
PHYSICS AWARE NEURAL NETWORKS : DENOISING FOR MAGNETIC NAVIGATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Magnetic-anomaly navigation, which leverages small-scale local variations in the Earth’s magnetic field, has emerged as a promising alternative for environments where GPS signals are unavailable or compromised. Airborne systems face a fundamental challenge in extracting the necessary geomagnetic field data: the aircraft itself induces magnetic noise. Although the classical Tolles-Lawson model addresses this, it inadequately handles stochastically corrupted magnetic data needed for operational navigation. To handle stochastic noise, we propose a novel approach using two physics-based constraints: divergence-free vector fields and $E(3)$ -equivariance. Our constraints guarantee that the generated magnetic field obeys Maxwell’s equations and ensure that output changes appropriately with sensor position/orientation. The divergence-free constraint is implemented by defining a neural network outputting vector potential A , with the magnetic field constructed as its curl. For $E(3)$ -equivariance, we use tensor products of geometric tensors that are representable using spherical harmonics with well-known rotational transformations. By enforcing physical consistency and constraining the space of admissible field functions, our formulation acts as an implicit regularizer, improving spatio-temporal performance. We conduct ablation studies evaluating these constraints’ individual and combined effects across CNNs, MLPs, Liquid Time Constant Models, and Contiformers. We note that continuous-time dynamics and long-term memory are critical for modelling magnetic time-series data; the Contiformer architecture, which inherently possesses both, surpasses state-of-the-art methods in our experiments. To handle data scarcity, we develop synthetic datasets by utilising the World Magnetic Model (WMM) in conjunction with time-series conditional GANs, generating realistic and temporally consistent magnetic field sequences spanning various trajectory patterns and environmental scenarios. Our experiments demonstrate that embedding these constraints significantly improves predictive accuracy and physical plausibility, outperforming the state-of-the-art across both classical and unconstrained deep learning approaches. Our code is available at <https://github.com/darksideofmoon123456/summerhouse>.

1 INTRODUCTION

Satellite based navigation systems, commonly referred to as Global Navigation Satellite System (GNSS), is now ubiquitous, underpinning activities from commerce to defense. However, today there is an increased threat of jamming and spoofing of GNSS signals, affecting civilian and military aircraft alike. Between Aug 21 - June 24, over 580,000 GNSS signal loss events on aircraft were reported IATA (2024). Fallback systems such as inertial sensors are limited due to sensor drift which increases over time Woodman (2007) Groves (2015).

A promising alternative is *magnetic-anomaly navigation*, which utilizes small-scale local anomalies in Earth’s magnetic field (which have well-defined geographical signatures) to overcome the limitations of traditional navigation systems Canciani & Raquet (2020). Simply put, with a known background magnetic field, sufficiently precise local measurements can be used for geolocation and navigation. However, magnetic-anomaly navigation is limited because the magnetic field sensors onboard the aircraft measure magnetic fields produced by the aircraft’s chassis, engines and instruments, which are essentially noise in the measurement Gnadt (2022a) Nerrise et al. (2023).

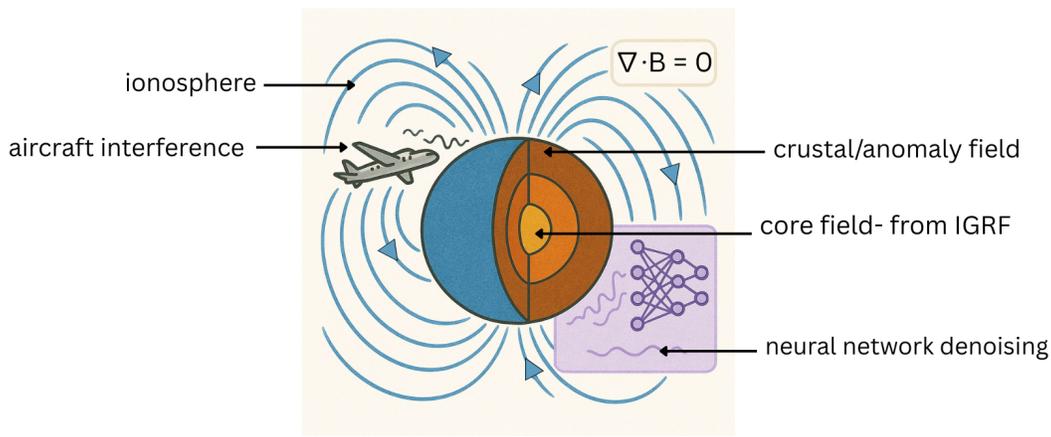


Figure 1: An illustrative diagram of our problem setting.

This limitation also applies to platforms other than aircraft, such as ships, unstaffed aerial vehicles (UAVs), etc. Extracting the real geomagnetic field data (i.e., the field that would have been measured if the aircraft did not exist) from onboard sensor measurements containing these superimposed magnetic components is difficult, especially for stochastic noise. Without denoising, navigation becomes error-prone and impractical.

Prior work in the area (detailed in the next section) has struggled to handle such noise, even with deep-learning based approaches. We aim to approach the problem from its fundamentals: Given the nature of the physical phenomenon we are measuring, what properties should an ideal network have to solve it?

