TRANSACTION ON MAGNETICS , VOL. XX, NO. XX, XXXX 2024

Utilizing Synthetic Data in Supervised Learning for Robust 5-DoF Magnetic Marker Localization

Mengfan Wu^{*}, Thomas Langerak^{*}, Otmar Hilliges, Juan Zarate

Tracking passive magnetic markers plays a vital role in advancing healthcare and robotics, offering the potential to significantly improve the precision and efficiency of systems. This technology is key to developing smarter, more responsive tools and devices, such as enhanced surgical instruments, precise diagnostic tools, and robots with improved environmental interaction capabilities. However, traditionally, the tracking of magnetic markers is computationally expensive due to the requirement for iterative optimization procedures. Moreover, these methods depend on the magnetic dipole model for their optimization function, which can vield imprecise outcomes due to the model's significant inaccuracies when dealing with short distances between non-spherical magnet and sensor. Our paper introduces a novel approach that leverages neural networks to bypass these limitations, directly inferring the marker's position and orientation to accurately determine the magnet's five degrees of freedom (5 DoF) in a single step without initial estimation. Although our method demands an extensive supervised training phase, we mitigate this by introducing a computationally more efficient method to generate synthetic, yet realistic data using Finite Element Methods simulations. Our novel method utilizes the rotational symmetry of axis-symmetric magnetic



markers to transform the 3D simulations into 2D. The benefits of fast and accurate inference significantly outweigh the offline training preparation. In our evaluation, we use different cylindrical magnets, tracked with a square array of 16 sensors. We perform the sensors' reading and position inference on a portable, neural networks-oriented single-board computer, ensuring a compact setup. We benchmark our prototype against vision-based ground truth data, achieving a mean positional error of 4 mm and an orientation error of 8 degrees within a 0.2x0.2x0.15 m working volume. These results showcase our prototype's ability to balance accuracy and compactness effectively in tracking 5 DoF.

Index Terms—Passive magnet localization, machine learning, position and orientation tracking, permanent magnets.

I. INTRODUCTION

TRACKING the three-dimensional position and orientation of objects hidden from view is crucial for numerous applications. For instance, in Human-Computer Interaction (HCI), where the users' own body might occlude their hands, it involves decoding gesture input and delivering force feedback [1]–[8]. In healthcare, tracking plays a vital role in guiding medical instruments such as capsules and catheters [9]–[13]. Similarly, in robotics, tracking is used for tasks such as magnetomicrometry [14] and tracking multi-arm robots [15]. Magnetic tracking enables the precise localization of objects hidden from view by utilizing magnetic markers. This is accomplished by measuring the magnetic fields around these markers with Hall sensors. The widespread interest in this technology is driven by its ability to permeate through various materials, including human tissues, without interference.

However, tracking magnetic markers is not a straightforward task. Existing research distinguishes between active and passive magnetic markers. Active markers emit a magnetic

* Equal Contribution

signal through an integrated power source like a battery, offering a higher signal-to-noise ratio and the advantage of frequency filtering. However, tracking active emitters requires either a wired connection to the tracked object or a battery-powered circuitry [3], [16]. Hashi et al. proposed a solution to overcome this limitation with semi-passive markers [17]. While the tracked elements were passive, they respond with a resonant signal when excited by a nearby source. In contrast, passive magnetic localization is simpler, requiring only the target to be equipped with a permanent magnet, eliminating the need for tethered connections or active electronics on the moving elements [5], [18], [19]. Our work primarily focuses on passive magnetic localization due to its simplicity and robustness; however, future work could potentially extend the ideas presented here to other types of magnetic markers.

To track a magnetic marker we need to map magnetic signals, measured by one or more Hall sensors, to a marker's 5-degree-of-freedom (5DoF) position and orientation (rotation). The most straightforward approach involves directly interpolating the reading intensity within the sensor array. In [20], [21], this approach was employed to track a stylus and a gaming object in a two-dimensional plane. In more complex configurations (e.g., [3], [4], [22]), multiple sensors with known spatial arrangements created an over-constrained system of equations, ensuring a unique solution for the magnet's state. These triangulation methods are fast but are not

Manuscript received March, 2024. This work was supported in part by grants from ERC (OPTINT StG2016-717054).

This work was conducted at the Department of Computer Science at ETH Zürich. Mengfan Wu is with TU Berlin and Huawei Research (email: mengfan.wu@campus.tu-berlin.de). All the other authors are with the Department of Computer Science at ETH Zürich. (email: first.last@inf.ethz.ch). All authors declare that they have no conflicts of interest.

This article has been accepted for publication in IEEE Transactions on Magnetics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TMAG.2024.3416729

robust to sensor noise and limited to positional degrees of freedom.

A more robust approach, that also infers rotation, involves minimizing the difference between a theoretical magnetic field model, based on a magnetic model, and the sensor readings. To achieve this, several non-linear optimization algorithms have been proposed [23], [24]. In [19], the authors implemented an analytical computation of gradients to expedite the optimization process. In [25], their algorithm decoupled the marker's orientation from its position, enhancing calculation speed and providing guarantees regarding global minima. Nonetheless, these iterative approaches come with significant challenges. Firstly, gradient-descent algorithms are computationally demanding, leading to a trade-off between tracking precision and frequency. Secondly, iterative non-convex optimization can be susceptible to the non-uniqueness of the solution, potentially converging to wrong local minima. These challenges make these methods highly dependent on their initialization.

Furthermore, these methods assume that the magnetic model, that predicts the theoretical magnetic field, is valid. All the mentioned methods employ the first term of the multipole series expansion derived from Maxwell's equations. The magnetic dipole approximation assumes all magnets to be spherical, providing the simplest explicit expression for the magnetic field concerning distance to the source. However, this approach has significant limitation. For non-spherical magnets, reliable results are only obtained when the magnets are far from the sensor [26]. Crucially, in gradient descent optimization process, a small magnetic field approximation error can lead to substantial positional discrepancies.

Recently, machine learning (ML) has emerged as a valuable approach to circumvent the computational burden of iterative methods. Machine learning, especially Neural Networks (NN), excels in approximating non-linearities through sequential multiply-addition operations. Additionally, it doesn't require initial estimates, often achieving reliable results with a single inference. While inference is rapid, creating a large and representative dataset for training to generalize well to unseen samples can be a challenging and time-consuming process.

