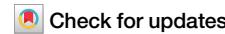




# NeuralMag: an open-source nodal finite-difference code for inverse micromagnetics



C. Abert<sup>1</sup> F. Bruckner<sup>1</sup>, A. Voronov<sup>1,2</sup>, M. Lang<sup>3,4</sup>, S. A. Pathak<sup>3,4</sup>, S. Holt<sup>3,4</sup>, R. Kraft<sup>1,2</sup>, R. Allayarov<sup>1</sup>, P. Flauger<sup>1</sup>, S. Koraltan<sup>5</sup>, T. Schrefl<sup>6</sup>, A. Chumak<sup>1</sup>, H. Fangohr<sup>3,4,7</sup> & D. Suess<sup>1</sup>

We present NeuralMag, a flexible and high-performance open-source Python library for micromagnetic simulations. NeuralMag leverages modern machine learning frameworks, such as PyTorch and JAX, to perform efficient tensor operations on various parallel hardware, including CPUs, GPUs, and TPUs. The library implements a novel nodal finite-difference discretization scheme that provides improved accuracy over traditional finite-difference methods without increasing computational complexity. NeuralMag is particularly well-suited for solving inverse problems, especially those with time-dependent objectives, thanks to its automatic differentiation capabilities. Performance benchmarks show that NeuralMag is competitive with state-of-the-art simulation codes while offering enhanced flexibility through its Python interface and integration with high-level computational backends.

Micromagnetic simulations are a fundamental tool in the study of magnetization dynamics and play a crucial role in understanding and designing magnetic materials and devices. These simulations model the behavior of magnetic and magnonic systems at the nanoscale, providing insight into phenomena such as domain wall motion, magnetization reversal, and spin wave propagation. The field relies on various computational methods, with finite-difference and finite-element schemes being widely used. Notable examples of established finite-difference codes include OOMMF<sup>1</sup> and fidimag<sup>2</sup> for CPU-based simulations and mumax3<sup>3</sup>, BORIS<sup>4</sup>, and magnum.np<sup>5</sup> for GPU-accelerated simulations. Finite-element-based methods, such as those implemented in NMag<sup>6</sup>, Tetramag<sup>7</sup>, FastMag<sup>8</sup>, FinMag<sup>9</sup>, and magnum.fe<sup>10</sup>, provide greater flexibility in handling complex geometries but can be computationally more expensive. More recently, the finite-element solver TetraX<sup>11</sup> has gained popularity in the magnonics community due to its efficient eigenmode solver in infinite geometries.

In addition to standard micromagnetic simulations, inverse problems have attracted considerable attention in recent years. These problems involve determining the optimal parameters, such as material properties, external fields, or device geometries, that lead to a desired magnetic configuration or device functionality. A significant body of work has focused on inverse modeling of the demagnetization field, a static inverse problem. This has been particularly useful in the context of magnetic 3D printing, where

topology optimization techniques are employed to design optimal material layouts, and the inverse modeling is used to infer the magnetization configuration of printed samples<sup>12–14</sup>.

More recently, research in the emerging field of inverse magnonics has gained momentum, focusing on optimizing the functionality of magnonic devices. Magnonics uses spin waves (magnons) for information processing, and designing efficient magnonic devices poses complex nonlinear optimization challenges. Inverse-design approaches have been increasingly applied to magnonics, allowing researchers to automate the design of devices by specifying a desired functionality and using computational algorithms to find the optimal configuration<sup>15–18</sup>.

In this paper, we present a novel discretization strategy for micromagnetic simulations, adjoint-state algorithms for efficiently solving time-dependent inverse problems, and the software design of NeuralMag, which integrates these advancements into a flexible and high-performance computational framework.

## Results

### Micromagnetics

The micromagnetic model provides a semi-classical continuum description of magnetization dynamics in ferromagnetic systems, as originally formulated by Brown<sup>19</sup>. The key governing equation is the

<sup>1</sup>Faculty of Physics, University of Vienna, Vienna, Austria. <sup>2</sup>Vienna Doctoral School of Physics, University of Vienna, Vienna, Austria. <sup>3</sup>Max Planck Institute for the Structure and Dynamics of Matter, Hamburg, Germany. <sup>4</sup>Center for Free-Electron Laser Science, Hamburg, Germany. <sup>5</sup>Institute of Applied Physics, TU Wien, Vienna, Austria. <sup>6</sup>Department for Integrated Sensor Systems, University for Continuing Education Krems, Krems, Austria. <sup>7</sup>University of Southampton, Southampton, UK. e-mail: [claas.abert@univie.ac.at](mailto:claas.abert@univie.ac.at)

Landau–Lifshitz–Gilbert (LLG) equation, which reads

$$\frac{\partial \mathbf{m}}{\partial t} = -\frac{\gamma}{1+\alpha^2} \mathbf{m} \times \mathbf{H}_{\text{eff}} - \frac{\alpha\gamma}{1+\alpha^2} \mathbf{m} \times (\mathbf{m} \times \mathbf{H}_{\text{eff}}) \quad (1)$$

with  $\mathbf{m}$  being the unit-vector field representation of the magnetization,  $\gamma$  being the reduced gyromagnetic ratio, and  $\alpha$  being a dimensionless damping parameter. The effective field  $\mathbf{H}_{\text{eff}}$  accounts for all relevant interactions within the system and derives from the total energy as

$$\mathbf{H}_{\text{eff}} = -\frac{1}{\mu_0 M_s} \frac{\delta E}{\delta \mathbf{m}} \quad (2)$$

with  $M_s$  being the saturation magnetization and  $\delta E/\delta \mathbf{m}$  denoting the variational derivative of the energy with respect to the magnetization<sup>20,21</sup>. When the energy  $E$  depends on spatial derivatives of the magnetization field  $\mathbf{m}$ , additional boundary conditions must be imposed to solve Eqs. (1) and (2). One such example is the micromagnetic exchange energy, which is defined as

$$E_{\text{ex}} = \int_{\Omega} A(\nabla \mathbf{m})^2 \, dx, \quad (3)$$

where  $A$  is the exchange stiffness constant. The variation of the exchange energy with respect to  $\mathbf{m}$  yields

$$\delta E_{\text{ex}}(\mathbf{m}, \delta \mathbf{m}) = \int_{\Omega} \underbrace{-2[\nabla \cdot (A \nabla \mathbf{m})]}_{=\delta E/\delta \mathbf{m}} \cdot \delta \mathbf{m} \, dx + \int_{\partial\Omega} \underbrace{2A \frac{\partial \mathbf{m}}{\partial \mathbf{n}} \cdot \delta \mathbf{m}}_{=B(\mathbf{n})} \, ds \quad (4)$$

leading to the exchange field definition

$$\mathbf{H}_{\text{ex}} = -\frac{1}{\mu_0 M_s} \frac{\delta E_{\text{ex}}}{\delta \mathbf{m}} = \frac{2}{\mu_0 M_s} \nabla \cdot (A \nabla \mathbf{m}). \quad (5)$$

The boundary term in Eq. (4) defines the appropriate exchange boundary condition. To satisfy equilibrium conditions in micromagnetics, the system must fulfill Brown's conditions, which require  $\mathbf{m} \times \delta E/\delta \mathbf{m} = 0$  for  $\mathbf{x} \in \Omega$ , and  $\mathbf{m} \times \mathbf{B} = 0$  for  $\mathbf{x} \in \partial\Omega$ .

A similar variational treatment at internal interfaces, where material parameters vary discontinuously, introduces additional interface conditions<sup>22</sup>. Assuming a continuous magnetization across such interfaces and dividing the domain into regions of continuous material parameters, the corresponding interface condition can be written as  $\mathbf{m}_1 \times \mathbf{B}_1(\mathbf{n}) = \mathbf{m}_2 \times \mathbf{B}_2(\mathbf{n})$ , where  $\mathbf{B}_1(\mathbf{n})$  and  $\mathbf{B}_2(\mathbf{n})$  represent the boundary terms on either side of the interface.