Contributions

1. We introduce two physics-aware constraints for magnetic navigation: (i) divergence-free vector fields and (ii) E(3)-equivariance. These ensure predictions obey Maxwell’s second law of electromagnetism ($\nabla \cdot \mathbf{B} = 0$) and respond correctly to changes in position and orientation. By enforcing physical consistency, we implicitly regularize the model, improving spatiotemporal performance.
2. We identify key properties for magnetic navigation—continuous-time dynamics and long-term memory suitable for spatiotemporal settings—and find appropriate architectures. Our ablation studies show how each constraint improves performance alone and in combination.
3. To address data scarcity, we generate synthetic magnetic field sequences using the World Magnetic Model and conditional GANs, covering diverse trajectories and environments.
4. We design, implement, and evaluate a novel system for denoising magnetic field sensor data for navigation, outperforming the state-of-the-art across both classical and unconstrained deep learning approaches.

2 RELATED WORK AND PRELIMINARIES

2.1 FILTERING BASED APPROACHES

Principal Component Analysis (PCA) decomposes flight lines or spatio-temporal data into principal components, isolating dominant noise trends - e.g., slowly varying diurnal baselines. Qiao et al. (2022) proposed an RPCA-based aeromagnetic compensation method, which improved outlier robustness and accuracy over traditional models. Simple moving-average filters better adapt to smooth, time-varying changes. Several finite impulse response (FIR) moving-average filters and an exponential moving average (IIR) for smoothing magnetometer signals in an unmanned ground vehicle were compared Pereira et al. (2024). Hybrid approaches such as wavelet–Wiener algorithms

also improve noise reduction, while adaptive Kalman filters can struggle with nonlinear coupling of aircraft interference Canciani (2022).

2.2 TOLLES-LAWSON MODEL

The linear aeromagnetic compensation model for scalar magnetometers was introduced by Tolles & Lawson (1950), with patents for associated hardware Tolles (1954; 1955). Enhancements using sinusoidal flight maneuvers improved observability of model coefficients Leliak (1961). Various extensions refined the model, but stochastic noise remains inadequately handled Leach (1979); Gu et al. (2013); Webb & White (2014); Wu et al. (2018); Gnatd et al. (2022). Shallow, feature-informed networks can improve performance while maintaining interpretability Gnatd (2022b).

2.3 DEEP-LEARNING BASED APPROACHES

Introduced neural networks for compensating aircraft magnetic interference, estimating the complete magnetic field as a combination of crustal, diurnal, and platform-generated components. Modern deep learning methods, including DeepMAD Xu et al. (2020), CNNs, recurrent networks, and Liquid Time-Constant (LTC) networks, capture nonlinear interference patterns more effectively than classical methods. Recently, there has been concurrent work on using similar physics-aware approaches to constraining for finding non-singularities in Navier-Stokes equations. Wang et al. (2025)

3 OUR APPROACH

We present a first-principles approach to magnetic field denoising that enforces physical laws and handles real-world magnetometer noise. First, we integrate physics-aware constraints directly into neural networks, for extraction of geomagnetic signals while preserving physical consistency. Second, we identify two key physical properties of magnetic fields essential for effective network design.

Magnetic fields in free space must satisfy (1) Maxwell’s equations, specifically: divergence-free constraint ($\nabla \cdot \mathbf{B} = 0$), which is depicted in Figure 2 and (2) E(3)-equivariance under rigid transformations. By embedding these constraints directly into the network architecture rather than as soft penalties, we constrain the space of admissible field functions to physically plausible solutions. One of the fundamental properties of magnetic data is that they are irregularly sampled, with measurement times t_i separated by $\Delta t_i = t_{i+1} - t_i$. Therefore, to correctly handle data, we model the hidden state $z(t)$ via a neural ODE $\dot{z}(t) = f_\theta(z(t), t)$ which naturally evolves in continuous time across each arbitrary interval Δt_i . Magnetometer noise exhibits long-term dependencies, as its auto-correlation $R(\tau) = \mathbb{E}[n(t)n(t+\tau)]$ decays slowly (e.g., $R(\tau) \sim e^{-\lambda|\tau|}$ for small λ), necessitating long-term memory for modeling and denoising scenarios. Our entire denoising pipeline is given in Figure 4.

3.1 DIVERGENCE-FREE CONSTRAINT

We parametrize \mathbf{B} via a vector potential \mathbf{A} such that $\mathbf{B} = \nabla \times \mathbf{A}$. We do this by defining a neural network $\mathcal{A}_\theta : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ that outputs the vector potential $\mathbf{A}_\theta(\mathbf{x})$. The magnetic field is then constructed as the curl of this potential:

$$\mathbf{B}_\theta(\mathbf{x}) := \nabla \times \mathbf{A}_\theta(\mathbf{x}) \tag{1}$$

This guarantees that $\nabla \cdot \mathbf{B}_\theta = 0$ identically, due to the vector calculus identity $\nabla \cdot (\nabla \times \mathbf{A}) = 0$ for all smooth vector fields \mathbf{A} . In 3D Euclidean space \mathbb{R}^3 , the magnetic field can be treated as a differential 2-form $\mathbf{B} \in \Omega^2(\mathbb{R}^3)$, and the vector potential as a 1-form $\mathbf{A} \in \Omega^1(\mathbb{R}^3)$. The curl operation corresponds to applying the exterior derivative d followed by the Hodge star operator \star :

$$\mathbf{B} = \star d\mathbf{A} \tag{2}$$

This geometric formulation further shows why $\nabla \cdot \mathbf{B} = 0$. In differential forms, the divergence of a vector field (as a 1-form) is:

$$\nabla \cdot \mathbf{B} = \star d \star \mathbf{B} \quad (3)$$

Since $\mathbf{B} = \star d\mathbf{A}$, we have:

$$\nabla \cdot \mathbf{B} = \star d \star (\star d\mathbf{A}) = \star dd\mathbf{A} = 0 \quad (4)$$

The last equality holds because $d^2 = 0$. This shows that our approach is compatible with differential forms used in recent works on neural PDEs and conservation laws Richter-Powell et al. (2022), while allowing direct incorporation of physical priors into model design.

