

~~TensorWeave~~ Tensorweave 1.0: Interpolating geophysical tensor fields with spatial neural networks

Akshay V. Kamath¹, Samuel T. Thiele¹, Hernan Ugalde², Bill Morris³, Raimon Tolosana-Delgado¹, Moritz Kirsch¹, and Richard Gloaguen¹

¹Helmholtz-Zentrum Dresden-Rossendorf, Helmholtz Institute Freiberg, Chemnitzer Str. 40, 09599 Freiberg, Germany.

²DIP Geosciences, Hamilton, ON Canada.

³Morris Magnetism Inc., Fonthill, ON Canada.

Correspondence: Akshay V. Kamath (a.kamath@hzdr.de)

Abstract. Tensor fields, as spatial derivatives of scalar or vector potentials, offer powerful insight into subsurface structures in geophysics. However, accurately interpolating these ~~measurements—such~~ measurements—such as those from full-tensor potential field ~~gradiometry—remains~~ gradiometry—remains difficult, especially when data are sparse or irregularly sampled. We present a physics-informed spatial neural network that treats tensors according to their nature as derivatives of an underlying scalar field, enabling consistent, high-fidelity interpolation across the entire domain. By leveraging the differentiable nature of neural networks, our method not only honours the physical constraints inherent to potential fields but also reconstructs the scalar and vector fields that generate the observed tensors. We demonstrate the approach on synthetic gravity gradiometry data and real full-tensor magnetic data from Geyer, Germany. Results show significant improvements in interpolation accuracy, structural continuity, and uncertainty quantification compared to conventional methods.

10 1 Introduction

Full-tensor gravity and magnetic gradiometry measurements capture the spatial derivatives of potential fields, offering rich detail about subsurface density and ~~magnetization~~ magnetisation variations. These tensor fields enhance geological imaging by encoding directional and gradient information that scalar fields do not straightforwardly provide (Brewster, 2011; Ugalde et al., 2024). However, gradiometry data are typically sparse and anisotropically sampled—often along sub-parallel flight
15 lines—posing significant challenges for downstream analysis, which relies on gridded representations.

Interpolating these tensor fields accurately is far from trivial. Conventional methods often treat tensor components as independent scalar fields, leading to noise amplification, loss of directional trends, and violations of physical constraints like symmetry and harmonicity. More advanced approaches, such as eigen-decomposition-based interpolation (~~Satheesh et al., 2023~~) (Fitzgerald et al., 2012; Satheesh et al., 2023), attempt to preserve tensor structure, but remain limited
20 in generalisability and scalability.

Recent advances in ~~neural fields~~ neural fields (Xie et al., 2022)—also known as coordinate-based or implicit neural representations—offer a promising alternative. These models learn continuous functions that map spatial coordinates to field values, ~~and crucially, they are differentiable. This.~~ Their differentiable nature allows them to incorporate gradient information directly into training—a key advantage for geophysical applications where tensor data often represents derivatives
25 of an underlying field (Raissi et al., 2019). However, standard multilayer perceptron (MLP) architectures suffer from spectral bias (Rahaman et al., 2019), meaning they struggle to capture high-frequency features common in geophysical signals (~~Rahaman et al., 2019~~). To address this, techniques like ~~random Fourier feature (RFF) mapping and~~ Random Fourier Feature mapping (Tancik et al., 2020), periodic activation functions (~~e.g., SIREN, Wavelet~~) (e.g. SiREN, Sitzmann et al., 2020), and wavelet activations (Saragadam et al., 2023) have been introduced, enabling neural fields to model fine-scale spatial
30 variations more effectively (~~Sitzmann et al., 2020; Saragadam et al., 2023~~).

In this paper, we introduce a physics-informed neural field approach tailored for interpolating geophysical tensor data. Our model learns a single scalar potential field from sparse tensor measurements, leveraging RFF mappings and embedded physical constraints (e.g., symmetry, Laplacian properties) to reconstruct consistent, physically meaningful tensor fields. We further introduce an ensemble strategy to quantify uncertainty in the interpolations, offering insights into data sensitivity and model
35 confidence. We demonstrate this framework on both synthetic gravity gradiometry data and real airborne magnetic gradiometry from Geyer, Germany, highlighting clear improvements over traditional methods in accuracy and structural continuity, as well as opening the door to uncertainty quantification.

2 Background

A tensor is an algebraic object that encodes multilinear relationships between sets of vectors and linear functionals (Lee, 2012).
40 A tensor field assigns a tensor to every point in space, allowing describing the local structure of a vector field or scalar potential

~~to be described~~ throughout a region. In geophysical applications, tensors naturally arise as derivatives of vector and scalar fields, extending classical multivariable calculus into field-based formulations.

2.1 Potential fields

Many measured geophysical quantities, such as gravitational acceleration $\mathbf{g}$ and the magnetic field $\mathbf{b}$, are conservative vector fields—i.e., they are gradients of scalar potential fields (Blakely, 1995). Within $\mathbb{R}^3$, a conservative vector field $\mathbf{v}$ is irrotational at all points (given by the position vector $\mathbf{r}$), satisfying

$$\mathbf{v} = \nabla \phi \quad \leftrightarrow \quad \nabla \times \mathbf{v} = 0$$

$$\mathbf{v} = \nabla \phi \quad \leftrightarrow \quad \nabla \times \mathbf{v} = 0 \tag{1}$$

50 Where $\nabla = \left[\frac{\partial}{\partial r_x}, \frac{\partial}{\partial r_y}, \frac{\partial}{\partial r_z} \right]$ is the gradient operator. For instance, the magnetic field can be expressed as the gradient of a scalar magnetic potential in regions free of electric currents—a condition typically met outside source distributions. Taking the gradient of $\mathbf{v}$ yields the Hessian tensor $\mathbf{H}$, a second-order tensor that captures the local curvature of the scalar potential

$$\mathbf{H} = \nabla \mathbf{v} = \nabla (\nabla \phi) = \frac{\partial^2 \phi}{\partial r_i \partial r_j} \equiv \partial_i \partial_j \phi \quad \forall i, j = \underline{1, 2, 3}.$$

$$\mathbf{H} = \nabla \mathbf{v} = \nabla (\nabla \phi) = \frac{\partial^2 \phi}{\partial r_i \partial r_j} \equiv \partial_i \partial_j \phi \quad \forall i, j = \underline{1, 2, 3} \tag{2}$$

In source-free regions, these fields are not only irrotational but also solenoidal—i.e., divergence-free. The divergence of $\mathbf{v}$ corresponds to the trace of the Hessian, which reflects the Laplace equation

$$\nabla^2 \phi \equiv \text{tr}(\mathbf{H}) = 0$$

$$\nabla^2 \phi \equiv \text{tr}(\mathbf{H}) = 0 \tag{3}$$

This implies that, outside source regions, scalar potentials are harmonic functions, and their Hessians are traceless. Additionally, since mixed partial derivatives commute (by Schwarz's theorem), the Hessians are symmetric and thus comprise five independent components.

2.2 Full Tensor Gradiometry

65 Direct measurements of second-order Hessian tensors—particularly gravity and magnetic gradient tensors—represent the current frontier in potential field surveying (Rudd et al., 2022; Stolz et al., 2021). Access to the full tensor enables ~~characterization~~ characterisation of scalar field curvature, aiding in tasks such as edge detection, structure delineation (Zuo and

Hu, 2015), and magnetic remanence ~~eharaeterization~~-characterisation (Ugalde et al., 2024). These measurements are typically acquired via airborne surveys, which are highly anisotropic in their sampling: dense along flight lines and sparse across them.

70 Vector fields are frequently reconstructed from tensor components using Fourier-domain transfer functions, which integrate the measured gradients into vector components while suppressing noise (Vassiliou, 1986). Since most downstream analyses, including Fourier-based reconstructions, require gridded tensor and vector fields, interpolation is a critical preprocessing step.

Rudd et al. (2022) note that, in practice, tensor components are often treated as separate scalar fields and interpolated using standard methods like minimum curvature or radial basis functions (RBFs). Several alternative approaches have been proposed

75 to enforce physical or geometric constraints during interpolation. For example, Brewster (2011) uses iterative Fourier-domain transformations, while Fitzgerald et al. (2012) suggest interpolating eigenvalues and eigenvectors separately—
, using a quaternion-based interpolation technique. In essence, the quaternion interpolation algorithm decomposes the process into two parts: interpolating the eigenvalues and the corresponding eigenvectors. Two of the three eigenvalues are interpolated using standard schemes (e.g., RBF or minimum curvature), with the third computed such that their sum is zero—a direct

80 consequence of the Hessian’s traceless-ness. The eigenvector matrix of any symmetric real matrix is guaranteed to be real and orthogonal, allowing it to be represented as a 3D orientation and encoded as a quaternion (Hamilton, 1844), provided some constraints on ordering and sign convention are imposed (Satheesh et al., 2023). These quaternions are then interpolated using Spherical Linear Interpolation or SLERP (Shoemake, 1985), which ensures smooth variation of orientation across space. While SLERP works for two quaternions, Markley et al. (2007) devised a scheme that works across a set of weighted quaternions.

