains challenging and requires an initial magnetization guess [9]. Recently, we introduced a hybrid strategy [11] that utilizes a U-Net to generate an initial guess, followed by micromagnetic relaxation. While this approach is very effective at recovering intrinsic material parameters, it relies on a multistep process that fundamentally separates the data-driven estimation from the final physics-based refinement. Here, we propose a powerful alternative to existing reconstruction strategies: a self-consistent physics-informed reconstruction framework that embeds micromagnetic physics directly into the optimization [Fig. 1(c)]. Compared to other variational methods that rely on mathematical regularization as a proxy for physics, e.g., using Tikhonov

*Contact author: alexander.setescak@univie.ac.at

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. Physics-informed reconstruction of magnetization textures from NV magnetometry. The schematic illustrates the workflow for overcoming the ill-posed nature of stray-field imaging. (a) NV magnetometry measures the stray field projection H^{meas} at an unknown effective distance d_{NV} . (b) Multiple distinct magnetization configurations can produce a nearly identical stray field, making direct inversion unreliable. (c) Our physics-informed framework solves this by coupling a differentiable forward model with a joint objective function. This objective balances data-fidelity with micromagnetic energy. Through iterative gradient-based optimization, the framework simultaneously reconstructs the magnetization $\mathbf{m}$ and refines the effective distance d_{NV} . This yields the final parameters $\mathbf{m}^*$ and d_{NV}^* as an energetically favorable configuration that remains consistent with the experimental observations.

regularization as a qualitative exchange interaction, we replace these heuristic regularizers with quantitative micromagnetic energy terms. This leaves us with a single regularization weight that balances the magnetic stray field residual against the system's total energy, and can be guided by the so-called L-curve method [12,13]. This yields a transparent, end-to-end optimization process that requires no training data and provides quantitative, model-constrained reconstruction of nanoscale magnetic textures along with critical experimental parameters, in particular the *effective distance* d_{NV} . Throughout this Letter, d_{NV} denotes the total separation between the magnetic material and the NV sensing center, including the NV implantation depth, any surface layers such as sample oxidation, the physical gap between sample and diamond surface, and any other separation between the two [14].

Forward model—Given the magnetization of the sample, our forward model essentially simulates what an NV magnetometer would detect at a given effective distance. It provides us with a differentiable map from a given magnetization configuration $\mathbf{m}$ and the effective distance d_{NV} to the expected signal. The required micromagnetic

calculations are performed using the open-source, finite-difference Python library NeuralMag [15], configured with its cell-averaged discretization and running on JAX (also verified in the PyTorch-based MAGNUM.NP [16]). Both back ends evaluate the stray field through FFT convolution with the demagnetization tensor kernel, resulting in a favorable computational complexity of $\mathcal{O}(N \log N)$.

All simulations presented in this Letter focus on an $800 \text{ nm} \times 800 \text{ nm}$ area of a $\text{Fe}_{3-x}\text{GaTe}_2$ flake of thickness 100 nm . The system is discretized with a cell size of $\Delta_x = \Delta_y = 1 \text{ nm}$ and $\Delta_z = 100 \text{ nm}$. The geometry consists of two layers: the bottom layer represents the magnetic material, while the top layer is set to vacuum to calculate stray fields above the sample surface. The measurement itself does not constrain depth variations, and for the chosen parameters the physics-informed regularizer favors a z -uniform state. Finer Δ_z would therefore not recover meaningful new information, and we adopt the single-layer geometry by construction.

The solver computes the stray field averaged over the simulation cell volume ($1 \text{ nm} \times 1 \text{ nm} \times 100 \text{ nm}$). While lateral spatial averaging is negligible, the vertical averaging

over the vacuum layer thickness Δ_z introduces significant deviations that require explicit correction to enable comparison with measurements performed at a specific effective distance d_{NV} . Fortunately, in the nonmagnetic region $z > z_0$ above the sample surface z_0 , the magnetic field satisfies Laplace's equation, permitting the derivation of an exact transfer function in 2D Fourier space. Starting with the averaged field $\langle \mathbf{H}^{\text{dem}} \rangle_{[z_0, z_0 + \Delta_z]}$ in the vacuum layer above the sample, this function recovers the magnetic field at the sample surface and extrapolates it to the effective distance d_{NV} :

