

# Scalable physical source-to-field inference with hypernetworks

Anonymous authors

Paper under double-blind review

## Abstract

We present a generative model that amortises computation for the field and potential around e.g. gravitational or electromagnetic sources. Exact numerical calculation has either computational complexity  $\mathcal{O}(M \times N)$  in the number of sources  $M$  and evaluation points  $N$ , or requires a fixed evaluation grid to exploit fast Fourier transforms. Using an architecture where a hypernetwork produces an implicit representation of the field or potential around a source collection, our model instead performs as  $\mathcal{O}(M + N)$ , achieves relative error of  $\sim 4\% - 6\%$ , and allows evaluation at arbitrary locations for arbitrary numbers of sources, greatly increasing the speed of e.g. physics simulations. We compare with existing models and develop two-dimensional examples, including cases where sources overlap or have more complex geometries, to demonstrate its application. The code for these models and experiments is available at this location.

## 1 Introduction

Physics-informed machine learning has proven to be a useful tool for modelling functions. In this work, we apply it to the problem of learning fields. This poses extra methodological challenges, since not only do we need to incorporate physical constraints, but we also require that we can compose the output functions linearly corresponding to the superposition of fields. Fields are the critical observable in most disciplines of physics. The gravitational field, magnetic and electric fields, as well as various quantum mechanical fields, have provided fundamental advances in the description of the physical world. Evaluating fields numerically is also crucial to the development of real-world technologies, e.g. in renewable or fusion energy systems. However, current simulation approaches are bounded by their scaling when faced with a high number of field sources (i.e. non-pointlike objects) or fine-grained field resolution, for which seminal numerical methods—the Fast Multipole Method (FMM) and Fast Fourier Transform (FFT) algorithms (Greengard & Rokhlin, 1987; Engheta et al., 1992)—either impose restrictive assumptions or fail to efficiently handle arbitrary source configurations.

We propose to tackle *source-to-field* inference—i.e. generating the field at any point around or inside an extended physical source, given its properties—using statistical learning, and in doing so, we uncover a set of distinct challenges that require innovative solutions. First, the model must exhibit exceptional flexibility, being capable of learning a function space that accurately represents fields generated by an arbitrary number of sources with varying properties. This demands a framework that goes beyond traditional grid-based learning approaches and that adapts to arbitrary source configurations. Second, we must rigorously maintain the *principle of superposition*. In this context, this means ensuring that each source contributes to the field in a manner that is both independent and linear. This requirement is not just a mathematical convenience but a physical necessity, reflecting the fundamental linear nature of how fields interact and combine. Lastly, the model must be informed by established principles of physics, for example Maxwell’s equations, which for magnetic fields guarantee a divergence-free field.

To address these challenges, we propose a general architecture encompassing two models and two baseline variants (cf. Table 1), designed to:

- generate continuous field functions at arbitrary locations using hypernetworks (Ha et al., 2017);

Property	Models		Baseline models		
	FOURIER	FC+ILR	FC+INR	Linear	Exact
$\mathcal{O}(M + N)$ scaling	✓	✓	✓	✓	×
Principle of superposition	✓	✓	×	✓	✓
Admits physical interpretation	✓	×	×	×	✓

Table 1: Comparison of desiderata across different explored models within the template architecture. The specification of models and baseline models is given in Section 5.2: FC+ILR is a fully-connected network with a hypernetwork-trained final layer (an ‘implicit linear representation’); FC+INR is a standard fully-connected implicit neural representation hypernetwork.

- achieve linear scaling with the number of sources and evaluation points by designing hypernetworks whose outputs are combined linearly; and
- incorporate relevant physical properties, such as enforcing divergence-free fields through novel Fourier hypernetworks.

By satisfying these criteria, our models significantly alleviate the computational cost of high-fidelity field simulations and open new directions for physical modelling through statistical learning. This approach advances the intersection of machine learning and physics, enhancing the applicability of generative models to real-world physical systems.

## 2 Background

### 2.1 Source-to-Field Inference

Our focus is on classical physical fields: the magnetic field generated by any electrical current or source with a magnetisation vector  $\mathbf{M}$ , or similarly an electrical field generated by electric charges; and the gravitational field generated by any object with a mass. In this work, we take as an example the magnetic field, but the results generalise directly to other fields.

It can be shown that the magnetic field  $\mathbf{H}$  at a position  $\mathbf{r}$  generated by an object with a magnetisation at position  $\mathbf{r}'$  is given by (Jackson, 1999)

$$\mathbf{H}(\mathbf{r}|\mathbf{M}, \mathbf{V}') = -\frac{1}{4\pi} \int \underbrace{\left( \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \nabla \cdot \right)}_{\equiv \mathbb{D}(\mathbf{r} - \mathbf{r}')} \mathbf{M}(\mathbf{r}') dV', \quad (1)$$

where the term in parentheses acts like an operator  $\mathbb{D}$  taking the divergence ( $\nabla \cdot$ ) of the source magnetisation  $\mathbf{M}$ , and where the integral is evaluated over the source (with volume  $V'$ ). The mathematical action of the operator is to take the magnetisation vector at each location  $\mathbf{r}'$  within the source and yield a rotated and scaled version of it at  $\mathbf{r}$ . In three dimensions, it is a real symmetrical  $3 \times 3$  tensor with components (Smith et al., 2010a, Eqs. A5–6)

$$D_{ij}(\mathbf{r} - \mathbf{r}') = \frac{\delta_{ij}}{|\mathbf{r} - \mathbf{r}'|^3} - \frac{3(r - r')_i(r - r')_j}{|\mathbf{r} - \mathbf{r}'|^5}, \quad (2)$$

where the indices  $i, j$  run pairwise over the Cartesian coordinates  $x, y, z$ . That is, the matrix  $\mathbb{D}$  is a purely *geometric* object. The total field at  $\mathbf{r}$  is then the superimposed contribution integrated over the spatial extent of the source.

For a *uniformly* magnetised source, like a permanent magnet or a magnetic domain, the magnetisation  $\mathbf{M}_0$  is constant across the object, and so the magnetic field vector at (a single)  $\mathbf{r}$  is given by

$$\mathbf{H}(\mathbf{r}|\mathbf{M}, \mathbf{V}') = -\left( \frac{1}{4\pi} \int \mathbb{D}(\mathbf{r} - \mathbf{r}') dV' \right) \cdot \mathbf{M}_0 = -\mathbb{N}(\mathbf{r}, \mathbf{V}') \cdot \mathbf{M}_0, \quad (3)$$

where by linearity  $\mathbb{N}$  remains a real, symmetric  $3 \times 3$  object known as the *demagnetisation tensor*, the mathematical operator  $(\cdot)$  is matrix multiplication and  $\mathbf{V}'$  is the geometry of the source (i.e. location  $\mathbf{r}'$ , volume, and shape). There exist closed-form expressions for the tensor  $\mathbb{N}$  for ellipsoids (Joseph & Schlömann, 1964), prisms (Aharoni, 1998), tetrahedra (Nielsen et al., 2019), and cylinders (Nielsen & Bjørk, 2020; Slanovc et al., 2022) among others. The same principles apply for the gravitational field, with the divergence of the magnetisation replaced by the density of the object.

Since most objects consist of uniformly magnetised domains, for calculations it is advantageous to consider the magnetic field at a point  $\mathbf{r}_n$  as being generated by a collection of  $M$  sources

$$\mathbf{H}(\mathbf{r}_n|\mathbf{M}, \mathbf{V}') = - \sum_{m=1}^M \mathbb{N}(\mathbf{r}_n, \mathbf{V}'_m) \cdot \mathbf{M}_m \quad \text{for } n = 1, \dots, N, \quad (4)$$

where  $m$  is indexing the individual magnetic sources, the index  $n$  specifies the evaluation point, and  $\mathbf{M}_m$  indicates the magnetisation of the  $m$ th source. If the field must be computed at  $N$  points, the overall cost scales as  $\mathcal{O}(M \times N)$ , since each of the  $N$  evaluations requires summing over the  $M$  sources with a distinct demagnetisation tensor termed  $\mathbb{N}_{mn}$ . This becomes particularly problematic in dynamical systems, where sources evolve over time and the field must be recomputed repeatedly. While the demagnetisation tensor  $\mathbb{N}$  only has to be computed once if the geometry remains constant, the calculation of this magnetic field is numerically the most time consuming task in simulations of e.g. the time evolution of a collection of magnetic spins (Abert et al., 2013), and various computational frameworks have been developed for evaluating the aggregate field from many sources at many points by geometric decomposition (Bjørk et al., 2021; Ortner & Coliado Bandeira, 2020; Liang et al., 2023).

## 2.2 Hypernetworks

The computational cost of evaluating Eq. 4 motivates the use of flexible machine-learning surrogates that can represent the field function directly, allowing field values to be queried anywhere without recomputing each demagnetisation tensor.

Hypernetworks provide precisely such a mechanism. Originally introduced as neural networks that generate the weights of another network (Ha et al., 2017), they enable model parameters to be predicted dynamically rather than stored as fixed quantities. This allows to model entire functions—such as a magnetic field—conditioned on physical context like magnetisation. Hypernetworks have powered advances in few-shot learning (Bertinetto et al., 2016), meta-learning (Munkhdalai & Yu, 2017), generative modelling (Dupont et al., 2022; Koyuncu et al., 2023; Peis et al., 2025), neural architecture search (Brock et al., 2018), and continuous scene representations, as in HyperNeRF (Park et al., 2021), where weights are conditioned on scene deformations.

Recent works have also demonstrated the potential of hypernetworks for modelling physical systems, where they can serve as fast surrogates for time-evolving or spatially distributed processes (Wan et al., 2023). In this work, we leverage hypernetworks to generate continuous field representations conditioned on a set of physical sources, enabling scalable and flexible source-to-field inference. Unlike traditional applications focused on discrete outputs or network weights, our approach produces entire continuous functions, making it well-suited to physical domains where evaluation at arbitrary spatial locations is required.