In case of the exchange energy being the only energy contribution introducing spatial derivatives and furthermore considering  $\mathbf{m} \perp \partial \mathbf{m} / \partial \mathbf{n}$ , this leads to the well-known exchange jump condition<sup>19</sup>

$$A_1 \frac{\partial \mathbf{m}_1}{\partial \mathbf{n}} = A_2 \frac{\partial \mathbf{m}_2}{\partial \mathbf{n}}. \quad (6)$$

In addition to satisfying equilibrium conditions, the boundary and interface conditions must be consistently fulfilled at all times when solving the LLG equation<sup>22</sup>.

NeuralMag implements a novel nodal finite-difference scheme described in section “Nodal finite-difference scheme” to accurately solve the micromagnetic equations. The code supports both PyTorch<sup>23</sup> and JAX<sup>24</sup> as computational backends, enabling efficient tensor operations and automatic differentiation on various hardware platforms. Details on the implementation are provided in section “Implementation”.

## Inverse micromagnetics

In addition to employing a nodal finite-difference scheme, NeuralMag is specifically designed to address inverse problems in both space and time domains. In this context, the computation of individual field terms or the solution of the LLG Eq. (1) is classified as a forward problem  $F$ . Given a vector of design variables  $\theta$ , which may include material properties or the initial magnetization configuration, these forward problems yield well-defined outputs  $y$ , such as effective field contributions or the resulting magnetization trajectory

$$F(\theta) = y. \quad (7)$$

An inverse problem is formulated to determine the design variables  $\theta$  that yield a specified result  $y$  from the forward problem. This task is often challenging, as inverse problems are typically ill-posed, and their solution vectors may encompass a large number of degrees of freedom. The most common strategy to solve such a problem is the reformulation in terms of a minimization problem that might be complemented by additional terms for regularization or smoothing purposes. In the case of a high-dimensional input  $\theta$  and a nonlinear function  $F$ , this problem is nontrivial. In such cases, iterative methods, typically based on the gradient of the functional  $\nabla_{\theta} \mathcal{L}$ , are commonly employed to find a solution. NeuralMag uses automatic differentiation<sup>25</sup> for static problems such as inverse strayfield calculations. In contrast to the adjoint method that has been used in previous works<sup>12,13</sup>, this approach performs the differentiation on the discrete level (discretize first). As for the adjoint method, the gradient computation requires a forward solve and a subsequent backward solve with the complexity of the backward solve being equivalent to that of the forward solve.

$$\min_{\theta} \mathcal{L}(\theta) \quad \text{with} \quad \mathcal{L}(\theta) = \| F(\theta) - y \| \quad (8)$$

For time-dependent problems, NeuralMag implements the adjoint-state method<sup>26</sup>. The adjoint-state method is a powerful tool for the solution of PDE-constrained optimization problems, also referred to as optimal-control problems. Given a forward problem

$$\frac{\partial \mathbf{m}}{\partial t} = f_{\theta}(t, \mathbf{m}), \quad (9)$$

with design variables  $\theta$ , we define an objective functional

$$\mathcal{L}(\theta) = L(\mathbf{m}(T; \theta), \mathbf{y}_{\text{target}}) \quad (10)$$

with  $\mathbf{m}(T; \theta)$  being the solution of Eq. (9) for a final time  $T$  and  $\mathbf{y}_{\text{target}}$  being the desired output of the forward problem. In order to compute the gradient of the objective functional with respect to the design variables  $\nabla_{\theta} \mathcal{L}(\theta)$ , the adjoint-state method requires two steps. In the first step, the forward problem (9) is solved for the given design variables  $\theta$ , which results in the output  $\mathbf{m}_{\text{output}} = \mathbf{m}(T)$ . In the second step, the so-called adjoint problem is solved, which is given by the following system of ODEs

$$\begin{cases} \frac{\partial \mathbf{m}}{\partial t} = f_{\theta}(t, \mathbf{m}) & \text{with} & \mathbf{m}(T) = \mathbf{m}_{\text{output}} \\ \frac{\partial \mathbf{a}}{\partial t} = -\frac{\partial f_{\theta}(t, \mathbf{m})}{\partial \mathbf{m}} \mathbf{a} & \text{with} & \mathbf{a}(T) = \nabla_y L(\mathbf{y}, \mathbf{y}_{\text{target}}), \mathbf{y} = \mathbf{m}_{\text{output}} \\ \frac{\partial \mathbf{u}}{\partial t} = -\frac{\partial f_{\theta}(t, \mathbf{m})}{\partial \theta} \mathbf{a} & \text{with} & \mathbf{u}(T) = \mathbf{0}. \end{cases} \quad (11)$$

with  $\mathbf{a}$  being the so-called adjoint variable. This system is solved backwards in time, starting from the final time  $T$  used in the forward pass. Successful integration of the system yields the output  $\mathbf{u}(0)$ , which can be identified as the desired gradient of the objective

$$\mathbf{u}(0) = \nabla_{\theta} \mathcal{L}(\theta). \quad (12)$$

While the objective (10) depends solely on the magnetization at the final time  $T$ , extending this method to objectives depending on multiple time points  $T_i$  can be done in a straightforward fashion by adding appropriate

terms depending on  $m(T_i, \theta)$  to (10). The computational and storage complexity of the adjoint system is comparable to that of the forward problem, yielding an exceptionally efficient strategy for the gradient computation of PDE-constrained optimization. This method is superior to the backpropagation method<sup>15,16</sup> with regard to the storage requirements that are similar to a regular forward pass. However, this advantage comes at the cost of reduced accuracy, which is caused by the backwards pass that reconstructs the magnetization trajectory by inverse integration instead of using the exact values from the forward pass.

## Validation and benchmarks

To validate the accuracy of NeuralMag, we solve two significant micromagnetic problems. These tests showcase NeuralMag's ability to handle both standard and advanced cases, verifying its precision and computational efficiency.

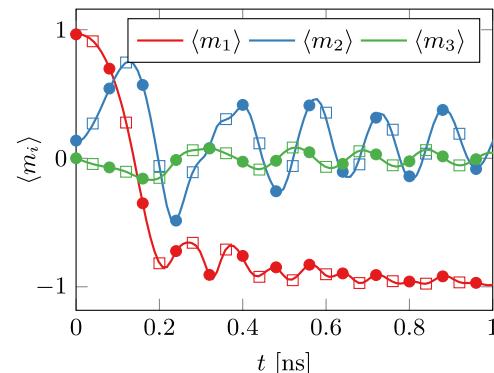
The first validation case is MuMag Standard Problem #4<sup>27</sup>, which simulates the dynamic behavior of a thin ferromagnetic film under an applied magnetic field tilted either by 170° or 190° toward the x-axis. The focus is on the time evolution of the averaged magnetization components. We solve this problem using a full 3D spatial discretization and compare the results to a 2D simulation as described in section "Low-Dimensional Geometries" of the paper. The results for the field tilted by 170°, displayed in Fig. 1, show excellent agreement with the reference solutions from the MuMag community, demonstrating the precision of NeuralMag in simulating the time dynamics of micromagnetic systems both with the 3D as well as 2D thin-film approximation. The second part of the standard problem #4, which simulates the switching under a field tilted by 190° is included in the demos that are accessible via the NeuralMag website<sup>28</sup>.

The second validation case involves solving the domain wall pinning problem proposed by Heistracher et al.<sup>29</sup>. This problem focuses on calculating the coercive field required to unpin a domain wall at the interface between two magnetic phases with varying material properties, such as exchange interaction, uniaxial anisotropy, and spontaneous magnetization. This problem is sensitive to discontinuities in these parameters, making it an ideal test for NeuralMag's handling of complex material boundaries.

In this validation, we compare the switching fields calculated by NeuralMag with the analytical results provided in Table 1 of the original paper. We varied the material parameters (exchange constant  $A$ , anisotropy constant  $K$ , and saturation magnetization  $M_s$ ) in different combinations across the two magnetic phases. Table 1 compares the switching fields obtained using NeuralMag with those presented in the reference paper. Our results closely match the analytical solutions, with minor deviations likely due to the time integration method and field rate used during the simulation. These successful validations confirm that NeuralMag correctly handles discontinuities at material interfaces and provides accurate predictions for complex micromagnetic systems.