In [11], [27], [28], ML was utilized to predict the locations of tracked magnets using input from magnetic sensors. The data for training the neural network was collected in-vivo by placing magnets at various known locations and gathering sensor readings. Generally, the data collection process is lengthy and not scalable to new scenarios and markers. [29] employed a neural network to track a magnetic marker, though they trained on synthetic data generated by the dipole model. While they eliminated the need for iterative gradient-based solutions, the work did not address the limitations of the underlying model, retaining inaccuracies for non-spherical or closely located magnets. Recently, authors in [16] and [30] demonstrated the use of neural networks trained on synthetic data to track active markers. Although we share goals with these works, their results applied exclusively to active emitters and did not consistently outperform methods based on iterative optimization.

To address the data limitation challenge, we employ Finite Element Methods (FEM) to simulate the complete set of



Fig. 1: Schematic pipeline overview. We use a Multi-Layer Perceptron (MLP) to output the location and orientation of the magnet. During the training phase, the MLP outputs are compared with the ground truth for generating synthetic readings. During the inference, the input to the MLP is the sensor data.

Maxwell's equations (see Fig. 1 for a method overview). FEM is computationally intensive, particularly when employing a fine mesh for accurate results. Therefore, it is unsuitable for real-time modeling. However, FEM can be leveraged to generate noise-free synthetic datasets for training machine learning models. In various other fields, the use of FEM-generated synthetic datasets has successfully trained neural networks, although not for magnetic tracking. Examples include mechanical deformations [31], elastoplasticity [32], material inspection [33], and nano-structures [34]. To further reduce computation time while increasing dataset size, we focus on axis-symmetric markers. We enable this approach by introducing a coordinate transformation algorithm that capitalizes on our markers' symmetric properties, converting FEM-simulation results from 2D to 3D.

To evaluate the efficacy of our system, we evaluate our method in both simulation and experiments. In simulation we find that our method outperforms iterative methods, in terms of accuracy, robustness and computational efficiency. We also compare our FEM-based approach, to a method that is trained on magnetic-dipole generated data, where our method significantly outperforms the baseline. In real-world experiments we assess various cylindrical magnets with 5 degrees of freedom using a lower-power portable Nvidia Jetson Nano, where sensing and tracking are performed in parallel. We achieve an averaged error of 4 mm and an orientation error of less than 8 degrees, with an interactive rate of 75 Hz when using 8 sensors.

In summary, our work contributes significantly to the magnetic tracking literature in five key ways: *i*) We propose a supervised learning approach that enables real-time tracking of arbitrary axis-symmetric magnets using Neural Networks, overcoming the computational limitations of previous iterative solutions by approximating the inverse function of the magnetic field. *ii*) To address the challenge of data limitation, we leverage Finite Element Methods (FEM) to simulate



Fig. 2: Coordinate system overview for Alg. 1

Maxwell's equations, generating noise-free synthetic datasets for training machine learning models. *iii*) We introduce a coordinate transformation algorithm that capitalizes on our markers' symmetric properties, converting FEM-simulation results from 2D to 3D, enabling accurate high-resolution magnetic field data for training neural networks. *iv*) Due to the computational efficiency of our approach, we enable lowpower and portable applications. We demonstrate this by using a Nvidia Jetson Nano. *v*) We demonstrate the validity of our approach with in-silico and real-world evaluations. These contributions collectively enable more accurate and efficient magnetic tracking solutions, offering the potential for the development of truly portable tracking devices.

II. METHOD

Our contribution revolves around a novel tracking method through supervised learning, specifically utilizing a Multi-Layer Perceptron (MLP). Our tracking pipeline consists of two distinct phases: the training phase involves comparing the MLP's output to the position and orientation used to generate synthetic sensor readings through Finite Element Methods (FEM), while during the inference phase the sensor readings are inputted into the MLP to determine the magnetic marker's position and orientation (as illustrated in Fig. 1).

In this section, we provide a detailed explanation of our approach, starting with the creation of a high-resolution synthetic dataset (Sec. II-A). Subsequently, we detail the architecture and training process of our neural network (Sec. II-B). Finally, we outline the hardware setup (Sec. II-C).

A. Synthetic Dataset

Our focus is on axisymmetric magnetic markers, encompassing any shape and size of magnets, provided they have rotational symmetry around their magnetization axis. This category encompasses cylinders, spheres, and toroids with arbitrary cross-sections, covering the most commonly used types of permanent magnets.

Instead of costly 3D simulations, we leverage the symmetry to conduct just one high-resolution 2D Finite Element Method (FEM) simulation for each magnet shape. From a single 2D FEM simulation, we generate synthetic sensor readings for any location and orientation of the magnet. We achieve the 3D volume by revolving this 2D cross-section around the magnet's principal axis. This approach of creating a synthetic dataset

TABLE I: Summary of Notations.

Variables	Description		
\mathbf{p}_i^C	Positional vector. For C and i see below.		
$C \in \{D,M\}$	Coordinate system on device (D) or magnet (M)		
$i \in \{s,m\}$	Positional vector of sensor (S) or magnet (m)		
$\mathbf{u},\mathbf{v},\mathbf{w}$	vectors in C_D aligned with the axes of the coordinate system C_M		
w	Magnetic moment direction in C_D , used as the third axis in C_M		
В	Magnetic flux density		
\mathbf{B}_{j}	Magnetic flux density detected at sensor j		
$B_j^d, \\ d \in \{x, y, z\}$	x, y, or z-directional component of mangetic flux density detected at sensor j		
n	number of sensors		

Algorithm 1 Synthethic data	set generation
-----------------------------	----------------

1:	Input \mathbf{p}_s^D	
2:	$\mathbf{p}_d^D \leftarrow \mathbf{p}_m^D - \mathbf{p}_s^D$	Magnet-sensor vector
3:	$\mathbf{v} \leftarrow \mathbf{w} imes \mathbf{p}_d^D$	
4:	while $\mathbf{v} = 0 \mathbf{d} \mathbf{o}$	\triangleright If w and \mathbf{p}_d^D are (anti-)parallel
5:	random q	
6:	$\mathbf{v} \gets \mathbf{q} \times \mathbf{w}$	
7:	$\mathbf{u} \leftarrow \mathbf{v} \times \mathbf{w}$	
8:	$d_w \leftarrow \mathbf{p}_d^D \cdot \frac{\mathbf{w}}{ \mathbf{w} }$	\triangleright the projection of \mathbf{p}_d^D on \mathbf{w}
9:	$d_u \leftarrow \sqrt{ \mathbf{p}_d^D ^2 - dw^2}$	\triangleright the projection of \mathbf{p}_d^D on \mathbf{u}
10:	$d_v \leftarrow 0$	
11:	$\mathbf{B}^M \leftarrow \mathbf{2D} - \mathbf{FEM}(d$	(u, d_w)
12:	$M \leftarrow \left[\begin{smallmatrix} \mathbf{u} \\ \mathbf{u} ^\top & \frac{\mathbf{v}}{ \mathbf{v} } ^\top & \frac{\mathbf{w}}{ \mathbf{w} } \end{smallmatrix} \right.$	
		Coordinate transformation matrix
13:	$\mathbf{B}^D \leftarrow \mathbf{B}^M \cdot \mathbf{M}$	

significantly enhances computational efficiency during training, striking an optimal balance between detailed simulation and minimal data storage requirements.

