

Structured Learning for Electromagnetic Field Modeling and Real-Time Inversion

António Bernardes^{*†}, Jasan Zughaibi^{*†}, Michael Muehlebach[‡], Bradley J. Nelson[†]

^{*}These authors contributed equally to this work

[†]Multi-Scale Robotics Lab, ETH Zürich, 8092 Zürich, Switzerland

[‡]Learning and Dynamical Systems Group, Max Planck Institute for Intelligent Systems, 72076 Tübingen, Germany
Email: costaa@ethz.ch, zjasan@ethz.ch, michael.muehlebach@tuebingen.mpg.de, bnelson@ethz.ch

Abstract—Precise magnetic field modeling is fundamental to the closed-loop control of electromagnetic navigation systems (eMNS) and the analytical Multipole Expansion Model (MPEM) is the current standard. However, the MPEM relies on strict physical assumptions regarding source symmetry and isolation, and requires optimization-based calibration that is highly sensitive to initialization. These constraints limit its applicability to systems with complex or irregular coil geometries. This work introduces an alternative modeling paradigm based on multi-layer perceptrons that learns nonlinear magnetic mappings while strictly preserving the linear dependence on currents. As a result, the field models enable fast, closed-form minimum-norm inversion with evaluation times of approximately 1 ms, which is critical for high-bandwidth magnetic control. For model training and evaluation we use large-scale, high-density datasets collected from the research-grade OctoMag and clinical-grade Navion systems. Our results demonstrate that data-driven models achieve predictive fidelity equivalent to the MPEM while maintaining comparable data efficiency, and we further assess their suitability for real-time magnetic control in a closed-loop tracking experiment running at 100 Hz. Furthermore, we demonstrate that straightforward design choices effectively eliminate spurious workspace ill-conditioning frequently reported in MPEM-based calibration. To facilitate future research, we release the complete codebase and datasets open source.

SUPPLEMENTARY MATERIAL

The complete codebase for model training, the experimental datasets, and the pre-trained models for evaluation are available at <https://github.com/antonio14bernardes/CalibrationForMagneticNavigation/tree/main>.

I. INTRODUCTION

Remote magnetic manipulation is rapidly emerging as a promising technology for contactless actuation across a wide range of length scales. It enables precise, responsive control of catheters, guidewires, and endoscopes for applications ranging from telesurgery to targeted drug delivery with micro- and nanorobots [1–4]. These capabilities are realized using magnetic navigation systems, which generate and shape the magnetic field using one or more magnetic sources: either permanent magnets with controllable pose [5, 6] or electromagnets with controllable coil currents [7, 8] (including, in some designs, moving electromagnets [9]). Electromagnetic navigation systems (eMNS) arrange multiple fixed electromagnets around a workspace and, given a sufficient number and suitable placement of the coils, can synthesize a desired




Fig. 1. Visualization of the positions where magnetic field data has been acquired. (a) Measured positions for the Octomag system (53248 total positions). (b) Measured positions for the clinically ready Navion system (7500 total positions). To the best of our knowledge, these represent the densest datasets recorded for such systems.

magnetic field at a target position by appropriately choosing the coil currents [10]. Accordingly, eMNS control hinges on a field model that is not only accurate, but also computationally efficient to support closed-loop operation. This includes a forward map from coil currents to the resulting magnetic field (and, when needed, its spatial gradients), as well as an inverse model to compute currents that realize a desired output. While our focus is medical eMNS, the need for real-time field prediction and model inversion is broadly shared across

electromagnetic platforms, including MRI field shaping [11], magnetic particle imaging [12], magnetic levitation / bearing actuation [13], and wireless power-transfer coil arrays [14]. More recently, researchers in the nuclear-fusion community have discussed whether the coils in fusion facilities could support new forms of remote handling and field-aware robotics—for example, using controlled magnetic forces for partial gravity compensation in long-reach maintenance systems [15].

In this work, we introduce a set of learning-based eMNS field models implemented with multilayer perceptrons (MLPs) that learn magnetic field mappings under different structure- and physics-aware representations, designed to preserve tractable, real-time inversion. We benchmark these models against the *de facto* standard, the Multipole Expansion Model (MPEM) [16], and against black-box regressors (standard MLPs and gradient-boosted linear-tree models) providing an empirical comparison of model-based and model-free approaches to magnetic field modeling. We additionally demonstrate, through a closed-loop control experiment, that the proposed structured models indeed retain the real-time invertibility required for feedback control. Finally, we investigate a repeatedly noted failure mode of the MPEM—spurious ill-conditioned regions in the workspace that can destabilize inversion and control [16, 17]—and show that straightforward design choices are sufficient for mitigation.

The analysis in this work considers two eMNS platforms: the OctoMag [8] shown in Figure 1(a), a research-grade eMNS with eight coils and an approximately 13 cm-diameter hemispherical workspace, and the Navion [7] shown in Figure 1(b), a clinical-grade eMNS with three coils that provides magnetic control over a large, human-scale workspace. For training and evaluation, we collected, to the best of our knowledge, the largest datasets reported to date for each system, with dense sampling and broad spatial coverage across their workspaces. This coverage enables accurate modeling and assessment of spatial gradients. We release both datasets open source as supplementary material, alongside our codebase for training and a suite of pretrained models for immediate evaluation.

A. Motivation for current-linear field models

Our key emphasis is on models that are linear (or affine) in the coil currents, targeting the system’s practical pre-saturation operating regime where the electromagnet-generated field scales approximately linearly with current. Preserving current-linearity is central to eMNS practice, enabling tractable workspace and actuation-capacity analyses [17–19], and supporting a fast, closed-form, minimum-norm inversion via the Moore–Penrose pseudoinverse [20]. Within this regime, the contributions of individual electromagnets superpose, and an eMNS with S electromagnets is typically described by a current-linear forward model: currents $\mathbf{i} \in \mathbb{R}^S$ produce at position $\mathbf{p} \in \mathbb{R}^3$ the magnetic field vector $\mathbf{b} \in \mathbb{R}^3$ as $\mathbf{b} = \mathcal{A}_b(\mathbf{p})\mathbf{i}$, where $\mathcal{A}_b(\mathbf{p}) \in \mathbb{R}^{3 \times S}$ is commonly referred to as the field actuation matrix.

In the present study, we augment the standard linear formu-

lation with a position-dependent affine term, leading to

$$\mathbf{b} = \mathcal{A}_b(\mathbf{p})\mathbf{i} + \mathbf{b}_0(\mathbf{p}), \quad (1)$$

which preserves the convenience of workspace analysis and inversion, while capturing background contributions such as sensor-current biases (e.g., from miscalibration) and remanent magnetization within the workspace. Note that although (1) is affine in the currents, both $\mathcal{A}_b(\mathbf{p})$ and $\mathbf{b}_0(\mathbf{p})$ may vary strongly/nonlinearly with position. In our learning-based approaches, we learn these terms either directly, or indirectly via a physics-informed scalar potential parameterization.