## 3 Method

The computational bottlenecks in classical field evaluation motivate the use of a learned surrogate model that can simultaneously i) encode an arbitrary number of sources into an efficient representation; and ii) impute the field at any  $\mathbf{r}_n$  for parameterised shapes, and ultimately from a learned embedding of sources. To address this, we propose a hypernetwork-based surrogate that maps source properties to continuous field functions, allowing evaluation at arbitrary locations and aggregation across sources.




Figure 1: In the template architecture, a basis function generator  $\psi$  expands the  $d$ -dimensional spatial coordinate  $\mathbf{r}_n$  in a (fixed or learnable) basis to specified order  $L$ . The coefficients of the expansion  $\mathbf{a}$  are learned as a hypernetwork of the source geometry and magnetisation, that is additive across the sources, giving the cumulative magnetic field or scalar potential at  $\mathbf{r}_n$ . It follows that once the basis weights  $\mathbf{a}$  are accumulated for all  $M$  sources, only  $\psi(\mathbf{r}_n)$  has to be recalculated for any other evaluation point.

For magnetic fields, the functions  $\mathbf{N}$  and  $\mathbf{H}$  are continuous and differentiable, and so are likely candidates for approximation by a network trained on data examples generated from the analytically known forms for  $\mathbf{H}(\mathbf{r}_n|\mathbf{M}, \mathbf{V}')$ .

The challenge, therefore, lies in breaking the quadratic coupling between sources and evaluation points: instead of explicitly computing the demagnetisation tensor at every point  $n$  for every source  $m$ , as in Eq. 4, we seek approximations for  $\mathbf{H}(\mathbf{r}|\mathbf{M}, \mathbf{V}')$  that condition on a permutation-invariant aggregation of arbitrary sources, such that the field can be inferred in  $\mathcal{O}(M + N)$ . To this end, we propose the approximation

$$\mathbf{H}(\mathbf{r}_n|\mathbf{M}, \mathbf{V}') \approx f\left(\mathbf{r}_n \left| \sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m) \right. \right) \quad \text{for } n = 1, \dots, N, \quad (5)$$

where  $f$  and  $g$  (cf. Fig 1) are functions to be specified. In our architecture,  $g$  is implemented as a hypernetwork that outputs the weights or parameters for a field representation model  $f$ , which evaluates the field at query locations. We interpret  $g$  as a representation for the sources, and it should be constructed so that a single representation of the source *collection*, i.e. the sum of the individual source embeddings, can accurately condition the approximating model  $f$  for the field. These arbitrary collections of magnetic sources, with each source input with features like source position, volume, shape, and magnetisation, generate a single additive fixed-length vector from an arbitrary number of sources, rather than applying the network many times in parallel to the features of each input source. Especially note that  $\mathbf{r}_n$  becomes independent of the summation index in Eq. 5.

Given access to a set of source configurations and their corresponding field observations, we train our model by minimizing the Huber loss via stochastic gradient descent:

$$\mathcal{L}_{\mathbf{H}} = \mathcal{L}_{\delta}\left(f\left(\mathbf{r}_n \left| \sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m) \right. \right) - \mathbf{H}(\mathbf{r}_n)\right), \quad \delta = 1, \quad (6)$$

where the loss is averaged over all observation locations  $\{\mathbf{r}_n\}_{n=1}^N$  and across mini-batches of source configurations. The use of the Huber loss maintains differentiability near zero residuals.

In the following sections, we explore different instantiations of this general framework. We begin by introducing a formulation that leverages the conservative nature of certain physical fields, enabling indirect field evaluation via scalar potentials. Next, we present additive hypernetworks, which enforce the principle of superposition by requiring  $f$  to be linear in the aggregated source representation  $g$ . We then propose a physically motivated variant of  $f$  based on Fourier expansions, before turning to fully connected alternatives that relax these inductive biases to achieve greater expressivity.



### 3.1 Direct and indirect evaluation of the field

An additional property for fields can be exploited to ease training. Because the gravitational field, as well as the magnetic field in the absence of currents, is conservative, it can be derived from scalar-valued functions  $\phi$  called *potentials*, where  $\mathbf{H} = -\nabla\phi$ . This observation allows us to model the field indirectly: instead of parameterising  $\mathbf{H}$  itself, we can parameterise a scalar potential  $\hat{\phi}$  and recover the field via automatic differentiation:

$$f\left(\mathbf{r}_n \left| \sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m) \right. \right) = -\nabla \hat{\phi}\left(\mathbf{r}_n \left| \sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m) \right. \right) \quad (7)$$

This ensures that the resulting field approximation is strictly conservative by construction, and may present a superior target for learning. Additionally, since our training data are generated using MagTense (Björk et al., 2021), which provides both the true field and its associated potential, we can incorporate this additional supervision into the loss function:

$$\mathcal{L} = \gamma_{\mathbf{H}} \cdot \mathcal{L}_{\mathbf{H}} + \gamma_{\nabla\phi} \cdot \mathcal{L}_{\nabla\phi} \quad (8)$$

where  $\mathcal{L}_{\mathbf{H}}$  and  $\mathcal{L}_{\nabla\phi}$  are Huber losses with  $\delta = 1$  applied to the field and potential predictions, respectively, and  $\gamma_{\mathbf{H}}, \gamma_{\nabla\phi} \in \mathbb{R}_+$  are scaling weights that balance their contribution.

Additional physical constraints could influence our choices for the loss function. Because the magnetic field is divergence-free, the scalar potential satisfies Laplace’s equation, i.e.

$$0 = \nabla \cdot \mathbf{H} = \nabla \cdot (-\nabla\phi) = -\nabla^2\phi; \quad (9)$$

everywhere except the boundary of the source, where  $\nabla^2\phi = -\nabla \cdot \mathbf{M}$  (Griffiths, 2013, Eq. 6.23). This knowledge could be incorporated via the loss (e.g., as in Pollok et al., 2023). However, in this work, we employ a loss function solely based on the known field / potential values in the provided data.

### 3.2 Additive hypernetworks for field inference

There are reasons to suppose that a representation like this might be possible. We know that physical sources obey the principle of superposition, contributing independently and linearly to the field at all locations. In the language of equivariance, this is to say linear additivity implies the sources will be permutation-invariant<sup>1</sup>. We provide additional discussion of this structure in Appendix B.2.

To guarantee the principle of superposition, then, we must require that  $f$  be linear in  $\sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m)$ , which is a significant constraint on possible architectures. This constraint ensures that the contribution of each source remains independent and additive in the latent space, faithfully mirroring the underlying physical laws. Geometrically, this implies that the space of source collections is a vector space, and that any  $g_m$  is a coordinate in the space, weighting an abstract basis function representation  $f$  of the field around source collections. The architecture set out in Fig. 1 implements this construction.

We implement  $g$  as a fully-connected hypernetwork that outputs the parameters of the basis function representation  $f$ . We leave open the specific parameterisation of  $f$ , and consider both a physically motivated choice and a more expressive, data-driven alternative to compare their performance.

### 3.3 Fourier Hypernetworks

An additive hypernetwork architecture guarantees the principle of superposition. In our initial setup, we focus further on the physical desirability of orthogonal decomposition of the potential, choosing a set of basis functions  $\psi_l = \{\cos(\mathbf{k}_l \cdot \mathbf{r}), \sin(\mathbf{k}_l \cdot \mathbf{r})\}$  corresponding to a Fourier expansion of the scalar potential in  $d$ -dimensions,

$$\hat{\phi}(\mathbf{r}_n) = \sum_{l=0}^{L-1} A_l \cos(\mathbf{k}_l \cdot \mathbf{r}_n) + B_l \sin(\mathbf{k}_l \cdot \mathbf{r}_n). \quad (10)$$

<sup>1</sup>Indeed Eq. 5 follows the form of a DeepSets architecture (Zaheer et al., 2017), however, we use  $f$  and  $g$  for the functions because  $\phi$  is needed for the scalar potential.

In this model,  $f$  corresponds to the Fourier basis evaluated at query locations, while  $\sum_{m=1}^M g(\mathbf{M}_m, \mathbf{V}'_m)$  produces the corresponding coefficients  $\{A_l, B_l\}$  for wavenumbers  $\mathbf{k}_l$  up to order  $(L-1)$  that modulate this basis.

To understand why this reduces the complexity, note that the terms  $\cos(\mathbf{k}_l \cdot \mathbf{r}_n)$  and  $\sin(\mathbf{k}_l \cdot \mathbf{r}_n)$  depend only on where the potential is to be evaluated, while all dependence on the configuration of the sources generating the potential, i.e. their geometry and magnetisation, is enclosed in the Fourier coefficients  $A_l$  and  $B_l$ .

As potentials are additive, and as the evaluation points stay constant in static problems, the Fourier coefficients  $A_l$  and  $B_l$  must be sums of the Fourier coefficients of each source. This means that if we compute these for every source, an operation of order  $M$ , and after adding them, we then only need to use these coefficients to evaluate the potential in the desired points, an operation of order  $N$ .

Rather than computing these coefficients via a Fourier integral, we amortise them through the hypernetwork, conditioned on the source configuration. We provide concrete examples of this construction in both 1-D and 2-D in Sec. 5.

Note that, while this is a non-linear function in the field location  $\mathbf{r}_n$ , it is linear in the weights as required to make the source collections a vector space. Unlike in a discrete Fourier transform, the wavenumbers  $\mathbf{k}$  need not be tied to the resolution or domain window where the field is evaluated, and there is no *requirement* that they form an integer sequence—they could e.g. be set as learnable parameters—however, choosing integrally spaced wavenumbers makes the basis functions orthogonal, which can help eliminate degeneracies in the training landscape. There is a geometric correspondence to the use of random Fourier features in fully-connected networks (Rahimi & Recht, 2007); the use of  $d_l$  random Fourier features projects the input coordinates into a high-dimensional random superspace spanned by  $d \leq d_l$  independent basis functions, while an explicit choice of  $d_l$  integrally-spaced wavenumbers projects to an orthogonal superspace of dimension  $d = d_l$ . Practical tuning considerations for this model are provided in Appendix C.2.