85 Another widely used approach for interpolating and transforming potential-field (and gradient) data is the equivalent-source / equivalent-layer method: one replaces the true 3D distribution of sources by a 2D layer of hypothetical monopoles or dipoles beneath the observation surface whose field exactly reproduces the measured data on that surface (Dampney, 1969; Blakely, 1995). In practice the infinite layer is discretised into a finite set of sources and the corresponding dense sensitivity matrix is solved—often with regularisation—to obtain source strengths

90 that honour the physical constraints of potential fields and enable stable continuation and derivative transforms (Hansen and Miyazaki, 1984; Blakely, 1995; Oliveira Junior et al., 2023). This formulation is powerful but computationally demanding for large surveys. Consequently, a substantial literature focuses on accelerating the method by exploiting data geometry and matrix structure: scattered equivalent-source gridding (Cordell, 1992); the “equivalent data” concept to reduce system size (Mendonça and Silva, 1994); wavelet compression and adaptive meshing to sparsify

95 sensitivities (Li and Oldenburg, 2010; Davis and Li, 2011); fast iterative schemes in the space/wavenumber domains (Xia and Sprowl, 1991; Siqueira et al., 2017); and scalable algorithms that leverage the block-Toeplitz Toeplitz-block (BTTB) structure of the sensitivity matrix (Piaulino et al., 2025). Recent machine-learning-inspired variants (e.g., gradient-boosted equivalent sources) further cut memory and runtime for continental-scale datasets (Soler and Uieda, 2021). Open-source implementations, notably *Harmonica*, provide practical tools for gravity and magnetic datasets using these ideas

100 (Fatiando a Terra Project et al., 2024).

However, these methods still have limitations: component-wise methods can be insensitive to the true shape of the tensor, whereas full-tensor schemes involve complex handling of 3D rotations, which are complicated due to the existence of indeterminate points and the need for shifting reference frames due to non-uniqueness of 3D rotations. Equivalent source methods offer a physically consistent approach, but suffer from the trade-off between computational expense and fidelity of the interpolated result (e.g., Piaulino et al., 2025). In this contribution, we propose a neural field method that interpolates the scalar potential field directly—constrained by physical laws and Hessian measurements—to produce consistent, ~~noise-minimizing~~ noise-minimising tensor and vector fields that respect observations and preserve geologically meaningful structures.

2.3 Neural fields

Neural fields—also known as implicit neural representations, or spatial neural networks—are models that represent continuous spatial functions using neural networks. Unlike traditional methods that store information in discrete grids or meshes, neural fields encode spatial structure within the weights and biases of a neural network, enabling resolution-independent, continuous representations of complex signals.

The application of spatial neural networks in geoscience dates back to Openshaw (1993), who used them for interpolating sparse spatial data and found their performance competitive with fuzzy logic and genetic algorithms, a conclusion also reached by Hewitson et al. (1994). More recently, neural fields have gained traction in computer vision—for example, in volumetric radiance field modelling (Mildenhall et al., 2020)—and in geoscience applications such as 3D geological modelling (Hillier et al., 2023) and potential field representation (Smith et al., 2025).

A key advantage of neural fields is their differentiability: they allow access not only to predicted signals but also to their spatial derivatives via automatic differentiation. This is especially useful when the scalar field itself is ~~unobservable-unmeasurable~~ or physically arbitrary, but its gradients are measurable—as is often the case in geological modelling using structural orientation data (Kamath et al., 2025; Thiele et al., 2025), or in reconstructing potential fields from tensor gradiometry data.

3 Methodology

This section outlines the key components of our proposed framework, including the use of random Fourier features, a harmonic feature embedding, model architecture, training regimen, and loss function. We also describe the methodology used to generate the synthetic dataset used in our study.

3.1 Random Fourier Feature mapping

A common challenge in implicit neural representations is the mismatch between low-dimensional input coordinates and the complex, high-frequency structure of the target signal. To address this, we employ *Random Fourier Feature(RFF)-mapping* (RFF) mapping—a kernel approximation technique introduced by Rahimi and Recht (2007) and adapted to deep learning by

130 Tancik et al. (2020). RFF mapping projects spatial coordinates into a higher-dimensional space, making it easier for the network to learn fine-scale spatial variation.

Given the position vector $\mathbf{r} \in \mathbb{R}^N$, we define a frequency (also called weights) matrix $\mathbf{W}$, of the dimension $M \times N$, with every component sampled from a standard Gaussian distribution

$$W_{ij} \sim \mathcal{N}(0, 1)$$

135 ~~Where M , where M is the number of Fourier features .Furthermore, to (i.e., frequencies). To~~ encode known signal frequency characteristics (e.g., the maximum possible frequency based on sampling resolution), we rescale the weights matrix using different length scales. ~~For every length scale $\ell_s \forall s = 1 \dots L$, we map the position vector onto~~ Therefore, for a 3D input, we get a $2M$ dimensional feature ~~space, giving us a feature-vector $\boldsymbol{\nu}$ with the components given by~~

$$\nu_i = \sin \left(2\pi \mathbf{W}_s \mathbf{r} + \delta_i \frac{\pi}{2} \right) \quad \forall i = 1 \dots 2M, \text{ where } \mathbf{W}_s = \frac{\mathbf{W}}{\ell_s}, \delta_i = \begin{cases} 0, & i \text{ is even} \\ 1, & i \text{ is odd} \end{cases}$$

140 $\boldsymbol{\nu}_s$ for every length scale ℓ_s given by

$$\boldsymbol{\nu}_s = [\sin(2\pi \mathbf{W}_s \mathbf{r}), \cos(2\pi \mathbf{W}_s \mathbf{r})], \text{ where } \mathbf{W}_s = \ell_s^{-1} \mathbf{W}^{M \times 3} \quad (4)$$

Where $\mathbf{r} = [x, y, z]^T$, $W_{ij} \sim \mathcal{N}(0, 1)$, $\sin(\mathbf{x}) := [\sin(x_i)]_i$, and $[:, :]$ represents concatenation along the feature axis. Hence, for L length scales, as the feature vectors are concatenated, we get a $2ML$ -dimensional feature vector. This feature vector is then fed into the subsequent multi-layer perceptron to get the scalar potential at the input coordinate.

145 The transformation enables the model to capture high-frequency details more effectively, while the random sampling of frequencies introduces a useful stochastic component. When followed by a linear MLP with no non-linear activations, the resulting mapping approximates a full Fourier reconstruction of the signal (Bracewell and Kahn, 1966). Non-linear activations help the model fit sparse data more flexibly (LeCun et al., 2015), albeit at the cost of simplicity, interpretability, and gradient stability.

150 3.2 Harmonic feature embedding

Applying Fourier features uniformly in all dimensions can ~~lead to poor~~ hinder convergence when modelling harmonic ~~functions. fields. By~~ Liouville's Theorem (Axler et al., 2001) ~~states that any bounded, periodic, harmonic function in, any bounded harmonic function on $\mathbb{R}^N$ must be constant. This can cause the network to collapse onto is constant, so naive periodic embeddings can bias the network toward~~ trivial solutions. ~~To avoid this, we partition the position vector into horizontal and~~

155 ~~vertical components~~

$$\mathbf{r} = \mathbf{r}_{xy} + r_z \hat{k} \quad \text{where } \mathbf{r}_{xy} \in \mathbb{R}^2$$

~~We apply the RFF mapping. We therefore introduce a *harmonic* Fourier mapping that uses Fourier features in the horizontal (xy) plane and encode a decaying term in the vertical (z) direction, to replicate the observed decay of scalar potentials and to encourage harmonicity. The modified feature vector ν therefore has its components given by~~

$$160 \quad \nu_i = \sin\left(2\pi \mathbf{W}_s \mathbf{r}_{xy} + \delta_i \frac{\pi}{2}\right) \odot e^{-\|\mathbf{W}_s\|_2 r_z} \quad \forall i = 1 \dots 2M, \text{ where } \mathbf{W}_s = \frac{\mathbf{W}}{\ell_s}, \delta_i = \begin{cases} 0, & i \text{ is even} \\ 1, & i \text{ is odd} \end{cases}$$

plane and an analytically motivated vertical decay. Since our model reconstructs the spatial domain signal from a combination of sinusoids, we use the frequencies of the sinusoids to encode the harmonicity condition into the mapping. Specifically, in the Fourier domain, if we couple Laplace’s equation with the standard separation-of-variables method for a scalar potential ϕ , we get an ordinary differential equation of the form

$$165 \quad (k_z^2 - \|\mathbf{k}_h\|_2^2) \mathcal{F}_{x,y}(\phi) = 0 \quad (5)$$

Where $\mathbf{W}$ is now an $M \times 2$ dimensional weights matrix, $\|\mathbf{W}_s\|_2$ denotes the row-wise Euclidean norm of the weights matrix, and $\odot$ denotes element-wise multiplication. This formulation helps the network respect the expected vertical decay of scalar potentials, aligning the embedding with $\mathcal{F}_{x,y}(\cdot)$ is the 2D Fourier transform operator, $\mathbf{k}_h = [k_x, k_y]$ represents the horizontal wavenumbers (i.e., frequencies), and k_z is the vertical wavenumber (Lacava, 2022). This implies that the vertical wavenumber is constrained by $k_z^2 = k_x^2 + k_y^2$, making the vertical dependence evanescent and not oscillatory. This is the classical half-space solution of Laplace’s equation. While the network output remains a non-linear combination of the features, this decay-aware encoding enhances generalization in under-sampled regions and theoretically supports and underpins upward/downward continuation in potential-field geophysics (Blakely, 1995; Parker, 1973). Our harmonic embedding scheme simply bakes in the same physics at a feature-level, helping with convergence while training on mostly co-planar datasets, and potentially allowing robust upward/downward continuation—though that is outside the scope of this study.