B. Why not rely on nonlinear forward models?

Prior work has applied learning-based models to eMNS primarily to extend accuracy beyond the linear regime, where higher currents drive ferromagnetic cores toward saturation and the MPEM’s linearity assumption, along with its modeling performance, breaks down. In particular, random forests, MLPs, and convolutional neural networks have been shown to achieve strong overall predictive accuracy across operating conditions, including regimes where saturation effects are non-negligible [21–23]. However, these gains can come with a trade-off: models optimized for global nonlinear fidelity do not necessarily preserve high accuracy in the near-linear operating range that is most relevant for fast, well-conditioned inversion. Nevertheless, nonlinear models have seen limited adoption in remote magnetic motion control, in part because they forfeit the algebraic structure available in the linear regime: the inverse problem is no longer a pseudoinverse and instead requires more expensive inversion techniques such as iterative nonlinear optimization [21, 24], formulation as a differential control problem [24], or a separately learned inverse map [23].

In addition, many eMNS applications operate in highly dynamic settings that demand high-bandwidth control [17, 25–29], making inverse-model latency a primary concern [30]. Iterative inversion schemes for nonlinear forward models are therefore difficult to accommodate at the required update rates.

C. Why not learn the inverse map directly?

While the inverse-MLP in [23] can, in principle, provide low-latency inference, fundamentally, learning an inverse map from fields (and gradients) to currents cannot, in general, guarantee energy-optimal actuation, even when regularized toward minimum-norm currents given a field/gradient target. Since many feedback control algorithms specify objectives in motion space (torques/forces), optimizing an inverse in this intermediate space (fields-to-currents) independently of the full composition does not yield a globally optimal solution to the end-to-end inverse relevant for control. The resulting currents can differ by an order of magnitude as shown in [17]. Furthermore, practical current inversion may need to be modified online to account for saturation, power-supply limits, and temperature-dependent weighting [31], which a fixed learned inverse cannot naturally accommodate.

In contrast, because the mapping from field/gradients to magnetic torque/force is linear, an affine field model preserves affinity from currents to the control input, enabling

fast, closed-form, minimum-norm inversion of the end-to-end mapping using a computationally inexpensive pseudoinverse operation.

D. Outline

The remainder of the paper is structured as follows. Section II introduces the experimental platforms and details the data collection procedure, including an empirical assessment of linearity and outlier removal. Section III describes the proposed learning-based architectures alongside the benchmark models. Section IV presents a comprehensive empirical evaluation of the models, covering predictive accuracy, data efficiency, physical consistency, computational performance, and performance in a closed-loop control task, while also investigating the sources of residual error and examining MPEM-induced ill-conditioning in field actuation maps. Finally, Section V summarizes the findings and discusses implications for real-time control.

II. DATA COLLECTION

Reliable learning-based field models require consistent measurements across currents and space; we therefore detail our acquisition protocol and the verification/cleaning steps applied to both datasets.

A. Experimental platforms

1) *OctoMag*: To collect magnetic-field data on the OctoMag, we used a $4 \times 4 \times 4$ array of Melexis® MLX90395 triaxial high-field magnetometers with 5 mm sensor spacing. The array was mounted in a rigid fixture and positioned at a wide range of poses, yielding a dense dataset spanning the full workspace with an effective spatial sampling of approximately 1.7 mm, as shown in Figure 1(a). To ensure accurate pose measurements, the cube was tracked with an infrared motion-capture system using reflective markers.

We collect data with currents in the range ± 4 A. Since we focus on the linear regime where superposition holds, each coil's contribution can be modeled independently. Accordingly, to streamline data collection, we ramped one coil at a time from -4 to 4 A in 1 A increments while holding all other coils at zero current.

2) *Navion*: For the Navion, we used a separate, larger sensor array arranged in a $5 \times 5 \times 5$ grid of Allegro® A1301 magnetometers with 5 cm spacing. The array was manually positioned at a large set of positions throughout the workspace, with denser sampling in regions closer to the magnetic sources. As in the OctoMag procedure, the sensor array was tracked using an infrared motion-capture system. The spatial extent of this dataset is seen in Figure 1(b).

Data were acquired using coil currents from 0 to 30 A in 5 A steps, with the same per-coil sweep strategy as for the OctoMag.

B. Linearity in currents

For both systems, we verified that the collected current ranges fall within their corresponding linear regime. For each sampled position, we fit an affine model relating the measured magnetic field to the commanded current in each electromagnet. The distributions of R^2 values from the per-position affine fits and those of the affine intercept magnitudes are shown in Figure 2. The consistently high R^2 values support linear behavior over the tested ranges, while the intercept magnitudes indicate a persistent constant field offset, likely originating from a miscalibration of the current sensors.




Fig. 2. Evaluation of magnetic field linearity through boxplots of fit quality (R^2) and field offset magnitudes ($\|b_0\|$) for the OctoMag and Navion. The consistently high R^2 values confirm that both systems operate within the linear-current regime over the tested ranges, while the nonzero offsets indicate persistent constant field contributions.

C. Outlier removal

To remove anomalous measurements, we use RANSAC [32], a robust estimator for structured models commonly used for outlier detection in the field of computer vision. Specifically, we exploit the observed affine-in-currents substructure and the per-coil sweep collection strategy of our data: for each position and active coil, we fit an affine model using RANSAC and classify as outliers those samples whose residuals exceed the inlier tolerance of the consensus fit.

A key design choice is the inlier criterion. A criterion based on the absolute residual, $\|b - \hat{b}\|$, tends to reject large-magnitude samples that are accurate in relative terms, whereas a purely relative residual, $\|b - \hat{b}\|/\|b\|$, over-penalizes smaller measurements—this is undesirable because low-magnitude samples encode the field's rapid decay as the distance from the source increases, an important concern for modeling. We therefore use a scale-aware relative residual,

$$r = \frac{\|b - \hat{b}\|}{\max(\|b\|, b_T)}, \quad (2)$$

with a denominator floor of b_T that prevents the relative error from being over-amplified in the low-field regime. We set the floor magnitude to the first quartile of the $\|b\|$ distribution, namely, $b_T = 2.66$ mT for the OctoMag dataset and $b_T = 3.27$ mT for the Navion dataset. A measurement is classified as an outlier if $r > 20\%$.

This procedure flagged approximately 5.03% of the OctoMag dataset as outliers and approximately 3.11% for the Navion. Note that, since the criterion is based on consistency with currents, it will not detect samples whose field direction

is systematically corrupted (e.g., due to a pose-tracking error) but that remain affine in the current command.