### 3.4 Networks without inductive biases

As an alternative, we relax the architectural constraints and allow the basis itself to be learned implicitly through a fully-connected network. In this setting,  $g$  no longer outputs coefficients for a fixed basis, but instead generates parameters that directly shape the representation capacity of  $f$ . To preserve the principle of superposition, we constrain  $\sum g_m$  to produce only the weights  $\mathbf{a}$  and bias  $b$  of the final linear layer of  $f$ . In this formulation, the activations of the penultimate layer, of dimension  $L$ , serve as learned basis functions  $\psi(\mathbf{r}_n)$ , and the potential can be expressed as

$$\hat{\phi}(\mathbf{r}_n) = \sum_{l=0}^{L-1} a_l \psi_l(\mathbf{r}_n) + b. \quad (11)$$

We refer to this construction as an *implicit linear representation* (FC+ILR), which preserves superposition while gaining expressivity through learnable basis functions. If we further relax this constraint,  $\sum g_m$  may generate the full parameter set of the network, resulting in a highly expressive model akin to *implicit neural representations* (INRs; Sitzmann et al., 2020), where a hypernetwork (Ha et al., 2017) outputs all weights and biases.

We compare three instantiations of the general hypernetwork architecture, each defined by a different choice of  $f$ : FOURIER, Eq. 10, where  $f$  is constructed from orthogonal basis functions and the hypernetwork outputs the corresponding Fourier coefficients; FC+ILR, Eq. 11, where  $f$  is a fully-connected network and the hypernetwork modulates only its final linear layer; and FC+INR, where the hypernetwork generates the full parameter set of  $f$ , albeit at the cost of losing superposition. Table 1 summarises the properties of these models, and in the experiments, we investigate their behaviour and performance.

## 4 Related Work

**Fast field evaluation and surrogate modelling.** A key contribution of our work is the scaling behaviour of our generative model, which operates with a computational complexity of  $\mathcal{O}(M + N)$ , compared to the

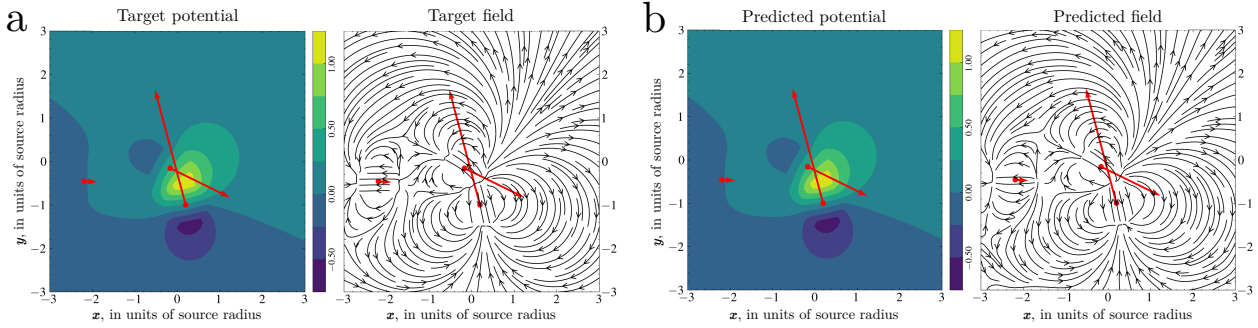


Figure 2: Magnetic potential and field for three finite circular sources (a) and their approximation by a small fully-connected network (b). Sources with locations  $\mathbf{r}'$  and direction and magnitude of the magnetisation  $\mathbf{M}$  shown by the red arrows, are randomly positioned within a  $[-3, 3] \times [-3, 3]$  domain, in units of the source radius, with the potential and field generated on a regular  $100^2$  grid.

standard  $\mathcal{O}(M \times N)$  in exact numerical simulation of fields such as the MagTense framework (Björk et al., 2021), or FFT-based methods requiring fixed evaluation grids. Unlike these methods, our model adopts amortised evaluation, significantly improving computational efficiency with only a small sacrifice in accuracy. This scaling perspective is also explored in the linearly constrained neural networks by Hendriks et al. (2021) for the problem of magnetic fields. Our work focuses on learning representations for physical sources, which can be used in a generative model, and this presents a different set of challenges and opportunities.

**Physics-informed machine learning and basis representations.** Our work differs from general physics-informed neural networks (PINNs; Raissi et al., 2019; Karniadakis et al., 2021), which embed PDE constraints into the loss but typically require retraining for each new source configuration. While related work, e.g. Schaffer et al. (2023), has explored neural surrogates for demagnetising fields, our method generalises across source geometries without retraining and enables continuous-space evaluation. We further depart from standard PINNs by explicitly embedding the principle of superposition for a collection of sources and optimising the linear coefficients of fixed (FOURIER) / learned (FC+ILR) basis functions. This makes our approach closer in spirit to classical Galerkin methods, while remaining end-to-end trainable.

**Conservative fields and architectural inductive biases.** Recent efforts to enforce physical laws through network design include divergence-free networks and conservation-driven approaches (Richter-Powell et al., 2022; Müller, 2023). These methods use differential forms or specialised architectural components to enforce symmetries. Our approach takes a complementary route: by constructing the magnetic field as the negative gradient of the scalar potential, we ensure a conservative field by design, which holds true in the absence of free currents, as in the present case of permanent magnets. Ghosh et al. (2023) similarly leverage harmonic representations in quantum systems; however, we focus on classical fields and emphasise linear additivity of sources rather than equivariant or gauge-theoretic formulations.

**Hypernetworks and physical models.** Our contribution extends the idea of hypernetworks (Ha et al., 2017), which have been increasingly applied in physical modelling (Pfaff et al., 2020; Cho et al., 2023; de Avila Belbute-Peres et al., 2021). We extend this direction by showing how additive hypernetworks can be designed to enforce superposition and enable scalable source-to-field inference. Unlike prior work using hypernetworks to model dynamical systems or PDE evolution, we target static field representations and develop a novel parameterisation (FOURIER, FC+ILR) that yields physically interpretable outputs with linear source aggregation.

## 5 Experiments

We implement experiments using finite magnetic sources, with fields and potentials derived from exact forms, either through the computational framework MagTense (Björk et al., 2021) or via direct implementation. Full details of the analytical expressions for spherical sources are given in Appendix B.1.

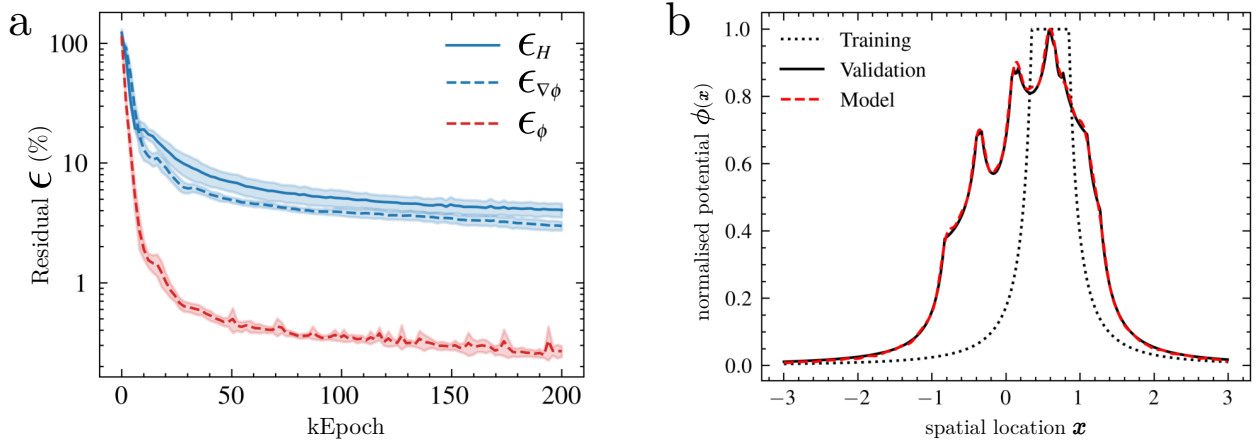


Figure 3: (a) Training curves for predicting magnetic potential (red) and field (blue), either directly (solid) or indirectly via the potential (dashed), shown as relative error percentages. (b) Output of a 1D FOURIER hypernetwork trained on single-source potentials (dotted black), successfully generalising to the combined potential of six randomly placed sources.

We express the performance of the model prediction ( $\hat{\phi}$  or  $\hat{\mathbf{H}}$ ) as a relative error

$$\epsilon_\phi := \text{median}_{n \in N} \left| \frac{\hat{\phi}(\mathbf{r}_n) - \phi(\mathbf{r}_n)}{\phi(\mathbf{r}_n)} \right|; \quad \epsilon_{\mathbf{H}} := \text{median}_{n \in N} \frac{\|\hat{\mathbf{H}}(\mathbf{r}_n) - \mathbf{H}(\mathbf{r}_n)\|}{\|\mathbf{H}(\mathbf{r}_n)\|}, \quad (12)$$

where  $\|\cdot\|$  is the vector norm; we report the median instead of the mean because relative errors diverge as the true field value approaches zero; the median offers a more robust and representative summary of model performance across the domain.

Our code is available at <https://anonymous.4open.science/r/hypermagnetics-TMLR/>.

### 5.1 Direct and indirect evaluation of the field

We first evaluate how well the architecture can represent the field around a single collection of sources, without requiring the model to learn source-conditioned representations—i.e., we train  $f$  while keeping  $g$  fixed. To demonstrate the ability of a small fully-connected network  $f$  to represent the potential and field, we use a single source configuration shown in Fig. 2. The sources (here in 2D represented as disks of equal radius) are randomly placed within a  $[-3, 3] \times [-3, 3]$  domain (in source-radius units), and the field and potential are evaluated on a  $100^2$  grid, with  $N = 100^2$  random points reserved for validation. An ensemble of fully-connected networks (width 32, depth 3) is trained for  $10^5$  epochs using the Adam optimiser (Kingma & Ba, 2014) with a learning rate of  $10^{-5}$ .