For a 3D input $\mathbf{r} = [x, y, z]^T$, we extract the horizontal coordinates $\mathbf{r}_{xy} = [x, y]^T$. As defined in the previous section, we generate a random weights matrix $\mathbf{W}^{M \times 2}$ with the entries independently drawn from a standard normal distribution $W_{ij} \sim \mathcal{N}(0, 1)$, and scale it with the length scale ℓ_s to acquire a scaled matrix $\mathbf{W}_s = \ell_s^{-1} \mathbf{W}$. For every length scale, we define a new vector, κ_s , such that

$$180 \quad (\kappa_s)_m = \|\mathbf{W}_{s,m}\|_2, \quad \kappa_s \in \mathbb{R}^M \quad (6)$$

Where m refers to the m th row of the $\mathbf{W}_s$ matrix, giving us a vector of length M i.e., one norm per row (per feature). With element-wise sine/cosine and Hadamard product $\odot$, the new feature vector ν_s for the scale ℓ_s is

$$\nu_s = [\sin(2\pi \mathbf{W}_s \mathbf{r}_{xy}) \odot e^{-\kappa_s z}, \cos(2\pi \mathbf{W}_s \mathbf{r}_{xy}) \odot e^{-\kappa_s z}] \quad (7)$$

where z is the vertical coordinate, and the exponential is applied element-wise, producing a $2M$ -dimensional vector. Concatenating across L scales yields a $2ML$ -dimensional embedding that encodes horizontal oscillations with physically

consistent vertical decay, aligning the features with solutions of Laplace’s equation and improving generalisation in under-sampled regions.

3.3 Synthetic dataset

To evaluate our method, we generated a synthetic gravity gradiometry dataset (Fig. 1) using SimPEG (Cockett et al., 2015).

190 The model consists of three density-contrast anomalies within a zero-density half-space:

1. *Torus*: $+1 \text{ g cm}^{-3}$, semi-major axis 450m, semi-minor axis 220m, cross-section radius 40m, lying in the xy plane and rotated 12° CCW from the y -axis.
2. *Dyke*: $+0.15 \text{ g cm}^{-3}$, 60m aperture, striking at 45° to the y -axis.
3. *Cube*: -0.2 g cm^{-3} , 400m side length, rotated 45° about the vertical (z) axis.

195 The simulation mesh has a voxel size of 20m. This geometry offers a challenging mix of sharp discontinuities and smooth curvature for testing interpolation. We generated a high-resolution ground truth dataset sampled at 25m spacing both along and across the lines, as well as a low-resolution airborne-style dataset with flight lines 200m apart in the y direction (perpendicular to the flight line), and sampled densely (15m) along the x direction (Fig. 1B-F). ~~Gaussian noise was added to simulate measurement uncertainty~~b-f). Furthermore, to make the data more realistic, We corrupt the full-tensor gradiometry data with
200 ~~additive white Gaussian noise at a prescribed signal-to-noise ratio (SNR). For each independent component of the Hessian H_k (where $k \in \{xx, xy, xz, yy, yz\}$), we estimate the power spectrum $P_k = \langle H_k^2 \rangle$, convert the target SNR from dB to linear units as:~~

$$\text{SNR}_k = 10^{\text{SNR}_{k, \text{dB}}/10} \quad (8)$$

Using this SNR, we set the noise variance of the signal as $\sigma_k^2 = \frac{P_k^2}{\text{SNR}_k}$. Independent samples $\varepsilon_{i,k} \sim \mathcal{N}(0, \sigma_k^2)$ are then added
205 to each datum, to acquire a noisy synthetic dataset ($\tilde{H}$), given by:

$$\tilde{H}_{i,k} = H_{i,k} + \varepsilon_{i,k} \quad \forall i = 1 \dots P \quad (9)$$

Where P is the number of points in the dataset.

To test robustness to data sparsity, we also computed 10 versions of the low-resolution dataset with line spacings varying from 80m to 560m. These were used to benchmark interpolation quality and information loss under varying acquisition
210 densities. Comparisons were made with a truncated RBF interpolator (250 nearest neighbours, smoothing factor 100), as well as results from the quaternion interpolation (QUAT, Fitzgerald et al. (2012)), (QUAT: Fitzgerald et al., 2012), combining RBF-interpolated eigenvalues with SLERP-interpolated quaternions. All results were evaluated on the same high-resolution grid using the R^2 (coefficient of determination), MSE (Mean Squared Error), and SSIM (Structural Similarity Index Measure).

3.4 Model architecture, ~~training regimen~~, and loss ~~function~~functions

215 Our model has two main components: a RFF mapping block followed by a sequence of fully-connected feed-forward layers that together produce a continuous scalar field representation (Fig. 2. ~~We found that the best-performing model for our usage used-). In our tests, we varied the number of Fourier features from anywhere between 16 to 64 RFFs and three hidden layers with 1024, 512, and 512 neurons each-,~~ depending on the complexity of underlying field. The specifications of each individual model showcased in this contribution can be found in the following sections.

220 The MLP block in our model uses non-linear activations for all layers except the output layer. As our framework involves computing second derivatives with AD, activation functions like *ReLU* (which do not satisfy the C^2 differentiability criterion) resulted in abrupt edges within the resultant interpolation. Notably, even within the activations that satisfy the aforementioned criterion, some functions performed better than the others. For example, the Hyperbolic Tangent activation function has extremely small second order derivatives which tend to get saturated, impeding convergence.

225 These activations are stable, but not ideal for our models. Among the various activation functions tested, *Swish* (Sigmoid Linear Unit, SiLU; Ramachandran et al., 2017) and *Mish* (Misra, 2019) activations provided the best results.

The loss function used to train our model involves two types of losses - a data loss and a Laplacian loss. The data loss is computed at the points of measurement between the measured tensor components and the ~~hessians~~-Hessians acquired from the predicted scalar field through automatic differentiation (AD) (Margossian, 2019). Since the model is built with ~~pytorch~~ ~~(Paszke et al., 2019)~~PyTorch (v2.8.0; Paszke et al., 2019), we use the inbuilt `torch.autograd` ~~autograd~~ engine to compute ~~hessians~~-Hessians from the scalar field output. For the predicted scalar field ϕ , the data loss term is given by

$$\mathcal{L}_d = |\partial_i \partial_j \phi - H_{ij}| \quad \forall i, j = 1, 2, 3$$

$$\mathcal{L}_d = \frac{1}{P} \sum_{p=1}^P |\partial_i \partial_j \phi^p - H_{ij}^p| \quad \forall i, j = 1, 2, 3 \quad (10)$$

235 Where $\partial_i(+)\partial_j\phi^p$ refers to the partial derivative with the respect to the i th input computed with AD ~~, and $\mathbf{H}$ is the~~at the p th location, and $\mathbf{H}^p$ is the corresponding measured hessian tensor. ~~Note that this reconstruction loss is computed using all the measurement points.~~The first, second, and third indices correspond to x (East-West), y (North-South) and z (Up-Down) axes respectively. Only the upper-triangular part of the Hessian is used for loss computation i.e., the losses for the off-diagonal components are only considered once per measurement. Hence we get a six-component data loss vector, consisting of the misfit

240 ~~between the xx , xy , xz , yy , yz and zz components.~~

The second term in the loss function is derived to encourage the predicted scalar field to conform with a partial differential equation defined across the whole domain of interest. Since the predicted field has to be harmonic not only at the points of measurement, but everywhere, we thus penalise non-zero traces of the predicted Hessian tensors, hereafter referred to as the Laplacian loss. During every training epoch, ~~hessians are computed for randomly sampled points within the domain of~~

245 interest(with the number of points usually set to be equal to the number of measurement points). The Laplacian loss for these predicted tensors is then computed as-

$$\mathcal{L}_l = |\text{tr}(\nabla(\nabla\phi))| = |\partial_i\partial_i\phi|$$

we use a Poisson disk sampling routine based on a hierarchical dart throwing approach (White et al., 2007), to select evenly spaced points within a predefined grid that covers the area of interest. The spacing for these points is evaluated using an exponentially decaying function that goes from a user specified large spacing (at the start of training) to a tighter spacing (towards the final epochs). For the Q sampled points in any epoch, the Laplacian loss is given by:

$$\mathcal{L}_l = \frac{1}{Q} \sum_{q=1}^Q |\text{tr}(\nabla(\nabla\phi^q))| = \frac{1}{Q} \sum_{q=1}^Q |\partial_i\partial_i\phi^q| \quad (11)$$

Here, the Einstein summation convention is used to represent the Laplacian, and the superscript refers to the point of evaluation. This loss penalises high values of the trace of the predicted hessian tensor outside the measured points, thereby encouraging harmonicity on the underlying scalar field within the domain of interest. Hence, for every epoch, we get a seven component combined loss vector, with the first six components corresponding to the data loss, and the seventh component referring to the global Laplacian loss.

When combined, the total loss that is optimised is thus-

$$\mathcal{L} = \alpha_d \mathcal{L}_d + \alpha_l \mathcal{L}_l$$

260 acquired from equations 10 and 11 is thus

$$\mathcal{L}_{total} = \sum_{i=1}^7 \alpha_i \mathcal{L}_i \quad (12)$$

The hyper-parameters α_d and α_l are (optionally) initialised as fractions of the respective initial losses, in order to help find an appropriately balanced weighting of the different loss terms. Where $\mathcal{L}_i$ is the i th component of the combined loss vector, $[\alpha_i]$ are the corresponding hyperparameters. Instead of manually fine-tuning these hyperparameters, we tested various multi-objective optimisation schemes for our framework. The best-performing scheme involved dynamically updating hyperparameters, such that every loss function was scaled by the real-time magnitude of the loss. Mathematically, we replace the $[\alpha_i]$ in Eq. 12 as follows:

$$\mathcal{L}_{total} = \sum_{i=1}^7 \frac{\mathcal{L}_i}{\tilde{\mathcal{L}}_i} \quad (13)$$

Where $\tilde{\mathcal{L}}_i$ refers to the magnitude of the loss, detached from the computational graph. This ensures that the gradients only flow through the numerator of the scaled loss, even when the real-time value of the loss is always equal to 1. This real-time

normalisation yields a scale-invariant objective whose update is approximately $\nabla_{\theta} \sum_i \log \mathcal{L}_i$ (where ∇_{θ} is the gradient with respect to the network parameters), encouraging proportional improvements across terms. Similar in spirit to uncertainty-based weighting (Kendall et al., 2017), *GradNorm* (Chen et al., 2017), and Density Weight Averaging (Liu et al., 2019), this variant requires no extra parameters or gradient-norm computations and worked reliably in our setting.