III. MODELING METHODS

Magnetic fields used for magnetic manipulation are well described by the quasi-static Maxwell equations. In the workspace, assumed to be current-free, these reduce to:

$$\nabla \cdot \mathbf{b} = 0, \quad \nabla \times \mathbf{b} = \mathbf{0}. \quad (3)$$

Thus, the workspace field is both divergence-free and curl-free. In eMNS control, a scalar potential Φ is commonly used to express the field as

$$\mathbf{b} = -\nabla \Phi. \quad (4)$$

Substituting (4) into the divergence-free condition yields Laplace's equation:

$$\nabla^2 \Phi = 0. \quad (5)$$

Conditions (3)–(5) provide a useful set of physical constraints for model formulation and physical-consistency evaluation of model-free methods.

A. The Multipole Expansion Model

The MPEM, introduced in [16], models each magnetic source by an azimuthally symmetric scalar potential Φ . Under this formulation, and considering a workspace external to the electromagnets, Laplace's equation (5) admits the truncated multipole expansion

$$\Phi(r, \cos \theta) = \sum_{n=1}^N \frac{B_n}{r^{n+1}} P_n(\cos \theta), \quad (6)$$

where N is the order of approximation, $P_n(\cdot)$ denotes the n -th Legendre polynomial, B_n is the corresponding multipole coefficient, r is the distance to the evaluation point, and θ is the polar angle from the source's zenith direction (i.e., its symmetry axis). The MPEM therefore satisfies Maxwell's equations (3) by construction.

As mentioned above, the MPEM is formulated under the assumption of linearity in currents, hence the field contribution of a single source is written as $\mathbf{b}(\mathbf{p}) = i \tilde{\mathbf{b}}(\mathbf{p})$, where the unit current field contribution $\tilde{\mathbf{b}}$ follows from (4); for brevity, we refer the reader to [16] for the resulting closed-form expression. To account for cross-magnetization, the MPEM augments each of the S electromagnets with F sources that model induced magnetization in nearby ferromagnetic components—namely the cores of the surrounding electromagnets. Additionally, to accommodate the non-negligible constant field offset observed in Section II, we augment the model by adding one source per electromagnet, each subject to a constant unit-current. The resulting field produced by currents vector $\mathbf{i} = [i_1, \dots, i_S]^\top$ at position $\mathbf{p}$ is:

$$\mathbf{b}(\mathbf{p}, \mathbf{i}) = \sum_{s=1}^S i_s \left(\tilde{\mathbf{b}}_s(\mathbf{p}) + \sum_{f=1}^F \tilde{\mathbf{b}}_{s,f}(\mathbf{p}) \right) + \mathbf{b}_{s,0}(\mathbf{p}), \quad (7)$$

where $\tilde{\mathbf{b}}_s$ and $\tilde{\mathbf{b}}_{s,f}$ are the unit-current field contributions of the s -th electromagnet by direct current excitation and through induced magnetization of ferromagnetic body f , respectively, and $\mathbf{b}_{s,0}$ (in T) is the offset term. Collecting the per-electromagnet contributions in (7) yields the affine map (1).

Calibration then amounts to estimating the B_n coefficients and the poses (location and zenith direction) of each source, iteratively solving a non-linear least squares problem using the Levenberg–Marquardt algorithm [16].

B. The ActuationNet

We have repeatedly emphasized the utility of the linear (or affine) structure applicable in the linear regime. Accordingly, we introduce the ActuationNet, an MLP f_{AN} that takes as input a position $\mathbf{p} \in \mathbb{R}^3$ in the workspace and predicts the local field actuation matrix and affine offset of (1):

$$(\hat{\mathbf{A}}_{\mathbf{b}}, \hat{\mathbf{b}}_0) = f_{\text{AN}}(\mathbf{p}), \quad \hat{\mathbf{A}}_{\mathbf{b}}(\mathbf{p}) \in \mathbb{R}^{3 \times S}, \quad \hat{\mathbf{b}}_0(\mathbf{p}) \in \mathbb{R}^3. \quad (8)$$

The ActuationNet and all MLPs introduced below use tanh as the activation function in all hidden layers.

C. The PotentialNet

While the MPEM, by construction, preserves the curl- and divergence-free conditions, the ActuationNet is free to learn an arbitrary affine mapping and thus does not necessarily respect Maxwell's equations (3). As an alternative, we introduce the PotentialNet, an MLP f_{PN} that maps a workspace position $\mathbf{p} \in \mathbb{R}^3$ to a set of scalar potentials: one unit-current potential per coil and a bias potential:

$$(\hat{\phi}, \hat{\phi}_0) = f_{\text{PN}}(\mathbf{p}), \quad \hat{\phi}(\mathbf{p}) \in \mathbb{R}^S, \quad \hat{\phi}_0(\mathbf{p}) \in \mathbb{R}, \quad (9)$$

which define the predicted scalar potential

$$\hat{\Phi}(\mathbf{p}, \mathbf{i}) = \hat{\phi}(\mathbf{p})^\top \mathbf{i} + \hat{\phi}_0(\mathbf{p}). \quad (10)$$

Taking the gradient of (10) with respect to $\mathbf{p}$ yields a predicted affine map of the form (1) which is curl-free by construction:

$$\begin{aligned} \hat{\mathbf{A}}_{\mathbf{b}}(\mathbf{p}) &= -\nabla \hat{\phi}(\mathbf{p}) \in \mathbb{R}^{3 \times S}, \\ \hat{\mathbf{b}}_0(\mathbf{p}) &= -\nabla \hat{\phi}_0(\mathbf{p}) \in \mathbb{R}^3. \end{aligned} \quad (11)$$

Note that this formulation guarantees a curl-free predicted field, but it does not enforce divergence-freeness. While, in principle, a learning-based method could enforce both curl- and divergence-free conditions by learning a parameterization over harmonic-gradient basis functions, such a choice would effectively amount to learning the parameters of a reduced, MPEM-like model class (a harmonic potential with a prescribed basis), and would therefore blur the distinction between model-based and model-free approaches that we seek to benchmark. Alternatively, we could encourage divergence-freeness during training via a Laplacian-based regularization term [23], but we do not apply such regularization in the present work.

D. Benchmark mappings

As additional baselines, we evaluate two models that directly learn the current-to-field map $\hat{\mathbf{b}} = f(\mathbf{p}, \mathbf{i})$, to quantify the value of the reduced representations used by the ActuationNet and the PotentialNet. In contrast to our structured models, these direct baselines impose no explicit assumption of affinity in currents. We consider the DirectNet, an MLP, and the DirectGBT, a gradient-boosted linear-tree regressor (i.e., each leaf predicts a $\hat{\mathbf{b}}$ contribution via a local affine function of the inputs). Because data are collected via per-coil sweeps, the dataset contains only one-hot current vectors, so these direct models may extrapolate poorly to simultaneous multi-coil actuation; we therefore consider them strictly as benchmarks in the current setting, not as deployable models without additional multi-coil training.