$$\tilde{\mathbf{H}}^{\text{dem}}(d_{\text{NV}}) = \langle \tilde{\mathbf{H}}^{\text{dem}} \rangle_{[z_0, z_0 + \Delta_z]} \cdot \left(\frac{k\Delta_z}{1 - e^{-k\Delta_z}} \right) e^{-k(d_{\text{NV}} - z_0)}, \quad (1)$$

where $k = \sqrt{k_x^2 + k_y^2}$ is the magnitude of the wave vector and $\tilde{\mathbf{H}}^{\text{dem}}$ denotes the stray field in the 2D Fourier domain. The multiplicative factors correspond to de-averaging over the cell thickness of the vacuum layer Δ_z and upward continuation toward the effective distance d_{NV} , respectively.

Following upward continuation, the vector field $\tilde{\mathbf{H}}^{\text{dem}}(d_{\text{NV}})$ is transformed back to real space via inverse FFT and projected onto the nitrogen-vacancy (NV) center quantization axis $\mathbf{n}_{\text{NV}} = (\sin 54.74^\circ, 0, \cos 54.74^\circ)$, the tetrahedral angle set by the diamond $\langle 111 \rangle$ lattice direction. This projection, $H^{\text{dem}} = \tilde{\mathbf{H}}^{\text{dem}} \cdot \mathbf{n}_{\text{NV}}$, yields a final scalar field map that can be directly compared to experimental field measurements.

Notably, our approach decouples the simulation's vertical discretization (Δ_z) from the measurement geometry. This allows for the modeling of thick magnetic samples using coarse discretization while accurately calculating the stray field at arbitrary effective distances. Furthermore, because the entire forward pipeline is implemented in a fully differentiable manner, it enables gradient-based optimization not only for the magnetization field $\mathbf{m}$ but also for auxiliary parameters such as the effective distance d_{NV} . See Supplemental Material (SM) Sec. I for the deaveraging and upward-continuation derivation, and SM Sec. II. C for the simulation geometry [17].

Physics-informed loss functional—To address the inverse problem of inferring a magnetization $\mathbf{m}$ from the measured field H^{meas} , we formulate an optimization problem by minimizing the loss functional $J(\mathbf{m}, d_{\text{NV}})$:

$$J(\mathbf{m}, d_{\text{NV}}) = \mathcal{L}_{\text{data}}^2(\mathbf{m}, d_{\text{NV}}) + \lambda \mathcal{L}_{\text{energy}}(\mathbf{m}). \quad (2)$$

Here $\mathcal{L}_{\text{data}} = \|\mathbf{H}^{\text{dem}}(\mathbf{m}, d_{\text{NV}}) - H^{\text{meas}}\| / \|H^{\text{meas}}\|$ is the relative field error between the modeled and measured stray field, and $\mathcal{L}_{\text{energy}}$ is the total micromagnetic energy $E_{\text{total}}(\mathbf{m})$ normalized to a uniformly out-of-plane magnetized state $\mathbf{m} = \mathbf{e}_z$.

Following the micromagnetic framework of Abert *et al.* [29], we define the total energy as the sum of five contributions: the exchange energy (E^{ex}), the demagnetization energy (E^{dem}), the uniaxial anisotropy energy (E^{aniu}), the interfacial Dzyaloshinskii-Moriya interaction (DMI) energy (E^{dmii}), and the Zeeman energy (E^{Zee}) from the applied bias field $\mathbf{H}^{\text{bias}}$. For the anisotropy and DMI formulations, the relevant axis and interface normal directions coincide with the out-of-plane direction $\mathbf{e}_z$. The material parameters used throughout this work are the interfacial DMI constant $D_i \approx -0.51 \text{ mJm}^{-2}$, saturation magnetization $M_s \approx 53 \text{ kAm}^{-1}$, uniaxial anisotropy constant $K_u \approx 0.302 \text{ MJm}^{-3}$, and exchange stiffness $A \approx 0.70 \text{ pJm}^{-1}$ [30]. The experimental bias field is $H^{\text{bias}} \approx 3.2 \text{ kAm}^{-1}$. The formulation is not tied to this particular Hamiltonian or geometry: any micromagnetic energy terms and sample topology supported by our differentiable inverse-micromagnetic libraries NeuralMag [15] and MAGNUM.NP [16] can be substituted.