275 We train the MLP parameters (weights and biases) with *Adam* (Kingma and Ba, 2014), while keeping the RFF encoder fixed after initialisation. The initial learning rate is set to 10^{-3} (occasionally 10^{-2} when the initial loss scale is large). We apply a plateau scheduler that reduces the learning rate by a factor of 0.8 whenever the loss fails to improve for 20 epochs. Optionally, we also optimise a set of learnable length-scale parameters that modulate the Fourier features; the log-values of these scales are stored as parameters and updated jointly with the MLP. However, the learnable nature of these length scales did not help
280 the model convergence greatly, reproducing results explored by Tancik et al. (2020), which suggested that neural fields fail to suitably optimise these length scale parameters.

To make the stopping criterion robust to small oscillations, we monitor an exponential moving average (EMA) of the pre-scaled loss:

$$\hat{\mathcal{L}}_n = \beta \hat{\mathcal{L}}_{n-1} + (1 - \beta) \mathcal{L}_n \quad \beta \in (0, 1) \quad (14)$$

285 Where $\hat{\mathcal{L}}_n$ denotes the smoothed loss on the n th epoch. This combination—*Adam* for fast, well-scaled updates, plateau-based learning-rate decay, and EMA-stabilised early stopping—follows common best practices for training smooth function approximators and has been shown to curb overfitting while maintaining convergence speed (Goodfellow et al., 2016; Prechelt, 1998; Bottou, 2012).

3.5 Uncertainty estimation

290 A key benefit of using RFF embeddings is that their stochastic nature allows for ensemble-based uncertainty estimation. As a result of the stochasticity, each initialisation of the RFF mapping induces a unique basis in the feature space, causing the neural network to converge on a solution that represents a random sample from a broader distribution of plausible scalar fields conditioned on the training data.

To exploit this property for uncertainty quantification, we generate an ensemble of model outputs by varying the random seed
295 used to sample the RFF projection matrix. Ensemble-based uncertainty quantification has a long and successful history in geophysics, particularly in subsurface modelling and inversion. In seismic full waveform inversion (FWI), ensembles have been used to assess the variability and reliability of recovered velocity models under data noise and model ambiguities (Fichtner et al., 2011). In reservoir geophysics, the Ensemble Kalman Filter (EnKF) has become a widely used tool to propagate uncertainty in dynamic reservoir simulation and history matching (Evensen, 2009). More recently, ensemble-based methods
300 have also been applied to probabilistic gravity and magnetotelluric inversion (Tveit et al., 2020; Giraud et al., 2023), demonstrating their utility in quantifying non-uniqueness and guiding data acquisition strategies.

In our implementation, each ensemble member corresponds to a different ~~realization~~realisation of the frequency space, leading to stochastically independent function approximations that depend, largely, on the degree to which the solution is constrained by the available data. This ensemble-based approach provides a Monte Carlo-style estimate of the model’s epistemic uncertainty.

305 Furthermore, because the scalar field is modelled continuously, we can propagate this ensemble approach to the field’s derivatives, helping us quantify uncertainty in derived physical quantities. Therefore, we showcase our results as the Ensemble Neural Field (ENF) method, which corresponds to the average prediction from an ensemble of ~~100 models~~models. We also compute results from the individual models within the ensemble (shown as the Neural Field or NF result), to ascertain the effect of averaging over multiple predictions.

310 4 Results

4.1 Synthetic Data

We first evaluate the Ensemble Neural Field (ENF) method on the synthetic gravity gradiometry dataset, comparing it against a Truncated Radial Basis Function (RBF) interpolator (Fig. 3). ~~Panels A and B~~The ensemble showcased here has 25 models, each with 16 Fourier features, three length scales of [200, 400, 1000], and two hidden layers with 256 neurons each. Each model within the ensemble was trained for 400 epochs. A predefined grid with a cell-size of 25 m was provided for evaluating the Laplacian loss, with the Poisson sampling radius going from 250 m to 80 m. Panels (a) and (b) show the residuals between predicted and true H_{xy} values for the RBF and ENF methods, respectively. The RBF output exhibits high-amplitude residuals (MSE = ~~4.698~~4.60 eotvos) between flight lines, indicating overfitting to sampled regions and poor generalisation across them. It also fails to preserve continuity in linear trends that lie at high angles to the flight direction. In contrast, the ENF method

315 yields spatially smoother residuals with significantly lower error (MSE = ~~0.825~~0.30 eotvos; improvement of $\approx 93.4\%$ over the RBF), suggesting homogeneous improved performance across the domain. Insets in both panels show 1:1 scatter kernel density plots, where the ENF predictions cluster more tightly along the identity line—further confirming its accuracy.

~~Panel C summarises~~For a quantitative measure of the improvement offered by our method, we plot the R^2 scores for each tensor component across ~~four~~three interpolation methods: RBF, and two neural field-based (NF and ENF) ~~-(Fig. 3c)~~. The NF method reflects the mean R^2 from ~~100~~25 independently trained models, with error bars showing standard deviation. The ENF method, by contrast, uses the averaged prediction across those same models. Both neural field approaches outperform classical methods, with ENF showing a slight edge—demonstrating that ensemble averaging reduces variance and enhances prediction stability. Fig. 3d shows the loss curves for the various losses for one of the models within the ensemble, as a function of epochs. The data loss terms reasonably plateau after reaching values of ≈ 1 eotvos, while the real-time updating hyperparameters help avoid overfitting to a single component. The Laplacian loss (dotted pink line; $\mathcal{L}_l$) keeps steadily decreasing as the sampling gets tighter and ever more points are sampled from the grid.

325

330

To further evaluate structural accuracy, we compute the Structural Similarity Index Measure (SSIM) between predicted and true tensor fields (Fig. 4). The ENF method achieves higher SSIM scores across all three components—~~0.89~~0.95 (H_{xx}), ~~0.90~~

0.97 (H_{xy}), and ~~0.87-0.96~~ (H_{xz})—compared to ~~0.79, 0.63, and 0.79~~ 0.78, 0.64, and 0.76 for RBF. The greatest improvement
 335 is seen in H_{xy} (improvement of $\approx 50.46\%$), where RBF results show structural distortion, over-smoothing, and “boudinage”
 artefacts along flight lines (Naprstek and Smith, 2019). ENF, on the other hand, preserves coherent anomalies and directional
 continuity even across sparsely sampled regions.

4.2 Rate of information loss

To assess robustness under sparse sampling, we compare the interpolation results for varying line spacings from 80m to 560m
 340 (Fig. 5). Classical methods (RBF and quaternion-based interpolation, or QUAT) show sharp drops in accuracy beyond 200m
 spacing. For example, the RBF method’s ~~average root-mean-squared~~ R^2 (computed over the components) drops to 0.54, and the
root-mean-squared SSIM plummets to ~~0.43 and SSIM to 0.21~~ 0.26 at 560m. In contrast, NF interpolation maintains relatively
 stable performance up to ≈ 400 m spacing, with a much gentler decline at wider gaps. At 560m, the NF model still achieves
~~an a root-mean-squared~~ R^2 of 0.87-0.91 and an SSIM of 0.580.65.

345 The MSE trends mirror this behaviour: classical methods exhibit steep error increases with sparser lines, while the NF
 model degrades more gracefully. QUAT offers minor improvements over component-wise interpolation but follows a similar
 performance trajectory. This suggests that the main bottleneck in full tensor interpolation lies in the eigenvalue interpolation
 step, which—like the component-wise case—relies on RBF methods.

4.3 Magnetic gradiometry from Geyer

350 Finally, we validated the method on real airborne magnetic gradiometry data from Geyer, located in Germany’s Erzgebirge
 region—part of the Central European Variscides. The area features high- and medium-pressure metamorphic units, orthogneiss
 domes, and post-orogenic granites (Kroner and Romer, 2013), with abundant ore-forming skarns containing magnetic minerals
 (Burisch et al., 2019; Lefebvre et al., 2019) as well as magnetite-rich quartzites and amphibolites that occur as intercalations
 within the metamorphic rocks. These rocks contribute to complex magnetic anomalies ideal for real-world evaluation.

355 We test the ENF method on a real airborne magnetic gradiometry dataset from Geyer (Fig. 6), acquired by Supracon AG in
 2016 as part of the E³ (ErzExploration Erzgebirge) project. As in the synthetic case, we compare ENF to RBF interpolation.
Due to the complex nature of the signal, we run a 50 model ensemble for the Geyer dataset. Each model uses 64 Fourier
features, with four length scales of [220, 400, 1000, 100000]. The number of hidden layers is increased to three, with 1024,
512 and 256 neurons respectively. Each model is trained for 600 epochs, with early stopping triggered after 30 epochs of no
 360 improvement. A predefined grid with a cell-size of 20 m is used to evaluate the Laplacian loss, with the Poisson sampling radius
starting at 500 m, and going to 200 m. Every fourth flight line is used for training, with the others reserved for ~~validation testing~~
the interpolation. Since ground-truth grids are unavailable, we assess accuracy using residual analysis and R^2 scores computed
~~along for points in~~ the withheld lines.