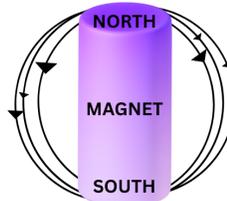
The network \mathcal{A}_θ is implemented using a standard feedforward or equivariant architecture. During inference and training, \mathbf{B}_θ is computed using automatic differentiation libraries to calculate the spatial derivatives needed for the curl:

$$[\nabla \times \mathbf{A}_\theta]_i = \epsilon_{ijk} \frac{\partial \mathcal{A}_\theta^k}{\partial x^j}, \quad (5)$$

where ϵ_{ijk} is the Levi-Civita symbol, and Einstein summation is implied. The vector potential representation introduces gauge freedom: for any scalar function ϕ , both \mathbf{A} and $\mathbf{A} + \nabla\phi$ yield the same magnetic field. This non-uniqueness might seem problematic, but it provides additional flexibility for optimization. Standard gradient descent naturally selects smooth solutions within the gauge equivalence class.

Total magnetic field passing through a given area or volume

Magnetic flux forms continuous closed loops. Flowing from the north pole to the south pole externally. Returning through the magnet internally.



The north and south poles of a magnet are inseparable, and thus no isolated magnetic charge can occur within the field.

Figure 2: Maxwell’s 2nd Law of Electromagnetism. Magnetic field lines have no beginning or end, but exist as loops or extend to infinity: magnetic monopoles cannot exist.

3.2 E(3)-EQUIVARIANCE

Vehicle-mounted magnetometers measure fields in their local reference frame, which changes with orientation. So, a physically consistent model must transform under rotations and translations of the sensor platform. Formally, if $g = (R, \mathbf{t}) \in E(3)$ represents a rigid transformation, then:

$$\mathbf{B}'(R \cdot \mathbf{r} + \mathbf{t}) = R \cdot \mathbf{B}(\mathbf{r}) \quad (6)$$

Or, for a neural network: Let G be a group with group elements $g \in G$. For transformations $\Psi_g : X \rightarrow X$ on input domain X , a mapping $\varphi : X \rightarrow Y$ satisfies G -equivariance when there exist transformations $\Phi_g : Y \rightarrow Y$ on output domain Y such that Satorras et al. (2021):

$$\varphi(\Psi_g(x)) = \Phi_g(\varphi(x)) \quad (7)$$

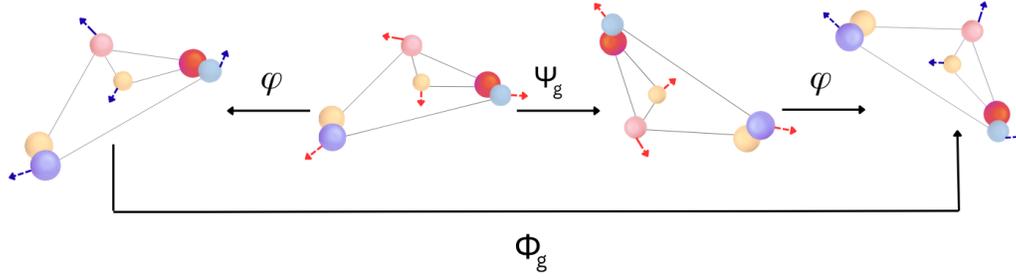


Figure 3: E(3)-equivariance demonstration: applying geometric transformations before or after the neural network φ yields equivalent results, ensuring consistent magnetic field predictions regardless of coordinate system orientation.

(See Figure 3). This ensures that our predictions remain consistent regardless of sensor pose, since orientation varies continuously during navigation. We implement E(3)-equivariance using geometric tensors, which can be represented using spherical harmonics:

$$\mathbf{T}^{(\ell)}(\mathbf{r}) = T(r)\mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}) \quad (8)$$

where $\mathbf{Y}^{(\ell)} = (Y_{\ell}^{-\ell}, Y_{\ell}^{-\ell+1}, \dots, Y_{\ell}^{\ell})^T$ and $T(r)$ is a radial function. This works because spherical harmonics are precisely the basis functions for irreducible representations (irreps) of $SO(3)$. A standard layer is given by: $\mathbf{y} = \sigma(W\mathbf{x} + \mathbf{b})$ – but this is not equivariant! If we rotate the input, the output does not rotate consistently. We can create a general E(3)-equivariant layer in a neural network through the following linear operation between geometric tensors:

$$\mathbf{T}_{\text{out}}^{(\ell_{\text{out}})} = \sum_{\ell_{\text{in}}} \sum_{\ell} W^{(\ell_{\text{out}}, \ell_{\text{in}}, \ell)} [\mathbf{T}_{\text{in}}^{(\ell_{\text{in}})} \otimes \mathbf{Y}^{(\ell)}]^{(\ell_{\text{out}})} \quad (9)$$

Input: $\mathbf{T}_{\text{in}}^{(\ell_{\text{in}})}$ is a geometric tensor of type (ℓ_{in})

Spherical Harmonics: $\mathbf{Y}^{(\ell)}$ provides the geometric information about relative positions

