

Polaris: Accurate, Vision-free Fiducials for Mobile Robots with Magnetic Constellation

Jike Wang[†], Yasha Iravantchi[◇], Alanson Sample[◇], Kang G. Shin[◇], Xinbing Wang[†], Dongyao Chen[†]

[†] Shanghai Jiao Tong University, [◇] University of Michigan, Ann Arbor
{jikewang,xwang8,chendy}@sjtu.edu.cn,{yiravan,apsample,kgshin}@umich.edu

ABSTRACT

Fiducial marking is indispensable in mobile robots, including their pose calibration, contextual perception, and navigation. However, existing fiducial markers rely solely on vision-based perception which suffers such limitations as occlusion, energy overhead, and privacy leakage.

We present Polaris, the first vision-free fiducial marking system, based on a novel, full-stack magnetic sensing design. Polaris can achieve reliable and accurate pose estimation and contextual perception, even in NLOS scenarios. Its core design includes: (1) a novel digital modulation scheme, Magnetic Orientation-shift Keying (MOSK) that can encode key information like waypoints and coordinates with passive magnets; (2) a robust and lightweight magnetic sensing framework to decode and localize the magnetic tags. Our design also equips Polaris with three key features: sufficient encoding capacity, robust detection accuracy, and low energy consumption. We have built an end-to-end system of Polaris and tested it extensively in real-world scenarios. The testing results have shown Polaris to achieve an accuracy of up to 0.58 mm and 1° in posture estimation with a power consumption of only 25.08 mW.

CCS CONCEPTS

• **Hardware** → **Sensor devices and platforms.**

KEYWORDS

Magnetic sensing; fiducial system; magnetometer

*Dongyao Chen is the corresponding author.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ACM MobiCom '24, November 18–22, 2024, Washington D.C., DC, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0489-5/24/09.

<https://doi.org/10.1145/3636534.3690711>




Figure 1: (A) shows the magnetic sensing array of Polaris on different robotic systems; (B) and (C) compare the legacy ground-mounted vision-based tags with Polaris tags. Note they have the same *encoding capacity* and *posture calibration capability*.

ACM Reference Format:

Jike Wang[†], Yasha Iravantchi[◇], Alanson Sample[◇], Kang G. Shin[◇], Xinbing Wang[†], Dongyao Chen[†]. 2024. Polaris: Accurate, Vision-free Fiducials for Mobile Robots with Magnetic Constellation. In *International Conference On Mobile Computing And Networking (ACM MobiCom '24)*, November 18–22, 2024, Washington D.C., DC, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3636534.3690711>

1 INTRODUCTION

Fiducial markers are vital for mobile robots [31, 33, 39, 47]. They serve two essential purposes: 1) estimating and calibrating the robot's posture and 2) identifying the marker information, e.g., the marker ID and concise contextual messages. Their efficacy has led to their wide deployment in various mobile "robots," including warehouse robots [20, 28, 29], household vacuum robot cleaners [18, 19].

However, detecting and tracking fiducials usually relies on vision sensors like cameras and LiDARs [45]. This leads to the following three limitations of modern fiducial systems. **Reliability.** Cameras suffer from occlusion that can severely limit their use in practice. This has been a major hindrance for fiducials [51, 66]. The visual markers are also susceptible to diverse disturbances to their visibility. For example, contamination on the factory floor incurs navigation errors to warehouse robots, thus severely limiting their efficiency [28].

Energy overhead. Vision-based fiducials require substantial computation resources, even with low-power cameras integrated [30]. This is a pressing issue as robotic systems are often powered by batteries and subject to diverse constraints [56]. For example, the localization of a microrobot (40 cm \times 7 cm) that employs vision-based fiducials consumes a power of 9 W [32]. In contrast, miniature robots usually operate with a power budget of $< 30\text{mW}$ for their perception [23, 35, 38, 50]. Due to their computation and energy constraints, they primarily depend on proximity sensors (e.g., infrared (IR), ultrasonic) to tackle critical tasks such as inspecting confined pipelines [40] and searching for survivors in disaster debris [22, 46], lacking the ability of pose estimation and contextual perception.

Privacy concerns. Detecting fiducial markers with a camera could become a major source of privacy leaks [65]. For example, today's household robots (e.g., Amazon Astro) need to frequently access private areas like bedrooms and bathrooms. Therefore, a robot equipped with an active camera can cause severe privacy threats to the users.