We plot the residual maps for H_{xy} on test lines (Fig. 7). While absolute R^2 scores are lower than in the synthetic case—owing
 365 to added geological complexity and noise—ENF still achieves ~~10-20~~ more than 30% better performance than RBF across most

tensor components, with a whopping increase of $\approx 157\%$ in the R^2 score for H_{yy} , and an average increase of approximately 57.27%. Both NF and ENF results are better than the RBF across all components. Residuals show that ENF (Fig. 7a) reduces systematic bias between lines and preserves anomaly shapes more faithfully. RBF (Fig. 7b), by contrast, displays patchy behaviour with abrupt shifts between lines—a well-known artifact of interpolating sparse or anisotropically sampled data (Hillier et al., 2014; Wittwer, 2009). The loss curves (Fig. 7d) show a similar trend to the synthetic case, with the data loss terms plateauing around 0.1 nT m^{-1} , and the Laplacian loss steadily decreasing.

To get a qualitative overview of the overall result, we plot the histogram-equalized, gridded visualisations of the H_{xx} , H_{xy} , and H_{xz} tensor components (Fig. 7). Panels A-C show the 8). We also compute a result from using all of the flight lines with an RBF interpolator (Fig. 8a-c), serving as our ground-truth. The RBF results from using every fourth line (Panels D-F Fig. 8d-f) reveal strong aliasing and inconsistent behaviour between flight lines—hallmarks of inadequate cross-line interpolation. In contrast, the ENF interpolations (Panels G-I Fig. 8g-i) exhibit smoother transitions and clearer structural trends, especially in directions orthogonal to flight lines. The ENF model successfully mitigates high-frequency striping and captures geologically meaningful features.

5 Discussion

5.1 Accurately reconstructing tensor fields

The proposed Neural Field (NF) Interpolator has shown remarkable success in interpolating tensor gradiometry data. Our results show that the additional information contained within the hessian tensor can help derive a more accurate reconstruction of the entire field as sampling gets sparser (Fig. 5), provided the interpolation algorithm can access the full tensor constraints. For equivalent inputs, the NF interpolation recovers a signal that better fits all the tensor components, while maintaining the integrability and physical properties inherent to a hessian tensor field.

We also see equivalent results from all methods when line spacings are tight (i.e., for line spacings of 80m, 100m and 120m in our synthetic tests). This suggests an oversampling with respect to the spatial frequencies in the signal, such that all the interpolation methods converge to the same (correct) result to yield high accuracy metrics. Results then diverge as line spacing increases to 200m, indicating the neural field interpolation is able to leverage information in the shape of the tensors to continue to derive accurate reconstructions, while the RBF and quaternion methods cannot.

The reason that the results converge with close spatial sampling could be attributed to the equivalence of SLERP and standard linear interpolation as the angle between the quaternions describing the orientations of the input data points goes to zero. Since a tighter line spacing ensures a smoother graduation of the eigenvector orientations (i.e., a smaller change in the angle between the corresponding quaternions), the resulting interpolation is closer to what one would achieve with standard linear interpolation of the components. But, under sparse sampling conditions, the differences seen in the results indicate that an

interpolation using neural field formulation better preserves the shape of interpolated tensors, without the need for cumbersome quaternion formalisms.

The interpolated tensor components for Geyer (Fig. 8) also showcase significant improvements over the component-wise interpolation of these tensors. The extension and continuation of the trend from the centre of the grid, towards the north-east is preserved in the ENF result, but is completely absent in the RBF result. Any interpretation of these grids would thus result in significantly different geological structures, highlighting the necessity for appropriate interpolation methods.

The Laplacian constraint is handled with an objective minimisation approach in our method. One could potentially enforce harmonicity by design, however this is challenging for 3D (i.e. geophysical potential) fields and difficult to enforce through the non-linear activation functions inherent to neural networks. In 2D, holomorphic functions (i.e., complex-differentiable functions of multiple variables) consist of real and imaginary parts that are harmonic functions, a fact that is utilised by Harmonic Neural Networks (e.g., PIHNNs; Calafà et al., 2024) to yield exactly harmonic outputs. These concepts do not directly extend to 3D, promoting an objective driven enforcement of the constraint. Vector potential based formulations (e.g. CurlNet; Ghosh et al., 2022) enforce divergence-free fields but fail to enforce the zero curl constraint. Furthermore, as our network consists of non-linear activations, and as non-linear compositions do not generally preserve harmonicity (Chen et al., 2010), we are further motivated to rely on our new mapping that has harmonic elements (see Section 3.2) and use an objective to constrain the Laplacian.

5.2 Recovery of vector fields

Many analysis methods applied to tensor gradiometry data require a domain-wide integral to estimate the underlying vector field. The simplest way of computing this integral is by ignoring everything but the last row of the gradiometry tensor, and using the H_{xz} , H_{yz} , and H_{zz} components to get vector components. Due to the Fourier domain properties, vector components are defined as a vertical integral in the Fourier domain (Mickus and Hinojosa, 2001). Similarly, the power spectrum of these signals can also be used to generate vector components, using transfer functions that fit all of the signals while minimising noise (Vassiliou, 1986). However, in our method, we can completely avoid this potentially complex integral. We can use automatic differentiation to acquire the vector field components from the predicted scalar potential as the neural field predicts scalar potential and not the gradiometry tensor itself. Importantly, we thus estimate the vector field components exclusively from real measurements, rather than from an integral over a regularly spaced (i.e. interpolated) grid that is already one-step removed from the data.

To test the recovery of vector components from our model, we compared it to the benchmark generated using the RBF interpolation on all flight lines and then applying Fourier domain transfer functions to compute the integral. We also use the transfer functions on the RBF interpolation results for our training data for a baseline comparison (Fig. 9). Comparing the resulting ~~bx~~ b_x (Fig. 9, Panels ~~A, D, and G~~a, d, g) components, we see that features present in both the ground truth and the ENF results are completely erased from the RBF result. Similarly, the shape of the anomaly at the top-right corner of the grid is distorted in the RBF result, but completely preserved within the ENF grid. Slight changes in trend directions (i.e., the

shift of the strike of the anomalies to having a smaller azimuth) also cannot be seen in the RBF results, which has prominent
430 “boudinage” artefacts along the flight lines that cause a loss of trend and directional information perpendicular to the flight line.
We suggest that these results highlight the ability of the neural field interpolation to extract sensible information (resembling
the ground truth) from data acquired at four times the line spacing.

5.3 Uncertainty analysis and ensemble models

We also used the stochastic nature of our feature embeddings to do a preliminary uncertainty analysis for the results from our
435 interpolator for the Geyer dataset (Fig. 10). The standard deviation plot shows higher variability in model predictions across
regions without data points (i.e., between the flight lines), which could be interpreted as an uncertainty measure. Interestingly,
the variance between flight lines seems to scale with the value of the underlying tensor component, leading to heteroscedasticity
in the predictions. This might need correction in future developments of our methodology. It is also worth noting that the NF
approach has parallels to the turning bands and spectral methods to simulate random fields (Mantoglou and Wilson, 1982),
440 suggesting that a deeper stochastic link to other Gaussian process methods may be possible. This link could be exploited to
better understand the variance of neural field ensembles or consider future modifications of the present NF algorithms towards
tuned frequency matrix distributions(Equation-??).

The variance of our ensemble model is generally higher for the components with two derivatives in the same dimension (i.e.,
 H_{xx} , H_{yy} , and H_{zz}), and for the derivatives involving the z component (i.e., H_{xy} seems to be the least uncertain). High same-
445 dimension double derivative uncertainties might reflect the propagation of uncertainty through differentiation, as uncertainties
in two variables have a chance of cancelling out, but are only amplified with multiple passes through the same derivative
operator Li and Oldenburg (1998)(Li and Oldenburg, 1998). The high uncertainty in the z components likely reflects the lack
of information in the z direction, as all of our training data are close to ~~eo~~planar~~co-planar~~. Furthermore, we also see that
the uncertainty in the recovered vector components (Fig. 10, Panels G, H, Ig, h, i) never goes to zero (even where we have
450 measurements of the tensor), reflecting the lack of information on the constant of integration.

Interpolated grids alter the observation error model: smoothing and continuation introduce spatially correlated errors that, if
ignored, can bias ensemble-based inversions (EnKF). Best practice is naturally to invert at the real measurement locations,
however when a grid is needed we suggest that our ENF ensemble could provide a mean and a sample covariance for the
pseudo-observations. It is possible (although untested) that this might be used as the observation-error covariance in the
455 inversion.

5.4 Challenges and future directions

We suggest that the proposed approach opens the door to using neural fields for potential field geophysics, and broader
applications involving tensor quantities (e.g., stresses and strains). However, further work and research is needed in several
areas. Firstly, our model is highly sensitive to the length scales chosen for the Fourier encoding,~~and the hyper-parameters.~~
460 As shown by Tancik et al. (2020), optimisation algorithms fail to tweak these scales, meaning they need to be selected with

careful empirical tuning. Furthermore, while we have utilised a real-time updating hyperparameter based on the magnitude of the loss function. This requires a lot of empirical tuning. Automatising, research into other possible avenues of automatising hyper-parameter tuning ~~would~~ could boost the usability of our method and help to ensure robust results. Furthermore, the inclusion of features with multiple length scales, while tested within our synthetic dataset, needs confirmation for effectiveness in larger areas.