To evaluate the performance of NeuralMag, we conducted a throughput benchmark, shown in Fig. 2, where we compare the time required for evaluating the right-hand side (RHS) of the LLG equation across different system sizes. Specifically, we measure the time for the integration of the full LLG, including the exchange and demagnetization field, and then divide by the number of field evaluations. This procedure can be easily applied to any micromagnetic code without the need to modify it and provides a robust measure of the overall performance at the same time. In this benchmark, NeuralMag is compared to two widely-used micromagnetic simulation tools: mumax3<sup>3</sup> and magnum.np<sup>5</sup>. mumax3 shows the best performance due to its highly optimized GPU implementation. However, NeuralMag, when using JAX as the backend, almost matches the performance of mumax3, being less than a factor of 2 slower.

Remarkably, NeuralMag maintains this competitive performance even for small system sizes, despite the computational overhead typically associated with a Python implementation. This performance can be attributed to the just-in-time (JIT) compilation feature of JAX, which optimizes the entire RHS of the LLG equation at runtime. Thanks to NeuralMag's architecture, JAX is able to analyze and compile the full computation into highly optimized machine code, reducing overhead and achieving near-optimal execution times. This demonstrates the strength of NeuralMag's design in



**Fig. 1 | MuMag standard problem #4.** The results are presented using both 2D and 3D discretizations as computed by NeuralMag. The reference solution, computed with OOMMF<sup>1</sup>, is depicted by solid lines for comparison. The NeuralMag solutions are illustrated using circles for the 3D discretization and squares for the 2D discretization.

leveraging modern machine learning frameworks to achieve high-performance computations while maintaining flexibility.

The remaining performance gap of approximately a factor of two compared to MuMax3 likely arises from the demagnetization field computation, as the current FFT interface in JAX is limited, preventing certain optimizations. However, as JAX's FFT capabilities expand, this gap could narrow significantly—or even vanish completely—in the future.

The CPU performance using the JAX backend is approximately 10 times slower than the single-precision GPU performance, which is remarkably fast. However, there is no notable difference in timings for single CPU usage compared to the OpenMP parallelization on 8 CPUs. For very small systems, the CPU implementation of JAX even outperforms the GPU implementation. This is most likely due to the missing overhead of kernel deployment and data transfer to the GPU, which leads to almost linear scaling down to very small system sizes. CPU computations with PyTorch are approximately 10 times slower than JAX computations.

Because the adjoint/autodiff gradient has the same asymptotic cost as one forward RHS evaluation, the timings in Fig. 2 also characterise a single gradient step in inverse-design workflows.

## Inverse problems

Listing 1. Gradient computation of a topology optimization problem.

```
state = State(...)

# Restrict the design region by masking the density rho
state.rho = (lambda rho_m: rho_m.at[:, :, 10:]).set(state.eps, "ccc", ())
state.rho_m = nm.CellFunction(state).fill(1.0)

# Register demag field and get field function depending on rho_m
demag = DemagField().register(state, "demag")
demag_func = state.resolve("h_demag", ["rho_m"])

# define loss and compute gradient of loss w.r.t. rho_m
def loss(rho):
    return -demag_func(rho*3)[10, 10, 12, 2] ** 2

grad_loss = jax.grad(loss)
```

A first, purely static inverse problem concerns the topology optimisation of a hard-magnetic cuboid of dimensions 100 nm × 100 nm × 50 nm that is uniformly magnetised along +z. Using the Solid-Isotropic Material with Penalisation (SIMP) approach<sup>30</sup>, the material density field  $\rho(\mathbf{x}) \in [0, 1]$  is restricted to the magnetic design region  $\Omega_m$ , see Listing 1). The objective

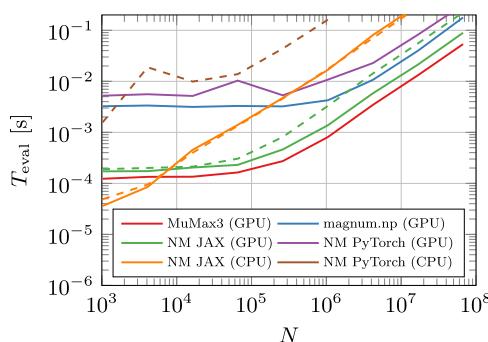
$$\mathcal{L}(\rho) = H_z[\rho^3(0, 0, M_z)]|_{\mathbf{x}=\mathbf{x}_{\text{center}}} \quad (13)$$

with the density

$$\rho(\mathbf{x}) = \rho_{\Omega_m}(\mathbf{x}) \mathbb{1}_{\Omega_m}(\mathbf{x}), \quad \text{with} \quad \rho_{\Omega_m}(\mathbf{x}) \in [0, 1] \quad (14)$$

**Table 1 | Depinning fields for a domain wall in a two-phase magnet, as defined in ref.<sup>29</sup>, computed with NeuralMag and compared to the analytical and numerical reference solutions computed with magnum.af**

Discontinuous parameters	Analytical [T]	magnum.af [T]	NeuralMag [T]
$A/K/M_s$	1.568	1.585	1.580
$A/K$	1.089	1.116	1.112
$A/M_s$	1.206	1.256	1.205
$A$	0.838	0.868	0.867
$K/M_s$	1.005	1.020	1.012
$K$	0.565	0.582	0.571



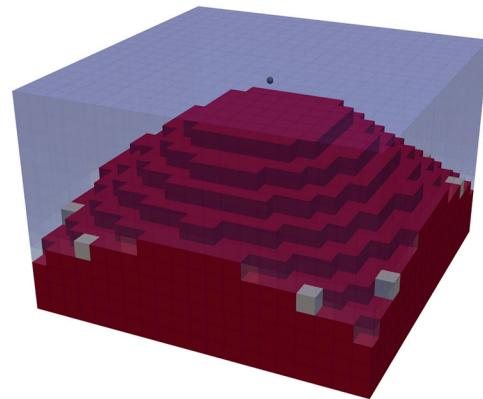
**Fig. 2 | Computational benchmark.** Comparison of the computation time for evaluating the right-hand side of the Landau-Lifshitz-Gilbert (LLG) equation, including both the exchange and demagnetization fields, across various system sizes  $N$ , with NeuralMag (NM) in comparison with other finite-difference codes. The legend indicates the code as well as the device that was used for the computation. As a GPU, we used an NVIDIA A100 card with 80 GB of RAM, and as a CPU, we used an Intel Xeon Gold 6326. For GPU timings, solid lines indicate single-precision computations while dashed lines indicate double-precision. CPU timings were all performed with double precision, with solid lines indicating single-core computations and dashed lines indicating the use of 8 cores.

maximises the z-component of the stray field at the probe point  $\mathbf{x}_{\text{center}}$  located 5 nm above the centre of the top surface. Because the demagnetization field in NeuralMag is differentiable with respect to  $\rho$ , its gradient  $\partial \mathcal{L} / \partial \rho$  is obtained in a single reverse pass and fed to any gradient-based optimiser that enforces  $0 \leq \rho \leq 1$ . The relevant part of the simulation script is shown in Listing 1. The final topology is shown in Fig. 3 in excellent agreement with published solutions<sup>5,14</sup>.