The regularization parameter $\lambda > 0$ controls the trade-off between data fidelity and physical stability. We visualize this trade-off with the L-curve method [12], plotting $\mathcal{L}_{\text{data}}$ against $\mathcal{L}_{\text{energy}}$ for varying λ and select a balanced operating point λ_{bal} at the “corner” of the resulting L-curve.

The loss $J(\mathbf{m}, d_{\text{NV}})$ is minimized by gradient-based optimization, jointly over $\mathbf{m}$ and d_{NV} , exploiting the differentiability of the forward model. The micromagnetic normalization constraint $|\mathbf{m}| = 1$ is enforced via a unit-norm parametrization of the magnetization rather than an additional penalty term. Full expressions for $\mathcal{L}_{\text{data}}$, $\mathcal{L}_{\text{energy}}$, and the five energy contributions, together with the parameter derivation and the optimizer, parametrization, and initialization choices, are given in SM Sec. II [17].

Reconstruction—We apply the framework to room-temperature experimental NV magnetometry measurements of $\text{Fe}_{3-x}\text{GaTe}_2$ flakes [30] (Fig. 2). Because optically detected magnetic resonance (ODMR) NV magnetometry probes the Zeeman splitting of the NV spin states, it is sensitive only to the field magnitude along the NV axis and cannot resolve its sign [31]. A bias field H^{bias} is applied to lift this degeneracy, with the experiment reporting the bias-subtracted field $H^{\text{meas}} = |H^{\text{bias}} + H^{\text{dem}}| - H^{\text{bias}}$ (where $H^{\text{dem}} = \tilde{\mathbf{H}}^{\text{dem}} \cdot \mathbf{n}_{\text{NV}}$). In theory, $H^{\text{meas}} = H^{\text{dem}}$ everywhere outside localized regions where the opposing demagnetizing field exceeds the bias ($H^{\text{dem}} < -H^{\text{bias}}$). There, the total projected field undergoes a sign inversion, creating ODMR sign-ambiguity artifacts (green contour in Fig. 2), which we exclude from the data-fidelity term $\mathcal{L}_{\text{data}}$ (see SM Sec. III [17]). Masking suppresses the spurious texture induced by these artifacts, even though for this measurement the effect is small.

The reconstruction jointly optimizes $\mathbf{m}$ and d_{NV} using the measured field map H^{meas} as input. We repeat this reconstruction independently for a range of regularization




FIG. 2. Experimental NV magnetometry measurements of a 100 nm thick $\text{Fe}_{3-x}\text{GaTe}_2$ thin film. Left: large-scale overview (2500 nm $\times$ 2500 nm, 100 $\times$ 100 pixels). Right: high-resolution scan (800 nm $\times$ 800 nm, 100 $\times$ 100 pixels) used for the quantitative magnetization reconstruction in this work. The green contour encloses the sign-ambiguity artifact regions, where the opposing projected demagnetizing field H^{dem} exceeds the bias H^{bias} , identified in practice by a data-driven threshold-and-fill criterion. These regions are excluded from the data-fidelity term $\mathcal{L}_{\text{data}}$. The exact effective distances are unknown.

strengths λ . Figure 3 shows (a) the L-curve and (b) the corresponding converged effective distances d_{NV}^* . We select a balanced operating point $\lambda_{\text{bal}} \approx 25$. We observe that d_{NV}^* rises with λ from ≈ 73 nm at $\lambda_{\text{low}} = 0.1$, then plateaus at ≈ 81 nm over an order of magnitude in λ around λ_{bal} , before climbing to ≈ 95 nm at $\lambda_{\text{high}} = 10^3$.

At $\lambda = 0$ there is no prior and the inverse problem is under-determined: upward continuation damps high- k stray fields as $e^{-kd_{\text{NV}}}$, so at the large effective distance d_{NV} the measurement is blind to subresolution structure and the reconstruction collapses into high-frequency noise that fits the data perfectly. A weak prior at $\lambda_{\text{low}} = 0.1$ suppresses this noise, yet the optimizer still overfits experimental imperfections such as sample inhomogeneities and local