In addition, while the recovery of integrated vector fields is a big advantage of our approach, these have arbitrary integration constants. This ambiguity means that, for every vector component, there is a constant that is unbounded in the other two dimensions. The same problem occurs when we use the Fourier domain transfer functions, as a fundamental lack on long wavelength information leads us to misrepresenting the baseline for the recovered vector field (Ugalde et al., 2024). However, in our methodology, this could be resolved with a few measurements of the vector components included as constraints on the neural field. Therefore, one additional future direction would be to include multiple datasets (e.g., TMI measurements for magnetic gradiometry, satellite or ground gravity measurements for gravity gradiometry) during the training process. Further research on the propagation of uncertainties through our model, as well as impact of ensembles during inversion, would help in improving the robustness of our proposed framework.

Finally, the inclusion of a harmonic decaying term in the feature mapping makes our method a possible contender for an innovative downward continuation scheme, and thus help with the problem of noise amplification in the ill-posed downward continuation of potential field anomalies. This application needs further research, with proper tuning of the weight matrices and data acquired at multiple elevations for validation.

6 Conclusion

We introduce an innovative ~~Neural Field (NF)~~ interpolation method tailored to tensor gradiometry data in potential field geophysics. This approach leverages the inherent physical relationships among tensor components by representing them as derivatives of an underlying scalar potential field. Our method clearly demonstrates advantages over conventional interpolation techniques, particularly in scenarios involving sparse and anisotropic data coverage, as are typical during aerial surveys.

Our method has shown substantial improvements in interpolation accuracy, structural fidelity, and robustness against data sparsity during evaluations on both synthetic gravity gradiometry data and a real-world magnetic gradiometry dataset from Geyer, Germany. Quantitative comparisons using metrics such as R^2 scores and Structural Similarity Index Measure (SSIM) highlights the NF interpolator’s performance across all tensor components, a preservation of geological trends that are typically used during interpretation, and a reduction of common ~~artifacts~~ artefacts caused by line-to-line inconsistencies.

Furthermore, by incorporating stochastic random Fourier features, our model likely opens the possibility to quantify uncertainty. Our analysis reveals heteroscedastic behaviour in the interpolations, and also highlights regions that require further data acquisition or refinement. Additionally, our approach seamlessly integrates vector and scalar field reconstructions through automatic differentiation, simplifying subsequent geophysical analyses and interpretations.

Overall, we argue that the proposed ~~NF~~-neural field interpolation method represents a significant advancement in processing tensor gradiometry data. Future developments should focus on larger scale applications, better understanding uncertainty of the model predictions, extended vertical interpolation capabilities (e.g., up- and downward continuations), and the integration of this approach into broader geophysical inversion and interpretation frameworks.

Acknowledgements. The authors gratefully acknowledge the Federal Institute for Geosciences and Natural Resources (BGR) and Supracon AG for providing the airborne full-tensor magnetic gradiometry dataset from Geyer, acquired in 2016 as part of the E³ (ErzExploration Erzgebirge) project. This research was supported by funding from the European Union's HORIZON Europe Research Council and UK Research and Innovation (UKRI) under grant agreement No. 101058483 (VECTOR). The authors ~~also~~ thank Ítalo Gonçalves, David Nathan, and one anonymous reviewer for their constructive comments which helped tremendously to improve the manuscript. AVK would also like to thank Vinit Gupta, Ralf Hielscher, and Parth Naik for their valuable insights and ~~constructive~~-delightful discussions, which significantly contributed to improving the clarity and depth of the manuscript.

Code and data availability. The code and datasets used in the study can be found at doi.org/10.5281/zenodo.16925858 (Kamath, 2025); Last access: 25 August, 2025.

Author contributions. AVK: Conceptualisation, Formal analysis, Investigation, Methodology, Visualisation, Writing – original draft; STT: Conceptualisation, Methodology, Writing – original draft, review & editing; HU: Writing – review & editing; BM: Writing – review & editing; RTD: Discussion, Writing - review & editing; MK: Writing – review & editing; RG: Writing – review & editing.

Competing interests. The authors declare that they have no conflict of interest.

References

- Axler, S., Bourdon, P., and Ramey, W.: Harmonic Function Theory, vol. 137 of *Graduate Texts in Mathematics*, Springer New York, New York, NY, ISBN 978-1-4419-2911-2 978-1-4757-8137-3, <https://doi.org/10.1007/978-1-4757-8137-3>, 2001.
- 515 Blakely, R. J.: Potential Theory in Gravity and Magnetic Applications, Cambridge University Press, 1 edn., ISBN 978-0-521-41508-8 978-0-521-57547-8 978-0-511-54981-6, <https://doi.org/10.1017/CBO9780511549816>, 1995.
- Bottou, L.: Stochastic Gradient Descent Tricks, p. 421–436, Springer Berlin Heidelberg, ISBN 9783642352898, https://doi.org/10.1007/978-3-642-35289-8_25, 2012.
- Bracewell, R. and Kahn, P. B.: The Fourier Transform and Its Applications, American Journal of Physics, 34, 712–712, <https://doi.org/10.1119/1.1973431>, 1966.
- 520 Brewster, J.: Description and evaluation of a full tensor interpolation method, in: SEG Technical Program Expanded Abstracts 2011, pp. 811–814, Society of Exploration Geophysicists, <https://doi.org/10.1190/1.3628199>, 2011.
- Burisch, M., Gerdes, A., Meinert, L. D., Albert, R., Seifert, T., and Gutzmer, J.: The essence of time – fertile skarn formation in the Variscan Orogenic Belt, Earth and Planetary Science Letters, 519, 165–170, <https://doi.org/10.1016/j.epsl.2019.05.015>, 2019.
- 525 Calafã, M., Hovad, E., Engsig-Karup, A. P., and Andriollo, T.: Physics-Informed Holomorphic Neural Networks (PIHNNs): Solving 2D linear elasticity problems, Computer Methods in Applied Mechanics and Engineering, 432, 117 406, <https://doi.org/10.1016/j.cma.2024.117406>, 2024.
- Chen, Z., Badrinarayanan, V., Lee, C.-Y., and Rabinovich, A.: GradNorm: Gradient Normalization for Adaptive Loss Balancing in Deep Multitask Networks, <https://doi.org/10.48550/ARXIV.1711.02257>, 2017.
- 530 Cockett, R., Kang, S., Heagy, L. J., Pidlisecky, A., and Oldenburg, D. W.: SimPEG: An open source framework for simulation and gradient based parameter estimation in geophysical applications, Computers & Geosciences, 85, 142–154, <https://doi.org/10.1016/j.cageo.2015.09.015>, 2015.
- Cordell, L.: A scattered equivalent-source method for interpolation and gridding of potential-field data in three dimensions, GEOPHYSICS, 57, 629–636, <https://doi.org/10.1190/1.1443275>, 1992.
- 535 Dampney, C. N. G.: The equivalent source technique, GEOPHYSICS, 34, 39–53, <https://doi.org/10.1190/1.1439996>, 1969.
- Davis, K. and Li, Y.: Fast solution of geophysical inversion using adaptive mesh, space-filling curves and wavelet compression, Geophysical Journal International, 185, 157–166, <https://doi.org/10.1111/j.1365-246x.2011.04929.x>, 2011.
- Evensen, G.: Data Assimilation: The Ensemble Kalman Filter, Springer, 2nd edn., 2009.
- Fatiando a Terra Project, Castro, Y. M., Esteban, F. D., Li, L., Oliveira Jr, V. C., Pesce, A., Shea, N., Soler, S. R., Souza-Junior, G. F., Tankersley, M., Uieda, L., and Uppal, I.: Harmonica v0.7.0: Forward modeling, inversion, and processing gravity and magnetic data, <https://doi.org/10.5281/ZENODO.3628741>, 2024.
- 540 Fichtner, A., Kennett, B. L. N., Igel, H., and Bunge, H.-P.: Full seismic waveform tomography for upper-mantle structure in the Australasian region using adjoint methods, Geophysical Journal International, 186, 1459–1473, publisher: Oxford University Press, 2011.
- Fitzgerald, D., Paterson, R., and Holstein, H.: Comparing Two Methods for Gridding and Honouring Gravity Gradient Tensor Data, Copenhagen, Denmark, <https://doi.org/10.3997/2214-4609.20148284>, 2012.
- 545 Ghosh, A., Gentile, A. A., Dagrada, M., Lee, C., Kim, S.-H., Cha, H., Choi, Y., Kim, B., Kye, J.-I., and Elfving, V. E.: Harmonic (Quantum) Neural Networks, <https://doi.org/10.48550/ARXIV.2212.07462>, 2022.