Fig. 3a shows the mean and variance of the training error across 10 runs. A model directly inferring the scalar potential ( $\epsilon_\phi$ ) performs with 0.31% error. The same model can be used to infer the field by taking the numerical gradient of the potential ( $\epsilon_{\nabla\phi}$ ), defined in the same way as  $\epsilon_H$  but notated differently to express that the field evaluation is indirect; this reaches an error of 3.70% and notably follows the same training shape. A separate model with the same hidden layer size that directly outputs the magnetic field from the fully-connected network ( $\epsilon_H$ ) performs comparably to the indirect field method, reaching 4.70%, though with an apparently slower convergence. This clearly demonstrates generalisation to arbitrary field locations—but not to arbitrary source collections.

It is unsurprising that, for a fixed network size, field prediction is less accurate than potential prediction due to the instability of numerical differentiation. Moreover, training scalar outputs is typically easier than training vector outputs. Given that direct field prediction does not offer accuracy advantages, we adopt the potential-based strategy in the remainder of our experiments, which implies a simpler network architecture and simultaneously enables direct evaluation of the potential, which is itself of practical interest.

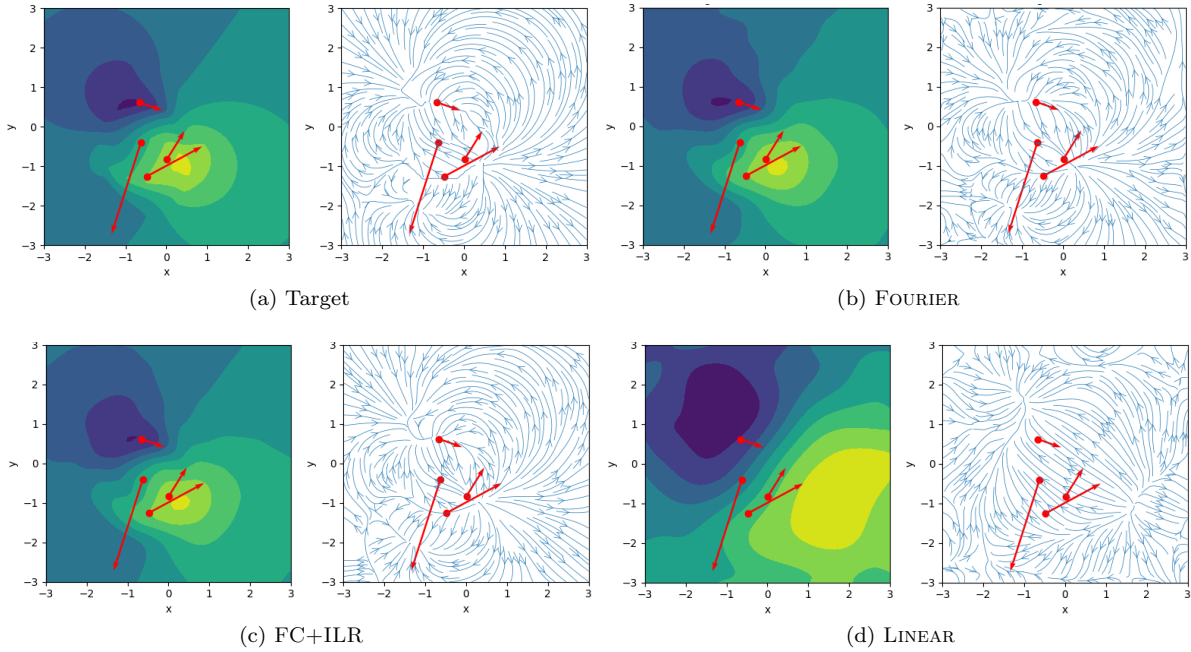


Figure 4: Multiple-source inference. (a) Ground-truth potential and field from four randomly placed circular magnets. (b-d) Predictions from models trained only on single-source examples, evaluated by first aggregating source representations before computing spatial values.

## 5.2 Evaluation of arbitrary source configurations

While the previous experiment demonstrated generalisation across spatial locations, achieving the desired  $\mathcal{O}(M + N)$  scaling requires that the model also generalise across arbitrary source configurations. To evaluate this, we generate an ensemble of up to  $10^4$  source collections, each with random positions and magnetisations within a  $[-3, 3] \times [-3, 3]$  domain.

To leverage the linearity of superposition, we train only on single-source examples. Each training example is evaluated at  $32^2$  uniformly sampled field points. At inference time, multiple source embeddings are aggregated before computing the field. For visualisation purposes, we use a fixed  $128^2$  grid in this section.

All models share a common hypernetwork  $g$ , implemented as a fully-connected network of depth 3 and width equal to the number of parameters in  $f$ . The networks are trained with a stepwise learning rate schedule using Adam (Kingma & Ba, 2014) for 5,000 epochs per step at learning rates  $10^{-\{3,4,5,6\}}$ . To ensure fair comparison, model sizes are kept similar—generally around 20M parameters—so that the overall number of parameter-epochs remains approximately constant. Detailed configurations are summarised in Table 4 in the Appendix.

### 5.2.1 The FOURIER network

For the FOURIER model, the ability to generalise to an arbitrary number of sources follows directly from the linearity of the hypernetwork outputs in the inference network. This is illustrated in Fig. 3b, where a 1D

Model	$M = 1$	$M > 1$
FOURIER	4.74 ( $\pm 0.38$ )	5.73 ( $\pm 0.45$ )
FC+ILR	4.38 ( $\pm 0.32$ )	4.76 ( $\pm 0.39$ )
FC+INR	3.68	—
LINEAR	70.0	71.2

Table 2: Out-of-sample relative error  $\epsilon_\phi$  (%) for predicting the potential generated by magnetised 2D disks of equal radius.

version of the model with 32 modes is trained solely on single-source examples. At inference time, weights from multiple sources are aggregated *before* evaluating the potential, demonstrating the use of superposition.

In two dimensions, the FOURIER model takes the form:

$$\begin{aligned} \hat{\phi}(\mathbf{r}_n) = \sum_{l=0}^{L_x-1} \sum_{p=0}^{L_y-1} & a_{lp} \cos(k_l r_{nx}) \cos(k_p r_{ny}) + b_{lp} \sin(k_l r_{nx}) \cos(k_p r_{ny}) \\ & + c_{lp} \cos(k_l r_{nx}) \sin(k_p r_{ny}) + d_{lp} \sin(k_l r_{nx}) \sin(k_p r_{ny}), \end{aligned} \quad (13)$$

where  $k_l, k_p$  denote the wavenumbers with  $k_0 = 0$  and  $k_l = 10^{-2.9+3.65*(l/(L_x-1))}$ , while  $L = L_x = L_y$  controls the expansion order. We set  $k_l = k_p$  to ensure a square basis, so the total number of Fourier coefficients—and hence hypernetwork outputs—is  $\sim L^2$ , comparable in size to a single linear layer of width  $L$ . Empirically, the expansion performs well with as few as 16 modes per dimension. For results in Table 2 and Fig. 4, we use  $L = 32$  modes.

### 5.2.2 The FC+ILR model

For the 2D setting, Fig. 4 illustrates the ability of the FC+ILR model to infer the potential and field around finite multiple sources. We emphasise both that i) the model is trained exclusively on single-source examples, and that ii) at inference time, the source representations are computed individually and aggregated *before* evaluating the field at any location.

### 5.2.3 LINEAR and FC+INR baselines

The FC+INR model lacks architectural constraints enforcing superposition, and thus fails to generalise to multi-source configurations when trained only on single-source examples. Table 2 summarises out-of-sample relative errors  $\epsilon_\phi$  across all models. While FC+INR performs best on single-source inputs (3.68% error), it cannot be evaluated on multi-source data.

As a further baseline, we implement a minimal LINEAR model, where both  $f$  and  $g$  are single linear layers scaled to match parameter count. Despite satisfying superposition and  $\mathcal{O}(M+N)$  scaling, it lacks sufficient expressiveness, yielding relative median errors around 70%.

## 6 Experiments using a larger collection of sources and scaling behaviour

To test the model’s applicability to more complex geometries, we extend our setup to prism-shaped sources, for which the magnetic scalar potential is known analytically (James et al., 2025). We use the FC+ILR model for this section, as it achieved the best overall trade-off in earlier experiments (see Table 2). The model is trained on  $2 \times 10^5$  2D potentials generated by individual prisms with varying side lengths ( $s_x = s_y \in [0.1, 1]$ ) in the  $xy$ -plane and fixed height ( $s_z = 10 \gg s_x, s_y$ ) along the  $z$ -axis in a  $[-1.25, 1.25] \times [-1.25, 1.25]$  domain. Each training example is evaluated at  $32^2$  uniformly sampled potential points.

To benchmark our model, we compare against the fast multipole method (FMM), in the open-source implementation by the Flatiron Institute (Cheng et al., 1999; Greengard et al., 2002; Greengard & Rokhlin, 1997; Greengard et al., 2002). The FMM method, which here is utilised in its 2D version, computes N-body interactions by compressing the interactions between clusters of source and target points at a hierarchy of scales. This means that the field can be computed as order  $\mathcal{O}((M+N)\log(1/\tau))$  where  $\tau$  is the desired relative precision, which we set to  $10^{-5}$  for all experiments.