We present *Polaris*, the first vision-free fiducial system. It achieves posture estimation and information encoding without using a camera. As shown in Fig. 1, *Polaris* consists of two core modules: (1) a novel *magnetic fiducial tag* using passive magnets with a compact form factor and high data-encoding capacity; (2) a robust *magnetic sensing framework* for interpreting fiducials with limited energy consumption. Upon detecting the magnetic tag by the sensing module, the robot can decode the embedded information and estimate the robot's orientation and position. *Polaris* is designed to provide three key features of a fiducial system: sufficient encoding capacity, robust detection accuracy, and low power consumption. To achieve these features, we must tackle two unique challenges as follows:

Compact and high-capacity tag design. The *Polaris* tag requires a compact design to support various robotic systems with different form factors, ranging from medium-sized robots (e.g., warehouse robots and vacuum robot cleaners) to miniature robots. To achieve this, *Polaris* employs a high-capacity tag design with a well-structured constellation of passive magnets (Sec. 5). In particular, a *Polaris* tag adopts a 2D tag layout similar to a chessboard, where each intersection point symbolizes a potential location for placing a magnet. To enable posture estimation, three magnets are positioned at three vertices to form a robust position-detection pattern. To embed rich information, we devise a unique digital modulation scheme called *Magnetic Orientation-shift Keying* (MOSK), which leverages magnets' polarity orientation and placement. Specifically, MOSK uses the polarity orientation of a diametrically magnetized disc magnet to encode data. With MOSK, a *Polaris* tag can achieve the same level of encoding capacity as existing visual fiducial systems with a

compact layout. For example, as shown in Fig. 1(C), a $3 \times 3\text{cm}^2$ *Polaris* tag consisting of 9 magnets can encode 36 bits of data. This encoding capacity is comparable to the tag family "36h11" in AprilTag 2 [60], recommended for most robotic applications. As we will elaborate in Sec. 7.3, such high encoding capacity also empowers the error correction capability of *Polaris* to ensure reliable detection, e.g., by adjusting the minimum Hamming distance [36] between two tags.

Robust and energy-efficient sensing framework. The sensing framework of *Polaris* must accurately interpret these compact magnetic fiducials, even in complex environmental conditions. Our sensing framework features a robust sensing pipeline and an energy-efficient hardware design. For the sensing pipeline, we propose a lightweight and accurate algorithm to effectively detect and localize each magnet with millimeter-level accuracy. We have implemented system-on-chip (SoC) support, i.e., integrate *Polaris* on low-power chips like ESP32, for embedded magnetic sensing. For the hardware design, we utilize low-cost, energy-efficient Hall-effect magnetometers. We also refine the hardware configuration (e.g., the circuit) and minimize the number of magnetometers to further reduce energy consumption. Our evaluation in Sec. 8.5 demonstrates that the sensor array of $2.5\text{cm} \times 1.2\text{cm}$, comprising three triaxial magnetometers, consumes a power of only 25.08 mW.

This work assesses *Polaris*' usability across two distinct robotic platforms varying in size and sensing modality, including a robot car and a miniature car, as shown in Fig. 1(A). We have built an end-to-end system of *Polaris* on the robot car and evaluated its performance under varying scenarios. We have thoroughly analyzed the resilience of *Polaris* against various real-world noises, including magnetic, non-magnetic, and ferromagnetic objects. The *Polaris* sensing framework has been incorporated into a low-power embedded chip (an ESP32-S3 SoC) to demonstrate its compatibility with miniature robots. The power consumption of *Polaris*' sensing array and computing unit have been analyzed, and these evaluation results demonstrate the practicality and resilience of *Polaris* as a fiducial system for mobile robots in real-world scenarios. To help reproduce the results of *Polaris*, we open-source all of its technical details in [17] for the research community.

In summary, this paper makes the following contributions:

- Presenting *Polaris*, the first vision-free fiducial system;
- Introducing MOSK — a new data encoding scheme that embeds rich information into a compact magnetic tag;
- A lightweight magnetic sensing module to decode the embedded information and estimate the robot's posture at millimeter-level accuracy; and
- Extensive experimental studies to demonstrate *Polaris*' performance in real-world settings.