We utilize COMSOL Multiphysics to acquire FEM data. The simulation centers on the magnet and due to the magnets symmetrical properties we constraint it to a single quadrant. We reconstruct the 3D magnetic B-field, using 2D FEM generated data and a sampled point. For generating training data and evaluating the neural network, we employ the transformation delineated in Alg. 1 (all variables are explicated in Tab. I) with coordinate systems, vectors, and scalars illustrated in Fig. 2. This algorithm converts 2D FEM data into 3D synthetic sensor readings as follows. Initially, it requires the magnet's current position and the locations of each sensor, all within a fixed coordinate system, C_D , centered at the array of sensors. Subsequently, we establish a coordinate system, C_M , centered on the magnet's current position and orientation. The three perpendicular axes of C_M are along u, v and w in C_D . u, v and w are then normalized to create coordinate transformation matrix M. The algorithm then use 2D - FEM to extract the magnetic flux density, \mathbf{B}^M in C_M , with the sensor's coordinates in C_M as $\langle du, 0, dw \rangle$, where the second coordinate is always 0 since we obtain FEM data in 2D. Ultimately, we transform \mathbf{B}^M back into the original global coordinates, C_D , to be used as features for training the MLP.

B. Tracking with Neural Networks

1) Multi-Layer Perceptron

We employ an MLP as our network architecture. The MLP takes in a 3n-element vector as input, which contains the (x, y, z) magnetic flux densities **B** from the *n* sensors. Its output is a 6-element tuple, comprising the magnet's position, $\mathbf{p} = [p_m^x, p_m^y, p_m^z]$, and orientation, $\mathbf{o} = [o_m^x, o_m^y, o_m^z]$. Therefore, we define the MLP as a non-linear mapping from the sensor readings to the tracking variables:

$$\mathcal{F}: \mathbb{R}^{3n} \to \mathbb{R}^6; \mathcal{F}(B_1^x, B_1^y, \dots, B_n^z) = [\mathbf{p}, \mathbf{o}]$$
(1)

It is important to note that, although we have a 5-Degree-of-Freedom (DoF) output, three in **p** and two in **o**, we represent the orientation vector in Cartesian coordinates. This approach helps avoid numerical discontinuity, particularly when the azimuth angle transitions from 359° to 0° . Our MLP architecture, depicted in Fig. 1, consists of a 3-layer perceptron with 2048 units per layer, excluding the input and output layers. We use ReLU as the activation functions.

2) Pre-processing

Crucially, the dipole model (Eq. 2) indicates that the magnetic field decreases with the cube of the distance $(1/r^3)$, where r represents the distance between the source and the sensors, and m denotes the dipole moment.

$$\mathbf{B}(\mathbf{p}_m) = \frac{\mu_0}{4\pi} \left[\frac{3\mathbf{r}(\mathbf{m} \cdot \mathbf{r})}{r^5} - \frac{\mathbf{m}}{r^3} \right],\tag{2}$$

We have empirically determined that system training fails to converge when using magnetic readings directly as inputs. This issue may arise because the input values change dramatically, by several orders of magnitude, as the magnet moves from near a sensor (approximately 10^{-2} Tesla) to the boundary of the working volume (around 10^{-6} Tesla), consistent with the predictions of the dipole model where $|\mathbf{B}| \propto 1/r^3$.

To address this challenge we re-scale the input signals by taking their cubic root, $f(B) = \sqrt[3]{B} \propto 1/r$. The alteration in distribution is visible in Fig. 3, where the values of $\sqrt[3]{B}$ is distributed on a wider range and more balanced. In general, an MLP benefits from having inputs that are more uniformly spread across the same order of magnitude.

3) Training

We train our neural networks using randomly sampled data within a cubic volume measuring $0.2 \times 0.2 \times 0.15 \text{m}^3$, where the sensor array covers the bottom face. The magnets' orientations are sampled as points on a unit sphere and paired with their locations to serve as training labels, as outlined in Alg. 1.

For our loss function, we employ a weighted sum of the positional and orientational differences:

$$\mathcal{L} = \|\mathbf{p}_{\text{true}} - \mathbf{p}_{\text{pred}}\|^2 + \eta \|\frac{\mathbf{o}_{\text{true}}}{\|\mathbf{o}_{\text{true}}\|} - \frac{\mathbf{o}_{\text{pred}}}{\|\mathbf{o}_{\text{pred}}\|}\|^2 \quad (3)$$

Our goal is to minimize both the positional and orientational discrepancies between the predictions and ground truth. We introduce the weight η to balance the differing scales of orientation and position error terms. After parameter tuning, we set $\eta = 10^{-5}$ as the standard in all experiments.



Fig. 3: Histogram of input values before and after adding feature-engineering function.



Fig. 4: Hall sensor array in a 4x4 grid. The center-to-center distance is 52mm. The array is connected to a Jetson Nano for the complete inference pipeline. The inset shows a rigid-tree to collect Optitrack groundtruth.

C. Implementation

1) Software

We implement the MLP in python with PyTorch. We use Adam [35] with an initial learning rate $\gamma = 10^{-4}$. We train for 40 epochs, and the learning rate decays by 0.98 after each epoch in training. We generate 10^6 random points per epoch to train the model with a batch size of 256. Using Alg. 1 we generate the training data independently and identically distributed. Data generation and training of the MLP takes ~ 1 hour on a standard desktop.

2) Hardware

Our hardware's foundation consists of 16 triaxial magnetometers (MLX90393, Melexis). These sensors feature a linear range up to 0.05 Tesla, aligning well with the magnetic fields we anticipate from our test magnets at close range. These sensors can achieve a data output rate of 716.9 Hz.