Tensor Product: $[\mathbf{T}_{\text{in}}^{(\ell_{\text{in}})} \otimes \mathbf{Y}^{(\ell)}]^{(\ell_{\text{out}})}$ combines them using Clebsch-Gordan coefficients

Weights: $W^{(\ell_{\text{out}}, \ell_{\text{in}}, \ell)}$ are scalar weights (learnable parameters)

The tensor product $[\mathbf{T}^{(\ell_1)} \otimes \mathbf{Y}^{(\ell_2)}]^{(\ell)}$ is computed as:

$$[\mathbf{T}^{(\ell_1)} \otimes \mathbf{Y}^{(\ell_2)}]_m^{(\ell)} = \sum_{m_1=-\ell_1}^{\ell_1} \sum_{m_2=-\ell_2}^{\ell_2} \langle \ell_1 m_1 \ell_2 m_2 | \ell m \rangle T_{m_1}^{(\ell_1)} Y_{\ell_2}^{m_2} \quad (10)$$

For a relative position vector $\mathbf{r}_{ij} = \mathbf{r}_j - \mathbf{r}_i$ we have $Y_{\ell}^m(\mathbf{r}_{ij}) = Y_{\ell}^m(\theta_{ij}, \phi_{ij})$, where (θ_{ij}, ϕ_{ij}) are the spherical angles of $\hat{\mathbf{r}}_{ij} = \mathbf{r}_{ij}/|\mathbf{r}_{ij}|$. For a node i with neighbors $\mathcal{N}(i)$:

$$\mathbf{T}_i^{(\ell_{\text{out}})} = \sum_{j \in \mathcal{N}(i)} \sum_{\ell_{\text{in}}} \sum_{\ell} W^{(\ell_{\text{out}}, \ell_{\text{in}}, \ell)} \times [\mathbf{T}_j^{(\ell_{\text{in}})} \otimes \mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}_{ij})]^{(\ell_{\text{out}})} \quad (11)$$

See Appendix A for proof of E(3)-equivariance. We modify this for the LTC and Contiformer architectures, as they operate on continuous-time dynamics governed by Neural ODEs. For the Contiformer, we embed the tensor product in the attention integral:

$$\alpha(\mathbf{r}, t; \mathbf{r}_i, t_i) = \frac{1}{t - t_i} \int_{t_i}^t \sum_{\ell_q, \ell_k, \ell} \mathbf{Q}^{(\ell_q)}(\mathbf{r}, \tau) \cdot \left[\mathbf{K}^{(\ell_k)}(\mathbf{r}_i, \tau) \otimes \mathbf{Y}^{(\ell)}(\widehat{\mathbf{r} - \mathbf{r}_i}) \right]^{(\ell_q)} d\tau \quad (12)$$

270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323

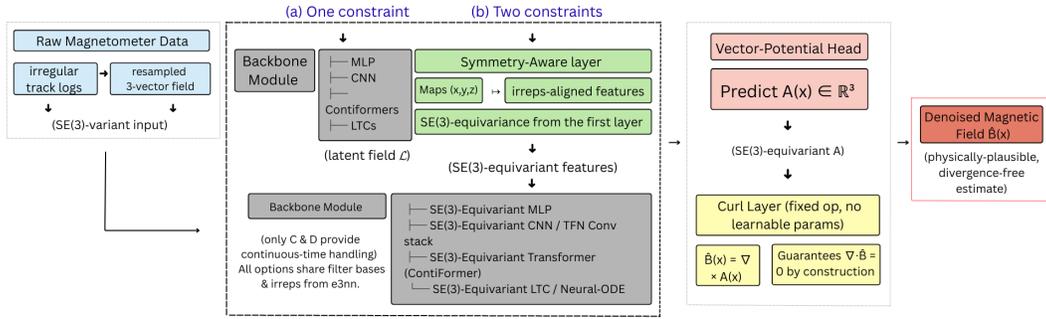


Figure 4: Denoising Pipeline with Divergence-Free and E(3)-Equivariant Constraints

and obtain the final output:

$$\mathbf{T}_{\text{out}}^{(\ell_{\text{out}})}(\mathbf{r}, t) = \sum_{i=1}^N \sum_{\ell_v, \ell_{\text{mix}}} W_{\text{out}}^{(\ell_{\text{out}}, \ell_v, \ell_{\text{mix}})} \times \left[\alpha(\mathbf{r}, t; \mathbf{r}_i, t_i) \mathbf{V}_{\text{exp}}^{(\ell_v)}(\mathbf{r}, t; \mathbf{r}_i, t_i) \otimes \mathbf{Y}^{(\ell_{\text{mix}})}(\widehat{\mathbf{r} - \mathbf{r}_i}) \right]^{(\ell_{\text{out}})} \quad (13)$$

For LTC, each hidden unit i is governed by the ODE Hasani et al. (2022):

$$\tau_i(x_t, h_t) \dot{h}_i(t) = -h_i(t) + \sigma(W_{x,i}x_t + W_{h,i}h_t + b_i) \quad (14)$$

Equation 14 is linear in h_i , which means it has a closed-form solution over an event span Δt :

$$h_i(t + \Delta t) = e^{-\Delta t / \tau_i} h_i(t) + \left(1 - e^{-\Delta t / \tau_i}\right) u_i(t), \quad u_i(t) = \sigma(\dots) \quad (15)$$

To achieve E(3)-equivariance in the LTC model, we treat the state vector h_t as a collection of geometric tensors: 64 copies of scalar irreps (0e) and vector irreps (1o) e3nn Developers (2025) corresponding to the magnetic field components. Since scalars are rotationally invariant and vectors transform as $\ell = 1$ geometric tensors, the ODE dynamics in Equation 15 automatically preserve E(3)-equivariance by evolving each irreducible representation independently. This approach never mixes channels of different tensor types (0e with 1o), maintaining the geometric structure throughout.