E. Training

All models are trained to minimize the mean-squared error (MSE) in the magnetic field predictions,

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N \|\mathbf{b}_i - \hat{\mathbf{b}}_i\|^2, \quad (12)$$

requiring the output of the ActuationNet to be converted with (1) and that of the PotentialNet with (11) and (1) prior to loss evaluation.

The full datasets are randomly split into 70% training, 15% validation, and 15% test sets based on the positional features. This means that for any fixed position $\mathbf{p}$, all corresponding samples across all applied currents $\mathbf{i}$ remain entirely within one split (train, validation, or test). Although a conventional random split over the full $(\mathbf{p}, \mathbf{i})$ feature set may appear appealing, it is inappropriate for the affine models (MPEM, ActuationNet, PotentialNet), whose outputs depend only on position. A sample-wise split would distribute identical positions across different splits, leading to information leakage and, moreover, may leave the training set with an insufficient diversity of current values at a given position to reliably estimate the affine parameters (slope and intercept).

We train the MPEM as described in Section III-A and stop when the validation RMSE, i.e., $\sqrt{\mathcal{L}_{\text{MSE}}}$, changes by less than 10^{-4} mT between successive iterations. Neural networks are trained with the Adam optimizer [33] (learning rate 10^{-3}) and stopped when the validation RMSE fails to improve by at least 10^{-3} mT over 10 consecutive epochs. DirectGBT training is stopped after 200 boosting iterations without validation improvement, using a shrinkage factor of 0.1 for each new tree. For all models, we retain the parameter set that achieves the lowest validation RMSE.

F. Gradient modeling

Spatial derivatives of the predicted field provide a useful diagnostic for physical consistency with Maxwell's equations (3). More importantly, obtaining spatial derivatives of the magnetic fields in real-time is critical for magnetic control, since magnetic forces depend on field gradients. In eMNS control,

the constraints in (3) allow for a reduced representation of the full 3×3 spatial Jacobian. By convention, we define the gradient vector as

$$\mathbf{g} := \left[\frac{\partial b_x}{\partial x}, \frac{\partial b_x}{\partial y}, \frac{\partial b_x}{\partial z}, \frac{\partial b_y}{\partial y}, \frac{\partial b_y}{\partial z} \right]^\top.$$

Affinity in coil currents is preserved under positional differentiation. Stacking the field and reduced gradient yields

$$\begin{bmatrix} \mathbf{b} \\ \mathbf{g} \end{bmatrix} = \begin{bmatrix} \mathcal{A}_b(\mathbf{p}) \\ \mathcal{A}_g(\mathbf{p}) \end{bmatrix} \mathbf{i} + \begin{bmatrix} \mathbf{b}_0(\mathbf{p}) \\ \mathbf{g}_0(\mathbf{p}) \end{bmatrix}, \quad (13)$$

where $\mathcal{A}_g(\mathbf{p}) \in \mathbb{R}^{5 \times S}$ is referred to as the gradient actuation matrix.

Although our datasets do not include direct measurements of magnetic field gradients, we can still obtain gradient estimates by exploiting the structure of each model. For the MPEM, spatial derivatives follow analytically from its closed-form expression. For the neural models, using tanh as the activation function, the forward mappings are inherently continuous and smooth, with gradients being obtained via automatic differentiation. For the DirectGBT, gradients are obtained by summing the derivatives of the active linear leaves across trees, yielding a piecewise-constant Jacobian, undefined only on split boundaries.

IV. RESULTS

In this section, we evaluate the proposed structured learning-based models against the physics-based MPEM and unstructured learning-based benchmarks (DirectNet, DirectGBT) using the collected datasets. We evaluate predictive accuracy relative to model complexity and data density, analyze the sources of residual error via geometric sensitivity analysis, assess physical consistency with Maxwell's equations, examine computational efficiency for real-time control, and demonstrate the applicability of the proposed models in a closed-loop tracking task. Finally, we propose practical methods to mitigate the MPEM-induced workspace ill-conditioning.

A. Field prediction and model complexity

We first evaluate field prediction accuracy across model classes. As shown in Figure 3(a) and (e), all methods achieve comparable test-set performance, converging to an RMSE of nearly 0.5 mT for the OctoMag dataset and 3 mT for the Navion dataset. To analyze the effect of model size/complexity, we define three complexity ranks (I, II, III) within each model class. For neural models, the ranks correspond to hidden-layer structures with increasing neuron counts, i.e. (256, 256), (256, 256, 256), and (512, 512, 512), respectively. For gradient-boosted tree models, ranks correspond to increasing per-tree leaf limits (32, 64, and 128 leaves). For the MPEM baseline, ranks correspond to increasing approximation/multipole order (first/dipole, second/quadrupole, third/octopole). Using these definitions, Figure 3(a) and (e) confirm that test RMSE is largely insensitive to complexity rank across model classes, suggesting an error floor that is not driven by model capacity.




Fig. 3. Field prediction accuracy, model complexity, and data-efficiency analysis for the OctoMag and the Navion. Panels (a) and (e) show test-set RMSE as a function of model complexity rank, while (b) and (f) report the corresponding generalization gap, defined as $\text{RMSE}_{\text{test}} - \text{RMSE}_{\text{train}}$. Test RMSE is largely insensitive to model complexity, indicating an error floor not primarily driven by model capacity. Panels (c) and (g) show test-set RMSE as a function of dataset percentage for the highest-capacity variant of each model class, and (d) and (h) show the corresponding generalization gap. For the OctoMag, the structured models, especially the PotentialNet, remain comparatively robust under reduced data availability, whereas unstructured baselines show stronger overfitting. For the Navion, where the baseline positional density is lower, model degradation occurs earlier in the sweep. The PotentialNet remains the most robust learning-based model, although it does not reach MPEM-level robustness at the lowest density.

Finally, to quantify overfitting, we report the generalization gap (Gap RMSE), defined as $\text{RMSE}_{\text{test}} - \text{RMSE}_{\text{train}}$, in Figure 3(b) and (f). With the exception of DirectGBT, which shows a higher tendency to overfit, the gap remains near zero for all models, indicating minimal overfitting; moreover, this behavior is again insensitive to model capacity.