2 BACKGROUND AND MOTIVATION

2.1 Fiducial Systems

Ground-mounted fiducial markers are essential for the localization of mobile robots. These markers can provide precise landmarks such as waypoints and coordinates. Robots use them to extract contextual information, calibrate their posture, and plan their paths. This capability differentiates fiducial systems from other perception techniques like Visual-Inertial Odometry (VIO), Global Positioning System (GPS), and Simultaneous Localization and Mapping (SLAM).

Vision-based fiducials have been widely used due to their accuracy [39]. A visual fiducial system typically comprises printed visual markers and a camera. One can embed vital contextual information, such as marker ID and waypoints, in a fiducial tag. By using the relative position of the fiducial with the camera reference frame, robots can then perform accurate posture estimation using computer vision methods. For example, upon detecting an AprilTag [47], a robot can estimate its posture by using homography estimation [37], i.e., finding the correlation between the tag coordinate system and the 2D image coordinate system. This feature is the cornerstone of various robotic applications. For example, Amazon’s warehouse robots [29, 34] navigate by utilizing ground-mounted 2D barcodes and handle billions of shipments every year [6]. Existing robotic systems also integrate visual fiducials to perform precise docking tasks for recharging their battery and transmitting data [10, 18, 19].

However, these visual fiducial systems are susceptible to occlusion and limited visibility. For example, the detection performance of AprilTag could decrease to 76%, with only a 4.38% occlusion rate according to [51]. We also assessed the detection performance of an AprilTag (ID “0”) across four occlusion scenarios with visual disturbances, e.g., dark dust, as shown in Fig. 2. The weights of dark dust used for light, moderate, and heavy occlusion were only 30 mg, 70 mg, and 100 mg, respectively. An Astra Plus camera [16] was stably mounted directly above the tag with a distance of 20 cm. We employed the standard AprilTag detection algorithm [8] with default parameters. The average detection accuracies are 1, 0.89, 0.42, and 0, respectively. Our results indicate that existing visual fiducials are highly susceptible to occlusion. Note that heavy occlusion (i.e., 0.1 g of dark dust) can fully paralyze the standard fiducial system despite the image features remaining visible to humans.

Current vision-based methods also require significant computational resources for detecting and localizing visual fiducials. For example, achieving 640×480 detection of AprilTag at 20 Hz on a Raspberry Pi 3 was feasible, yet the CPU usage exceeded 150% according to [39]. This limits their practicability for diverse robotic configurations, e.g., resource-constrained robots. SOTA ultra-low power cameras, such as




Figure 2: Occlusion test for AprilTag.

the HM01B0 [12], usually offer limited resolution (320×320). Our results in Fig. 2 demonstrated that even high-resolution cameras, i.e., 1920×1080 , failed to detect visual fiducials at a close range. Moreover, using an always-on camera to detect fiducial markers accompanies privacy risks [65]. The above limitations undermine the performance and usability of legacy fiducial systems in the real world. Thus, a vision-free and energy-efficient fiducial system is crucial for enabling the robust and precise localization of mobile robots.

2.2 Why Magnetic Sensing?

Magnetic sensing uses MEMS magnetometers to sense the magnetic field. These sensors leverage the Hall effect [49] to measure the strength and orientation of the magnetic field. Compared with other sensing modalities, magnetic sensing has the following advantages:

(1) **Reliability.** Unlike vision-based methods, the magnetic field is resilient to the NLOS occlusion. Moreover, the magnetic field can penetrate water and concrete, making it well-suited for harsh environments. Passive magnets are also durable, allowing for deployment in challenging environments for critical tasks [61]. (2) **Energy-efficiency.** Hall-effect sensors incur low power consumption. For example, COTS magnetometers typically have a current consumption of less than 1 mA [9, 14]. (3) **Privacy protection.** Unlike cameras, magnetic sensing relies on detecting and measuring the magnetic field. This makes it well-suited for sensing tasks where privacy protection is critical.