3.3 CONTINUOUS-TIME HANDLING

Magnetometer data is irregularly sampled, with measurement times t_1, t_2, \dots, t_N and unpredictable intervals $\Delta t_i = t_{i+1} - t_i$ arising from sensor jitter, asynchronous system clocks, and platform motion. This irregular sampling violates the fixed-step assumption. Additionally, real-world magnetic measurements are corrupted by structured, time-varying noise: e.g., eddy-current noise arises from induced currents in nearby conductors and obeys diffusion-type ODEs, leading to non-stationary distortions; temperature-driven bias and external electromagnetic interference introduce further drift and slowly varying biases that cannot be modelled as white noise. To address irregular sampling and ODE-governed noise, we model the latent state $z(t) \in \mathbb{R}^d$ via a continuous-time neural ODE, $\frac{dz}{dt} = f_\theta(z(t), t)$, where f_θ is a neural network. At each t_i , the model maps $z(t_i)$ to observation space as $\hat{x}(t_i) = g_\phi(z(t_i))$ and penalizes the discrepancy $\|\hat{x}(t_i) - x_{\text{meas}}(t_i)\|$ in the loss.

3.4 LONG-TERM MEMORY

Magnetic noise sources like eddy currents, sensor drift, and thermal bias often have long-lasting effects.

Table 1: CNN and MLP results.

		CNN				MLP			
		Base	Divergence-Free	E(3)-Equivariance	Both	Base	Divergence-Free	E(3)-Equivariance	Both
RMSE (nT)	Training	55.36	36.58	33.27	30.19	65.23	40.12	37.45	25.67
	Testing	60.29	41.14	38.61	35.48	70.47	45.28	42.19	30.34
SNR (dB)	Training	45.18	48.78	49.60	50.45	43.75	47.98	48.57	51.85
	Testing	44.44	47.76	48.31	49.04	43.08	46.93	47.54	50.40

Table 2: LTC and Contiformer results. RMSE ranges from 13.09 to 42.31 nT and SNR ranges from 47.51 to 57.70 dB.

		LTC				Contiformer			
		Base	Divergence-Free	E(3)-Equivariance	Both	Base	Divergence-Free	E(3)-Equivariance	Both
RMSE (nT)	Training	38.74	30.85	28.43	24.92	19.04	16.16	15.22	13.09
	Testing	42.31	34.67	32.16	28.55	21.57	18.62	17.75	15.09
SNR (dB)	Training	48.28	50.26	50.97	52.11	54.45	55.87	56.40	57.70
	Testing	47.51	49.24	49.90	50.93	53.37	54.64	55.06	56.47

For example, eddy currents can last minutes, while temperature bias evolves over hours. These slow dynamics cause long-range autocorrelation, with noise at time t still correlated with much earlier noise.

Sensor offsets and low-frequency interference accumulate rather than reset. Thus, we must maintain a latent state $z(t)$ capable of storing long-term context, so past noise remains accessible for accurate prediction.

3.5 SYNTHETIC DATA GENERATION USING CONDITIONAL GAN

MagNav datasets are limited, with only a handful of publicly available flight recordings. The data comprises multichannel magnetometer sequences recorded at 10 Hz aboard fixed-wing aircraft; the limited available datasets constrain supervised learning models for magnetic denoising. To expand this corpus, we train a conditional Generative Adversarial Network (cGAN) that learns $p(\mathbf{x} | c)$, the joint distribution of windowed sensor readings $\mathbf{x} \in \mathbb{R}^{T \times C}$ conditioned on a discrete context label c (e.g. flight ID, attitude bin). The generator G_θ accepts latent noise $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, I)$ and an embedding \mathbf{e}_c and produces synthetic sequences $\tilde{\mathbf{x}} = G_\theta(\mathbf{z}, \mathbf{e}_c)$; the discriminator D_ϕ receives $(\mathbf{x}, \mathbf{e}_c)$ and outputs two heads: $D_{\text{adv}} \in [0, 1]$ for real/fake and $D_{\text{cls}} \in \Delta^{K-1}$ for class prediction. The GAN design is as follows: Generator: 2-layer LSTM ($h = 256$) with linear expansion of $[\mathbf{z}; \mathbf{e}_c]$ to T time steps, output projected to $C = 4$ channels and bounded by \tanh . Discriminator: matching LSTM depth, spectral normalization on recurrent weights, final hidden state feeds dual heads.

Furthermore, training minimizes the discriminator loss

$$\mathcal{L}_D = -\mathbb{E}_{(\mathbf{x}, c)} \log D_{\text{adv}} + \mathbb{E}_{(\mathbf{z}, c)} \log(1 - D_{\text{adv}}) + \lambda \mathcal{L}_{\text{cls}}^D \quad (16)$$

and generator loss

$$\mathcal{L}_G = -\mathbb{E}_{(\mathbf{z}, c)} \log D_{\text{adv}} + \lambda \mathcal{L}_{\text{cls}}^G, \quad (17)$$

where D_{adv} denotes the adversarial head output, $\mathcal{L}_{\text{cls}}^D$ and $\mathcal{L}_{\text{cls}}^G$ are cross-entropy classification losses on real and synthetic samples, respectively, and $\lambda = 1$. We employ AdamW optimizer ($\eta = 10^{-4}$, $\beta_{1,2} = (0.5, 0.999)$) with gradient clipping at 1.0 and one-sided label smoothing of 0.9.