B. Dependence on positional resolution

To evaluate how performance degrades with reduced data resolution, we retrain the highest-capacity variant (complexity rank III) of each model class using progressively smaller random subsamples of the training and validation sets, while keeping the test set fixed. Given that high-resolution acquisition is resource-intensive, this analysis aims to identify the minimum sampling density required to maintain predictive fidelity. Subsampling is performed at the positional level as detailed in Section III-E. We report, for each system, the test-set RMSE in Figure 3(c) and (g) as well as the generalization gap in Figure 3(d) and (h) as a function of dataset percentage.

For the OctoMag (Figure 3(c) and (d)), test RMSE and generalization remain largely stable under data reduction down to 20%. At 5%, both unstructured baselines show clearer signs of overfitting (increased generalization gap), whereas the structured models retain a smaller gap. At 1%, the ActuationNet begins to noticeably overfit while the PotentialNet retains performance approaching MPEM-level robustness, suggesting that imposing model structure is beneficial for mitigating overfitting without sacrificing test RMSE.

For the Navion (Figure 3(g) and (h)), similar trends are observed, although, given the lower baseline positional density of the Navion dataset, performance degradation occurs earlier in the data reduction sweep. As with the OctoMag results,

PotentialNet remains the most robust learning-based model under reduced data availability. At the lowest data percentage, however, it does not reach the MPEM’s performance.

C. RMSE floor and uncertainty analysis

As shown in Figure 3, all model classes, even with increasing capacities, converge to a similar RMSE of approximately 0.5 mT for the OctoMag dataset and 3 mT for the Navion when provided sufficient data. This convergence suggests error floors determined not by model capacity, but by physical limitations. Accordingly, we now investigate the hypothesis that this residual error is dominated by geometric uncertainties—specifically, small deviations in sensor position and orientation.

The sensor array used for the OctoMag comprises four flexible PCBs mounted in a 3D-printed fixture that also carries the motion-capture markers, while for the Navion each sensor is manually assembled onto the fixture. Manufacturing and assembly tolerances inevitably introduce mechanical deviations, causing fixed offsets between the actual sensor poses and the marker-defined frame. Although these geometric offsets are static, they induce varying field errors as the arrays are re-oriented throughout the workspace during data collection. We model this geometric uncertainty as an isotropic translational error $\delta \mathbf{p} \sim \mathcal{N}(\mathbf{0}, \sigma_p^2 \mathcal{I})$ and a small-angle rotational error $\delta \theta \sim \mathcal{N}(\mathbf{0}, \sigma_\theta^2 \mathcal{I})$, where $\mathcal{I} \in \mathbb{R}^{3 \times 3}$ is the identity matrix. Using a first-order Taylor expansion, the resulting field-error covariance $\Sigma \in \mathbb{R}^{3 \times 3}$ for a given sample is:

$$\Sigma = \sigma_p^2 \nabla \mathbf{b} \nabla \mathbf{b}^\top + \sigma_\theta^2 [\mathbf{b}]_\times [\mathbf{b}]_\times^\top, \quad (14)$$

where $[\mathbf{b}]_\times$ is the skew-symmetric cross-product matrix of the field vector. Using a dipole-order MPEM to approximate the




Fig. 4. Position-orientation uncertainty combinations that reproduce the observed RMSE floor under the first-order geometric uncertainty model. Each curve shows the $\text{RMSE}_{\text{floor}}$ iso-contour corresponding to the empirical error floor of the respective system: 0.5 mT for the OctoMag and 3 mT for the Navion. This analysis demonstrates that sub-millimeter position error and degree-scale orientation error reproduce the observed error floors. The larger uncertainty levels required for the Navion are consistent with its manually assembled sensor array, whereas the OctoMag array is PCB-based and therefore more geometrically certain.

gradients, we compute the expected RMSE induced by these tolerances over the N samples of each test set:

$$\text{RMSE}_{\text{floor}} = \sqrt{\frac{1}{N} \sum_{i=1}^N \text{trace}(\Sigma_i)}. \quad (15)$$

The 0.5 mT and 3 mT iso-contours for the OctoMag and Navion, respectively, in Figure 4 demonstrate that realistic combinations of sub-millimeter position error and degree-scale orientation error reproduce the observed error floors. This supports the hypothesis that geometric uncertainty, rather than model capacity, is the dominant source of residual error.

D. Physical consistency

Beyond prediction accuracy, we evaluate each model’s physical consistency by checking adherence to the quasi-static Maxwell’s equations (3). Focusing on the highest-capacity variants (rank III) trained on the full datasets, we compute the divergence and curl magnitude using the spatial derivatives on each test set. The MPEM is excluded from this analysis as it satisfies both conditions in (3) by construction.

Figure 5 shows that, for both systems, the structured models (PotentialNet and ActuationNet) exhibit substantially smaller Maxwell residuals than the direct current-to-field baselines, therefore achieving better physical consistency. Additionally, as intended from its scalar-potential formulation, the PotentialNet successfully produces a curl-free magnetic field model.

E. Evaluation time

Magnetic real-time control hinges on the rapid evaluation of the actuation matrix which needs to be recomputed at each control cycle. We assess the computational efficiency of our control-ready models by averaging runtimes over 10^4 evaluations (Figure 6). All computed times remain sufficiently small (<1.5 ms), confirming that our proposed models are suitable for high-frequency real-time magnetic control.




Fig. 5. Evaluation of physical consistency through boxplots showing the distribution of curl magnitude (top) and divergence (bottom) computed on each system’s test set. The structured learning-based models (ActuationNet, PotentialNet) demonstrate significantly better adherence to Maxwell’s equations compared to the unstructured baselines (DirectNet, DirectGBT). Notably, the PotentialNet exhibits zero curl error, confirming that it satisfies the curl-free condition by construction. The MPEM baseline is excluded from this plot as it inherently satisfies both conditions by construction.




Fig. 6. Computational efficiency: Mean evaluation time versus model complexity rank for the structured learning-based models. Solid lines denote the computation of the field actuation matrix ($\mathcal{A}_b$), while dashed lines correspond to the joint evaluation of both field and gradient actuation matrices ($\mathcal{A}_{b,g}$). All evaluation times remain below 1.5 ms, indicating suitability for high-frequency real-time control. In accordance with real-time control usage, the neural models run on the CPU (Apple M3 (8-core CPU), 16 GB unified memory). The MPEM is not included since, owing to its lightweight closed-form analytical formulation, its evaluation time is negligible on the scale of this comparison.

F. Closed-loop control application

To assess the suitability of the proposed models for closed-loop control, we consider a circular trajectory-tracking task involving a magnetically actuated rod with two angular degrees of freedom, α and β , following the methodology of [17]. Since the system exhibits inherent Lyapunov instability, achieving stabilization requires a sufficiently high sampling frequency, thereby confirming our models’ capacity for real-time inversion. The rod carries a rigidly attached stack of permanent




Fig. 7. Closed-loop circular trajectory tracking of the Lyapunov-unstable rod in the (α, β) plane using the MPEM, ActuationNet, and PotentialNet field models. The circular reference trajectory is denoted by a dashed gray line. The results show successful tracking for all three models.

magnets, modeled as a point dipole with magnetic moment $\mathbf{m} \in \mathbb{R}^3$ located at its geometric center, position $\mathbf{p}_m \in \mathbb{R}^3$. Consequently, both the dipole position $\mathbf{p}_m$ and the orientation of the moment $\mathbf{m}$ are functions of α and β .