3 RELATED WORKS

3.1 Fiducial Marking System

Fiducial markers are used to provide position references and posture estimation for mobile robots. Existing fiducial marker systems commonly use cameras thanks to their availability at low cost. 2D barcodes, e.g., Data Matrix [7] and QR Code [55], have been employed as fiducial markers for robots in industrial settings [6, 34]. ARTag [33] achieves robust marker detection even under partial occlusion, providing a data payload of 11 bits. AprilTag [47], building upon the framework of ARTag, provides enhanced robustness against occlusions and distortions using a graph-based segmentation algorithm. It offers a data payload ranging from 4 to 16 bits.

However, these visual fiducial systems rely heavily on onboard cameras and computer vision techniques for marker




Figure 3: System overview of Polaris.

detection. As we have elaborated in Sec. 2.1, they suffer from real-world issues, including reliability, energy efficiency, and privacy concerns. In contrast, Polaris is the first vision-free fiducial marker system that provides rich data encoding capability and accurate posture estimation for robots.

3.2 Magnetic Sensing

Sensing magnetic fields has been used extensively for object positioning and tracking applications. For example, Finexus [26] utilizes a triangulation-based algorithm to localize electromagnets nearby. However, this approach relies on each electromagnet generating oscillatory magnetic fields, thus requiring a battery pack. Recently, the Levenberg–Marquardt (LM) algorithm [25, 54] has been employed to track passive magnets over short distances. However, the iterative optimization process incurs quadratic computational load.

Magnetic sensing is also used for robot navigation. For example, automated guided vehicles (AGVs) navigate along a predetermined path constructed with magnetic tapes by using an onboard magnetometer ruler to detect the magnetic field [21]. However, this approach cannot provide contextual information, e.g., the robot’s precise location and orientation. To address this problem, magnetic spots are employed to enhance precise localization and enable flexible movement paths [28]. However, these systems are constrained by a limited localization accuracy (i.e., around 11 cm) and are ill-suited for encoding rich context information.

4 GOALS OF POLARIS

We explore the feasibility of employing magnetic sensing in a vision-free and robust fiducial system for mobile robots. Polaris tag offers the same key functionalities as classic fiducial markers, i.e., information encoding and posture estimation, with magnetic sensing. Polaris’ unique advantages, i.e., robustness, energy efficiency, and privacy protection, make it an ideal candidate for various challenging scenarios. Specifically, it can complement existing visual markers in challenging scenarios (e.g., occlusion/contamination on the floor of heavy-duty factories). For resource-constrained robotic systems, Polaris can be the standalone module for providing fiducial capability. Polaris is privacy-preserving in various real-world settings. For example, it can help robotic

vacuums [19] dock precisely to their charging stations without using camera(s).

Information encoding. As we will elaborate in Sec. 5, a Polaris tag is designed in a 2D plane, where each intersection symbolizes a potential position for a magnet. By varying both the polarity orientation and the spatial placement of each magnet, Polaris tag can encode rich information. For example, with just four magnets, Polaris can encode 20,480 unique messages, e.g., marker ID “0” to “20479”, similar to existing fiducial mechanisms. To enrich the context information, each ID can denote unique (navigation) instructions, e.g., “0” is “move slowly”. Upon detecting the tag, robots can decode the ID and perform the actions.

Posture estimation. Upon detecting the tag and reconstructing its layout, a robot equipped with Polaris’ sensing array can determine its posture relative to the tag, i.e., the x, y, z coordinates and the heading angle. Our experiment results in Sec. 8.4 demonstrate that Polaris can achieve an accuracy of 0.58 mm and 1° in position and orientation estimation, respectively.

Fig. 3 shows a system overview of Polaris. To facilitate the construction of the magnetic tag (Sec. 5), Polaris allows users to determine the tag’s layout — the number, spacing, and layout of passive magnets. Polaris’ sensing module (Sec. 6) performs three tasks: detect magnets, recognize the polarity and localize magnets.

5 DESIGN OF POLARIS TAG

We first introduce MOSK for information encoding. Next, we analyze Polaris’ encoding capacity.

5.1 Two-dimensional Tag Layout

To support the ground-mounted usability, i.e., reliable posture estimation and rich encoding capacity, we propose a 2D, chessboard tag layout. As shown in Fig. 4(a), each intersection point of the grid serves as a potential position for magnet placement. Here, K denotes the tag order, and d is the distance between two adjacent intersections.