4 EXPERIMENTS

We evaluate our denoising system with three goals: (i) Noise removal, (ii) Preservation of physical properties – divergence-free structure and E(3) equivariance, and (iii) Verifying that models with long-term memory and continuous-time handling outperform those without.

4.0.1 DATA PRE-PROCESSING

While raw magnetometer data can be used directly, we implemented physics-based corrections to remove known interference. Transient Level compensation eliminated aircraft-induced noise. IGRF

correction removed temporal variations from the Earth’s core field, isolating crustal signatures. Diurnal corrections minimized time-dependent fluctuations and measurement noise.

4.0.2 TRAINING/TESTING

We tested 16 models, combining four backbone architectures (MLP, 1D CNN, LTC, and ContiFormer) with 4 layers of constraint (Plain, Div-Free, E3NN, and E3NN + Div-Free) applied orthogonally. Each model was trained on an NVIDIA RTX A6000 GPU using consistent parameters, including 10-second windows, batch size of 128, 40 epochs, Adam optimizer with exponential learning rate decay, and early stopping based on validation RMSE.

4.0.3 BASELINES AND METRICS

A significant bottleneck in evaluating magnetic denoising and compensation techniques is due to the lack of standardized datasets and metrics. Image-based methods often use PSNR, SSIM, and RMSE, while sensor-based approaches favor STD, NRF, and IR Pereira et al. (2024); Wang et al. (2022). Datasets vary from images, to time-series, to UAV surveys. For our MagNav flight data and similar synthetic datasets, traditional metrics like SPE, PSNR, and SSIM are ill-suited due to the non-stationary, time-series nature. Instead, we use time-domain metrics: Root Mean Squared Error (RMSE) to measure reconstruction accuracy, and Signal-to-Noise Ratio (SNR) to evaluate denoising quality.

$$\text{RMSE} = \sqrt{\frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{B}}(t) - \mathbf{B}(t)\|_2^2}$$

while SNR measures relative noise suppression: Higher SNR values indicate better performance - for eg., an SNR of 100 dB means our signal is 10^{10} times stronger than the error, while 80 dB means it’s 10^8 times stronger, so even a 20 dB difference represents a 100-fold improvement in signal quality.

$$\text{SNR} = 10 \cdot \log_{10} \left(\frac{\sum_{t=1}^T \|\mathbf{B}(t)\|_2^2}{\sum_{t=1}^T \|\hat{\mathbf{B}}(t) - \mathbf{B}(t)\|_2^2} \right)$$

$\mathbf{B}(t)$ = true field; $\hat{\mathbf{B}}(t)$ = field generated by the NN.

5 RESULTS

We evaluated four architectures: MLP, CNN, LTC, and ContiFormer on the Flt1005, Flt1006, Flt1007, 2017, 2015, 2016 from MagNav and FltS005, FltS006, FltS007 from the conditional GAN synthetic dataset. Without constraints, performance varied widely: MLP performed worst (Train RMSE: 65.23; Test: 70.47), while ContiFormer was best (Train: 25.38; Test: 28.76). CNN and LTC achieved intermediate results (Test RMSEs: 60.29 and 42.31), consistent with prior work Nerrise et al. (2024). (See Table.) Adding physics-aware constraints consistently improved performance. The combination of divergence-free and equivariance constraints yielded the largest gains: MLP test RMSE dropped by 57%, and ContiFormer improved by 30%. These results highlight the benefit of embedding physical priors, especially in data-scarce or noisy settings.

6 DISCUSSION

Our experiments show that neural networks with divergence-free and equivariance constraints consistently outperform unconstrained architectures for magnetic field estimation. Though the reasons are more nuanced than they first appear. While magnetic fields inherently obey these constraints, real-world measurements are corrupted by sensor noise, platform interference, and calibration errors, which introduce physically invalid divergences. Magnetic fields can be decomposed into solenoidal, gradient, and harmonic components via the Helmholtz-Hodge decomposition Bhatia et al. (2013).

432 *True magnetic fields lie entirely in the solenoidal subspace* due to Gauss’s law, while the gradi-
433 ent component captures noise that violates this constraint. Thus, forcing network outputs to be
434 divergence-free effectively removes this non-physical gradient noise. The $E(3)$ -equivariance con-
435 straint complements this by averaging over rotations, suppressing noise components that do not
436 transform correctly under coordinate changes. Since most disturbances arise from hardware orienta-
437 tion, not the field itself, this symmetry-based filtering enhances robustness while preserving the true
438 rotational structure of the magnetic field.

439

440 7 CONCLUSION

441

442 We address the challenge of magnetic anomaly navigation, highlighting the importance of denois-
443 ing magnetometer data for downstream performance. We develop neural architectures tailored to
444 the problem’s structure—marked by continuous-time dynamics and long-range dependencies—and
445 introduce two physics-aware constraints: divergence-freeness and $E(3)$ -equivariance. Our experi-
446 ments show that combining these priors with suitable architectures yields significant gains over the
447 state-of-the-art. This line of research demonstrates that physics constraining can be effectively ap-
448 plied for other purposes as well, utilizing different architectures for various spatio-temporal systems
449 beyond magnetic anomaly navigation. The incorporation of fundamental physical and mathemati-
450 cal principles into network architectures provides a generalizable framework for diverse dynamical
451 systems.