variations in saturation magnetization. It can use fine domain patterns to fit the measurement very well at little energy cost, due to the short intrinsic length scale set by the strong out-of-plane anisotropy that keeps domain walls narrow, combined with the DMI that lowers the energy cost of creating them. At smaller d_{NV} the sensor would resolve these structures directly and prevent this overfitting. At $\lambda_{\text{high}} = 10^3$, the energy term dominates and the optimizer favors a strictly physical solution that conflicts with the data. To compensate, d_{NV} is artificially increased to blur and dampen the simulated stray field, producing a field noticeably weaker than the measurement. At λ_{bal} , the reconstructed magnetization $\mathbf{m}^*$ balances these extremes, minimizing the micromagnetic energy while still reproducing the measured field map fairly well (relative field error $\mathcal{L}_{\text{data}} \approx 14.5\%$; see Fig. 4).

A key result is the convergence of the effective distance to $d_{\text{NV}}^* \approx 81$ nm, with a plateau-like region in λ around λ_{bal} similar to, but less pronounced than, the plateau observed in reconstructions from synthetic data. Since the contributions to d_{NV} are individually unknown in typical experiments, the ability to estimate this composite quantity directly from the measurement data is a significant practical advantage.

For the present setup, we estimate these contributions from the fabrication and setup parameters. The 6 keV nitrogen implantation places the NV center 10 to 20 nm below the diamond surface, the oxidation layer of the uncapped, air-sensitive $\text{Fe}_{3-x}\text{GaTe}_2$ surface adds 10 to 30 nm, the mechanical contact gap adds a few nanometers, and because the NV can lie anywhere across the 200 to 400 nm wide tip, a tip-sample tilt can lift it by a further 0 to 40 nm. These contributions sum to a range that contains the extracted $d_{\text{NV}}^* \approx 81$ nm, consistent with the 26 to 128 nm effective distances reported for comparable scanning-NV probes [14]. The full budget is given in SM Sec. VI [17].




FIG. 3. L-curve analysis for experimental measurement data. (a) The L-curve illustrates the trade-off between the data-fidelity term $\mathcal{L}_{\text{data}}$ and the normalized regularization energy $\mathcal{L}_{\text{energy}}$. The red marker denotes the balanced operating point $\lambda_{\text{bal}} \approx 25$, where d_{NV}^* plateaus (b). (b) Converged effective distance d_{NV}^* as a function of the regularization parameter λ .




FIG. 4. Influence of regularization strength λ on reconstructed magnetization and field maps. Comparison between reconstructed magnetizations $\mathbf{m}^*$ and their corresponding demagnetizing fields $H^{\text{dem}}(\mathbf{m}^*, d_{\text{NV}}^*)$ across varying λ . The first column shows the experimental measurement and the second the unregularized $\lambda = 0$ inversion, which overfits the data into an unphysical, high-energy texture. While higher regularization yields a less diffuse $\mathbf{m}^*$, it results in a poorer fit to the experimental field. At high λ values, the field strength is significantly underestimated, indicating a decoupling from the experimental data in favor of the regularization term.

Limitations and outlook—The first class of limitations is set by the information content of stray-field measurements themselves rather than by the reconstruction. Certain magnetization configurations are inherently invisible to any $\mathbf{H}$ -based method. A classic example is a perfectly flux-closed in-plane magnetic vortex, which generates zero external stray field. Additionally, upward continuation acts as a low-pass filter on the stray field [32], exponentially suppressing high-frequency components as $e^{-kd_{\text{NV}}}$, so any magnetization texture concentrated at high wave vectors k is encoded only weakly in the measurement. In our $\text{Fe}_{3-x}\text{GaTe}_2$ thin film, the strong uniaxial out-of-plane anisotropy K_u confines the in-plane magnetization components to narrow domain walls, making their stray-field signatures hard to detect at large distances. The reconstructed wall type and chirality are then determined by the micromagnetic prior rather than by the data. Where the measurement is uninformative, the choice of physics determines the solution.

A second class of limitations stems from the physical model itself: the reconstruction is only as good as the assumed model. Its parameters may be off: a $\pm 20\%$ uncertainty in the saturation magnetization M_s , for instance, propagates into a 10% to 14% shift in the optimized effective distance d_{NV}^* . It may also be incomplete, neglecting variation of the magnetization along the film thickness z , surface roughness, pinning sites, or additional energy terms. Such inhomogeneities are not excluded in principle: once characterized, they can be built into the same differentiable energy and forward model. Two internal consistency checks can flag such effects when they remain unmodeled, although they cannot identify the cause of an inconsistency. The first is the shape of the L-curve itself.