With this 2D layout, Polaris first introduces a robust position-detection pattern tailored for accurate posture estimation. To achieve this, Polaris utilizes the three vertices of the 2D plane. As shown in Fig. 4(b), the bottom-left vertex of a 2D plane is labeled as v_1 , and the subsequent vertices, moving in a clockwise direction, are labeled as v_2, v_3 , and v_4 . The three magnets can be placed at v_1, v_2 , and v_3 . This design resembles the QR code’s finder pattern [55]. We can estimate position by detecting and localizing these three magnets using the sensing pipeline in Sec. 6.2. The triangulation from these three magnets also enables us to determine the heading angle between the robot and the position pattern.




Figure 4: 2D grid layout of Polaris tag.




Figure 5: Polarity configuration of disc magnets.

5.2 Magnetic Orientation-shift Keying

We build a novel encoding scheme called *magnetic orientation-shift keying* (MOSK), which utilizes and augments the polarity configurations of a diametrically magnetized disc magnet to embed information.

5.2.1 Rotational polarity orientation. Encoding information with magnets presents significant challenges owing to their limited data encoding capacity. Conventional approaches typically depend on the North/South (N/S) polarity of a magnet to represent binary bits, i.e., “0” and “1”, such as magnetic stripe card [52]. However, this encoding scheme severely limits the amount of information that can be encoded. For example, eight magnets can only encode eight bits of data.

How can we reliably embed rich information into the tag using a minimal number of magnets? Our key innovation is using the *fine-grained polarity orientation* of passive magnets. Specifically, different magnetic orientations would incur changes in the magnetometer readings, thus ensembling different information. To achieve this goal, we first determine the form factor of passive magnets and then leverage the polarity orientation of each magnet. Considering the deployability and durability of Polaris tag, we use disc-shaped cylindrical magnets. There are two types of disc magnets based on the direction of magnetization – diametrically and axially magnetized magnets – as shown in Fig. 5(a). Compared to the axially magnetized magnet, the diametrically magnetized magnet offers a unique advantage in encoding information: the variability of its polarity orientation. As shown in Fig. 5(b), we can alter the polarity of a diametrically magnetized disc magnet to different orientations to represent different information.

Hence, in MOSK, the polarity orientation of a disc magnet is shifted to represent the digital data. The angular constellation graph in Fig. 6(a) demonstrates *the simplest form of MOSK*. That is, a polarity orientation at 0 degrees (i.e., forward orientation) represents a binary ‘1’, while an orientation




Figure 6: Constellations of MOSK modulation.




Figure 7: Orientation setting and an exemplary tag.

at 180 degrees denotes a binary ‘0’. Figs. 6 (b) and (c) show that by employing more orientation shifts, MOSK can encode more bits of information. Thus, we integrate these orientation configurations into the 2D layout. As shown in Fig. 7(a), the polarity orientation of a disc magnet is denoted as 0° aligned with the orientation of the position pattern. Then, we can rotate the disc magnet clockwise to have different polarity orientations. For example, Fig. 7(b) shows a 3 × 3 tag of eight magnets, each with a unique orientation. This 3-order tag configuration can embed 24-bits based on Fig. 6(c). The encoding capacity is significantly improved compared to the eight bits encoded solely using the N/S pole.

5.2.2 Spatial permutation of magnets. Relying solely on the polarity orientation for encoding may restrict the usability of the Polaris tag. Specifically, a $K \times K$ tag requires $K^2 - 1$ magnets for maximum encoding capacity. This demands an excessive amount of magnets as K increases. Although magnets are low-cost, i.e., less than \$0.015 in our tag prototype (Sec. 7), using unnecessary magnets can lead to extra fabrication costs. To solve this problem, MOSK introduces the *spatial permutation* of magnets in the 2D plane. This encoding factor enhances the encoding capacity, allowing Polaris to encode rich information while using a minimal number of magnets. Specifically, three vertices are fixed as the position-detection pattern (Fig. 4(b)). Now, one can only manipulate the polarity orientation of each magnet in the vertex. The remaining $K^2 - 4$ intersection points can allocate a maximum of $K^2 - 4$ magnets. Thus, for a specific number of magnets denoted by M , there are a total of $\binom{K^2-4}{M-3}$ placement permutations for the remaining $M - 3$ magnets. Using this scheme, Polaris can achieve a customizable tag design, depending on the required information capacity.