However, it remains necessary to correct the potential computed by FMM within each source, as the FMM method assumes point dipoles or collections of these, which within each source provides an incorrect potential as shown with the “Pure FMM” method in Fig. 5b. This correction cannot be applied in a manner that produces an overall correct potential, as the far-field dipole potential cannot be matched with an internally correct potential computed using the known analytical formulas for the potential generated by a rectangular prism. However, ignoring this, we apply the correction to FMM to get the correct potential within each source. This can be done in two ways: i) If the field is evaluated, and this is done in the center of each

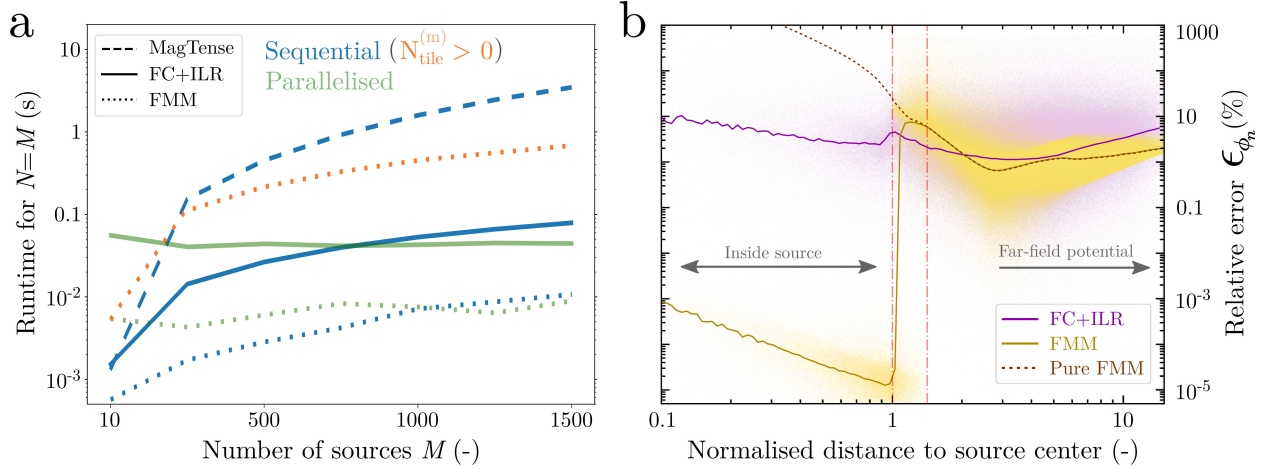


Figure 5: (a) Runtime as a function of the number of sources  $M$  and keeping  $N = M$ , e.g. evaluation of  $\phi$  in the center of each source. The blue coloured lines show the sequential evaluation time to make the underlying computational complexity of each method apparent. The green lines show the runtimes when parallelising the workload on the GPU (NVIDIA GeForce RTX 5090) with FC+ILR and on 96 threads of the CPU (Intel® Xeon® Gold 6248R) in case of FMM. The increased runtime visualised by the orange line originates from the additional correction required for evaluation points lying inside other sources in case sources overlap or  $\phi$  is evaluated on a regular grid, which leads to  $N_{\text{tile}}^{(m)} > 0$ . Notice that this only affects the evaluation time of FMM. (b) Relative error  $\epsilon_{\phi_n}$  for single evaluation points from single-source configurations ( $M = 1$ ) grouped by distance to the source center normalised by the side length  $s_x$  of the source. Here, the data consists of  $10^3$  2D potentials generated by single prism-shaped source configurations with varying side lengths ( $s_x = s_y \in [0.12, 0.48]$ ) in a  $[-1.2, 1.2] \times [-1.2, 1.2]$  domain. Each validation sample is evaluated at  $32^2$  uniformly sampled potential points. The FMM method used as a benchmark is shown in yellow, while we depict in brown the "pure" FMM without any corrections for the physical extension of the magnetic source. FC+ILR, depicted in purple, outperforms FMM close to the edges of the sources.

source (i.e.  $\mathbf{r}_n = \mathbf{r}'_m$ ), the field can easily be computed from the magnetometric demagnetization factor  $N_{mn}$ , which for the field in 2D for symmetric objects is simply  $1/2$  and the field is derived to  $\mathbf{H}_n = -1/2\mathbf{M}_m$ , and in 3D this factor is  $1/3$  for symmetric objects. When evaluating the potential only in the center, then FMM outputs the correct value ( $\phi_n = 0$ ) for symmetric objects, and no correction is required. ii) However, when the field or potential is evaluated across a grid including points within each tile, i.e.  $N_{\text{tile}}^{(m)} > 0$ , this simple magnetometric correction is not applicable. Instead here the field or potential must be calculated analytically using the known expressions (Smith et al., 2010b; James et al., 2025) within each point in the tile, which will add another  $\sum_{m=1}^M N_{\text{tile}}^{(m)}$  operations to the runtime of FMM as shown with the orange line of Fig. 5a.

When comparing the runtimes in Fig. 5a, we can see the expected quadratic complexity  $\mathcal{O}(M \times N)$  for the classical evaluation method done using MagTense and a scaling of  $\mathcal{O}(\lambda_m M + \lambda_n N)$  for the other methods, however with different scaling factors:  $\lambda_{m,\text{FMM}} = \log(1/\tau) + N_{\text{tile}}$  and  $\lambda_{n,\text{FMM}} = \log(1/\tau)$ , and  $\lambda_{m,\text{FC+ILR}} = L$  and  $\lambda_{n,\text{FC+ILR}} = L + 1$ . When parallelising the workload for FC+ILR, we obtain a flat runtime which becomes beneficial for around  $M + N > 2 \cdot 750$ , as long as the problem fits on GPU memory. For larger-scale problems, we expect the runtime to follow the theoretical  $\mathcal{O}(M + N)$  profile. Similarly, FMM exhibits a flat runtime for parallelisation across the CPU.

In Fig. 5b, we compare the performance of the method proposed here to FMM in configurations of prism-shaped magnetic sources. The relative errors of each evaluation point  $\epsilon_{\phi_n}$  are grouped by the normalised distance to the source center to highlight the peculiarities of each method. FC+ILR relative error rates are centered between 1% and 10% for the whole distance range. The best performance can be seen for distances between  $1.5s_x$  and  $6s_x$  from the source center, while the relative error goes up to 6% for the far-field potential.



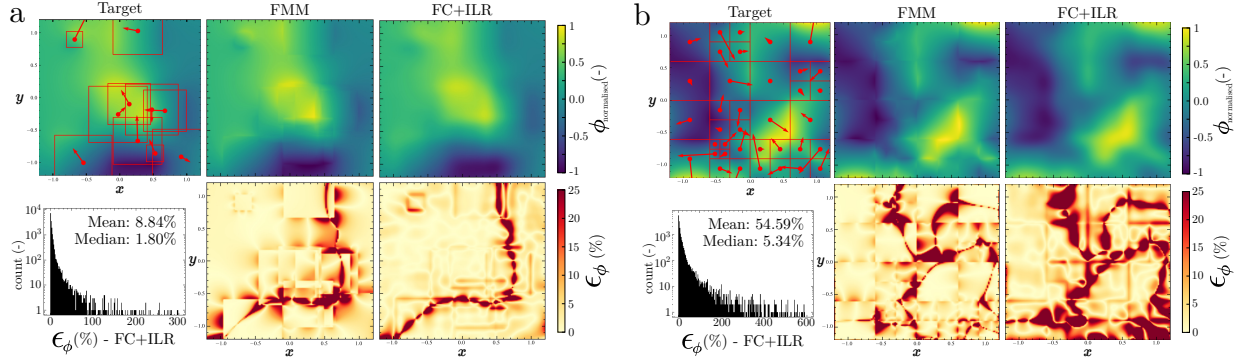


Figure 6: Scalar potential from sampled configurations of magnetic sources used in our experiments, comparing ground truth  $\phi$  with the predictions from FMM,  $\hat{\phi}_{\text{FMM}}$ , and our FC+ILR model,  $\hat{\phi}_{\text{FC+ILR}}$ . The prism-shaped sources are depicted in red with the arrows indicating their magnetisation  $\mathbf{M}$ . At the bottom, a histogram of the relative error of  $\hat{\phi}_{\text{FC+ILR}}$  is shown along with  $\epsilon_\phi$  spatially resolved for  $\hat{\phi}_{\text{FMM}}$  and  $\hat{\phi}_{\text{FC+ILR}}$ . (a) Normalised scalar potential of 10 overlapping sources. (b) Normalised scalar potential of a quadtree structure with 37 sources.

Interestingly, the error rates approach high error rates of 10% when reaching the center of the sources. This is arguably due to the potential approaching 0 close to the source center, hence small errors become relatively large. FMM, corrected by MagTense inside the sources, reaches relative error rates of around 10% close to the edges of the magnetic source, when higher order poles come into play. For the far-field potential, when the potential from a dipole becomes a better approximation,  $\epsilon_\phi$  grouped by the normalised distance stays below 2%. Nevertheless, the relative error becomes larger when increasing the normalised distance beyond 3, which is because absolute error diminishes at a slower rate than the  $1/|\mathbf{r}|$  scaling of the 2D potential values.

Further, we consider two different scenarios for comparing the overall performance of our model to FMM, namely one with overlapping sources in Sec. 6.1 and one with quadtree structures in Sec. 6.2.

### 6.1 Overlapping sources

It is of great interest to be able to model overlapping sources, as this is often the only way of modelling complex geometries without resorting to discretising the geometry to a large degree. Because fields and potentials are additive, a geometry with e.g. a hollow cylinder can be realised by starting with a cylinder and then on this overlapping a smaller cylinder with an opposite magnetisation compared to the larger cylinder. This will result in a cancellation where the cylinders overlap, and this produces the correct field with only two sources. The alternative strategy of discretising the hollow cylinder with e.g. prisms or tetrahedra, which would have resulted in a much larger number of sources. Therefore, it is computationally desirable to be able to correctly model overlapping sources.

In this experiment, we calculate the magnetic scalar potential from a varying number of sources ( $M = 10, 50, 250, 1000$ ) on a regular grid of  $32^2$  evaluation points. The median relative error in the scalar potential  $\epsilon_\phi$  (vs. the exact solution computed using MagTense) is presented in Table 3. An example with  $M = 10$  sources is depicted in Fig. 6a. Note that here the evaluation points are a regular grid of  $128^2$  values, which is chosen for visualisation purposes. Each entry represents an average over 10 runs, with its standard deviation in brackets. Notably, the potential error remains below 5% for this scenario with FC+ILR, regardless of the number of sources—in line with the small-scale experiments with spherical sources in Table 2. The reason for the slightly large error of the FMM method is exactly that the dipole approximation is not applicable in the proximity of a physical magnetic source, which means that error accumulates when multiple sources overlap, exemplified at the bottom of Fig. 6a.