When the energy prior and the measurement cannot be simultaneously satisfied, the corner softens and d_{NV}^* no longer settles onto a clear plateau. Comparison against a synthetic baseline or other datasets makes this easier to recognize (SM Sec. IV. B [17]). The second check is dynamic relaxation of the reconstructed state using the Landau-Lifshitz-Gilbert (LLG) equation [16,33,34], which simulates the temporal relaxation of the spin system toward an energy minimum. Because the reconstruction minimizes the joint loss rather than the energy alone, some relaxation is expected even for a correctly specified model. Rapid evolution nevertheless shows that the reconstructed state is not an equilibrium of the assumed model, but it cannot distinguish whether the model is incomplete or whether artifacts and noise imprinted spurious structure on the state (SM Sec. IV. C [17]).

Finally, the underlying reconstruction of the magnetization $\mathbf{m}$ from the magnetic field $\mathbf{H}$ remains fundamentally ill-posed. Even with physics constraints, we are not guaranteed to recover the true sample state. Moreover, due to the highly non-convex nature of the micromagnetic energy landscape, gradient-based optimization can converge to different metastable configurations depending on the initialization. Nevertheless, this framework allows us to find significantly better, physically plausible solutions compared to unregularized inversion, and the $(\mathcal{L}_{\text{data}}, \mathcal{L}_{\text{energy}})$ plane in which the L-curve lives provides a physics-grounded criterion for selecting a best guess among these candidate solutions: at equal data fit, prefer the lowest-energy configuration. At equal energy, prefer the better data fit. This allows informed choices about initialization, optimizer, and other hyperparameters (SM Secs. II. E and V [17]), which the data fit alone cannot rank meaningfully.

Looking forward, the same sensitivity to the physical assumptions that bounds the reconstruction is also a source of physical insight. Where the measurement is informative, comparing how different energy models stabilize the reconstructed states at a given data fit can help point to the dominant physical mechanisms. Where it is not, the prior still ensures that the recovered state is a physical solution consistent with the data. This variational approach is also easily extendable: multiheight or multiorientation NV acquisitions could add complementary constraints that further reduce the non-uniqueness of the inversion in realistic noisy, finite-field-of-view measurements. Moreover, because the framework relies on a differentiable forward model, additional auxiliary experimental parameters (such as relative shifts in multidistance acquisitions) or even parameters like the saturation magnetization M_s (when d_{NV} is independently known) can in principle be co-optimized jointly with the magnetization. Ultimately, this combination of a differentiable forward model and a differentiable micromagnetic energy formulation serves as a template for adapting physics-informed reconstructions to other magnetic imaging modalities.

Acknowledgments—This research was funded in whole, or in part, by the Austrian Science Fund (FWF) DOI: 10.55776/PIN1629824, DOI: 10.55776/P34671, and DOI: 10.55776/PAT3864023. The computational results have been achieved using the Austrian Scientific Computing (ASC) infrastructure.

A.S. conceptualized the specific physics-informed reconstruction framework, implemented the method, performed the synthetic and experimental evaluations, and wrote the original manuscript. This work builds upon the foundational concept of inverse micromagnetics established by F.B. and C.A., who developed the core inverse micromagnetic simulation software MAGNUM.NP and NeuralMag. C.A., D.S., and F.B. also directly contributed to the development of the methodology. On the experimental side, the samples were provided by H.Y. and C.Z. Data acquisition was performed by Y.-G.C., H.B., L.B., C.D., and U.V. The NV magnetometry infrastructure is overseen by C.D. and U.V. Additionally, Y.-G.C., U.V., and C.A. supported the application of the method to the experimental data, including the configuration of simulation parameters and the interpretation of the results. All authors engaged in scientific discussions and reviewed the manuscript.

Data availability—The NV measurement maps, the synthetic ground truth, and every reconstruction shown in the main text and the Supplemental Material [17] are archived at Zenodo [35]. The reconstruction code is available on GitHub and archived at Zenodo [36].

[1] F. Casola, T. van der Sar, and A. Yacoby, Probing condensed matter physics with magnetometry based on nitrogen-vacancy centres in diamond, *Nat. Rev. Mater.* **3**, 17088 (2018).