Magnetic torques and forces exerted on the dipole (i.e. the magnet stack), are given by $\mathbf{t} = \mathbf{m} \times \mathbf{b}$ and $\mathbf{f} = (\mathbf{m} \cdot \nabla)\mathbf{b}$, respectively, and can be expressed as the linear mapping:

$$\begin{bmatrix} \mathbf{t} \\ \mathbf{f} \end{bmatrix} = \begin{bmatrix} \mathcal{M}_b(\alpha, \beta) & 0 \\ 0 & \mathcal{M}_g(\alpha, \beta) \end{bmatrix} \begin{bmatrix} \mathbf{b} \\ \mathbf{g} \end{bmatrix} = \mathcal{M}(\alpha, \beta) \begin{bmatrix} \mathbf{b} \\ \mathbf{g} \end{bmatrix}, \quad (16)$$

where, defining the magnetic dipole moment as $\mathbf{m} = [m_x, m_y, m_z]^\top$,

$$\begin{aligned} \mathcal{M}_b(\alpha, \beta) &= \text{skew}(\mathbf{m}) = \begin{bmatrix} 0 & -m_z & m_y \\ m_z & 0 & -m_x \\ -m_y & m_x & 0 \end{bmatrix}, \\ \mathcal{M}_g(\alpha, \beta) &= \begin{bmatrix} m_x & m_y & m_z & 0 & 0 \\ 0 & m_x & 0 & m_y & m_z \\ -m_z & 0 & m_x & -m_z & m_y \end{bmatrix}. \end{aligned} \quad (17)$$

The resulting torque about the rod's pivot point is $\mathbf{t}_c = \mathbf{t} + \tilde{\mathcal{J}}(\alpha, \beta)\mathbf{f}$, we refer the reader to [17] for the definition of the mechanical Jacobian $\tilde{\mathcal{J}}(\alpha, \beta) \in \mathbb{R}^{3 \times 3}$. By introducing the matrix $\Lambda := [\mathcal{I} \quad \tilde{\mathcal{J}}] \mathcal{M}$, where $\mathcal{I} \in \mathbb{R}^{3 \times 3}$ denotes the identity matrix, this torque can be expressed as an affine function of the input currents:

$$\mathbf{t}_c = \Lambda(\alpha, \beta) \left(\begin{bmatrix} \mathcal{A}_b(\mathbf{p}_m) \\ \mathcal{A}_g(\mathbf{p}_m) \end{bmatrix} \mathbf{i} + \begin{bmatrix} \mathbf{b}_0(\mathbf{p}_m) \\ \mathbf{g}_0(\mathbf{p}_m) \end{bmatrix} \right). \quad (18)$$

Given a desired torque $\mathbf{t}_c^*$ computed by the feedback controller, the required setpoint currents $\mathbf{i}^*$ are then calculated as:

$$\mathbf{i}^* = \left(\Lambda \begin{bmatrix} \mathcal{A}_b \\ \mathcal{A}_g \end{bmatrix} \right)^\dagger \left(\mathbf{t}_c^* - \Lambda \begin{bmatrix} \mathbf{b}_0 \\ \mathbf{g}_0 \end{bmatrix} \right), \quad (19)$$

where the dependence on α , β , and $\mathbf{p}_m$ has been omitted for compactness and $\dagger$ denotes the Moore-Penrose pseudoinverse.

Control is implemented using the LQRI methodology described in [17], running at 100Hz with the considered field models. As shown in Figure 7, like the MPEM, both proposed models enable successful tracking, demonstrating suitability for real-time inversion.

G. Workspace ill-conditioning

As reported in [16, 17] and illustrated in Figure 8, the actuation map derived from the MPEM can suffer from significant ill-conditioning. We note that this phenomenon is not exclusive to the MPEM: similar ill-conditioned regions emerge with the PotentialNet, although they are not observed with the ActuationNet. Our analysis focuses on the MPEM, as it serves as the current *de facto* standard model and its ill-conditioning has been explicitly documented in the aforementioned literature.

We observe an empirical trend that calibration samples placed closer to the coils yield markedly better-conditioned actuation maps, whereas emphasizing far-field samples tends to exacerbate ill-conditioned regions. To visualize this behavior, we partition the measured positions of the Navion dataset (Figure 1(b)) into three radial segments according to their distance, in the horizontal plane, from the Navion electromagnets; the boundary radii are chosen such that all segments contain the same number of positions. We then train a dipole-order MPEM separately on Segment 1 (nearest), Segment 2 (intermediate), and Segment 3 (farthest), and visualize in Figure 8(a) the conditioning of the resulting field actuation map over a horizontal slice of the workspace. Consistent with the observed trend, near-source sampling substantially mitigates the ill-conditioned regions, whereas training predominantly on farther samples exacerbates them. Intuitively, because higher-order terms of the magnetic field decay rapidly with distance from the source, far-field samples are inherently less informative.

Accordingly, we also evaluate the effect of increasing the model's representational capacity by allowing it to capture the contributions of higher-order multipole terms, now using the full Navion dataset. As shown in Figure 8(b), moving from dipole to higher-order MPEM models further reduces ill-conditioning, demonstrating that limited model expressiveness also contributes to spurious conditioning artifacts.

V. CONCLUSION

In this work, we presented a comprehensive framework for learning-based magnetic field modeling in electromagnetic navigation systems (eMNS), introducing the ActuationNet and PotentialNet architectures. By rigorously benchmarking these data-driven approaches against the physics-based Multipole Expansion Model (MPEM) and unstructured regressors, we demonstrated that learning-based models can effectively replace analytical formulations without sacrificing the algebraic structure required for real-time control.

Our empirical analysis on the OctoMag and Navion systems reveals that the proposed neural models achieve predictive fidelity equivalent to the MPEM, while maintaining comparable data efficiency. Crucially, we identified that the performance of all evaluated models converges to a common RMSE floor for each system. Our sensitivity analysis suggests this limit is governed not by model capacity, but by irreducible geometric uncertainties in the ground-truth acquisition setup. Using the Navion, we also demonstrated that appropriate sampling and




Fig. 8. Analysis of workspace ill-conditioning for the MPEM. Heatmaps displaying the logarithmic condition number of the field actuation matrix across a horizontal workspace slice. (a) Effect of training data distribution: Dipole-order MPEMs trained on radial data segments of increasing distance (Segment 1: nearest, Segment 3: farthest). Training on near-field samples (Segment 1) significantly mitigates the ill-conditioned regions (yellow arcs) compared to training on far-field samples. (b) Effect of approximation order: Conditioning maps for MPEMs of increasing expansion order (Dipole, Quadrupole, Octopole), illustrating how increasing the expansion order mitigates ill-conditioning.

sufficient model order effectively eliminate the workspace ill-conditioning frequently reported in MPEM-based calibration.