Model	$M = 1$	$M$ overlapping sources				Quadtree structure with $M$ sources		
		10	50	250	1000	10	50	250
FMM	<b>1.14</b> ( $\pm 0.28$ )	<b>2.21</b> ( $\pm 0.78$ )	4.00 ( $\pm 0.84$ )	6.14 ( $\pm 2.12$ )	5.32 ( $\pm 2.30$ )	<b>2.75</b> ( $\pm 0.49$ )	<b>2.69</b> ( $\pm 0.47$ )	<b>2.39</b> ( $\pm 0.45$ )
FC+ILR	2.21 ( $\pm 1.48$ )	2.43 ( $\pm 0.54$ )	<b>3.05</b> ( $\pm 0.88$ )	<b>4.40</b> ( $\pm 1.53$ )	<b>4.21</b> ( $\pm 1.51$ )	6.40 ( $\pm 2.79$ )	5.70 ( $\pm 1.14$ )	-

Table 3: Relative error  $\epsilon_\phi$  (%) averaged over 10 runs for an increasing number of prism-shaped sources. The best performing method is marked in bold. The experiment with  $M = 1$  refers to the single-source configurations used in Fig. 5b. For FC+ILR, the last two values of the quadtree experiment can not be evaluated properly as the sizes of the magnetic sources in these setups become much lower than the sizes of the sources, which FC+ILR had seen during training.

## 6.2 Quadtree structures

It is also of interest to consider the sources arranged in a way such that they span a volume of space in a non-overlapping way, and with more sources located at volumes where a high resolution is desired. Such a structure could be a quadtree structure, which we term a structure where space is subdivided into increasingly smaller prisms by splitting each prism into four smaller prisms, with the same aspect ratio, as depicted in Fig. 6b for  $M = 37$ . Such a configuration is important in e.g. micromagnetic simulations where a simulation domain can be meshed using rectangular prisms, with the smaller prisms resolving finer features such as grain boundary layers (Björk & Insinga, 2023).

In this experiment, we calculate the magnetic potential similarly from a varying number of sources ( $M = 10, 50, 250$ ) and evaluate it on a regular grid of  $32^2$  points. The results in Table 3 clearly show that FMM outperforms our model for this scenario with  $\epsilon_{\phi, \text{FMM}} < 3\%$ . All evaluation points lie within a source and, as shown in Fig. 5b, FMM greatly benefits from the applied correction in this area. However, as  $N_{\text{tile}}^{(m)} \gg 0$ , the runtime of FMM is around one order of magnitude larger than FC+ILR.

Our trained model stays below 10% relative error for  $M = 10, 50$  and can not be evaluated for the scenarios with  $M = 250$  as the side lengths of some prism-shaped sources  $s_x, s_y \ll 0.1$ , which is simply outside the range FC+ILR has been trained for. The increased error compared to the scenario of overlapping sources arises from the prevalence of small sources with  $s_x < 0.2$  as can be seen in Fig. 6b. This is further emphasised in Fig. 8a in the Appendix, where the relative error of single-source configurations is plotted using the FC+ILR model. The reason for the large error of the method proposed here is that it is hard for the neural network to learn the potential far from a source, where the potential is close to zero. The dynamic range of potential values for small prisms is simply too large for the model to learn them properly, which means that the predictions for sources with small side lengths become incorrect. As can be seen from Fig. 8b, it is the smallest sources that result in a large error in the prediction over the whole domain.

## 7 Limitations & Discussion

We finish by delineating what this architecture should be capable of *in principle*, against what we have concretely demonstrated through these experiments. This work develops models that provide inference from arbitrary collections of sources, but our experiments demonstrate this only in a limited range of evaluation domains, source volumes, and magnetisations. For practical results in (e.g.) magnetostatics, it is necessary to develop models in the same architecture trained with the maximum expected parameter range in such applications, while overcoming the difficulties to learn a large dynamic range of magnetic field values as discussed in the Appendix D. When replacing the magnetisation with the mass as input to  $g$ , the same architecture can perform inference for the field and potential around gravitational sources. Further, developing these models as amortising replacements for infeasible numerical simulation requires extension to three dimensions. All models developed in this paper can be extended in this way; the FOURIER computation will then contain  $(2L)^3$  terms, which makes the parameter requirements for the FC+ILR model seem more scalable. We have not explicitly treated equivariance beyond noting a connection between the principle of superposition and permutation invariance of sources.

This work develops models for static configurations. This architecture could be integrated iteratively into a time-varying problem, using the model output to update properties of the sources and then re-evaluating the field. However, it would be more natural to consider training neural operator(s) to evolve the system. This model is most effectively used with collections of sources with spatial extent; for particle-like sources the FMM is a proven non-statistical algorithm with  $O(M + N)$  scaling (Greengard & Rokhlin, 1987). Lastly, our particular choice of orthogonal decomposition (Eq. 10), while physically motivated, may be improved by explicit weighting so that  $\lim_{r \rightarrow \infty} \phi = 0$ ; more physical information known from Maxwell’s equation, i.e.  $\nabla \cdot (\mathbf{H} + \mathbf{M}) = 0$ , may further improve training the proposed architecture; with defining  $g(\mathbf{M}_m, \mathbf{V}'_m) = |\mathbf{M}_m| \cdot g(\mathbf{I}_m, \mathbf{V}'_m)$ , where  $\mathbf{I}_m$  is the unit vector pointing in the direction of  $\mathbf{M}_m$ , as we know from Eq. 3 that the field is linear in the strength of the magnetisation or directly applying rotations and translations in the vector space spanned by the abstract basis representation  $f$ , may be another viable route to improve the proposed method. These limitations and ideas can be addressed with future work, which will be of interest to both the statistics and physics communities.

## 8 Conclusions

In this work, we addressed the problem of amortising field computation around physical sources by introducing a general hypernetwork-based architecture and three model variants, each balancing scalability, accuracy, and physical fidelity. Modelling via the scalar potential improves training efficiency and simplifies network design. Among the proposed approaches, the FOURIER and FC+ILR models achieve  $\sim 5\%$  error while enabling field or potential inference for arbitrary source collections through the principle of superposition. While the physically motivated FOURIER model offers interpretability through its orthogonal decomposition, it does not outperform the more flexible FC+ILR model and incurs greater architectural complexity and tuning overhead.

## References

- Claas Abert, Lukas Exl, Florian Bruckner, et al. magnum.fe: A micromagnetic finite-element simulation code based on fenics. *Journal of Magnetism and Magnetic Materials*, 345:29–35, 2013. ISSN 0304-8853. doi: <https://doi.org/10.1016/j.jmmm.2013.05.051>. URL <https://www.sciencedirect.com/science/article/pii/S0304885313004022>.
- Amikam Aharoni. Demagnetizing factors for rectangular ferromagnetic prisms. *Journal of Applied Physics*, 83(6):3432–3434, March 1998. ISSN 0021-8979, 1089-7550. URL <https://doi.org/10.1063/1.367113>.
- Luca Bertinetto, João F Henriques, Jack Valmadre, Philip Torr, and Andrea Vedaldi. Learning feed-forward one-shot learners. *Advances in neural information processing systems*, 29, 2016.
- R Bjørk and M d’Aquino. Accuracy of the analytical demagnetization tensor for various geometries. *Journal of Magnetism and Magnetic Materials*, 587:171245, 2023.
- R. Bjørk, E. B. Poulsen, K. K. Nielsen, et al. MagTense: A micromagnetic framework using the analytical demagnetization tensor. *Journal of Magnetism and Magnetic Materials*, 535:168057, October 2021. URL <https://www.sciencedirect.com/science/article/pii/S0304885321003334>.
- Rasmus Bjørk and Andrea Roberto Insinga. Explaining Browns paradox in NdFeB magnets from micromagnetic simulations. *Journal of Magnetism and Magnetic Materials*, 571:170510, April 2023. URL <https://www.sciencedirect.com/science/article/pii/S0304885323001592>.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2024. URL <http://github.com/google/jax>.
- Andrew Brock, Theo Lim, JM Ritchie, and Nick Weston. Smash: One-shot model architecture search through hypernetworks. In *International Conference on Learning Representations*, 2018.