Moreover, we calibrate our system using the ellipsoid-fitting method referenced in [36]. The arrangement of the 16 sensors follows a 4 x 4 grid pattern, maintaining a 52 mm interval between the centers of adjacent sensors, as depicted in Fig. 4.

To maintain minimal instrumentation and facilitate portable applications, such as prosthetics, we have implemented the tracking inference on an AI-oriented single-board computer (Jetson Nano, NVIDIA). This device weighs only 138 grams



Fig. 5: In the iterative method, the magnet is initialized at \mathbf{p}_{init} with magnet momentum direction aligned with vector \mathbf{o}_{init} , while the true position and direction are \mathbf{p}_{true} and \mathbf{o}_{true} . The initial offset is then $p_{\text{offset}} = |\mathbf{p}_{\text{init}} - \mathbf{p}_{\text{true}}|$, $\theta_{\text{offset}} = \cos^{-1}[\mathbf{o}_{\text{init}} \cdot \mathbf{o}_{\text{true}}/(|\mathbf{o}_{\text{init}}| \cdot |\mathbf{o}_{\text{true}}|)]$. Similarly, we can also compute the init. positions given the true (sampled positions) and an offset. Substituting \mathbf{p}_{init} and \mathbf{o}_{init} with $\mathbf{p}_{\text{optimized}}$ and $\mathbf{o}_{\text{optimized}}$, we can also get the error e_p and e_{θ} .

and also supports direct sensor readings using the I2C communication protocol. The minimum time required to read 16 sensors sequentially is approximately 24 milliseconds.

III. IN-SILICO EVALUATION

In this section, we compare the performance of our neural network-based tracking method with that of an iterative, gradient-descent-based technique. This comparison utilizes simulated sensor readings, allowing us to control experimental variables such as the initialization of the iterative, optimization-based algorithm. We first introduce how we sample data points in order to evaluate both methods. Then we compare our method versus the iterative baseline in terms of accuracy as function of initialization, as well as computational time. Finally, we compare our method trained on FEM data versus trained on data generated with the dipole model.

Sampling Method: For the in-silico evaluation, we randomly sample magnet positions and orientations and compute the corresponding magnetic flux densities. We either use the magnetic dipole model Eq. 2 to sample from or use FEMgenerated data. On one hand, using dipole-model-generated data ensures that the characteristics of the synthetic signals matches the internal physical model used in the optimization method. On the other hand, the FEM-generated data is more realistic. We initialize the optimization method by offsetting the sampled position and orientation. Offsetting the position and orientation means we initializate the iterative methods with the ground truth plus the offset (see Fig. 5). We expect that the a larger offset results in a worse performance for the iterative methods.. We configure the iterative algorithm to stop upon reaching the maximum number of iterations. This process is similar to a real-world application, where we would initialize with the last known position and the newly measured magnetic flux densities, and we have a limited computational budget.

TABLE II: Accuracy results for the iterative method after 50 Iterations. For comparison, we include results obtained via our MLP tracking on the same evaluation data. Our method is not an iterative approach, and does not rely on an initialization. Percentage is with regard to the size of the sensor array (200mm).

		e_p in mm		e_{θ} in degrees	
		median	3 rd quart.	median	3 rd quart
		(%)	- (%)		-
Iter. Op	otimization				
Initial	Initial				
e_p in	e_{θ} in				
mm	de-				
	grees				
80	10	0.65 (0.3)	11.90 (6.0)	1.17	9.98
80	30	2.06 (1.0)	16.91 (8.5)	4.49	29.76
80	90	35.12 (17.6)	86.65 (43.3)	86.00	90.00
80	180	86.27 (43.1)	147.30 (73.7)	179.75	179.99
30	45	0.53 (0.3)	8.50 (4.3)	1.04	28.33
60	45	2.96 (1.5)	22.38 (11.2)	5.78	44.40
90	45	5.82 (2.9)	33.88 (16.9)	12.85	44.82
120	45	10.37 (5.2)	65.97 (33.0)	25.20	45.03
	all	26.80 (13.4)	95.24 (47.6)	29.96	91.34
MLP (d	ours)				
n of se	nsors used				
	4	1.93 (1.0)	2.88 (1.4)	4.96	7.78
	8	1.62 (0.8)	2.25 (1.1)	4.28	6.54
	12	1.42(0.7)	1.89(0.9)	3.38	5.11

A. Influence of initialization on the optimization methods

1.87 (0.9)

3.34

5.06

1.40 (0.7)

16

In our simulation study, we conducted a comparison between our method and an optimization-based technique by analyzing 400 data points, sampled as outlined above. For the optimization process, we utilized PyTorch, configuring the magnet's position and orientation as adjustable parameters. This setup leveraged automatic differentiation for efficient computation, following the approach detailed in [37]. The optimization algorithm was L-BFGS, a quasi-Newton method. L-BFGS optimizes step size through a line search mechanism [38]. For the optimization-based approach we employed the magnetic dipole model as our internal physical model, integrating it as our loss function in a manner consistent with [19]. Our results are presented in two parts: firstly, we assess the accuracy of our method compared to the optimization baseline for various initialization errors; secondly, we compare the tracking performance of our method trained on FEM data versus data generated by the dipole model.

a) Accuracy: The full results of our optimization method are detailed in Fig. 6, showcasing the tracking errors relative to the maximum iteration count across different initialization of position (Fig. 6.a) and orientation (Fig. 6.b). Similar results are shown for the angular error (Fig. 7) For a detailed analysis, we have also provided statistical analysis of the tracking outcomes after 50 iterations in Tab. II. Note that our MLP method is a single inference step, and does not require multiple iterations or an initialization.

Since the non-converging cases in optimization method yield a long tail in error distribution, the average errors of optimization methods are not representative of the perforThis article has been accepted for publication in IEEE Transactions on Magnetics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TMAG.2024.3416729

TRANSACTION ON MAGNETICS , VOL. XX, NO. XX, XXXX 2024



Fig. 6: Positional errors of the iterative method in simulation for different number of iterations. (a) We vary the initial orientation mismatch, keeping a fixed distance of 80 mm to the truth. (b) We vary the initial positional mismatch, keeping a fixed orientation difference of 45° . As our method does not rely on iterations or initialization, it is a single value. The standard deviation of our method is too small to show in the graph.