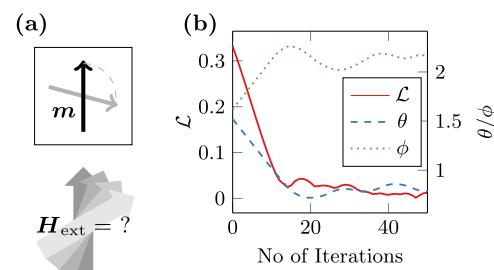
Listing 2. Simulation script for inverse problem.

```
state = State(...)
state.angles = [jnp.pi / 2, jnp.pi / 2]
h_ext = lambda angles: jnp.stack(
    [
        Hc / 2 * jnp.sin(angles[0]) * jnp.cos(angles[1]),
        Hc / 2 * jnp.sin(angles[0]) * jnp.sin(angles[1]),
        Hc / 2 * jnp.cos(angles[0]),
    ]
)
...
llg = nm.LLGSolver(state, parameters={"angles"})
m_target = nm.VectorFunction(state).fill((0.5**0.5, 0, 0.5**0.5)).tensor
def loss(angles, args):
    m_pred = llg.solve(state.tensor([0.0, 0.05e-9]), angles).ys[-1]
    return jnp.mean((m_target - m_pred) ** 2)
solver = optax.BFGS(1e-3, 1e-3, optax.max_norm)
result = optax.minimise(loss, solver, state.angles)
```

After the static topology-optimisation example above, we now turn to a genuinely dynamic inverse problem. Specifically, we aim to optimize the direction of an external magnetic field to align the magnetization of a single-domain particle with a target configuration  $\mathbf{m}_{\text{target}}$  after a given time  $T$ , see Fig. 4(a). The optimization minimizes the objective function  $\mathcal{L}$  with respect



**Fig. 3 | Topology optimization.** Optimization result of a permanent magnetic sample with magnetization  $\mathbf{M} = (0, 0, M_z)$  that maximizes  $H_z(\mathbf{M})$  at a single point above the sample marked by the grey sphere.



**Fig. 4 | Dynamic inverse problem.** Simple inverse micromagnetic problem for the optimization of the external field direction in order to align the magnetization of a single-domain particle in a given direction. **a** Sketch of the problem setup. **b** Convergence of the objective function  $\mathcal{L}$  and the optimized field angles  $\theta$  and  $\phi$ .

to the field angles  $\theta$  and  $\phi$ , as defined by the system

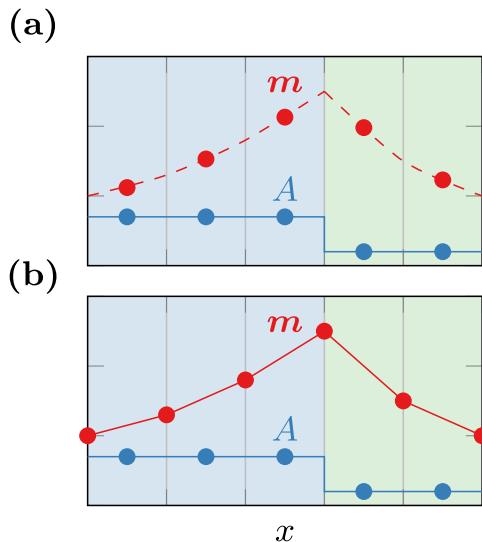
$$\mathbf{H}_{\text{ext}}(\theta, \phi) = H_c \begin{pmatrix} \sin(\theta) \cos(\phi) \\ \sin(\theta) \sin(\phi) \\ \cos(\theta) \end{pmatrix}, \quad (15)$$

$$\mathbf{H}_{\text{eff}} = \mathbf{H}_{\text{aniso}} + \mathbf{H}_{\text{exchange}} + \mathbf{H}_{\text{ext}}, \quad (16)$$

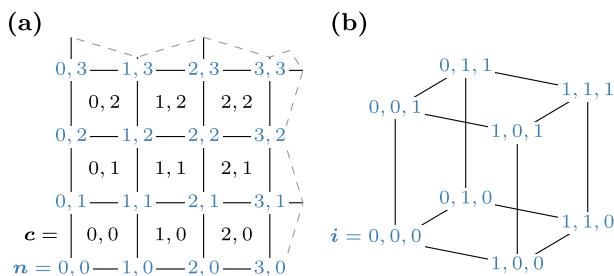
$$\mathcal{L}(\theta, \phi) = \int_{\Omega} \|\mathbf{m}(T) - \mathbf{m}_{\text{target}}\| \, d\mathbf{x} \quad (17)$$

with  $\mathbf{m}(t)$  being constrained by the LLG (1). A shortened code listing demonstrating the setup for this inverse problem is provided in Lst. 2. In this simple optimization, convergence is achieved after 30–50 gradient-descent steps, see Fig. 4b. NeuralMag computes the gradient of the objective function by performing one forward and one backward simulation of the dynamic problem.

Finally, a recently published application study by Voronov et al.<sup>31</sup> demonstrates that NeuralMag also scales to highly complex, fully dynamic topology-magnetization tasks: the framework tackled (i) the inverse design of a Stoner-Wohlfarth nanoparticle that converged in  $\approx 250$  adjoint iterations and (ii) a  $1 \mu\text{m} \times 1 \mu\text{m}$  magnonic demultiplexer, parameterised by 400 radial-basis functions on a  $512 \times 64 \times 1$  mesh, which achieved more than an order-of-magnitude spin-wave-contrast between 2.6 GHz and 2.8 GHz channels after only  $\approx 100$  optimization steps—underscoring the robustness and scalability of our approach for real-world, nonlinear device geometries.



**Fig. 5 | Discretization strategies.** Illustration of the discretization of the magnetization  $\mathbf{m}$  and the material parameter  $A$  for a one-dimensional representation of a two-phase magnetic system, using different numerical schemes: **a** Standard finite differences: Both the material parameter and the magnetization are discretized with a single value per simulation cell. The magnetization degrees of freedom are treated as sample points of a continuous function. **b** Finite elements: The material parameters are discretized using piecewise constant functions, while the magnetization is represented as piecewise affine, with degrees of freedom located at the vertices.



**Fig. 6 | Numbering of degrees of freedom.** Cell and vertex numbering using multiindices for nodal finite differences. **a** Two-dimensional representation of global cell and node indices denoted by  $\mathbf{c}$  (black) and  $\mathbf{n}$  (blue). **b** Three-dimensional representation of local vertex numbering denoted by index  $i$  (blue).

## Methods

### Nodal finite-difference scheme

Existing micromagnetic simulation software usually employs either a finite-difference discretization on regular grids<sup>22,23</sup> or a finite-element discretization on irregular grids<sup>20,22</sup>. The use of regular cuboid grids in the case of finite-difference micromagnetics allows for a very efficient computation of the demagnetization field by means of an FFT-accelerated convolution. On the other hand, the finite-element method allows for the accurate modeling of complex structures due to the use of irregular meshes.

Moreover, finite-element micromagnetics provides a more subtle but sometimes highly relevant advantage over finite-difference micromagnetics: In finite-element micromagnetics, the magnetization is usually explicitly defined on each mesh-vertex, whereas standard finite-difference tools store one magnetization vector per simulation cell, which is typically taken to be the magnetization in the center of this cell. While this difference appears to be insignificant for the micromagnetic modeling in the bulk, it plays a crucial role when considering material interfaces where the magnetization is subject to boundary and jump conditions. Consider e.g., the exchange jump condition (6), which prescribes a discontinuity in the first spatial derivative of the magnetization across material interfaces. Choosing the degrees of freedom of the magnetization in the cell centers, as illustrated in Fig. 5a,

requires a careful treatment of the boundary conditions that are defined on the vertices<sup>29</sup>. Similar considerations apply to interfacial energy contributions such as the RKKY coupling between two ferromagnetic layers<sup>33</sup>. Inaccurate modeling of the boundary conditions can lead to a loss of convergence order and consequently, introduce significant numerical errors. Introducing more energy contributions depending on surface integrals or spatial derivatives of  $\mathbf{m}$  results in more complex boundary conditions<sup>22</sup> that become unfeasible to handle in standard finite-difference micromagnetics.

In contrast, the finite-element method allows for the choice of tailored function spaces for the magnetization and material parameters, as shown in Fig. 5b. Moreover, the inherently variational nature of the finite-element method allows to solve for the effective-field contributions by directly considering the variation of the energy<sup>22</sup>, resulting in the weak form

$$-\int_{\Omega} \mu_0 M_s \mathbf{H}(\mathbf{m}) \cdot \mathbf{v} \, dx = \delta E(\mathbf{m}, \mathbf{v}) \quad \forall \mathbf{v} \in V \quad (18)$$

with  $V$  being a sufficiently smooth function space referred to as the test space. By a proper choice of function spaces for the material parameters and fields, this procedure does not require explicitly accounting for the boundary conditions at all.