5.2.3 Minimal inter-magnet distance. To achieve a compact tag design, we aim to determine the distance of two adjacent




Figure 8: Coupling effect of two adjacent magnets.

intersections, i.e., the value of d in Fig. 4(a). This value determines the minimal distance between two adjacent magnets in a tag. The key challenge is: reducing d introduces undesired magnetic field interference due to the *coupling effect* between adjacent magnets. We elaborate on the interference incurred by the coupling effect in Fig. 8. Two diametrically magnetized disc magnets (1 mm height, 3 mm diameter) are positioned 5 cm apart. We use a magnetometer that moves across the two magnets at a speed of 10 cm/s and a height of 2 cm. Next, we reduce the inter-magnet distance to 2cm to demonstrate the coupling effect and measure the raw data in the same setting. Fig. 8(a) shows the magnetometer reading without coupling effect. Peaks incurred by two magnets can be differentiated. In contrast, as shown in Fig. 8(b), the coupling issue produces undesired peak overlapping and signal distortion, leading to erroneous interpretation of magnets.

To address the coupling issue, we determine the optimal value of d via simulation. With a specified magnet configuration (e.g., the size and magnetization), we can model the magnetic field created by two adjacent disc magnets using MagPylib [48]. Then, we can determine the minimal inter-magnet distance d_{\min} when no magnetic fields overlap between the two magnets. Next, we explore the feasibility of achieving a more compact tag design by reducing the inter-magnet distance. Our empirical exploration suggests that the minimal inter-magnet distance can be reduced to $0.8d_{\min}$.

5.3 Encoding Capacity Analysis

The encoding capacity of MOSK depends on three factors: the tag order K , the number of rotational polarity configurations, denoted as P , and the number of disc magnets M . Specifically, with M magnets and P orientations of each magnet, we can encode P^M unique messages. The spatial permutations of $M-3$ magnets among the K^2-4 intersection points in the 2D plane can encode $\binom{K^2-4}{M-3}$ messages. Thus, the total encoding capacity can be calculated as:

$$C = P^M \cdot \binom{K^2-4}{M-3}. \quad (1)$$

Based on Eq. (1), the encoding capacity increases **quadratically** with the tag order K . The effects of P and M are plotted in Fig. 9. The extended number of orientations and magnets




Figure 9: Encoding capacity with different factors.




Figure 10: The bar-shaped sensor array.

can **exponentially** enhance the encoding capability. For example, with $P = 16$, a 4-order tag can encode more than 41 bits using only eight magnets, showing the potential of encoding rich information with Polar is.

6 SENSING MODULE DESIGN

6.1 Magnetometer Array

We use COTS Hall-effect magnetometers [9, 14] to detect the magnetic field strength of magnets. Magnetometers are usually low-cost components with a price of \$1.5 for each unit. We integrate several magnetometers into a bar-shaped sensor array for detecting and locating the ground-mounted tag (Fig. 10). This linear form factor enables easy installation on robots. The overall length of the sensor array is denoted by L , and the inter-magnetometer distance is d_s .

6.2 Sensing Pipeline

Polar is first detects each magnet and determines each magnet's polarity orientation. Then, it localizes the relative position of each magnet to the sensor array. Based on these results, Polar is can reconstruct and interpret the 2D tag.

6.2.1 Detect magnet(s). As shown in Fig. 8(a), we can detect the presence of a magnet by detecting the peak signal from the raw data using the existing amplitude-based peak-detection algorithm [58]. The raw magnetometer readings may vary due to environmental factors such as the Earth's magnetic field or the soft/hard iron effect caused by surrounding ferromagnetic objects [62]. To ensure reliable and accurate detection, we employ the first derivative signal of the raw peak [61] to further eliminate false detection(s).

6.2.2 Detecting the polarity orientation. Polar is needs to derive the polarity orientation of each magnet relative to the position-detection pattern of the 2D tag, i.e., the value of θ_m^t in Fig. 11(a). With θ_m^t , Polar is can decode the embedded bits based on the constellation shown in Fig. 6.




Figure 11: The analytical model of polarity orientation.