Fig. 7: Angular errors of the iterative method in simulation for different number of iterations. (a) We vary the initial orientation mismatch, keeping a fixed distance of 80 mm to the truth. (b) We vary the initial positional mismatch, keeping a fixed orientation difference of 45° . While the median angular errors are decreasing in general the difference between the 1st quartile and 3rd quartile are increasing with the number of iterations, especially for scenarios with large initial offsets. This might be due to an increased probability to converge to local minima. We observe that the positional error shows a faster convergence than the angular counterpart, showing that even after 50 iterations, some trajectories might have not converged yet. As our method does not rely on iterations or initialization, it is a single value. The standard deviation of our method is too small to show in the graph.

mance. Instead, we evaluate the results with the median value of errors. Our approach showed improved performance with a median positional error of $e_p = 1.40$ mm and a median orientation error of $e_{\theta} = 3.34^{\circ}$, significantly outperforming the baseline's median positional error of $e_p = 26.80$ mm and orientation error of $e_{\theta} = 29.96^{\circ}$ under nearly all test conditions. As expected, the accuracy of the optimization based approach improved with more iterations. This trend was also observed in relation to the quality of the initialization; more accurate initial conditions led to more precise optimization outcomes. However, our method was only surpassed in scenarios with the highest number of iterations (50) and the optimal initial conditions, specifically at 80 mm of initial positional error and 10° of initial orientation error, along with 30 mm of initial positional error and 45° of initial orientation error.

Decreasing the number of iterations led to a decrease in performance. With ten iterations, the median positional error rose above 200 mm in certain instances, especially when the initial orientation significantly deviated from the target. We observed that even with initial orientations closely aligned to the target (with an offset of 45°), the resulting positional errors on par with scenarios with a large initialization offset.

b) Running Time: Our Multilayer Perceptron (MLP) model delivers outcomes after a single inference step, which includes feature engineering, along with additions and multiplications within its hidden layers. On the other hand, optimization-based approaches necessitate the calculation of second-order gradients relative to the estimated positions at every iteration, proceeding until convergence is achieved or

TRANSACTION ON MAGNETICS , VOL. XX, NO. XX, XXXX 2024



Fig. 8: Inference time comparison between MLP and L-BFGS implementations for a different number of iterations. Iterative methods are strongly dependent on initialization and it takes between 20 and 50 iterations (i.e. > 10 ms) to achieve results similar to those given by a single MLP inference (0.8 ms).

the maximum number of iterations achieved. We evaluated the speed differences between these two methodologies. To ensure a fair comparison, we executed both methods on the same hardware, specifically a laptop powered by an Intel i7-7500U CPU.

Fig. 8 illustrates the comparison between the total processing time of our data-driven MLP approach and the optimization-based method across various maximum iteration counts. A single inference operation for 5-degree-of-freedom (DoF) tracking using the MLP is completed in just 0.8 milliseconds (ms), this time frame includes the feature engineering stage. In contrast, the L-BFGS optimization process demands approximately 1 ms per iteration. As shown in Fig. 6, achieving satisfactory outcomes with the optimization method often requires dozens of iterations, depending on how well the process is initialized. Consequently, the optimization approach can take tens of milliseconds to converge. Thus, it is evident that our method surpasses the iterative optimization baseline in terms of speed.

B. Neural Networks trained with FEM vs Dipole Model

Previous works train neural networks on data generated with the approximated analytical model [29], [39]. In contrast, we propose using FEM to generate a dataset and take full advantage of the powerful representation of Neural Networks for non-linear systems. To ensure a fair comparison, we distinguish two scenarios for training the MLP: i) using FEM data, and ii) using the magnetic dipole model directly, which is similar with the model used in the optimization method.

Our evaluation encompasses six different magnet shapes, similar to those in our practical experiments (Sec. IV). Including, a variety of shapes from cylinders to disks, all magnetized along their principal axis, and we also assess performance on a spherical magnet, where the dipole model theoretically provides an exact solution. For each magnet shape, we generated a dataset consisting of 1000 samples.

Fig. 9a presents a comparison of the accuracy between MLPs trained on two synthetic datasets: one created from FEM simulations, as explained in Sec. II-A, and the other based on the magnetic dipole approximation (Eq. 2). Additionally, Fig. 9b illustrates the angular error measurements. Both FEM and

dipole trained methods perform well in terms of the rotational error (c.f. Fig. 7). Overall this is an indication that our FEMbased approach generalizes well across all shapes.

To ensure the validity of our data analysis, we first conducted a Shapiro-Wilk test to confirm the normal distribution of our data. Subsequently, we applied a Student's T-Test to identify significant differences between MLPs trained with dipole- and FEM-generated datasets.

The results clearly demonstrate that our FEM-based approach significantly outperforms the dipole model-based baseline in positional accuracy for cylindrical and disk magnets across all tests (with all p < .005), while showing no significant difference for the spherical magnet (p = 0.95). These results are as expected, since there is no difference between the dipole model and the FEM data for a sphere. We observed greater differences in the results as the shape of the magnet distances itself from the sphere, thereby, reinforcing our hypothesis of the superiority of our method for magnets not shaped like spheres.

IV. EXPERIMENTAL EVALUATION

We compare the results of our MLP tracking method to experimental ground truth data collected with OptiTrack. We not only evaluate positional and orientation accuracy, but also evaluate the computational performance on a lightweight portable computer; enabling application such as prosthetics. We experiment with different numbers of Hall sensors in our system (4, 8, 12, or 16), to evaluate the computational speed versus accuracy trade-off.

A. Experimental Setup

We equipped the permanent magnet with OptiTrack markers for precise tracking, illustrated in the inset of Fig. 4. The setup also features ten OptiTrack cameras, and the cylindrical magnets utilized are identical to those described in Section Sec. III-B for simulation purposes.

In our experiments, we manipulate the magnet above the sensor grid, allowing for free movement and tilting, while ensuring optimal visibility of the optical markers to the cameras. This movement is executed at a pace similar to joystick gaming, providing a dynamic test environment. To address any discrepancies in timing between the magnetic and optical tracking systems, we implement a calibration process that involves adjusting a time-offset variable within the magnetic signal data. This adjustment aims to minimize tracking errors. Furthermore, we utilize Piecewise Cubic Hermite Interpolating Polynomial (PCHIP) interpolation [40], to reconcile any discrepancies in sampling frequencies between the two systems. This approach ensures the synchronization and accuracy of our tracking data across both modalities.