- Hongwei Cheng, Leslie Greengard, and Vladimir Rokhlin. A fast adaptive multipole algorithm in three dimensions. *Journal of Computational Physics*, 155(2):468–498, 1999.
- Woojin Cho, Kookjin Lee, Donsub Rim, and Noseong Park. Hypernetwork-based meta-learning for low-rank physics-informed neural networks. *Advances in Neural Information Processing Systems*, 36:11219–11231, 2023.
- Filipe de Avila Belbute-Peres, Yi-fan Chen, and Fei Sha. Hyperpinn: Learning parameterized differential equations with physics-informed hypernetworks. *The symbiosis of deep learning and differential equations*, 690, 2021.
- Emilien Dupont, Yee Whye Teh, and Arnaud Doucet. Generative models as distributions of functions. In *International Conference on Artificial Intelligence and Statistics*, pp. 2989–3015. PMLR, 2022.
- N. Engheta, W.D. Murphy, V. Rokhlin, et al. The fast multipole method (fmm) for electromagnetic scattering problems. *IEEE Transactions on Antennas and Propagation*, 40(6):634–641, 1992. doi: 10.1109/8.144597.
- Atiyo Ghosh, Antonio A. Gentile, Mario Dagrada, et al. Harmonic neural networks. In *Proceedings of the 40th International Conference on Machine Learning, Icm1’23*. JMLR.org, 2023.
- L Greengard and V Rokhlin. A fast algorithm for particle simulations. *Journal of Computational Physics*, 73(2):325–348, 1987. ISSN 0021-9991. doi: [https://doi.org/10.1016/0021-9991\(87\)90140-9](https://doi.org/10.1016/0021-9991(87)90140-9). URL <https://www.sciencedirect.com/science/article/pii/0021999187901409>.
- Leslie Greengard and Vladimir Rokhlin. A new version of the fast multipole method for the laplace equation in three dimensions. *Acta numerica*, 6:229–269, 1997.
- Leslie Greengard, Jingfang Huang, Vladimir Rokhlin, and Stephen Wandzura. Accelerating fast multipole methods for the helmholtz equation at low frequencies. *IEEE Computational Science and Engineering*, 5(3):32–38, 2002.
- D.J. Griffiths. *Introduction to Electrodynamics*. Pearson, 2013. ISBN 9780321856562. URL [https://books.google.dk/books?id=AZx\\_zwEACAAJ](https://books.google.dk/books?id=AZx_zwEACAAJ).
- David Ha, Andrew M. Dai, and Quoc V. Le. Hypernetworks. In *International Conference on Learning Representations*, 2017. URL <https://openreview.net/forum?id=rkpACe1lx>.
- Johannes Hendriks, Carl Jidling, Adrian Wills, et al. Linearly Constrained Neural Networks, April 2021. URL <http://arxiv.org/abs/2002.01600>.
- J. D. Jackson. *Classical Electrodynamics*, 3rd ed. Hoboken, NJ, USA: Wiley, 1999.
- Berian James, Stefan Pollok, Jes Frelsen, and Rasmus Bjørk. The magnetic scalar potential and demagnetization vector for a prism. *arXiv preprint arXiv:2501.18604*, 2025.
- R. I. Joseph and E. Schlömann. Demagnetizing Field in Nonellipsoidal Bodies. *Journal of Applied Physics*, 36(5):1579–1593, November 1964. ISSN 0021-8979. URL <https://doi.org/10.1063/1.1703091>.
- George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, et al. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, May 2021. ISSN 2522-5820. doi: 10.1038/s42254-021-00314-5. URL <https://www.nature.com/articles/s42254-021-00314-5>.
- Patrick Kidger and Cristian Garcia. Equinox: neural networks in JAX via callable PyTrees and filtered transformations. *Differentiable Programming workshop at Neural Information Processing Systems 2021*, 2021.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Batuhan Koyuncu, Pablo Sanchez Martin, Ignacio Peis, Pablo M Olmos, and Isabel Valera. Variational mixture of hypergenerators for learning distributions over functions. In *International Conference on Machine Learning*, pp. 17660–17683. PMLR, 2023.
- Ting-Ou Liang, Yan Hao Koh, Tie Qiu, et al. Magtetriz: A simulator for fast magnetic field and force calculation for permanent magnet array designs. *Journal of Magnetic Resonance*, 352:107463, 2023.
- Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *International conference on machine learning*, pp. 2554–2563. PMLR, 2017.
- Eike Hermann Müller. Exact conservation laws for neural network integrators of dynamical systems. *Journal of Computational Physics*, 488:112234, 2023. ISSN 0021-9991. doi: <https://doi.org/10.1016/j.jcp.2023.112234>. URL <https://www.sciencedirect.com/science/article/pii/S0021999123003297>.
- K. K. Nielsen and R. Bjørk. The magnetic field from a homogeneously magnetized cylindrical tile. *Journal of Magnetism and Magnetic Materials*, 507:166799, August 2020. URL <https://www.sciencedirect.com/science/article/pii/S0304885319342155>.
- Kaspar Kirstein Nielsen, Andrea Roberto Insinga, and Rasmus Bjørk. The Stray and Demagnetizing Field of a Homogeneously Magnetized Tetrahedron. *IEEE Magnetics Letters*, 10:1–5, 2019.
- Michael Ortner and Lucas Gabriel Coliado Bandeira. Magpylib: A free python package for magnetic field computation. *SoftwareX*, 2020. doi: 10.1016/j.softx.2020.100466.
- Keunhong Park, Utkarsh Sinha, Peter Hedman, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Ricardo Martin-Brualla, and Steven M Seitz. Hypernerf: a higher-dimensional representation for topologically varying neural radiance fields. *ACM Transactions on Graphics (TOG)*, 40(6):1–12, 2021.
- Ignacio Peis, Batuhan Koyuncu, Isabel Valera, and Jes Frellsen. Hyper-transforming latent diffusion models. In *International Conference on Machine Learning*, 2025.
- Tobias Pfaff, Meire Fortunato, Alvaro Sanchez-Gonzalez, and Peter Battaglia. Learning mesh-based simulation with graph networks. In *International conference on learning representations*, 2020.
- Stefan Pollok, Nataniel Olden-Jørgensen, Peter Stanley Jørgensen, et al. Magnetic Field Prediction Using Generative Adversarial Networks. *Journal of Magnetism and Magnetic Materials*, 571:170556, April 2023. URL <http://arxiv.org/abs/2203.07897>.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In J. Platt, D. Koller, Y. Singer, and S. Roweis (eds.), *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007. URL <https://proceedings.neurips.cc/paper%5Ffiles/paper/2007/file/013a006f03dbc5392effeb8f18fda755-Paper.pdf>.
- M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, February 2019. ISSN 0021-9991. doi: 10.1016/j.jcp.2018.10.045. URL <https://www.sciencedirect.com/science/article/pii/S0021999118307125>.
- Jack Richter-Powell, Yaron Lipman, and Ricky T. Q. Chen. Neural Conservation Laws: A Divergence-Free Perspective, December 2022. URL <http://arxiv.org/abs/2210.01741>. arXiv:2210.01741 [cs].
- Sebastian Schaffer, Thomas Schrefl, Harald Oezelt, et al. Physics-informed machine learning and stray field computation with application to micromagnetic energy minimization. *Journal of Magnetism and Magnetic Materials*, 576:170761, June 2023. URL <https://www.sciencedirect.com/science/article/pii/S0304885323004109>.
- Vincent Sitzmann, Julien Martel, Alexander Bergman, David Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in neural information processing systems*, 33: 7462–7473, 2020.

Florian Slanovc, Michael Ortner, Mohssen Moridi, Claas Abert, and Dieter Suess. Full analytical solution for the magnetic field of uniformly magnetized cylinder tiles. *Journal of Magnetism and Magnetic Materials*, 559:169482, 2022.

A. Smith, K. K. Nielsen, D. V. Christensen, et al. The demagnetizing field of a nonuniform rectangular prism. *Journal of Applied Physics*, 107(10):103910, May 2010a. URL <https://doi.org/10.1063/1.3385387>.

Anders Smith, Kaspar Kirstein Nielsen, DV Christensen, Christian Robert Haffenden Bahl, Rasmus Bjørk, and J Hattel. The demagnetizing field of a nonuniform rectangular prism. *Journal of Applied Physics*, 107(10), 2010b.

Zhong Yi Wan, Leonardo Zepeda-Nunez, Anudhyan Boral, and Fei Sha. Evolve smoothly, fit consistently: Learning smooth latent dynamics for advection-dominated systems. In *The Eleventh International Conference on Learning Representations*, 2023.

Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, et al. Deep sets. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper%5Ffiles/paper/2017/file/f22e4747da1aa27e363d86d40ff442fe-Paper.pdf>.

## A Broader Impact Statement

Magnets are a critical energy and infrastructure technology because they convert between motion and electricity. The goal is to create permanent magnets that are as strong and as stable as possible: the stronger the field, the more efficient the conversion of energy; the more stable the magnet, the more versatile its applications across a range of temperatures and environments. Despite a sound analytical basis for atomic-scale physics, it is not possible to derive the magnetic properties of new materials or solve an inverse problem to design a material with a desired field. Our research could be used to provide much more rapid exploration of magnetic material and configuration possibilities, accelerating the development of these much-needed energy technologies.

The scaling improvements from our architecture enable compute savings for numerical work with data that satisfy the general physical principle of superposition. We hope that our research will enable highly accurate modelling with vastly less energy expenditure at inference time. Additionally, our research supports the goal of explainable inference for physical systems by pursuing amortising models grounded in physical theory. This should be seen as a desirable contingency against a purely empirical model, which may exhibit hidden performance biases resulting in poor or unpredictable performance for physical systems.

Nevertheless, there are risks from unintentional inaccuracy or imprecision in model output: at a minimum, the resource and opportunity cost of pursuing experimental and manufacturing effort in materials and configurations that cannot exhibit the performance the model claims. More seriously, successful model performance accelerating the development of energy technologies can lead to negative societal outcomes if those new technologies are not used responsibly or with appropriate safeguards. Certainly, this has been the case for energy technologies in the past, and in general, social and regulatory structures have emerged that can mitigate against such disasters.

## B Theoretical Foundations and Extensions

### B.1 Details of Analytical Field Computation

For spherically-symmetric sources, the dipole field at  $\mathbf{r}$  from a collection of  $M$  point-like sources with magnetic moments  $\mathbf{m}_m$ , radii  $d_m$  at positions  $\mathbf{r}'_m$  is computed via the scalar potential (Jackson, 1999, p. 196)

$$\mu_0 \mathbf{H}_\odot(\mathbf{r}, \mathbf{M}, \mathbf{V}') = -\nabla \sum_{m=1}^M \underbrace{\frac{1}{4\pi|\mathbf{r} - \mathbf{r}'_m|^2}}_{\text{scalar potential } \phi(\mathbf{r}, \mathbf{M}_m, \mathbf{V}'_m)} \underbrace{\frac{\mathbf{m}_m \cdot (\mathbf{r} - \mathbf{r}'_m)}{|\mathbf{r} - \mathbf{r}'_m|}}_{\text{dipole term}} \quad \text{or} \quad -\nabla \sum_{m=1}^M \frac{1}{4\pi d_m^2} \frac{\mathbf{m}_m \cdot (\mathbf{r} - \mathbf{r}'_m)}{d_m} \quad (14)$$

corresponding to outside and inside the source, respectively. The distinguishing feature of the dipole term is the  $1/r$  dependence; higher multipole terms ( $1/r^2$ ,  $1/r^3$ ) will be present for specific shapes of sources, but the dipole term will dominate at larger  $r$  (Björk & d'Aquino, 2023), although higher multipole terms are required near the source.