**Local field terms.** The nodal finite-difference scheme proposed in this work applies the finite-element method for local field contributions on a regular cuboid grid. This enables the use of an FFT-accelerated demagnetization-field computation as in standard finite-difference micromagnetics, see section “Demagnetization Field”, while providing the rigorous and accurate handling of material interfaces for all local field contributions due to finite-element modeling. In order to address the cells and nodes of the regular grid, we introduce multi-indices  $\mathbf{c}$ ,  $\mathbf{n}$ , and  $\mathbf{i}$  as

$$\mathbf{c} = (c_1, c_2, c_3) \in \{0, \dots, N_1 - 1\} \times \{0, \dots, N_2 - 1\} \times \{0, \dots, N_3 - 1\} \quad (19)$$

$$\mathbf{n} = (n_1, n_2, n_3) \in \{0, \dots, N_1\} \times \{0, \dots, N_2\} \times \{0, \dots, N_3\} \quad (20)$$

$$\mathbf{i} = (i_1, i_2, i_3) \in \{0, 1\} \times \{0, 1\} \times \{0, 1\} \quad (21)$$

with  $N_1$ ,  $N_2$ , and  $N_3$  being the number of simulation cells in the respective mesh dimension. The indices  $\mathbf{c}$  and  $\mathbf{n}$  are used to address simulation cells and nodes, respectively, according to the numbering introduced in Fig. 6a. The index  $\mathbf{i}$  either acts as a local vertex number in a simulation cell according to Fig. 6b or, more generally, as a relative index to address neighborships.

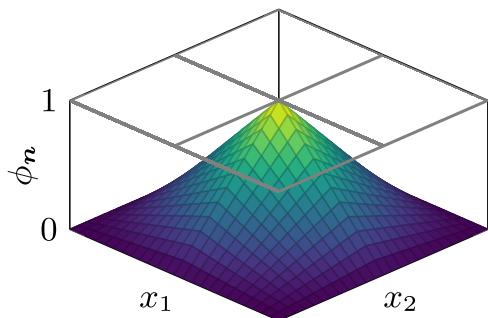
We discretize all continuous fields appearing in weak forms with standard piecewise polynomial and globally continuous basis functions  $\phi_n$  that form a nodal basis on the cuboid mesh. Each basis function  $\phi_n$  is defined per simulation cell in terms of reference basis functions  $\hat{\phi}_i$  as

$$\phi_n(\mathbf{x}) = \sum_i \hat{\phi}_i \left[ \begin{array}{c} x_1 / \Delta x_1 - n_1 + i_1 \\ x_2 / \Delta x_2 - n_2 + i_2 \\ x_3 / \Delta x_3 - n_3 + i_3 \end{array} \right] \quad (22)$$

with  $\Delta x_k$  being the simulation-cell size in the  $k$ -th dimension. The reference basis functions  $\hat{\phi}_i$  are defined on the reference unit cell  $\Omega_{\text{ref}} = [0, 1] \times [0, 1] \times [0, 1]$  as

$$\hat{\phi}_i(\mathbf{x}) = \mathbb{1}_{\Omega_{\text{ref}}}(\mathbf{x}) \begin{bmatrix} 1 - i_1 + (2i_1 x_1 - x_1) \\ 1 - i_2 + (2i_2 x_2 - x_2) \\ 1 - i_3 + (2i_3 x_3 - x_3) \end{bmatrix}. \quad (23)$$

where  $\mathbb{1}_{\Omega_{\text{ref}}}$  denotes the characteristic function of  $\Omega_{\text{ref}}$ , which evaluates to 1 if  $\mathbf{x} \in \Omega_{\text{ref}}$  and to 0 else. This restricts the support of the reference basis



**Fig. 7 | Nodal basis function.** Two-dimensional representation of a basis function  $\phi_n$  in nodal finite differences with a support spanning 4 simulation cells.

functions  $\hat{\phi}_i$  to the reference cell  $\Omega_{\text{ref}}$ . A 2D representation of a basis function is visualized in Fig. 7. Furthermore, we introduce vector basis functions as

$$\phi_{n,j} = \phi_n e_j \quad (24)$$

with  $e_j$  being the unit vector in direction  $j \in \{1, 2, 3\}$ . Continuous vector fields such as the magnetization  $\mathbf{m}$  and the effective field  $\mathbf{H}_{\text{eff}}$  are then discretized as

$$\mathbf{m} \rightarrow \mathbf{m}^h = \sum_{n,j} m_{n,j} \phi_{n,j} \quad (25)$$

with the superscript  $h$  denoting the discretized version of a field and coefficients  $m_{n,j}$  being the nodal values of the vector field  $\mathbf{m}$ .

For material parameters, such as the saturation magnetization  $M_s$ , we choose a piecewise constant function space in order to allow for the accurate modeling of rapid material interfaces. Namely, we define these parameters per simulation cell resulting in the following discretization

$$M_s \rightarrow M_s^h = \sum_c M_{s,c} \vartheta_c \quad (26)$$

with basis functions

$$\vartheta_c = \mathbb{1}_{\Omega_c}. \quad (27)$$

Replacing all fields with their discretized counterparts in the weak form (18) and testing with individual basis functions instead of arbitrary test functions yields the discretized weak form

$$-\int_{\Omega} \mu_0 M_s \mathbf{H}^h(\mathbf{m}^h) \cdot \phi_{n,j} \, d\mathbf{x} = \delta E(\mathbf{m}^h, \phi_{n,j}) \quad \forall \quad n, j. \quad (28)$$

For a given node  $\mathbf{n}$ , we split the variation  $\delta E(\mathbf{m}^h, \phi_{n,j})$  into its contributions from the eight simulation cells that share node  $\mathbf{n}$  and we address these cells by the local index  $\mathbf{i} \in \{0, 1\}^3$ . In general, the variation over a single simulation cell depends on the magnetization values of all nodes of this cell. Considering the three components of the magnetization, the contribution of the cell  $\mathbf{i}$  to the variation can be written as

$$\delta E_j^{*i} = F_{i,j}(m_{i,j'}) \quad \text{for} \quad \mathbf{i}' \in \{0, 1\}^3 \text{ and } j' \in \{1, 2, 3\} \quad (29)$$

where  $m_{i,j'}$  denotes all nodal values of the magnetization in cell  $\mathbf{i}'$ . If the energy  $E$  is quadratic in  $\mathbf{m}$ , the function  $F$  is linear in  $m_{i,j'}$  and can be described by a  $24 \times 24$  matrix considering the  $2^3 \cdot 3$  degrees of freedom defined by the index pairs  $\mathbf{i}, j$  and  $\mathbf{i}', j'$ . In the finite-element context, this matrix is usually referred to as the element matrix of the weak form.

In order to compute the variation at all nodes, we introduce the vector  $\delta \mathbf{E}$  with components  $\delta E_{n,j} = \delta E(\mathbf{m}^h, \phi_{n,j})$  and the auxiliary vectors  $\delta \mathbf{E}^{*i}$  containing the cell-wise variations according to (29) for all nodes.

Considering the node and cell numbering introduced in Fig. 6, the global node index is given by the global cell index and the relative node index as  $\mathbf{n}, \mathbf{i} = \mathbf{c} + \mathbf{i}$  resulting in

$$\delta E_{c+i,j}^{*i} = F_{i,j}(m_{c+i,j'}) \quad \text{for} \quad \mathbf{i}' \in \{0, 1\}^3 \text{ and } j' \in \{1, 2, 3\} \quad (30)$$

$$\delta \mathbf{E} = \sum_i \delta \mathbf{E}^{*i}. \quad (31)$$

Note that  $F$  only depends on the relative index  $\mathbf{i}$  and the component  $j$ . Eqs. (30) and (31) deliver a straight-forward strategy for a parallel evaluation over the cell index  $\mathbf{c}$ , see section “Form compilation”.

If the energy  $E$  depends on further fields, such as an external field or material parameters, the mapping function can be easily extended by adding additional arguments

$$F_{i,j}(m_{c+i,j'}, a_{c+i}^1, a_{c+i}^2, \dots, b_c^1, b_c^2, \dots) \quad (32)$$

where the variables  $a_{i+j}^l$  are the coefficients of arbitrary scalar fields discretized with nodal basis functions (22) and the variables  $b_i^l$  are the coefficients of arbitrary scalar fields discretized with cell basis functions (27). Since  $F_{j,k}$  does not explicitly depend on the cell index  $\mathbf{i}$ , it is fully determined by the integrand of the weak form (28) and the dimensions of a single simulation cell  $\Omega_i$ .