B. Results

a) Accuracy: Fig. 10 presents a comparison of a single trajectory as tracked by both our MLP method and the OptiTrack system, aligned within the same time frame. This specific trajectory involves a magnet with dimensions of 10mm

This article has been accepted for publication in IEEE Transactions on Magnetics. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TMAG.2024.3416729

TRANSACTION ON MAGNETICS , VOL. XX, NO. XX, XXXX 2024



Fig. 9: Comparison of a) positional errors and b) angular errors for MLP neural networks when trained with FEM- or dipole model-generated datasets.

in diameter and 20mm in height. Throughout the tracking period, the deviation between the two trajectories remained below the 4mm threshold for more than half of the duration, with the largest discrepancies occurring during rapid maneuvers but not exceeding 10mm. It's important to note that the discontinuities observed in the OptiTrack trajectories can be attributed to partial occlusion of the infrared markers.

We show (Fig. 11) the positional error as a function of distance from the center of the sensor array, demonstrating the robustness of our tracking method even at extended ranges. Furthermore we see that the rotation error is stable for all distances, once again highlighting the robustness of our method.

In Fig. 12, we detail the tracking performance statistics for five different cylindrical magnets and number of sensors. Consistent with expectations, the use of a greater number of sensors results in reduced tracking errors for all magnets tested. With the entire array of 16 sensors activated, the average errors in both position and orientation typically fall below 4mm and 8° respectively, with the exception of the pole-shaped magnet (with dimensions of 5mm in diameter and 25mm in height).

b) Computational Time: Finally, we evaluated the responsiveness of our tracking method by comparing the performance of inferences executed on the CPU versus the GPU of the Jetson Nano. It's crucial to highlight that sensor data collection was consistently handled by the CPU through the I2C communication protocol, taking an average time of 1.75 ms for each sensor read.

When focusing on the inference aspect, running the process on the CPU took approximately 28 ms, and this duration remained fairly consistent regardless of increases in the input feature size (i.e., number of sensors). On the other hand, executing the inference on the GPU showed a variation in processing times, starting at 10 ms with a setup of 4 sensors and extending up to 15ms when using the full array of 16



Fig. 10: Comparison of trajectories obtained by OptiTrack and MLP. The difference is generally smaller than 4mm



Fig. 11: (a) Positional tracking errors versus distance between the magnet position and the center of the sensor array. (b) Orientational tracking errors versus distance between the magnet position and the center of the sensor array.

sensors. This variation underscores the GPU's ability to efficiently manage larger datasets, albeit with a modest increase in processing time as the number of sensors—and consequently, the dimensionality of the input data—grows.

TRANSACTION ON MAGNETICS, VOL. XX, NO. XX, XXXX 2024



Fig. 12: Experimental positional and orientation errors obtained with MLP tracking, for different numbers of input sensor and magnet shapes.

V. APPLICATION EXAMPLE

To show applicability of our approach we implemented it on a novel haptic in- and output device (Fig. 13) [5]. This implementation consists of eighth Hall sensor configured in a circle on two different level-planes. The positions and angles are tracked directly on the Jetson Nano. The latency is below 40ms and the tracking frequency reaches up to 83Hz, which is more than sufficient for most interactive applications. We refer to the supplementary material for a video of the demonstration.

VI. DISCUSSION

In Sec. III-A, we identified a known limitation of iterative methods: tracking accuracy heavily depends on the initial estimation, particularly the magnet's initial orientation. The optimization often becomes trapped in a local minimum if the initial orientation estimate significantly deviates from the actual value.

However, as Fig. 6 demonstrates, with a reasonably accurate initial orientation (mismatch less than 45°), the convergence is less susceptible to other perturbations. The results after 50 iteration steps, as seen in Tab. II, reveal that the third quartile errors are as significant as the initialization errors. This finding suggests that the optimization method might not reach the global minimum even after 50 epochs.

In contrast, tracking with neural networks is independent of initialization. Tab. II shows that the MLP can surpass the optimization-based method in all cases except those with the most accurate initial estimations and the maximum number of iterations. Furthermore, both the third quartile positional



Fig. 13: A user using a permanent magnet as input device for a haptic system [5]. It is easy to imagine this magnet being embedded in an otherwise passive tool. The sphere on the device houses an haptic electromagnetic actuator.

and orientation errors are consistently within acceptable limits when using MLP, underscoring its stability.

Remarkably, we were able to directly apply the models trained with simulations in experimental tests without adjusting any hyperparameters. However, we observe a simulationto-real gap where performance declines when training on synthetic data and evaluating on real data. This is likely due to sensor and background noise, or because our realworld magnetic markers are not perfect. For future work, we consider incorporating background and sensor noise during neural network training, akin to approaches used in other magnetic tracking systems [41], [42]. Furthermore, we can fine-tune on real-world data. Finally, future work could include investigating additional inputs to the neural network that describe the properties of the realworld magnet (such as a magnetization scalar). These adaptations will likely diminish the simulation-to-reality gap, enhancing the accuracy of the MLP method.

In Sec. III-B, we examined the impact of utilizing training datasets generated either through FEM or the magnetic dipole model. As anticipated, models trained with the dipole method showed performance more closely aligned with FEM-trained models when the magnet's shape was closer to or exactly a sphere. For cylindrical magnets, models trained with FEM simulations improved tracking accuracy by 0.2 mm to 1.2 mm. However, we also noted that the tracking performance of the MLP degraded as the magnet shapes deviated further from a spherical form (see Fig. 12). The bigger errors in pole-shaped magnet might be because of the relatively small magnetic momentum. The volume of the magnet is smaller, so the magnetic field around it is weaker. Consider the same level of noise for sensors, the signal-noise-ratio (SNR) for the pole-shaped magnet is smaller and leads to worse performance.

We observed that the total time for one-shot inference in MLP was comparable to each iteration of the many required by iterative methods. Notably, the MLP can be activated sporadically on demand, without the need to continuously track and lock the target to ensure correct convergence within a few iteration steps. We also demonstrated the feasibility of implementing an MLP tracking algorithm on a portable, energy-constrained device, such as the Jetson Nano. We found that the inference time using the GPU was about the same as

reading 8 sensors via the I2C protocol on the CPU. Thus, the sensor reading process currently limits the refresh rate of our prototype. Alternative protocols like SPI could potentially alleviate this bottleneck.