### B.2 On the principle of superposition

The principle of superposition implies that the field generated by a collection of sources can be obtained by summing the individual contributions of each source. In our architecture, this property is realised through a latent representation that is additive across sources. Specifically, each source embedding produced by the hypernetwork  $g$  contributes linearly to the field representation  $f$ , such that the aggregation of embeddings corresponds to the aggregation of fields.

$$\begin{array}{ccc} \text{Set}\{\mathbf{H}_\mathbf{r}(\mathbf{M}_m, \mathbf{V}'_m)\} & \xrightarrow{\Sigma_{m=1}^M} & \mathbf{H}_\mathbf{r}(\mathbf{M}, \mathbf{V}') \\ \uparrow \mathbf{H}_\mathbf{r}(\cdot) \quad \forall m \in M & & \uparrow \mathbf{H}_\mathbf{r} \\ (\mathbf{M}, \mathbf{V}') = \text{Set}\{(\mathbf{M}_m, \mathbf{V}'_m)\} & \xrightarrow{\bigcup_{m=1}^M} & \bigcup (\mathbf{M}_m, \mathbf{V}'_m) \end{array}$$

Figure 7: Commutative diagram for the construction of the field around a source collection via a fixed-length vector encoding.  $\mathbf{H}_\mathbf{r}$  describes here the evaluation of the magnetic field at point  $\mathbf{r}$ .

This structure is naturally expressed using the commutative diagram of Fig. 7, where the field produced by aggregating source embeddings is equivalent to aggregating the fields of individual sources.

## C Experimental and Implementation Details

### C.1 Hyperparameters

We have implemented our models using JAX 0.4.25 (Bradbury et al., 2024) and Equinox (Kidger & Garcia, 2021), with training performed on the EuroHPC Karolina GPU cluster. Table 4 collates the model parameters used for the experiments in Sec. 5 and Sec. 6.

The output size of a hypernetwork is the number of parameters consumed by the main (inference) network, and this varies across the models as follows:

- The FOURIER model is parameterised by the order  $L$  of the expansion, and in two dimensions its inference ‘network’ requires  $4L^2$  parameters;
- In the FC+ILR model, only the final layer of size  $L$  of the inference network is set by the hypernetwork, so the total number of learnable parameters is the remaining layers of the inference network plus the hypernetwork size;

	FOURIER	FC+ILR(Sec. 5)	FC+ILR(Sec. 6)	FC+INR
Width	$4 \times 32 \times 32$	400	400	20
Depth	-	3	3	3
HNet Width*	0.25	1.5	2	1.0
HNet Depth	3	3	3	2
<i>Total Parameters</i>	<i>6.3M</i>	<i>1.5M</i>	<i>2.1M</i>	<i>1.7M</i>
Optimiser	Adam			
Average Learning Rate	5E-4	1E-4	7.5E-4	1E-5
$\gamma_{\mathbf{H}}$	1	1	0.25	1
Steps (thousands)	20	25	20	20
GPUs	NVIDIA A100-SXM4-40GB ( $\diamond$ )			
Time (hours)	12	6	40	-

Table 4: Collation of model parameters used for the experiments in Sec. 5 and Sec. 6 with the architecture developed in this work. If not specified otherwise, the same value is used for all experiments. A single entry in a row means the same value for all models. (\*) The hypernetwork (hidden layer) width is expressed as a multiple of its output size, which is the total number of learned parameters in the inference network. ( $\diamond$ ) For experiments in Sec. 6, we have instead used NVIDIA GeForce RTX 5090-32GB as GPU.

- In the FC+INR model, the total learnable parameters are those of the hypernetwork, which will quickly have a large output size as the inference network grows. However, because his model does not have the property of linear additivity that allows inference to generalise to multi-source examples, we have not tried to train a large network of this kind.

Consequently, we have found it easiest to express the hypernetwork width in units of its output size. The total number of parameters across both networks is then listed separately.

## C.2 Tuning tactics for the Fourier model

To coax strong performance from the FOURIER architecture, we use the following tactics, which run contrary to intuitions about Fourier series for sampled data:

**Selection of modes.** Our intuition for reconstructing a signal from sampled data would follow the constraints on wavelengths for the FFT—the shortest mode is set by the Nyquist sampling frequency and the longest mode is set by the domain window size. In fact, the domain size is irrelevant here *because the underlying function is not periodic*, and the modes should be significantly longer. In practice, we let the upper and lower bounds for the modes be free parameters, and use a logarithmically-spaced sequence (of length the order of the expansion), affixing the zero-frequency mode to automatically include the Fourier bias terms.

**Irregularity of data sampling.** The periodic Fourier modes mean that training on a regular grid encourages overfitting: the reconstruction oscillates wildly even a short distance from the node points. Good practice would in any event suggest that training with data that are irregularly sampled over the domain would be more reliable, and this is the case.

**Joint fitting of potential and field.** Given the practical interest in the field resulting from the model output, we find improved performance using a loss with a term for the target potential and an equally-weighted term fit to the target field, evaluated via the gradient of the potential.

## D High error rates from small sources in FC+ILR

The source size has a substantial influence on the metric  $\epsilon$  used in our experiments in Sec. 6. To analyse the root cause of this, we plot  $\epsilon_\phi$  against the prism-shaped source side lengths  $s_x, s_y$  in Fig. 8a. The data

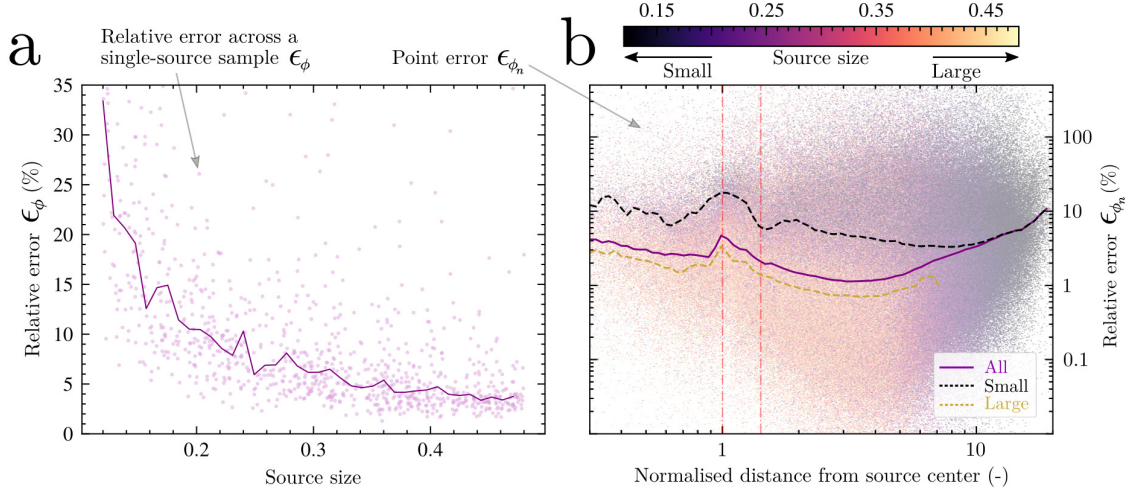


Figure 8: (a) Relative error  $\epsilon_\phi$  as a function of prism-shaped source side lengths. The 2D potentials generated by these configurations ( $M = 1$ ) is the same data as used in Fig. 5b. (b) Relative error  $\epsilon_{\phi_n}$  for single evaluation points from single-source configurations ( $M = 1$ ) grouped by distance to the source center normalised by the side length  $s_x$  of the source. The overall performance of FC+ILR is depicted in purple. The dotted lines show the grouped median relative error for small, i.e.  $s_x < 0.2$ , and large, i.e.  $s_x > 0.4$ , source sizes in black and yellow, respectively.

points correspond to the median relative error of  $10^3$  2D potentials generated by single prism-shaped source configurations ( $M = 1$ ) with varying side lengths ( $s_x = s_y \in [0.12, 0.48]$ ) in a  $[-1.2, 1.2] \times [-1.2, 1.2]$  domain. Each validation sample is evaluated at  $32^2$  uniformly sampled potential points. It can be clearly seen that when  $s_x < 0.2$ , then  $\epsilon_\phi > 10\%$ , while we obtain consistently low relative error for larger source sizes. To further investigate where this behaviour originates, we plot the relative error  $\epsilon_{\phi_n}$  for single evaluation points from single-source configurations ( $M = 1$ ) grouped by distance to the source center normalised by the side length  $s_x$  of the source in Fig. 8b. Across the whole range of the normalised distance, the relative error  $\epsilon_{\phi_n}$  of sources with  $s_x < 0.2$  is almost an order of magnitude larger compared to the error of sources with  $s_x > 0.4$ .

When we look at four random samples with side lengths  $s_x < 0.2$  from this dataset in Fig. 9, we spot two mechanisms responsible for large error rates: i) For small sources lying at the edge or corner of the sampling domain, the normalised distance to their source center becomes large and the dynamic range becomes considerably larger than in other samples in the dataset, which is apparently difficult for the neural network to learn, and ii) the relative errors become generally large in the axis perpendicular to the magnetisation, where the potential approaches values close to 0. The latter effect is also apparent for samples with larger sources size, but far less pronounced.






Figure 9: (a-d) Four samples of single prism-shaped source configurations ( $M = 1$ ) with side lengths  $s_x < 0.2$ . (e-f) Two samples of single prism-shaped sources with a side length of  $s_x > 0.4$ . Along with the target and source configuration, the prediction from FC+ILR, the respective spatially-resolved relative error  $\epsilon_\phi$  are presented on a regular grid with  $32^2$  evaluation points.