452

453

454

455

456

457

458

459

460

461

462

463

464

465

466

467

468

469

470

471

472

473

474

475

476

477

478

479

480

481

482

483

484

485

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

REFERENCES

- Harsh Bhatia, Gregory Norgard, Valerio Pascucci, and Peer-Timo Bremer. The helmholtz-hodge decomposition—a survey. *IEEE Transactions on Visualization and Computer Graphics*, 19(8): 1386–1404, 2013. doi: 10.1109/TVCG.2012.316.
- Aaron Canciani and John F. Raquet. Positioning using magnetic fields. In Y. T. Jade Morton, Frank van Diggelen, James J. Spilker, Bradford W. Parkinson, Sherman Lo, and Grace Gao (eds.), *Position, Navigation, and Timing Technologies in the 21st Century*, pp. 1521–1540. Wiley, 2020. doi: 10.1002/9781119458555.ch48. Chapter 48.
- Aaron J. Canciani. Magnetic navigation on an f-16 aircraft using online calibration. *IEEE Transactions on Aerospace and Electronic Systems*, 58(1):420–434, 2022. doi: 10.1109/TAES.2021.3101567.
- e3nn Developers. Irreps. Online, e3nn Documentation, docs.e3nn.org, 2025.
- Albert Gnad, Allan Wollaber, and Aaron Nielsen. Derivation and extensions of the tolles-lawson model for aeromagnetic compensation. *arXiv preprint arXiv:2212.09899*, 2022. URL <https://arxiv.org/abs/2212.09899>.
- Albert R. Gnad. *Advanced Aeromagnetic Compensation Models for Airborne Magnetic Anomaly Navigation*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, United States, May 2022a. Ph.D. Thesis, MIT; URL: <https://dspace.mit.edu/bitstream/handle/1721.1/145137/gnad-gnad-phd-aeroastro-2022-thesis.pdf>.
- Albert Reuben Gnad. *Advanced Aeromagnetic Compensation Models for Airborne Magnetic Anomaly Navigation*. PhD thesis, Massachusetts Institute of Technology, 2022b.
- Paul D. Groves. Navigation using inertial sensors [tutorial]. *IEEE Aerospace and Electronic Systems Magazine*, 30(2):42–69, 2015. doi: 10.1109/MAES.2014.130191.
- Bin Gu, Qingli Li, and Hongying Liu. Aeromagnetic compensation based on truncated singular value decomposition with an improved parameter-choice algorithm. In *2013 6th International Congress on Image and Signal Processing (CISP)*, pp. 1545–1551, 2013. doi: 10.1109/CISP.2013.6743921.
- Ramin Hasani, Mathias Lechner, Alexander Amini, Lucas Liebenwein, Aaron Ray, Max Tschaikowski, Gerald Teschl, and Daniela Rus. Closed-form continuous-time neural networks. *Nature Machine Intelligence*, 4(11):992–1003, 2022. doi: 10.1038/s42256-022-00556-7.
- IATA. Global navigation satellite system (gnss) radio frequency interference: Safety risk assessment, version 4, 2024. URL https://www.iata.org/contentassets/c8e90fe690ce4047a8edfa97f4824890/iata_safety_risk_assessment_gnss_interference.pdf. Accessed: 2025-09-24.
- B. W. Leach. Automatic aeromagnetic compensation. Technical report, 1979. URL <https://doi.org/10.4224/40003373>.
- Paul Leliak. Identification and evaluation of magnetic-field sources of magnetic airborne detector equipped aircraft. *IRE Transactions on Aerospace and Navigational Electronics*, ANE-8(3):95–105, 1961. doi: 10.1109/TANE3.1961.4201799.
- Favour Nerrise, Andrew Sosanya, and Patrick Neary. Physics-informed calibration of aeromagnetic compensation in magnetic navigation systems using liquid time-constant networks. In *Advances in Neural Information Processing Systems 36 (NeurIPS 2023)*, New Orleans, Louisiana, USA, 2023. Curran Associates, Inc. URL <https://neurips.cc/media/neurips-2023/Slides/76255.pdf>.
- Favour Nerrise, Andrew Sosa Sosanya, and Patrick Neary. Physics-informed calibration of aeromagnetic compensation in magnetic navigation systems using liquid time-constant networks. *arXiv preprint arXiv:2401.09631*, 2024. URL <https://arxiv.org/abs/2401.09631>.