A limitation of data-driven methods is the requirement to retrain the neural network for each new condition, such as different numbers and placements of sensors or changes in the magnet's shape. Although our training process, including data generation, takes only about 1 hour, this requirement could hinder applications that necessitate online optimization of sensor location. The training process could be facilitated with certain trade-off on data variety by storing and reusing generated data points in each epoch of training, rather than generating new data. Other promising directions for future research include employing neural networks to track multiple magnets, using neural network predictions to initialize optimization-based methods, and exploring the use of recurrent neural networks to enhance temporal consistency.

VII. CONCLUSION

In this paper, we demonstrated the accuracy and efficiency of using neural networks to predict the location and orientation of magnets directly. We combined 2D FEM-simulated data with a coordinate transformation algorithm to generate synthetic training data on demand for any type of axis-symmetric magnet. The tracking performance of neural networks was stable and did not experience the convergence issues often seen in optimization-based tracking methods. Our experiments also showed that it is feasible to move the tracking algorithms to energy-restricted devices, thereby enabling portable interactive magnetic applications.

REFERENCES

- X. Han, H. Seki, Y. Kamiya, and M. Hikizu, "Wearable handwriting input device using magnetic field," in *SICE Annual Conference* 2007. IEEE, 2007, pp. 365–368.
- [2] Y. Ma, Z.-H. Mao, W. Jia, C. Li, J. Yang, and M. Sun, "Magnetic hand tracking for human-computer interface," *IEEE Transactions on Magnetics*, vol. 47, no. 5, pp. 970–973, 2011.
- [3] K.-Y. Chen, S. N. Patel, and S. Keller, "Finexus: Tracking precise motions of multiple fingertips using magnetic sensing," in *Proceedings* of the 2016 CHI Conference on Human Factors in Computing Systems, 2016, pp. 1504–1514.
- [4] J. McIntosh, P. Strohmeier, J. Knibbe, S. Boring, and K. Hornbæk, "Magnetips: Combining fingertip tracking and haptic feedback for around-device interaction," in *Proceedings of the 2019 CHI Conference* on Human Factors in Computing Systems, 2019, pp. 1–12.
- [5] T. Langerak, J. J. Zárate, D. Lindlbauer, C. Holz, and O. Hilliges, "Omni: Volumetric sensing and actuation of passive magnetic tools for dynamic haptic feedback," 2020.
- [6] D. Chen, M. Wang, C. He, Q. Luo, Y. Iravantchi, A. Sample, K. G. Shin, and X. Wang, "Magx: Wearable, untethered hands tracking with passive magnets," in *Proceedings of the 27th Annual International Conference* on Mobile Computing and Networking, 2021, pp. 269–282.
- [7] T. Langerak, J. J. Zárate, V. Vechev, D. Lindlbauer, D. Panozzo, and O. Hilliges, "Optimal control for electromagnetic haptic guidance systems," in *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, 2020, pp. 951–965.
- [8] A. Abler, J. Zárate, T. Langerak, V. Vechev, and O. Hilliges, "Hedgehog: Handheld spherical pin array based on a central electromagnetic actuator," in 2021 IEEE World Haptics Conference (WHC). IEEE, 2021, pp. 133–138.
- [9] S. Krueger, H. Timinger, R. Grewer, and J. Borgert, "Modality-integrated magnetic catheter tracking for x-ray vascular interventions," *Physics in Medicine & Biology*, vol. 50, no. 4, p. 581, 2005.

- [10] M. Salerno, R. Rizzo, E. Sinibaldi, and A. Menciassi, "Force calculation for localized magnetic driven capsule endoscopes," in 2013 IEEE International Conference on Robotics and Automation. IEEE, 2013, pp. 5354–5359.
- [11] Z. Sun, L. Maréchal, and S. Foong, "Passive magnetic-based localization for precise unterhered medical instrument tracking," *Computer methods and programs in biomedicine*, vol. 156, pp. 151–161, 2018.
- [12] A. Z. Taddese, P. R. Slawinski, M. Pirotta, E. De Momi, K. L. Obstein, and P. Valdastri, "Enhanced real-time pose estimation for closed-loop robotic manipulation of magnetically actuated capsule endoscopes," *The International journal of robotics research*, vol. 37, no. 8, pp. 890–911, 2018.
- [13] X. Luo, C. Shi, H.-Q. Zeng, H. C. Ewurum, Y. Wan, Y. Guo, S. Pagnha, X.-B. Zhang, Y.-P. Du, and X. He, "Evolutionarily optimized electromagnetic sensor measurements for robust surgical navigation," *IEEE Sensors Journal*, vol. 19, no. 22, pp. 10859–10868, 2019.
- [14] C. R. Taylor, S. S. Srinivasan, S. H. Yeon, M. O'donnell, T. Roberts, and H. M. Herr, "Magnetomicrometry," *Science robotics*, vol. 6, no. 57, p. eabg0656, 2021.
- [15] S. Song, H. Ge, J. Wang, and M. Q.-H. Meng, "Real-time multi-object magnetic tracking for multi-arm continuum robots," *IEEE Transactions* on *Instrumentation and Measurement*, vol. 70, pp. 1–9, 2021.
- [16] F. S. Parizi, E. Whitmire, and S. Patel, "Auraring: Precise electromagnetic finger tracking," *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 3, no. 4, pp. 1–28, 2019.
- [17] S. Hashi, M. Toyoda, S. Yabukami, K. Ishiyama, Y. Okazaki, and K. I. Arai, "Wireless magnetic motion capture system for multi-marker detection," *IEEE transactions on magnetics*, vol. 42, no. 10, pp. 3279– 3281, 2006.
- [18] D. Son, S. Yim, and M. Sitti, "A 5-d localization method for a magnetically manipulated unterhered robot using a 2-d array of halleffect sensors," *IEEE/ASME Transactions on Mechatronics*, vol. 21, no. 2, pp. 708–716, 2015.
- [19] C. R. Taylor, H. G. Abramson, and H. M. Herr, "Low-latency tracking of multiple permanent magnets," *IEEE Sensors Journal*, vol. 19, no. 23, pp. 11 458–11 468, 2019.
- [20] H.-C. Kuo, R.-H. Liang, L.-F. Lin, and B.-Y. Chen, "Gaussmarbles: spherical magnetic tangibles for interacting with portable physical constraints," in *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, 2016, pp. 4228–4232.
- [21] R.-H. Liang, K.-Y. Cheng, C.-H. Su, C.-T. Weng, B.-Y. Chen, and D.-N. Yang, "Gausssense: attachable stylus sensing using magnetic sensor grid," in *Proceedings of the 25th annual ACM symposium on User interface software and technology*, 2012, pp. 319–326.
- [22] K.-Y. Chen, K. Lyons, S. White, and S. Patel, "utrack: 3d input using two magnetic sensors," in *Proceedings of the 26th annual ACM symposium* on User interface software and technology, 2013, pp. 237–244.
- [23] C. Hu, M. Q. Meng, and M. Mandal, "Efficient magnetic localization and orientation technique for capsule endoscopy," in 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005, pp. 628–633.
- [24] T. D. Than, G. Alici, H. Zhou, and W. Li, "A review of localization systems for robotic endoscopic capsules," *IEEE Transactions on Biomedical Engineering*, vol. 59, no. 9, pp. 2387–2399, 2012.
- [25] M. Yousefi, H. Nejat Pishkenari, and A. Alasty, "A fast and robust magnetic localization technique based on elimination of the orientation variables from the optimization," *IEEE Sensors Journal*, vol. 21, no. 19, pp. 21 885–21 892, 2021.
- [26] A. J. Petruska and J. J. Abbott, "Optimal permanent-magnet geometries for dipole field approximation," *IEEE Transactions on Magnetics*, vol. 49, no. 2, pp. 811–819, 2013.
- [27] N. Sebkhi, N. Sahadat, S. Hersek, A. Bhavsar, S. Siahpoushan, M. Ghoovanloo, and O. T. Inan, "A deep neural network-based permanent magnet localization for tongue tracking," *IEEE Sensors Journal*, vol. 19, no. 20, pp. 9324–9331, 2019.
- [28] C. Watson and T. K. Morimoto, "Permanent magnet-based localization for growing robots in medical applications," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2666–2673, 2020.
- [29] S. Su, S. Yuan, M. Xu, H. Gao, X. Yang, and H. Ren, "Amagposenet: Real-time 6-dof magnet pose estimation by dual-domain few-shot learning from prior model," *IEEE Transactions on Industrial Informatics*, 2023.
- [30] A.-i. Sasaki and E. Ohta, "Magnetic-field-based position sensing using machine learning," *IEEE Sensors Journal*, vol. 20, no. 13, pp. 7292– 7302, 2020.