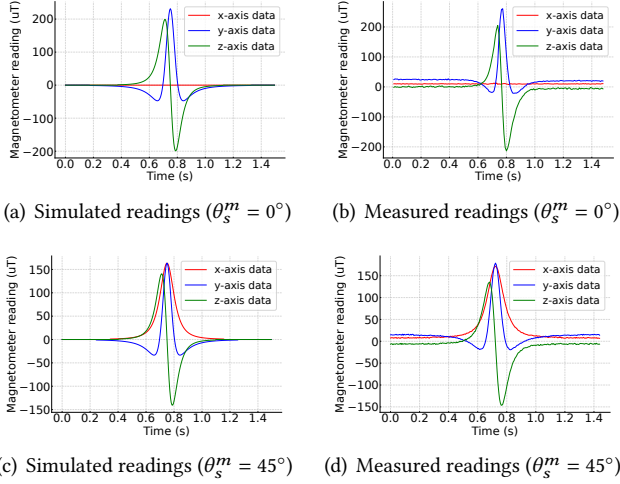


Figure 12: Simulated vs. real-world sensor readings at different headings.

However, calculating θ_m^t is challenging due to the arbitrary headings of a mobile robot. To overcome this challenge, we use two key insights: the relative orientation of the sensor array to the tag (denoted as θ_s^t) and the relative orientation of the sensor array to the magnet (denoted as θ_s^m). As shown in Figs. 11 (b) and (c), once θ_s^t and θ_s^m are derived, we can have $\theta_m^t = \theta_s^t + \theta_s^m$.

How can we accurately derive θ_s^t and θ_s^m ? Note that θ_s^t is the relative heading angle of the tag to the sensor array, which can be determined by the posture estimation process in Sec. 6.2.4. To obtain θ_s^m , we leverage the mathematical model of disc magnets and propose a unique analytical scheme. Specifically, we discover a strong correlation between the theoretical magnetic field derived from the mathematical model [24] and the real-time measured data. Note that the mathematical model has been integrated into widely used simulation tools such as MagPylib [48]. Interested readers can refer to [24] for more details.

We now analyze the correlation between the theoretical magnetic field and the raw measured data using simulation. Figs. 12 (a) and (b) show the simulated readings via MagPylib and the real-time measured data under the same settings, i.e., the same magnet configuration and moving heading. In Fig. 12(b), we observe an offset in each axis between the




Figure 13: The magnetic field template can be generated by varying lateral distances and orientations.

synthesized and measured data caused by the environmental magnetic field, i.e., the earth's magnetic field. By leveraging derivatives, this offset does not affect the template alignment result. This feature helps enable a calibration-free scheme since it is resilient against the magnetometer's intrinsic drift [62], i.e., the data offset and scale distortion in a sensor. Figs. 12 (c) and (d) show the results with a 45° heading, demonstrating that the data exhibits distinct features at different movement headings. Based on this insight, we can derive θ_s^m by aligning the measured data with a template consisting of synthesized data series.

The data series alignment can be implemented using existing algorithms, such as dynamic time warping (DTW) [2], convolutional neural networks (CNNs) [44], and recurrent neural networks (RNNs) [63]. However, DTW can introduce data misalignment due to bias [2, 59] since the magnetic field data varies with location or orientation. CNNs and RNNs require substantial computational resources, thus limiting their efficiency on resource-constrained platforms. To achieve *lightweight* and *highly accurate* data series alignment, we employ a derivative dynamic time warping (DDTW) algorithm [41] and utilize three-axis magnetometer readings. Compared to DTW, the DDTW algorithm uses the derivative of the magnetic field, thus overcoming the offset issue in the sensor reading as shown in Fig. 12, which can provide robust features and eliminate data bias. Combining three-axis data allows the alignment algorithm to differentiate fine-grained polarity orientations (Sec. 8.3.5). Our approach consists of template construction and data series alignment.

Template construction. Based on the analytical model, Polaris first constructs a template by synthesizing the magnetometer data in the offline stage. The synthesis process requires two inputs: the configuration of the magnet (e.g., magnetization and size) and the ground clearance of the robot. Note that these are a *one-time effort* that users can efficiently make in real-world applications. Since robots can be highly mobile, the template must cover various dynamic magnetic patterns to enable accurate alignment. However, a large template requires more computing time, thus limiting its efficiency. To efficiently emulate real-world data traces when a magnet impinges the sensing array, we consider two key parameters, i.e., the lateral distance and the relative heading θ_s^m as shown in Fig. 13.