From a control perspective, this study emphasizes the value of preserving current-linearity. Unlike general non-linear approximations that require iterative optimization for inversion, our affine formulations support closed-form, minimum-norm inversion via the pseudoinverse. This yields evaluation times in the range of approximately 1 ms, ensuring that computational latency remains sufficiently low for high-bandwidth feedback loops in real-time magnetic control. Moreover, we experimentally evaluated that the proposed models can indeed support real-time feedback control.

In addition, while analytical models like the MPEM rely on strict assumptions of source symmetry and magnetic isolation, our data-driven models are agnostic to both. This flexibility makes them particularly suitable for applications where idealized analytical assumptions fail—for instance in remote maintenance in nuclear fusion facilities, where complex coil geometries render classical expansions insufficient.

Finally, the open-source release of our complete codebase and high-density datasets for both clinical- and research-scale systems provides a valuable foundation for future research.

ACKNOWLEDGMENTS

The authors would like to thank Jonathan Fiene, Felix Grueniger, Arjen Houweling, Rob Faulkner, and Julian Martus for the technical support. Further thanks go to the Max Planck ETH Center for Learning Systems for the financial support. Michael Muehlebach thanks the German Research Foundation for the financial support.

REFERENCES

- [1] Fabian C. Landers, Lukas Hertle, Vitaly Pustovalov, Derrick Sivakumaran, Cagatay M. Oral, Oliver Brinkmann, Kirstin Meiners, Pascal Theiler, Valentin Gantenbein,

Andrea Veciana, Michael Mattmann, Silas Riss, Simone Gervasoni, Christophe Chautems, Hao Ye, Semih Sevim, Andreas D. Flouris, Josep Puigmartí-Luis, Tiago Sotto Mayor, Pedro Alves, Tessa Lühmann, Xiangzhong Chen, Nicole Ochsenbein, Ueli Moehrlen, Tilman Schubert, Zsolt Kulcsar, Philipp Gruber, Miriam Weisskopf, Quentin Boehler, Salvador Pané, and Bradley J. Nelson. Clinically ready magnetic microrobots for targeted therapies. *Science*, 390(6774):710–715, 2025. URL <https://www.science.org/doi/10.1126/science.adx1708>.

- [2] Yoonho Kim, Emily Genevriere, Pablo Harker, Jaehun Choe, Marcin Balicki, Robert W. Regenhart, Justin E. Vranic, Adam A. Dmytriw, Aman B. Patel, and Xuanhe Zhao. Telerobotic neurovascular interventions with magnetic manipulation. *Science Robotics*, 7(65):eabg9907, 2022. URL <https://www.science.org/doi/10.1126/scirobotics.abg9907>.
- [3] Alexandre Mesot, Michelle Mattille, Quentin Boehler, Nina Schmid, Sean Lyttle, Florian Heemeyer, Shannon Melissa Chan, Philip Wai Yan Chiu, and Bradley James Nelson. Teleoperated magnetic endoscopy: A case study and perspective. *Advanced Intelligent Systems*, 7(10):2400522, 2025. URL <https://advanced.onlinelibrary.wiley.com/doi/10.1002/aisy.202400522>.
- [4] Florian Heemeyer, Quentin Boehler, Minsoo Kim, Bernard R. Bendok, Evelyn L. Turcotte, H. Hunt Batjer, Ryan D. Madder, Vitor M. Pereira, and Bradley J. Nelson. Telesurgery and the importance of context. *Science Robotics*, 10(99):eadq0192, 2025. URL <https://www.science.org/doi/abs/10.1126/scirobotics.adq0192>.
- [5] Giovanni Pittiglio, Lavinia Barducci, James W. Martin, Joseph C. Norton, Carlo A. Avizzano, Keith L. Obstein, and Pietro Valdastrì. Magnetic Levitation for Soft-Tethered Capsule Colonoscopy Actuated With a Single