In order to determine the discretized effective field  $\mathbf{H}^h$ , the weak form requires the solution of a linear mass system defined by the left-hand side of Eq. (28). To avoid this costly procedure, we employ mass lumping to the left-hand side of Eq. (18) as described in Abert<sup>22</sup> resulting in

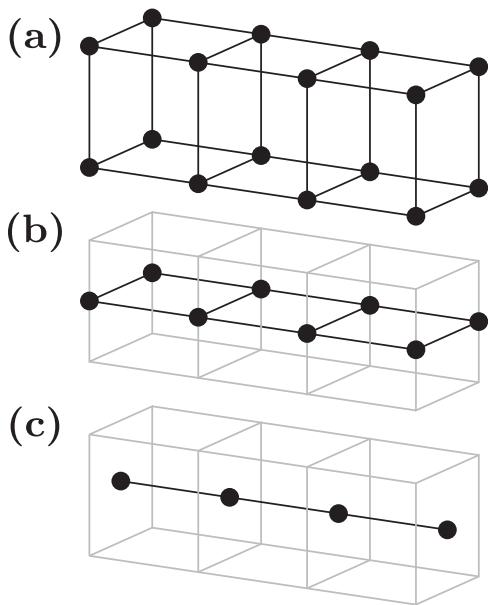
$$H_{n,j} = - \left[ \int_{\Omega} \mu_0 M_s^h \phi_n \, d\mathbf{x} \right]^{-1} \delta E_{n,j} \quad (33)$$

where the saturation magnetization  $M_s$  is discretized cell-wise according to Eq. (26).

The proposed method is applicable to any energy contribution whose density depends solely on the magnetization and its first-order spatial derivatives, such as Zeeman energy, crystalline anisotropies, and both symmetric and antisymmetric exchange interactions. Due to the regularity of the cuboidal grid, a matrix-free implementation of the presented scheme is straightforward. The local support of the basis functions results in a computational complexity of  $\mathcal{O}(N)$  for the evaluation of any local field term, with  $N$  being the number of simulation cells.

**Demagnetization field.** To compute the demagnetization field, we employ the well-established FFT-accelerated method commonly used in standard finite-difference micromagnetic simulations<sup>32</sup>. This algorithm calculates the demagnetization field generated by homogeneously magnetized cuboidal simulation cells arranged on a regular grid through fast convolution. Since this method requires both the magnetization and the resulting field to be specified for each simulation cell, we introduce a straightforward pre- and post-processing step. This procedure averages the values to transition between nodal and cell-centered discretizations efficiently. FFT-accelerated methods that operate directly on node-wise discretized magnetizations have been proposed in previous studies<sup>34,35</sup>. However, we opt for the standard method based on homogeneously magnetized cuboids due to its advantages in memory efficiency and computational performance, specifically for 2D computations where the FFT also reduces to two dimensions.

**Low-dimensional geometries.** Discretizing a mesh with  $N_1 \times N_2 \times N_3$  cells results in  $(N_1 + 1) \times (N_2 + 1) \times (N_3 + 1)$  degrees of freedom when using a nodal basis for the function discretization. In bulk system simulations, this introduces only a negligible overhead in comparison to standard finite-difference schemes, where the degrees of freedom are



**Fig. 8 | Low-dimensional discretization.** Representation of the degrees of freedom for a square-shaped rod using the following methods: **a** full three-dimensional discretization, **b** two-dimensional discretization with basis functions that are constant along the third dimension, and **c** one-dimensional discretization with basis functions that are constant along both the second and third dimensions.

equal to the number of simulation cells. However, a significant application area for micromagnetic simulations involves magnetic thin films, which are often discretized with just a single layer of simulation cells. In such cases, the 3D nodal discretization introduces a notable overhead, roughly doubling the computational cost, because it requires separate descriptions for the top and bottom surfaces of the thin film. This contrasts with standard finite differences, where the problem effectively reduces to a 2D formulation. By transitioning to 2D basis functions while maintaining full 3D integration in the weak form (28), the nodal finite-difference scheme can accurately describe magnetic thin films. This approach reduces the degrees of freedom to  $(N_1 + 1) \times (N_2 + 1) \times 1$ , making it more efficient for thin film simulations. Namely, the 2D basis function on the reference cell are chosen as

$$\hat{\phi}_i(\mathbf{x}) = \mathbb{1}_{\Omega_{\text{ref}}}(\mathbf{x}) [1 - i_1 + (2i_1 x_1 - x_1)] \\ [1 - i_2 + (2i_2 x_2 - x_2)] \quad (34)$$

with a 2D multiindex  $i = (i_1, i_2) \in \{0, 1\}^2$ . As illustrated in Fig. 8 this approach can also be generalized to 1D problems, leading to basis functions

$$\hat{\phi}_i(\mathbf{x}) = \mathbb{1}_{\Omega_{\text{ref}}}(\mathbf{x}) [1 - i + (2ix_1 - x_1)] \quad (35)$$

with a scalar index  $i \in \{0, 1\}$ .

### Implementation

NeuralMag is a Python library designed specifically for micromagnetic simulations, with a focus on high-performance tensor computations. A key feature of NeuralMag is its ability to operate with either PyTorch<sup>23</sup> or JAX<sup>24,36</sup> as a computational backend, allowing users to select the framework that best suits their needs. By leveraging these modern machine learning frameworks, NeuralMag achieves efficient computations on a variety of parallel hardware, including CPUs, GPUs, and TPUs. This versatility is complemented by the advantages these frameworks offer, such as optimized performance for large-scale simulations and built-in support for automatic differentiation, which simplifies solving inverse problems. Through the modular

design, the software is prepared to simplify the use of other computational backends in the future.

The use of either PyTorch or JAX as backends allows NeuralMag to fully exploit the unique strengths of each framework. PyTorch's `torch.compile()` feature enables JIT compilation, optimizing the computational workflow by reducing operation overhead and enabling kernel fusion for faster execution on compatible hardware. However, PyTorch currently has limitations when compiling complex functions, such as those involving the demagnetization field, which means `torch.compile()` can only be applied to certain field terms.

In contrast, JAX's `jit()` function can be applied to the entire right-hand side of the LLG equation. This capability allows JAX to significantly reduce Python overhead and leads to notable performance gains, particularly for smaller systems where the overhead would otherwise be a bottleneck.

Both backends support single- and double-precision computations, enabling NeuralMag to offer users flexibility in balancing computational speed with numerical accuracy according to the requirements of each simulation. The dual-backend approach ensures that NeuralMag can adapt to the user's preferred ecosystem while maintaining high computational efficiency and flexibility.

**Form compilation.** Listing 3. Symbolic definition of the exchange energy (3) in NeuralMag.

```
def e_expr(m, dim):
    A = Variable("material_A", "c" * dim)
    return (A * (
        m.diff(N.x).dot(m.diff(N.x)) +
        m.diff(N.y).dot(m.diff(N.y)) +
        m.diff(N.z).dot(m.diff(N.z))
    ) * dV(dim)
    )
```

Listing 4. Automatically generated code for the computation of the exchange field.

```
def h(dx, m, material_A, material_Ms, rho):
    h = torch.zeros_like(m)
    h[:-1, :-1, :-1, 0] += material_A[...]*rho[...]*(
        m[:-1, :-1, :-1, 0]*(
            0.2222222222222222*dx[0]*dx[1]/dx[2] +
            0.2222222222222222*dx[0]*dx[2]/dx[1] +
            0.2222222222222222*dx[1]*dx[2]/dx[0]
        ) +
        m[:-1, :-1, 1, 0]*(
            ...
        )
    )
    h[:-1, :-1, :-1, 1] += ...
    h[:-1, :-1, :-1, 2] += ...
    h[:-1, :-1, 1, 0] += ...
    ...
    return h / mass
```

At the heart of NeuralMag is a form compiler that translates a symbolic representation of a finite-element weak form into efficient tensor operations tailored to the chosen backend. For symbolic computation, NeuralMag leverages the Python library SymPy<sup>37</sup>. SymPy provides a powerful framework for representing the mathematical structures involved in micromagnetic simulations. Specifically, NeuralMag introduces custom SymPy symbols to represent functions that are discretized either node-wise or cell-wise, as described in the section “Nodal finite-difference scheme” of this paper. Users can define the weak form of the micromagnetic problem using SymPy's symbolic language, allowing them to work in an intuitive mathematical formulation.