[2] Y. Xu, W. Zhang, and C. Tian, Recent advances on applications of NV[−] magnetometry in condensed matter physics, *Photonics Res.* **11**, 393 (2023).

[3] G. E. Backus, Non-uniqueness of the external geomagnetic field determined by surface intensity measurements, *J. Geophys. Res.* **75**, 6339 (1970).

[4] F. Bruckner, C. Abert, G. Wautischer, C. Huber, C. Vogler, M. Hinze, and D. Suess, Solving large-scale inverse magnetostatic problems using the adjoint method, *Sci. Rep.* **7**, 40816 (2017).

[5] C. B. Clement, J. P. Sethna, and K. C. Nowack, Reconstruction of current densities from magnetic images by Bayesian inference, *arXiv:1910.12929*.

[6] H. J. Hug, B. Stiefel, P. J. A. van Schendel, A. Moser, R. Hofer, S. Martin, H.-J. Güntherodt, S. Porthun, L. Abelmann, J. C. Lodder, G. Bochi, and R. C. O’Handley, Quantitative magnetic force microscopy on perpendicularly magnetized samples, *J. Appl. Phys.* **83**, 5609 (1998).

[7] Y. Feng, A.-O. Mandru, O. Yildirim, and H. J. Hug, Quantitative magnetic force microscopy: Transfer-function method revisited, *Phys. Rev. Appl.* **18**, 024016 (2022).

[8] C. Yao, Y. Yu, Y. Shi, J.-I. Jung, Z. Váci, Y. Wang, Z. Liu, C. Zhang, S. Tikoo-Schantz, and C. Zu, Universal reconstruction of complex magnetic profiles with minimal prior assumptions, *Phys. Rev. Appl.* **24**, 064020 (2025).

[9] D. A. Broadway, M. Flaks, A. E. E. Dubois, and P. Maletinsky, Reconstruction of nontrivial magnetization textures from magnetic field images using neural networks, *Phys. Rev. Appl.* **23**, 044012 (2025).

[10] A. E. E. Dubois, D. A. Broadway, A. Stark, M. A. Tschudin, A. J. Healey, S. D. Huber, J.-P. Tetienne, E. Greplova, and P. Maletinsky, Untrained physically informed neural network for image reconstruction of magnetic field sources, *Phys. Rev. Appl.* **18**, 064076 (2022).

[11] D. Suess, A. Setescak, J. Smith, F. Bruckner, H. J. Hug, and C. Abert, Reconstruction of magnetic structures and material parameters with convolutional neural network and bias field-constrained micromagnetic relaxation, *Sci. Rep.* **15**, 42867 (2025).

[12] P. C. Hansen and D. P. O’Leary, The use of the L-curve in the regularization of discrete ill-posed problems, *SIAM J. Sci. Comput.* **14**, 1487 (1993).

[13] D. Calvetti, S. Morigi, L. Reichel, and F. Sgallari, Tikhonov regularization and the L-curve for large discrete ill-posed problems, *J. Comput. Appl. Math.* **123**, 423 (2000).

[14] Z. Xu, M. L. Palm, W. Huxter, K. Herb, J. M. Abendroth, K. Bouzehouane, O. Boulle, M. S. Gabor, J. Urrestarazu Larranaga, A. Morales, J. Rhensius, G. Puebla-Hellmann, and C. L. Degen, Minimizing sensor-sample distances in scanning nitrogen-vacancy magnetometry, *ACS Nano* **19**, 8255 (2025).

[15] C. Abert, F. Bruckner, A. Voronov, M. Lang, S. A. Pathak, S. Holt, R. Kraft, R. Allayarov, P. Flauger, S. Koraltan, T. Schrefl, A. Chumak, H. Fangohr, and D. Suess, NeuralMag: An open-source nodal finite-difference code for inverse micromagnetics, *npj Comput. Mater.* **11**, 193 (2025).