- Giraud, J., Seillé, H., Lindsay, M. D., Visser, G., Ogarko, V., and Jessell, M. W.: Utilisation of probabilistic magnetotelluric modelling to constrain magnetic data inversion: proof-of-concept and field application, *Solid Earth*, 14, 43–68, <https://doi.org/10.5194/se-14-43-2023>, 2023.
- Goodfellow, I., Bengio, Y., and Courville, A.: Deep learning, Adaptive computation and machine learning, The MIT press, Cambridge, Mass, ISBN 978-0-262-03561-3, 2016.
- Hamilton, W. R.: II. *On quaternions; or on a new system of imaginaries in algebra*, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 25, 10–13, <https://doi.org/10.1080/14786444408644923>, 1844.
- Hansen, R. O. and Miyazaki, Y.: Continuation of potential fields between arbitrary surfaces, *GEOPHYSICS*, 49, 787–795, <https://doi.org/10.1190/1.1441707>, 1984.
- Hewitson, B. C., Crane, R. G., and Tietze, W., eds.: Neural Nets: Applications in Geography, vol. 29 of *The GeoJournal Library*, Springer Netherlands, Dordrecht, ISBN 978-94-010-4490-5 978-94-011-1122-5, <https://doi.org/10.1007/978-94-011-1122-5>, 1994.
- Hillier, M., Wellmann, F., De Kemp, E. A., Brodaric, B., Schetselaar, E., and Bédard, K.: GeoINR 1.0: an implicit neural network approach to three-dimensional geological modelling, *Geoscientific Model Development*, 16, 6987–7012, <https://doi.org/10.5194/gmd-16-6987-2023>, 2023.
- Hillier, M. J., Schetselaar, E. M., De Kemp, E. A., and Perron, G.: Three-Dimensional Modelling of Geological Surfaces Using Generalized Interpolation with Radial Basis Functions, *Mathematical Geosciences*, 46, 931–953, <https://doi.org/10.1007/s11004-014-9540-3>, 2014.
- Kamath, A., Thiele, S., and Gloaguen, R.: (Auto) Differentiating geology: Geological modelling with random Fourier features and neural fields, <https://doi.org/10.5194/egusphere-egu25-9746>, 2025.
- Kamath, A. V.: k4m4th/tensorweave: tensorweave-1.0, <https://doi.org/10.5281/ZENODO.16941182>, 2025.
- Kendall, A., Gal, Y., and Cipolla, R.: Multi-Task Learning Using Uncertainty to Weigh Losses for Scene Geometry and Semantics, <https://doi.org/10.48550/ARXIV.1705.07115>, 2017.
- Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, <https://doi.org/10.48550/ARXIV.1412.6980>, 2014.
- Kroner, U. and Romer, R.: Two plates — Many subduction zones: The Variscan orogeny reconsidered, *Gondwana Research*, 24, 298–329, <https://doi.org/10.1016/j.gr.2013.03.001>, 2013.
- Lacava, F.: Separation of Variables in Laplace Equation, p. 137–177, Springer International Publishing, ISBN 9783031050992, https://doi.org/10.1007/978-3-031-05099-2_8, 2022.
- LeCun, Y., Bengio, Y., and Hinton, G.: Deep learning, *Nature*, 521, 436–444, <https://doi.org/10.1038/nature14539>, 2015.
- Lee, J. M.: Introduction to Smooth Manifolds, vol. 218 of *Graduate Texts in Mathematics*, Springer New York, New York, NY, ISBN 978-1-4419-9981-8 978-1-4419-9982-5, <https://doi.org/10.1007/978-1-4419-9982-5>, 2012.
- Lefebvre, M. G., Romer, R. L., Glodny, J., Kroner, U., and Roscher, M.: The Hämmerlein skarn-hosted polymetallic deposit and the Eibenstock granite associated greisen, western Erzgebirge, Germany: two phases of mineralization—two Sn sources, *Mineralium Deposita*, 54, 193–216, <https://doi.org/10.1007/s00126-018-0830-4>, 2019.
- Li, Y. and Oldenburg, D. W.: 3-D inversion of gravity data, *GEOPHYSICS*, 63, 109–119, <https://doi.org/10.1190/1.1444302>, 1998.
- Li, Y. and Oldenburg, D. W.: Rapid construction of equivalent sources using wavelets, *GEOPHYSICS*, 75, L51–L59, <https://doi.org/10.1190/1.3378764>, 2010.
- Liu, S., Johns, E., and Davison, A. J.: End-To-End Multi-Task Learning With Attention, in: 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), p. 1871–1880, IEEE, <https://doi.org/10.1109/cvpr.2019.00197>, 2019.

- 585 Mantoglou, A. and Wilson, J. L.: The Turning Bands Method for simulation of random fields using line generation by a spectral method, *Water Resources Research*, 18, 1379–1394, <https://doi.org/10.1029/WR018i005p01379>, 1982.
- Margossian, C. C.: A Review of automatic differentiation and its efficient implementation, <https://doi.org/10.1002/WIDM.1305>, arXiv:1811.05031 [cs], 2019.
- Markley, F. L., Cheng, Y., Crassidis, J. L., and Oshman, Y.: Averaging Quaternions, *Journal of Guidance, Control, and Dynamics*, 30,
590 1193–1197, <https://doi.org/10.2514/1.28949>, 2007.
- Mendonça, C. A. and Silva, J. B. C.: The equivalent data concept applied to the interpolation of potential field data, *GEOPHYSICS*, 59, 722–732, <https://doi.org/10.1190/1.1443630>, 1994.
- Mickus, K. L. and Hinojosa, J. H.: The complete gravity gradient tensor derived from the vertical component of gravity: a Fourier transform technique, *Journal of Applied Geophysics*, 46, 159–174, [https://doi.org/10.1016/S0926-9851\(01\)00031-3](https://doi.org/10.1016/S0926-9851(01)00031-3), 2001.
- 595 Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., and Ng, R.: NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis, in: *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 405–421, Springer, 2020.
- Misra, D.: Mish: A Self Regularized Non-Monotonic Activation Function, <https://doi.org/10.48550/ARXIV.1908.08681>, 2019.
- Naprstek, T. and Smith, R. S.: A new method for interpolating linear features in aeromagnetic data, *GEOPHYSICS*, 84, JM15–JM24, <https://doi.org/10.1190/geo2018-0156.1>, 2019.
- 600 Oliveira Junior, V. C., Takahashi, D., Reis, A. L. A., and Barbosa, V. C. F.: Computational aspects of the equivalent-layer technique: review, *Frontiers in Earth Science*, 11, <https://doi.org/10.3389/feart.2023.1253148>, 2023.
- Openshaw, S.: Modelling spatial interaction using a neural net, in: *Geographic Information Systems, Spatial Modelling and Policy Evaluation*, edited by Fischer, M. M. and Nijkamp, P., pp. 147–164, Springer Berlin Heidelberg, Berlin, Heidelberg, ISBN 978-3-642-77502-4 978-3-642-77500-0, https://doi.org/10.1007/978-3-642-77500-0_10, 1993.
- 605 Parker, R. L.: The Rapid Calculation of Potential Anomalies, *Geophysical Journal International*, 31, 447–455, <https://doi.org/10.1111/j.1365-246x.1973.tb06513.x>, 1973.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Köpf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S.: PyTorch: An Imperative Style, High-Performance Deep Learning Library, <https://doi.org/10.48550/ARXIV.1912.01703>, version Number: 1, 2019.
- 610 Piaulino, L. S., Oliveira Junior, V. C., and Barbosa, V. C. F.: A Scalable and Consistent Method for Multi-Component Gravity-Gradient Data Processing, *Applied Sciences*, 15, 8396, <https://doi.org/10.3390/app15158396>, 2025.
- Prechelt, L.: Early Stopping - But When?, p. 55–69, Springer Berlin Heidelberg, ISBN 9783540494300, https://doi.org/10.1007/3-540-49430-8_3, 1998.
- Rahaman, N., Baratin, A., Arpit, D., Draxler, F., Lin, M., Hamprecht, F. A., Bengio, Y., and Courville, A.: On the Spectral Bias of Neural
615 Networks, in: *Proceedings of the 36th International Conference on Machine Learning (ICML)*, vol. 97 of *Proceedings of Machine Learning Research*, pp. 5301–5310, 2019.
- Rahimi, A. and Recht, B.: Random features for large-scale kernel machines, in: *Proceedings of the 21st International Conference on Neural Information Processing Systems, NIPS’07*, p. 1177–1184, Curran Associates Inc., Red Hook, NY, USA, ISBN 9781605603520, 2007.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E.: Physics-informed neural networks: A deep learning framework for solving
620 forward and inverse problems involving nonlinear partial differential equations, *Journal of Computational Physics*, 378, 686–707, <https://doi.org/10.1016/j.jcp.2018.10.045>, 2019.
- Ramachandran, P., Zoph, B., and Le, Q. V.: Searching for Activation Functions, <https://doi.org/10.48550/ARXIV.1710.05941>, 2017.