- Permanent Magnet: A Dynamic Control Approach. *IEEE Robotics and Automation Letters*, 4(2):1224–1231, 2019. URL <https://ieeexplore.ieee.org/document/8624310/>.
- [6] Seunguk Kim, Suhong Bae, Wonseo Lee, and Gunhee Jang. Magnetic Navigation System Composed of Dual Permanent Magnets for Accurate Position and Posture Control of a Capsule Endoscope. *IEEE Transactions on Industrial Electronics*, 71(1):739–748, 2024. URL <https://ieeexplore.ieee.org/document/10049799/>.
- [7] Simone Gervasoni, Norman Pedrini, Tarik Rifai, Cedric Fischer, Fabian C. Landers, Michael Mattmann, Roland Dreyfus, Silvia Viviani, Andrea Veciana, Enea Masina, Buse Aktas, Josep Puigmartí-Luis, Christophe Chautems, Salvador Pané, Quentin Boehler, Philipp Gruber, and Bradley J. Nelson. A Human-Scale Clinically Ready Electromagnetic Navigation System for Magnetically Responsive Biomaterials and Medical Devices. *Advanced Materials*, 36(31):2310701, 2024. URL <https://onlinelibrary.wiley.com/doi/10.1002/adma.202310701>.
- [8] Michael P. Kummer, Jake J. Abbott, Bradley E. Kratochvil, Ruedi Borer, Ali Sengul, and Bradley J. Nelson. OctoMag: An Electromagnetic System for 5-DOF Wireless Micromanipulation. *IEEE Transactions on Robotics*, 26(6):1006–1017, 2010. URL <http://ieeexplore.ieee.org/document/5595508/>.
- [9] Xingzhou Du, Lidong Yang, Jiangfan Yu, Kai Fung Chan, Philip Wai Yan Chiu, and Li Zhang. Robomag: A magnetic actuation system based on mobile electromagnetic coils with tunable working space. *International Conference on Advanced Robotics and Mechatronics*, pages 125–131, 2020. URL <https://ieeexplore.ieee.org/document/9195280>.
- [10] Andrew J. Petruska and Bradley J. Nelson. Minimum Bounds on the Number of Electromagnets Required for Remote Magnetic Manipulation. *IEEE Transactions on Robotics*, 31(3):714–722, 2015. URL <http://ieeexplore.ieee.org/document/7104159/>.
- [11] Christoph Juchem, Omar M. Nalhass, Terence W. Nixon, and Robin A. de Graaf. Multi-slice mri with the dynamic multi-coil technique. *NMR in Biomedicine*, 28(11):1526–1534, 2015. URL <https://analyticalsciencejournals.onlinelibrary.wiley.com/doi/abs/10.1002/nbm.3414>.
- [12] Matthias Gräser, Florian Thieben, Patryk Szwargulski, Franziska Werner, Nadine Gdaniec, Marija Boberg, Florian Griesse, Martin Möddel, Peter Ludewig, D Van De Ven, et al. Human-sized magnetic particle imaging for brain applications. *Nature Communications*, 10(1):1936, 2019. URL <https://www.nature.com/articles/s41467-019-09704-x>.
- [13] Spasoje Mirić, Rosario Giuffrida, Dominik Bortis, and Johann Walter Kolar. Dynamic electromechanical model and position controller design of a new high-precision self-bearing linear actuator. *IEEE Transactions on Industrial Electronics*, 68(1):744–755, 2020. URL <https://ieeexplore.ieee.org/document/9091931>.
- [14] Takuya Sasatani, Alanson P Sample, and Yoshihiro Kawahara. Room-scale magnetoquasistatic wireless power transfer using a cavity-based multimode resonator. *Nature Electronics*, 4(9):689–697, 2021. URL <https://www.nature.com/articles/s41928-021-00636-3>.
- [15] Rob Buckingham and Antony Loving. Remote-handling challenges in fusion research and beyond. *Nature Physics*, 12(5):391–393, 2016. URL <https://www.nature.com/articles/nphys3755>.
- [16] Andrew J. Petruska, Janis Edelmann, and Bradley J. Nelson. Model-Based Calibration for Magnetic Manipulation. *IEEE Transactions on Magnetics*, 53(7):1–6, 2017. URL <http://ieeexplore.ieee.org/document/7817832/>.
- [17] Jasan Zughabi, Denis von Arx, Maurus Derungs, Florian Heemeyer, Luca A Antonelli, Quentin Boehler, Michael Muehlebach, and Bradley J Nelson. Expanding the workspace of electromagnetic navigation systems using dynamic feedback for single-and multi-agent control. *arXiv preprint: 2511.18486*, 2025. URL <https://arxiv.org/abs/2511.18486>.
- [18] Quentin Boehler, Simone Gervasoni, Samuel L. Charreyron, Christophe Chautems, and Bradley J. Nelson. On the Workspace of Electromagnetic Navigation Systems. *IEEE Transactions on Robotics*, 39(1):791–807, 2023. URL <https://ieeexplore.ieee.org/document/9868240/>.
- [19] Ashkan Pourkand and Jake J. Abbott. A Critical Analysis of Eight-Electromagnet Manipulation Systems: The Role of Electromagnet Configuration on Strength, Isotropy, and Access. *IEEE Robotics and Automation Letters*, 3(4):2957–2962, 2018. URL <https://ieeexplore.ieee.org/document/8383988/>.
- [20] R. Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(3):406–413, 1955. URL https://www.cambridge.org/core/product/identifier/S0305004100030401/type/journal_article.
- [21] Samuel L. Charreyron, Quentin Boehler, Byungsoo Kim, Cameron Weibel, Christophe Chautems, and Bradley J. Nelson. Modeling Electromagnetic Navigation Systems. *IEEE Transactions on Robotics*, 37(4):1009–1021, 2021. URL <https://ieeexplore.ieee.org/document/9329101/>.
- [22] Ruoxi Yu, Samuel L. Charreyron, Quentin Boehler, Cameron Weibel, Christophe Chautems, Carmen C. Y. Poon, and Bradley J. Nelson. Modeling electromagnetic navigation systems for medical applications using random forests and artificial neural networks. *IEEE International Conference on Robotics and Automation*, pages 9251–9256, 2020. URL <https://ieeexplore.ieee.org/document/9197212>.
- [23] Pascal Ernst, Simone Gervasoni, Derick Sivakumaran, Enea Masina, David Sargent, Bradley J. Nelson, and Quentin Boehler. A physics-informed neural network for the calibration of electromagnetic navigation systems. *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5743–5749, 2025. URL <https://ieeexplore.ieee.org/document/11246707>.
- [24] Jake J. Abbott, Eric Diller, and Andrew J. Petruska. Mag-

- netic Methods in Robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):57–90, 2020. URL <https://www.annualreviews.org/doi/10.1146/annurev-control-081219-082713>.
- [25] Jasan Zughaibi, Bradley J. Nelson, and Michael Muehlebach. Dynamic Electromagnetic Navigation. *IEEE Robotics and Automation Letters*, 10(6):6095–6102, 2025. URL <https://ieeexplore.ieee.org/document/10972059/>.
 - [26] Neelaksh Singh, Jasan Zughaibi, Denis von Arx, Bradley J. Nelson, and Michael Muehlebach. Remote Magnetic Levitation Using Reduced Attitude Control and Parametric Field Models. *arXiv preprint: 2512.15207*, 2025. URL <http://arxiv.org/abs/2512.15207>.
 - [27] Peter Berkelman and Michael Dzadovsky. Magnetic Levitation Over Large Translation and Rotation Ranges in All Directions. *IEEE/ASME Transactions on Mechatronics*, 18(1):44–52, 2013. URL <http://ieeexplore.ieee.org/document/5985533/>.
 - [28] Muneaki Miyasaka and Peter Berkelman. Magnetic levitation with unlimited omnidirectional rotation range. *Mechatronics*, 24(3):252–264, 2014. URL <https://linkinghub.elsevier.com/retrieve/pii/S0957415814000294>.
 - [29] Peter Berkelman and Steven Kang. Multiple Magnet Independent Levitation and Motion Control using a Single Coil Array. *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, pages 537–542, 2023. URL <https://ieeexplore.ieee.org/document/10196212/>.
 - [30] Fengqiu Xu, Yun Lv, Xianze Xu, and Venkata Dinavahi. FPGA-Based real-time wrench model of direct current driven magnetic levitation actuator. *IEEE Transactions on Industrial Electronics*, 65(12):9635 – 9645, 2018. URL <http://ieeexplore.ieee.org/document/8305657/>.
 - [31] Ashkan Pourkand and Jake J. Abbott. Magnetic Actuation With Stationary Electromagnets Considering Power and Temperature Constraints. *IEEE Robotics and Automation Letters*, 5(4):6964–6971, October 2020. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2020.3025512.
 - [32] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. URL <https://doi.org/10.1145/358669.358692>.
 - [33] Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. *arXiv preprint: 1412.6980*, 2017. URL <http://arxiv.org/abs/1412.6980>.