In addition to defining weak forms symbolically, NeuralMag leverages SymPy to automatically perform the variation of a symbolic energy expression, allowing it to derive the corresponding weak form. This capability streamlines the process of converting complex energy functionals into their weak form representations. For instance, in Lst. 3, the exchange energy is defined symbolically using SymPy, demonstrating how users can express physical energy terms within the framework.

NeuralMag's form compiler processes the symbolic weak form and transforms it into the discrete mapping function  $F_{i,j}$  as defined in Eq. (30).

This transformation is achieved by applying Gauss quadrature to integrate over the finite elements, converting the weak form into a set of tensor operations—primarily multiplications and summations—that can be efficiently executed by the selected backend. The role of the relative cell index, as discussed in Eqs. (30) and (31) are handled by tensor slicing. This involves slicing along specific tensor dimensions by removing either the first `[1 : ]` or the last `[ : -1]` value of the tensor in that dimension, which is necessary for handling the spatial relationships between adjacent cells in the discretized domain. This systematic conversion of symbolic expressions into backend-specific tensor operations is key to NeuralMag’s high-performance computational capabilities. An example code snippet for the PyTorch backend, generated from the exchange energy defined in Lst. 3, is shown in Lst. 4. The generated function is highly optimized, as it operates solely on raw tensor objects without introducing any loops or conditional statements. This structure ensures that the function is ideally suited for optimization by the JIT compilers of both PyTorch and JAX. By avoiding control flow statements, the generated code can be compiled into efficient low-level machine instructions, maximizing performance on parallel hardware architectures.

**Dynamic attributes.** Listing 5. Example usage of dynamic attributes in NeuralMag.

```
>>> state = State(...)  
>>> state.a = 1.0  
>>> state.b = lambda a: 2.0 * a  
>>> state.c = 5.0  
>>> state.d = lambda b, c: b + c  
>>> print(state.d)  
7.0
```

Listing 6. Automatically generated function for the evaluation of the dynamic attribute `d`.

```
def lmda(a, c):  
    b = __b(a)  
    return __lmda(b, c)
```

NeuralMag introduces the concept of dynamic attributes through its state object, which allows attributes to be either tensors or functions that depend on tensors and return tensors. This flexible design enables dynamic relationships between attributes, where some can be defined as functions of others, with NeuralMag automatically managing these dependencies. For example, consider the code in Lst. 5: attributes `a`, `b`, `c`, and `d` are defined, where `b` depends on `a`, and `d` depends on both `b` and `c`. When `state.d` is accessed, NeuralMag resolves these dependencies, and the output is `7.0` because `d` is computed as the sum of `b` (which is `2 * a = 2.0`) and `c` (which is `5.0`). Importantly, instead of scalar values, any tensor can be used as an attribute, allowing for more complex operations on multidimensional data.

When defining such dynamic attributes, NeuralMag analyzes the function signatures to identify all dependencies in a recursive manner. It then generates a new Python function at runtime that only relies on pure tensors and eliminates any control structures, such as loops or conditionals, ensuring the function remains optimal for high-performance tensor computation. In the case of the example from Lst. 5, the dynamically created function looks like Lst. 6, where `d` depends on `b`. Although `b` is not explicitly listed in the function arguments, its dependency on `a` is automatically resolved within the body of the function. This approach simplifies the handling of complex dependencies while maintaining the computational efficiency needed for the PyTorch and JAX backends.

As an example, a material parameter such as the exchange constant `A` can be defined either as a regular, constant attribute or as a dynamic attribute, depending on other state variables such as the temperature and the time in order to simulate the magnetic response to a heat pulse. Either way, the value of `A` can be accessed by `state.material.A` and the code responsible for the computation of the exchange field does not have to be changed to account for a dynamic exchange constant.

**Automatic differentiation and time integration.** In the context of inverse problems, NeuralMag leverages automatic differentiation and efficient time integration to solve complex optimization tasks. Both PyTorch and JAX offer powerful automatic differentiation capabilities, which are crucial for computing gradients with respect to parameters in inverse problems. For time integration, NeuralMag integrates with `torchdiffeq`<sup>38</sup> (for PyTorch) and `diffra`<sup>39</sup> (for JAX), both of which provide support for solving ordinary differential equations (ODEs).

Time integration is essential in dynamic micromagnetic problems, where the system’s evolution must be accurately tracked. Both libraries support a variety of numerical schemes for time stepping, including Euler methods, Runge–Kutta methods (such as RK4), and adaptive solvers like the Dormand–Prince method. These methods ensure that NeuralMag can flexibly adapt to different accuracy and performance requirements in dynamic simulations.

For gradient-based optimization in inverse problems, NeuralMag supports both the adjoint method<sup>26</sup> and traditional backpropagation. The adjoint method is particularly well-suited for problems with long time horizons or large state spaces, as it computes gradients more efficiently by solving an adjoint ODE backward in time. Both `torchdiffeq` and `diffra` support the adjoint method for time integration, offering an efficient way to compute gradients when optimizing over dynamic systems. At the same time, they also allow for direct backpropagation through the time integration process, which can be more straightforward for shorter time intervals or simpler problems.

By combining automatic differentiation with advanced time integration techniques, NeuralMag can effectively tackle inverse problems in micromagnetic simulations, allowing users to optimize parameters while ensuring accurate numerical solutions over time.

## Discussion

In this paper, we have introduced NeuralMag, an open-source Python library for micromagnetic simulations that leverages modern machine learning frameworks such as PyTorch and JAX to achieve high performance. NeuralMag implements a novel nodal finite-difference discretization scheme, which provides a rigorous numerical description of continuous fields such as the magnetization as well as discontinuous material parameters. This approach is particularly useful for the accurate modeling of material interfaces while maintaining the same computational complexity as standard finite-difference schemes. Its performance is competitive with state-of-the-art micromagnetic simulation codes, yet it offers unparalleled flexibility due to its Python-based interface and support for optimized tensor operations on a variety of hardware platforms.

NeuralMag is especially well-suited for solving inverse problems, particularly those with time-dependent objectives, thanks to its ability to seamlessly compute gradients using automatic differentiation. This makes it a powerful tool for a wide range of optimization and simulation tasks in micromagnetics. NeuralMag is freely available<sup>28</sup>, making it accessible to the broader research community for further development and application.

## Data availability

The code of NeuralMag is publicly available<sup>28</sup>. The datasets generated and analysed during the current study can be reproduced by running the respective demo scripts provided in the NeuralMag repository.

## Code availability

The source code of NeuralMag is publicly available under the MIT License on GitLab at <https://gitlab.com/neuralmag/neuralmag><sup>28</sup>. Comprehensive documentation, an API reference, and tutorials are available at <https://neuralmag.gitlab.io/>. NeuralMag can be installed via standard package managers such as pip or conda. Users and community members are encouraged to contribute to the codebase, tutorials, and documentation. Continuous integration workflows are set up using GitLab CI/CD to automatically run tests after every code change. These tests are run with both the PyTorch and the JAX backend and include unit, integration, and system

tests, covering both the standard problems and benchmarks. The repository includes all numerical problems discussed in this paper, as well as the code to reproduce the benchmarks.