- Rudd, J., Chubak, G., Larnier, H., Stolz, R., Schiffler, M., Zakosarenko, V., Schneider, M., Schulz, M., and Meyer, M.: Commercial operation of a SQUID-based airborne magnetic gradiometer, *The Leading Edge*, 41, 486–492, <https://doi.org/10.1190/tle41070486.1>, 2022.
- 625 Saragadam, V., LeJeune, D., Tan, J., Balakrishnan, G., Veeraraghavan, A., and Baraniuk, R. G.: WIRE: Wavelet Implicit Neural Representations, <https://doi.org/10.48550/ARXIV.2301.05187>, 2023.
- Satheesh, A., Schmidt, C. P., Wall, W. A., and Meier, C.: Structure-preserving invariant interpolation schemes for invertible second-order tensors, *International Journal for Numerical Methods in Engineering*, 125, <https://doi.org/10.1002/nme.7373>, 2023.
- Shoemake, K.: Animating rotation with quaternion curves, *ACM SIGGRAPH Computer Graphics*, 19, 245–254, <https://doi.org/10.1145/325165.325242>, 1985.
- 630 Siqueira, F. C. L., Oliveira Jr., V. C., and Barbosa, V. C. F.: Fast iterative equivalent-layer technique for gravity data processing: A method grounded on excess mass constraint, *GEOPHYSICS*, 82, G57–G69, <https://doi.org/10.1190/geo2016-0332.1>, 2017.
- Sitzmann, V., Martel, J. N. P., Bergman, A. W., Lindell, D. B., and Wetzstein, G.: Implicit Neural Representations with Periodic Activation Functions, *CoRR*, abs/2006.09661, <https://arxiv.org/abs/2006.09661>, 2020.
- 635 Smith, L. T., Horrocks, T., Akhtar, N., Holden, E.-J., and Wedge, D.: Implicit neural representation for potential field geophysics, *Scientific Reports*, 15, 9799, <https://doi.org/10.1038/s41598-024-83979-z>, 2025.
- Soler, S. R. and Uieda, L.: Gradient-boosted equivalent sources, *Geophysical Journal International*, 227, 1768–1783, <https://doi.org/10.1093/gji/ggab295>, 2021.
- Stolz, R., Schmelz, M., Zakosarenko, V., Foley, C. P., Tanabe, K., Xie, X., and Fagaly, R.: Superconducting sensors and methods in geophysical applications, *Superconductor Science and Technology*, <https://doi.org/10.1088/1361-6668/abd7ce>, 2021.
- 640 Tancik, M., Srinivasan, P. P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J. T., and Ng, R.: Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains, <https://doi.org/10.48550/ARXIV.2006.10739>, 2020.
- Thiele, S., Kamath, A., and Gloaguen, R.: Less is more: Weakly supervised interpolation using geological neural fields, <https://doi.org/10.5194/egusphere-egu25-8735>, 2025.
- 645 Tveit, S., Mannseth, T., Park, J., Sauvin, G., and Agersborg, R.: Combining CSEM or gravity inversion with seismic AVO inversion, with application to monitoring of large-scale CO₂ injection, *Computational Geosciences*, 24, 1201–1220, <https://doi.org/10.1007/s10596-020-09934-9>, 2020.
- Ugalde, H., Morris, B., Kamath, A., and Parsons, B.: Full-tensor magnetic gradiometry: Comparison with scalar total magnetic intensity, processing and visualization guidelines, *Geophysical Prospecting*, 73, 303–314, <https://doi.org/10.1111/1365-2478.13629>, 2024.
- 650 Vassiliou, A.: Numerical techniques for processing airborne gradiometer data, <https://doi.org/10.11575/PRISM/23539>, 1986.
- White, K. B., Cline, D., and Egbert, P. K.: Poisson Disk Point Sets by Hierarchical Dart Throwing, in: *2007 IEEE Symposium on Interactive Ray Tracing*, p. 129–132, IEEE, <https://doi.org/10.1109/rt.2007.4342600>, 2007.
- Wittwer, T.: Regional gravity field modeling with radial basis functions, vol. 72 of *Publications on Geodesy*, Nederlandse Commissie voor Geodesie, ISBN 978-90-6132-315-0, <https://doi.org/10.54419/hboxky>, 2009.
- 655 Xia, J. and Sprowl, D. R.: Correction of topographic distortion in gravity data, *GEOPHYSICS*, 56, 537–541, <https://doi.org/10.1190/1.1443070>, 1991.
- Xie, Y., Takikawa, T., Saito, S., Litany, O., Yan, S., Khan, N., Tombari, F., Tompkin, J., Sitzmann, V., and Sridhar, S.: Neural Fields in Visual Computing and Beyond, *Computer Graphics Forum*, 41, 641–676, <https://doi.org/10.1111/cgf.14505>, 2022.

660 Zuo, B. and Hu, X.: Edge detection of gravity field using eigenvalue analysis of gravity gradient tensor, *Journal of Applied Geophysics*, 114, 263–270, <https://doi.org/10.1016/j.jappgeo.2015.01.013>, 2015.

Figures




Figure 1. Synthetic subsurface model and corresponding gravity gradiometry data. (A) Horizontal cross-section of the synthetic geological model at a depth of 140m, with high-resolution observation points shown as black dots. (B-F) The five independent components of the gravity gradiometry tensor generated via forward modelling using SimPEG are also shown (b-f). Each panel displays both the high-resolution dataset (grey scale; cell size of 25m) and the low-resolution dataset (colour; 200m cross-line spacing and 15m inline spacing) for the corresponding tensor component.




Figure 2. Neural field-Fourier Field model architecture. The orange-blue block projects the input position vector into a feature space and passes it through the fully-connected layers (orange block) to acquire the scalar potential. The red arrows signify the use of automatic differentiation to acquire the first (gradient) and second (curvature) spatial derivatives of the potential.




Figure 3. Quantitative comparison of interpolation performance for the synthetic dataset. (A–B) Spatial distribution of residuals between the true and predicted H_{xy} tensor component using (Aa) the Truncated Radial Basis Function (RBF) method and (Bb) the Ensemble Neural Field (ENF) approach (with 25 models in the ensemble). Insets show 1:1 parity kernel density estimate-plots comparing predicted and true values. (Cc) R^2 scores for each tensor component (H_{xx} , H_{xy} , H_{xz} , H_{yy} , H_{yz} , H_{zz}) across three interpolation methods: RBF, the mean of the individual Neural Field (NF) scores from the models within the ensemble, and ENF. The ENF and NF models consistently achieve higher accuracy across all components, while RBF exhibits reduced performance, particularly. The loss curves (d) for off-diagonal terms various components of the loss show the data fitting parts (solid lines) plateau while the Laplacian part (dotted line) keeps decreasing owing to increased sampling with each progressive epoch.




Figure 4. Comparison of gravity gradiometry tensor components derived from two interpolation methods applied to the synthetic dataset. (A–C) Gridded, histogram-equalised The ground truth H_{xx} , H_{xy} , and H_{xz} components generated using a (a–c) are compared with the results from the Truncated Radial Basis Function (RBF) (d–f) interpolation with 250 nearest neighbours and a smoothing factor of 100. (D–F) Corresponding 100, and corresponding results produced by the Ensemble Neural Field (ENF) method (g–i) with 25 models in the ensemble. All values range from 0 to 1. Black lines in panel-interpolated results (Ed–i) indicate the input flight lines used for interpolation.




Figure 5. Accuracy metrics as a function of increasing line spacing for the synthetic dataset. (A) R^2 Score and (B) Structural Similarity Index Measure (SSIM) (C) and Mean Squared Error (MSE) (c) were computed between the ground truth and the gridded results from the interpolation methods. The Radial Basis Function (RBF) used 250 nearest neighbours, with a smoothing factor of 100, and the Neural Field (NF) model used the same architecture as discussed in Section 3.3. The full tensor interpolation algorithm from Fitzgerald et al. (2012) (QUAT) (QUAT; Fitzgerald et al., 2012) was also included for comparison, using the aforementioned RBF for the eigenvalue interpolation, and SLERP for rotational interpolation. The shaded regions show the minimum and maximum metric across all the components, and the plotted line shows the root-mean-squared metric computed across the components.




Figure 6. Case study site near Geyer, Erzgebirge, Germany. (A) Flight lines from a subset of the airborne magnetic gradiometry survey (a), with every fourth line (red) used as input for interpolation and the remaining lines (black) reserved for validation. (Bb) Spatial distribution of the measured zz -component of the magnetic gradiometry tensor across the survey region.




Figure 7. Quantitative comparison of interpolation performance for the Geyer dataset. **(A–B)** Spatial distribution of residuals between the true and predicted H_{xy} tensor component along the test flight lines using **(A)** the Truncated Radial Basis Function (RBF) **method and** **(Ba)** **method and** the Ensemble Neural Field (ENF) **(b)** **approach (with 50 models in the ensemble)**. Insets show 1:1 parity kernel density estimate plots comparing predicted and true values. **(Cc)** R^2 scores for each tensor component (H_{xx} , H_{xy} , H_{xz} , H_{yy} , H_{yz} , H_{zz}) across three interpolation methods: RBF, mean of the individual Neural Field (NF) scores from the models within the ensemble, and ENF. The ENF and NF models consistently achieve higher scores across all components, while RBF exhibits reduced performance, **particularly**. **The loss curves (d)** for **off-diagonal terms** **various components of the loss show similar characteristics to the synthetic loss curve.**




Figure 8. Comparison of magnetic gradient tensor components interpolated onto a uniform grid (cell size $=25\text{m} \times 20\text{m}$) using two methods. (A–C) Gridded, histogram-equalised H_{xx} , H_{xy} , and H_{xz} components obtained using the Truncated Radial Basis Function (RBF) interpolation method, with 250 nearest neighbours and a smoothing factor of 100 for all of the flight lines. (D–F) Corresponding are used as the ground truth. We compare the ground truth with the corresponding components interpolated with RBF using every fourth flight line. (G–I) Corresponding, and the corresponding components interpolated using the Ensemble Neural Field (ENF) approach (g–i) with 50 models in the ensemble. Each column visualises a distinct tensor-component of the tensor. All values range from -0.8 to 0.8. Black lines in panels B, E and H within the plots indicate the locations of the input flight lines used in the interpolation process.




Figure 9. Comparison of recovered vector magnetic field components from two interpolation methods, evaluated against a high-resolution reference model. (A–C) Reference vector components b_x , b_y , and b_z computed using Fourier domain transfer functions applied to magnetic tensor components gridded via the Truncated Radial Basis Function (RBF) from all available flight lines. (D–F) Reconstructed as our reference, Fourier domain reconstruction of the vector components obtained using the RBF method on tensor data from the training set of flight lines. (G–I) Corresponding, and the corresponding results computed from the spatial derivatives of the scalar field predicted by the Ensemble Neural Field (ENF) model (g–i) are shown. The black lines in each panel represent the flight lines used to generate the corresponding component. Each panel shows the histogram-equalised spatial distribution of the respective vector component across the subset of the Geyer survey area, mapped from 0 to 1.




Figure 10. Uncertainty maps for the ~~100-model~~ 50-model ensemble. The standard deviation computed across ~~100~~ 50 models for the ~~A)~~ H_{xx} ~~B)~~ H_{xy} , ~~C)~~ H_{xz} , ~~D)~~ H_{yy} , ~~E)~~ H_{yz} , ~~F)~~ H_{zz} (f) tensor components, and the recovered ~~G)~~ components b_x , b_y (h), and b_z (i) vector components magnetic field.