-
- 540 Tiago Pereira, Victor Santos, Tiago Gameiro, Carlos Viegas, and Nuno Ferreira. Evaluation of
541 different filtering methods devoted to magnetometer data denoising. *Electronics*, 13(11):2006,
542 2024. doi: 10.3390/electronics13112006. URL [https://www.mdpi.com/2079-9292/
543 13/11/2006](https://www.mdpi.com/2079-9292/13/11/2006).
- 544 Zhong-Kun Qiao, Peng Yuan, Lin-Fei Wang, Zhi-Hou Zhang, Yu-Sen Huang, Jia-Jun Zhang, Lin-
545 Ling Li, Zong-Yu Zhang, Bin Wu, and Qiang Lin. Research on aeromagnetic compensation of
546 a multi-rotor uav based on robust principal component analysis. *Journal of Applied Geophysics*,
547 206:104791, 2022. doi: <https://doi.org/10.1016/j.jappgeo.2022.104791>. URL [https://www.
548 sciencedirect.com/science/article/pii/S0926985122002622](https://www.sciencedirect.com/science/article/pii/S0926985122002622).
- 549 Jack Richter-Powell, Yaron Lipman, and Ricky T. Q. Chen. Neural conservation laws: a divergence-
550 free perspective. In *Advances in Neural Information Processing Systems 36*, pp. 2759–2772, 2022.
- 551 Víctor Garcia Satorras, Emiel Hooeboom, and Max Welling. E(n) equivariant graph neural net-
552 works. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139,
553 pp. 9323–9332, 2021. URL [http://proceedings.mlr.press/v139/satorras21a/
554 satorras21a.pdf](http://proceedings.mlr.press/v139/satorras21a/satorras21a.pdf).
- 555 Walter E Tolles. Compensation of aircraft magnetic fields. US Patent 2,692,970, 1954. URL
556 <https://patents.google.com/patent/US2692970A>.
- 557 Walter E Tolles. Magnetic field compensation system. US Patent 2,706,801, 1955. URL <https://patents.google.com/patent/US2706801A>.
- 558 Walter E Tolles and John D Lawson. Magnetic compensation of mad equipped aircraft. *Airborne
559 Instruments Lab. Inc., Mineola, NY, Rept*, pp. 201–1, 1950.
- 560 Weilin Wang, Keyan Li, Zhihao Yang, Jun Chen, Linliang Miao, Jun Ouyang, and Xiaofei Yang. An
561 adaptive alternating magnetic interference suppression (aais) algorithm for geomagnetic vector
562 measurement. *Sensors*, 22(10):3642, 2022. doi: 10.3390/s22103642. URL [https://www.
563 mdpi.com/1424-8220/22/10/3642](https://www.mdpi.com/1424-8220/22/10/3642).
- 564 Yongji Wang, Mehdi Bennani, James Martens, Sébastien Racanière, Sam Blackwell, Alex
565 Matthews, Stanislav Nikolov, Gonzalo Cao-Labora, Daniel S. Park, Martin Arjovsky, Daniel
566 Worrall, Chongli Qin, Ferran Alet, Borislav Kozlovskii, Nenad Tomašev, Alex Davies, Pushmeet
567 Kohli, Tristan Buckmaster, Bogdan Georgiev, Javier Gómez-Serrano, Ray Jiang, and Ching-Yao
568 Lai. Discovery of unstable singularities, 2025. URL [https://arxiv.org/abs/2509.
569 14185](https://arxiv.org/abs/2509.14185).
- 570 Michael L Webb and Aaron C White. Environmental noise reduction for magnetometry. US Patent
571 8,786,277, 2014. URL <https://patents.google.com/patent/US8786277>.
- 572 Oliver J. Woodman. An introduction to inertial navigation. Research report 696, University of Cam-
573 bridge, August 2007. URL [http://citeseerx.ist.psu.edu/viewdoc/summary?
574 doi=10.1.1.63.7402](http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.63.7402).
- 575 Peilin Wu, Qunying Zhang, Luzhao Chen, Wanhua Zhu, and Guangyou Fang. Aeromag-
576 netic compensation algorithm based on principal component analysis. *Journal of Sen-
577 sors*, 2018(1):5798287, 2018. doi: <https://doi.org/10.1155/2018/5798287>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1155/2018/5798287>.
- 578 Xin Xu, Ling Huang, Xiaojun Liu, and Guangyou Fang. Deepmad: Deep learning for magnetic
579 anomaly detection and denoising. *IEEE Access*, 8:121257–121266, 2020. doi: 10.1109/ACCESS.
580 2020.3006795.

581 A APPENDIX

582 We claim that the expression given in equation 11 is E(3)-equivariant. Let us formally prove this.
583 Consider a transformation $g = (R, \mathbf{t}) \in E(3)$.
584
585
586
587
588
589

594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647

1. Under the transformation:
 - Positions: $\mathbf{r}'_i = R\mathbf{r}_i + \mathbf{t}$
 - Relative positions: $\mathbf{r}'_{ij} = \mathbf{r}'_j - \mathbf{r}'_i = R(\mathbf{r}_j - \mathbf{r}_i) = R\mathbf{r}_{ij}$
 - Unit vectors: $\hat{\mathbf{r}}'_{ij} = R\hat{\mathbf{r}}_{ij}$
2. Spherical harmonics transform as:

$$\mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}'_{ij}) = \mathbf{Y}^{(\ell)}(R\hat{\mathbf{r}}_{ij}) = D^{(\ell)}(R)\mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}_{ij}) \quad (18)$$

3. Input tensors transform as:

$$\mathbf{T}_j^{(\ell_{in})} = D^{(\ell_{in})}(R)\mathbf{T}_j^{(\ell_{in})} \quad (19)$$

4. The tensor product preserves equivariance:

$$\begin{aligned} [\mathbf{T}_j^{(\ell_{in})} \otimes \mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}'_{ij})]^{(\ell_{out})} &= D^{(\ell_{out})}(R) \\ &\times [\mathbf{T}_j^{(\ell_{in})} \otimes \mathbf{Y}^{(\ell)}(\hat{\mathbf{r}}_{ij})]^{(\ell_{out})} \end{aligned} \quad (20)$$

5. Since weights are scalars, the final result is:

$$\mathbf{T}_i^{(\ell_{out})} = D^{(\ell_{out})}(R)\mathbf{T}_i^{(\ell_{out})} \quad (21)$$

This proves E(3)-equivariance.