[16] F. Bruckner, S. Koraltan, C. Abert, and D. Suess, MAGNUM.NP: A PyTorch based GPU enhanced finite difference

- micromagnetic simulation framework for high level development and inverse design, *Sci. Rep.* **13**, 12054 (2023).
- [17] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/nc5g-3yw9> for the upward-continuation derivation, simulation setup, sign-ambiguity analysis, synthetic validation and diagnostics, sensitivity study, and the physical composition of the extracted effective distances, which includes Refs. [18–28].
- [18] R. Bodenberger and A. Hubert, Zur bestimmung der blochwandenergie von Einachsigen ferromagneten, *Phys. Status Solidi (a)* **44**, K7 (1977).
- [19] A. Hubert and R. Schäfer, *Magnetic Domains: The Analysis of Magnetic Microstructures*, softcover reprint of the hardcover 1st ed. 1998, corrected printing 2000 ed. (Springer, Berlin, Heidelberg, 2014).
- [20] S. Rohart and A. Thiaville, Skyrmion confinement in ultrathin film nanostructures in the presence of Dzyaloshinskii-Moriya interaction, *Phys. Rev. B* **88**, 184422 (2013).
- [21] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, Y. Katariya, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of Python+NumPy programs (2018), <http://github.com/google/jax>.
- [22] M. Blondel, Q. Berthet, M. Cuturi, R. Frostig, S. Hoyer, F. Llinares-López, F. Pedregosa, and J.-P. Vert, Efficient and modular implicit differentiation, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2022), Vol. 35, pp. 5230–5242, [10.52202/068431-0378](https://doi.org/10.52202/068431-0378).
- [23] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [24] G. Bécigneul and O.-E. Ganea, Riemannian adaptive optimization methods, [arXiv:1810.00760](https://arxiv.org/abs/1810.00760).
- [25] M. Kochurov, R. Karimov, and S. Kozlukov, Geoopt: Riemannian optimization in PyTorch, [arXiv:2005.02819](https://arxiv.org/abs/2005.02819).
- [26] MAGNUM.NP developers, Inverse Problems—MAGNUM.NP 2.2.0 documentation, https://magnum.np.gitlab.io/magnum.np/inverse_problems.html (2026), accessed: 2026-02-18.
- [27] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2019), Vol. 32.
- [28] W. Xie, J. Zhang, Y. Bai, Y. Liu, H. Wang, P. Yu, J. Li, H. Chang, Z. Wang, F. Gao, G. Wei, W. Zhao, and T. Nie, Air stability and composition evolution in van der Waals Fe_3GeTe_2 , *APL Mater.* **12**, 031102 (2024).
- [29] C. Abert, Micromagnetics and spintronics: Models and numerical methods, *Eur. Phys. J. B* **92**, 120 (2019).
- [30] C. Zhang, Z. Jiang, J. Jiang, W. He, J. Zhang, F. Hu, S. Zhao, D. Yang, Y. Liu, Y. Peng, H. Yang, and H. Yang, Above-room-temperature chiral skyrmion lattice and Dzyaloshinskii-Moriya interaction in a van der Waals ferromagnet $\text{Fe}_{3-x}\text{GaTe}_2$, *Nat. Commun.* **15**, 4472 (2024).
- [31] L. Rondin, J.-P. Tetienne, T. Hingant, J.-F. Roch, P. Maletinsky, and V. Jacques, Magnetometry with nitrogen-vacancy defects in diamond, *Rep. Prog. Phys.* **77**, 056503 (2014).
- [32] R. J. Blakely, *Potential Theory in Gravity and Magnetic Applications* (Cambridge University Press, Cambridge, England, 1995).
- [33] L. D. Landau and E. M. Lifshitz, Theory of the dispersion of magnetic permeability in ferromagnetic bodies, *Phys. Z. Sowjetunion* **8**, 153 (1935).
- [34] T. L. Gilbert, A phenomenological theory of damping in ferromagnetic materials, *IEEE Trans. Magn.* **40**, 3443 (2004).
- [35] A. Setescak, F. Bruckner, D. Suess, Y.-G. Choi, H. Binger, L. Boer, C. Zhang, H. Yang, C. Donnelly, U. Vool, and C. Abert, Data for: A Fourier-space approach to physics-informed magnetization reconstruction from nitrogen-vacancy measurements (2026), Zenodo, [10.5281/zenodo.22674823](https://doi.org/10.5281/zenodo.22674823).
- [36] A. Setescak, nvrecon: physics-informed magnetization reconstruction from NV magnetometry maps (2026), Zenodo, [10.5281/zenodo.22675268](https://doi.org/10.5281/zenodo.22675268); code: <https://github.com/ajjaybob/nv-mag-recon>.