TRANSACTION ON MAGNETICS , VOL. XX, NO. XX, XXXX 2024

- [31] A. Mendizabal, P. Márquez-Neila, and S. Cotin, "Simulation of hyperelastic materials in real-time using deep learning," *Medical image analysis*, vol. 59, p. 101569, 2020.
- [32] E. Haghighat, M. Raissi, A. Moure, H. Gomez, and R. Juanes, "A deep learning framework for solution and discovery in solid mechanics," *arXiv* preprint arXiv:2003.02751, 2020.
- [33] G. Psuj, "Multi-sensor data integration using deep learning for characterization of defects in steel elements," *Sensors*, vol. 18, no. 1, p. 292, 2018.
- [34] I. Malkiel, M. Mrejen, A. Nagler, U. Arieli, L. Wolf, and H. Suchowski, "Plasmonic nanostructure design and characterization via deep learning," *Light: Science & Applications*, vol. 7, no. 1, pp. 1–8, 2018.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2017.
- [36] T. Pylvänäinen, "Automatic and adaptive calibration of 3d field sensors," *Applied Mathematical Modelling*, vol. 32, no. 4, pp. 575–587, 2008.
- [37] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, "Automatic differentiation in machine learning: a survey," *Journal of machine learning research*, vol. 18, 2018.
- [38] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Mathematical programming*, vol. 45, no. 1, pp. 503–528, 1989.
- [39] S. Song, B. Li, W. Qiao, C. Hu, H. Ren, H. Yu, Q. Zhang, M. Q.-H. Meng, and G. Xu, "6-d magnetic localization and orientation method for an annular magnet based on a closed-form analytical model," *IEEE Transactions on Magnetics*, vol. 50, no. 9, pp. 1–11, 2014.
- [40] K. Brodlie and S. Butt, "Preserving convexity using piecewise cubic interpolation," *Computers & Graphics*, vol. 15, no. 1, pp. 15–23, 1991.
- [41] G. Shao, Y. Tang, L. Tang, Q. Dai, and Y.-X. Guo, "A novel passive magnetic localization wearable system for wireless capsule endoscopy," *IEEE Sensors Journal*, vol. 19, no. 9, pp. 3462–3472, 2019.
- [42] Q. Li, Z. Shi, Z. Li, H. Fan, G. Zhang, and T. Li, "Magnetic object positioning based on second-order magnetic gradient tensor system," *IEEE Sensors Journal*, vol. 21, no. 16, pp. 18 237–18 248, 2021.

Mengfan Wu received his B.E. degree in Telecommunication from Tongji University, Shanghai, China in 2018 and the M.S. degree in Electrical Engineering and Information Technology from ETH Zürich, Switzerland in 2021. He is currently with the Advanced Wireless Technology Lab of Huawei Munich Research Center and is pursuing a Dr.Ing. degree at the Telecommunication Networks Group of TU Berlin.

Thomas Langerak received the B.Sc. in Industrial Design from Eindhoven University of Technology, Eindhoven, the Netherlands in 2016 and the doubledegree M.Sc. in Computer Science from Aalto University, Helsinki, Finland, and Twente University, Enschede, the Netherlands in 2018. He is currently a Ph.D. candidate in the Advanced Interactive Technologies group at ETH Zürich, Zürich, Switzerland. His research interests include HCI and haptics.

Otmar Hilliges was awarded the Diplom in Computer Science from Technische Universität München and the PhD in Computer Science from LMU München, Germany in 2009. Then he joined the Microsoft Research Cambridge, in the I3D group, first as a postdoc and later as a researcher. Since 2013 he is with ETH Zürich. He is a Full Professor in the Computer Science Department, leads the AIT lab and serves as the head of the Institute of Intelligent Interactive Systems. His research covers machine perception, computer vision and machine learning.

Juan Zarate completed his M.Sc. and Ph.D. in physics at Instituto Balseiro, Bariloche, Argentina, working on microfabricated infrared sensors. In 2014, he joined the Soft Transducers Lab at EPFL, Switzerland, to work on haptic displays. Since 2018, he is part of the Advanced Interactive Technologies group at ETH Zürich, Zürich, Switzerland. His current research covers volumetric video, haptics, and HCI.