Received: 19 November 2024; Accepted: 2 June 2025;

Published online: 21 June 2025

## References

1. Donahue, M. J. & Porter, D. G. OOMMF User's Guide, Version 1.0. Report No. 6376 <https://doi.org/10.6028/NIST.IR.6376> (Accessed June 5, 2025) (National Institute of Standards and Technology, 1999).
2. Bisotti, M.-A. et al. Fidimag—a finite difference atomistic and micromagnetic simulation package. *J. Open Res. Softw.* **6**, 22 (2018).
3. Vansteenkiste, A. et al. The design and verification of MuMax3. *AIP Adv.* **4**, 107133. <https://doi.org/10.1063/1.4899186> (2014).
4. Lepadatu, S. Boris computational spintronics—high performance multi-mesh magnetic and spin transport modeling software. *J. Appl. Phys.* **128**, 243902 <https://doi.org/10.1063/5.0024382> (2020).
5. Bruckner, F., Koraltan, S., Abert, C. & Suess, D. magnum.np: a PyTorch-based GPU-enhanced finite difference micromagnetic simulation framework for high-level development and inverse design. *Sci. Rep.* **13**, 12054 (2023).
6. Fischbacher, T., Franchin, M., Bordignon, G. & Fangohr, H. A systematic approach to multiphysics extensions of finite-element-based micromagnetic simulations: nmag. *IEEE Trans. Magn.* **43**, 2896–2898 (2007).
7. Kakay, A., Westphal, E. & Hertel, R. Speedup of FEM micromagnetic simulations with graphical processing units. *IEEE Trans. Magn.* **46**, 2303–2306 (2010).
8. Chang, R., Li, S., Lubarda, M., Livshitz, B. & Lomakin, V. FastMag: fast micromagnetic simulator for complex magnetic structures. *J. Appl. Phys.* **109**, 07D358 (2011).
9. Bisotti, M.-A. et al. FinMag: finite-element micromagnetic simulation tool (Version 1.0). *Zenodo* (2018).
10. Abert, C., Exl, L., Bruckner, F., Drews, A. & Suess, D. magnum.fe: a micromagnetic finite-element simulation code based on FEniCS. *J. Magn. Magn. Mater.* **345**, 29–35 (2013).
11. Körber, L., Quasebarth, G., Otto, A. & Kákay, A. Finite-element dynamic-matrix approach for spin-wave dispersions in magnonic waveguides with arbitrary cross section. *AIP Adv.* **11**, 095006 (2021).
12. Bruckner, F. et al. Solving large-scale inverse magnetostatic problems using the adjoint method. *Sci. Rep.* **7**, 40816 (2017).
13. Abert, C. et al. A fast finite-difference algorithm for topology optimization of permanent magnets. *J. Appl. Phys.* **122**, 113904 (2017).
14. Huber, C. et al. Topology optimized and 3D printed polymer-bonded permanent magnets for a predefined external field. *J. Appl. Phys.* **122**, 053904 (2017).
15. Papp, A., Porod, W. & Csaba, G. Nanoscale neural network using non-linear spin-wave interference. *Nat. Commun.* **12**, 6422 (2021).
16. Wang, Q., Chumak, A. V. & Pirro, P. Inverse-design magnonic devices. *Nat. Commun.* **12**, 2636 (2021).
17. Yan, Z., Xing, Y. & Han, X. Inverse design of magnonic filter. *J. Magn. Magn. Mater.* **563**, 169976 (2022).
18. Zenbaa, N. et al. A universal inverse-design magnonic device. *Nat. Electron.* **8**, 106–115 (2025).
19. Brown, W. F. Micromagnetics, domains, and resonance. *J. Appl. Phys.* **30**, S62–S69 (1959).
20. Schrefl, T. et al. Numerical micromagnetics (Finite element method). *Handb. Magn. Adv. Magn. Mater.* **2**, 765–795 (2007).
21. Skomski, R. *Simple Models of Magnetism*. <https://doi.org/10.1093/acprof:oso/9780198570752.001.0001> (Oxford University Press, 2008).
22. Abert, C. Micromagnetics and spintronics: models and numerical methods. *Eur. Phys. J. B* **92**, 1–45 (2019).
23. Paszke, A. et al. Pytorch: an imperative style, high-performance deep learning library. In *Proc. 33rd International Conference on Neural Information Processing Systems*. 8026–8037 (Curran Associates Inc., 2019).
24. Bradbury, J. et al. Jax: Autograd and XLA. *Astrophysics Source Code Library* ascl-2111 <https://ui.adsabs.harvard.edu/abs/2021ascl.soft11002B> (2021).
25. Li, S. et al. PyTorch distributed. *Proc. VLDB Endow.* **13**, 3005–3018 (2020).
26. Chen, R. T., Rubanova, Y., Bettencourt, J. & Duvenaud, D. K. Neural ordinary differential equations. In: S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (eds.) *Advances in Neural Information Processing Systems*. Vol. 31 (NeurIPS, 2018).
27. μMAG standard problem #4. <https://www.ctcms.nist.gov/~rdm/std4/spec4.html> (2024).
28. NeuralMag: an open-source nodal finite-difference code for inverse micromagnetics. <https://gitlab.com/neuralmag/neuralmag> (2024).
29. Heistracher, P., Abert, C., Bruckner, F., Schrefl, T. & Suess, D. Proposal for a micromagnetic standard problem: domain wall pinning at phase boundaries. *J. Magn. Magn. Mater.* **548**, 168875 (2022).
30. Rozvany, G. The SIMP method in topology optimization—theoretical background, advantages and new applications. In *Proc. 8th Symposium on Multidisciplinary Analysis and Optimization*, Vol. 4738. <https://doi.org/10.2514/6.2000-4738> (American Institute of Aeronautics and Astronautics, 2000).
31. Voronov, A. A. et al. Inverse-design topology optimization of magnonic devices using level-set method. *npj Spintronics* **3**, 1–8 (2025).
32. Miltat, J. E & Donahue, M. J. Numerical micromagnetics: finite difference methods. In *Handbook of Magnetism and Advanced Magnetic Materials*. (eds Kronmülle, H., & Parkin, S.) 742–764 (John Wiley & Sons, 2007).
33. Suess, D., Koraltan, S., Slanovc, F., Bruckner, F. & Abert, C. Accurate finite-difference micromagnetics of magnets including RKKY interaction: analytical solution and comparison to standard micromagnetic codes. *Phys. Rev. B* **107**, 104424 (2023).
34. Berkov, D. V., Ramstöck, K. & Hubert, A. Solving micromagnetic problems: towards an optimal numerical method. *Phys. Status Solidi* **137**, 207–225 (1993).
35. Ramstöck, K., Leibl, T. & Hubert, A. Optimizing stray field computations in finite-element micromagnetics. *J. Magn. Magn. Mater.* **135**, 97–110 (1994).
36. Frostig, R., Johnson, M. J. & Leary, C. Compiling machine learning programs via high-level tracing. *Syst. Mach. Learn.* **4**, 9 (2018).
37. Meurer, A. et al. SymPy: symbolic computing in Python. *PeerJ Comput. Sci.* **3**, e103 (2017).
38. Kidger, P., Chen, R. T. & Lyons, T. J. "Hey, that's not an ODE": faster ODE adjoints via seminorms. In: *Proc. ICML*, 5443–5452 (PMLR, 2021).
39. Kidger, P. *On Neural Differential Equations*. Ph.D. thesis (University of Oxford, 2021).

## Acknowledgements

This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/P34671, 10.55776/I6068, 10.55776/PAT3864023, and 10.55776/PIN1434524. For open access purposes, the author has applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission. This work has been supported by the Horizon Europe research and innovation program through MaMMoS grant agreement No 101135546 and Marie Skłodowska Curie grant agreement No 101152613. We gratefully acknowledge the wedding of the Koraltans for fruitful discussions and a great time, which led to this publication.

## Author contributions

C.A. conceived the project and developed the software architecture as well as the form compiler of NeuralMag. F.B. and R.A. contributed to coding

efforts and performance optimization. R.K., P.F., T.S., and D.S. provided key input in the development of the numerical formalism. A.V., A.C., and S.K. contributed to the development and testing of the framework's inverse problem-solving capabilities. S.A.P., S.H., M.L., and H.F. enhanced the codebase and established automatic code formatting and testing workflows. C.A. wrote the manuscript with contributions from all authors. All authors discussed the results and gave feedback on the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to C. Abert.

**Reprints and permissions information** is available at

<http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025