Figure 14: Movement analysis for localizing a magnet.

Specifically, the template size, i.e., the number of synthesized trajectories are related to the resolution of the lateral distance (denoted as res_d) and the headings (denoted as res_θ). As shown in Fig. 13(a), when detecting a magnet in a 2D tag, the value of lateral distance ranges from $-\frac{d}{2}$ to $\frac{d}{2}$ since the minimal inter-magnet distance is d (Sec. 5). The value of θ_s^m ranges from 0° to 360° (Fig. 13(b)). Thus, the template size can be calculated by:

$$(d/res_d + 1) \cdot (360/res_\theta + 1). \quad (2)$$

According to Eq. (2), fine-grained res_d and res_θ can generate a template that contains detailed trajectories and enhances the alignment accuracy of θ_s^m . However, this setting can be demanding for the computational resource. To tackle this issue, we allow users to adjust the resolution of Polaris for specific robotic systems. In particular, our experimental results indicate that a lateral distance resolution of $d/4$ is sufficient for the alignment process thanks to the compact tag layout. For the value of res_θ , as we will elaborate in Sec. 8, 10° resolution can differentiate tags encoded with 8-MOSK, which leads to encode 2^{24} unique messages with a 3-order tag of eight magnets. At a 1° resolution, Polaris can support 20-MOSK – it enables encoding 2^{35} unique messages.

Data series alignment. We perform precise time-series alignment between the raw data sequence and the template. The number of data points in each data series can be determined by the value of d . Specifically, combining the robot’s velocity v and magnetometer’s sampling rate f , the number of data points $n = f \frac{d}{v}$. To further reduce the time consumption, we implement the DDTW-based alignment algorithm in C. Our experimental results in Sec. 8 indicate that Polaris can decode a 8-MOSK tag with $res_\theta = 10^\circ$, $res_d = 4/d$ and $n = 40$, which only incurs a time delay of <30 ms.

6.2.3 Localize magnets. To estimate the posture and decode information, Polaris needs to localize each magnet in a 2D tag, i.e., the coordinates of each magnet relative to the sensor array. Existing localization schemes in magnetic sensing usually use iterative optimization, such as the Levenberg–Marquardt (LM) algorithm [25, 54]. These approaches are not feasible due to the requirements of quadratic computation complexities. In Polaris, we present a lightweight and highly accurate localization algorithm comprised of

three key steps: movement modeling, lateral position estimation, and longitudinal position estimation.

Movement modeling. We first analyze and model the movement when a sensor array detects a Polaris tag. As shown in Fig. 14(a), we present the geometry between the magnet and Polaris’ sensing array when detecting a magnet, with the y-axis of the coordinate aligned with the magnet’s N pole. The sensor height, i.e., the ground clearance of the robot, is represented as h .

We present an effective method for locating a magnet with two adjacent magnetometers, s_i and s_j . We first determine the hardware requirements for this localization scheme by analyzing the worst-case scenario: localizing two adjacent magnets simultaneously. As shown in Fig. 14(b), s_1 and s_2 are used to localize “Mag 1”, while s_3 , s_4 localize “Mag 2”. We need to determine the inter-sensor distance d_s to ensure no magnetic field superposition on s_2 and s_3 . To solve this problem, we leverage the minimal inter-magnet distance d (Sec. 5). Specifically, for each magnet, the maximal distance that a sensor can sense is $d/2$ based on the minimal inter-magnet distance d . Thus, the maximal inter-sensor distance d_s should be $d/2$. For a K -order tag, $2K$ sensors are needed to simultaneously locate at most K magnets.

Next, we need to determine the position of s_i and s_j relative to the magnet’s coordinate. Here, we only need to determine the 2D positions (x and y axis) of s_i and s_j . This insight helps us convert the 3D localization problem to a 2D position estimation task using two adjacent sensors.

Lateral position estimation. We first derive the lateral position, i.e., the position in the x -axis of the magnet as shown in Fig. 14(a). Specifically, the lateral position of the magnet is determined by the values of d_s^i and d_s^j , which are subject to the constraint $d_s^i + d_s^j = d_s$. The key challenge is the computation of d_s^i and d_s^j with arbitrary headings, as shown in Fig. 14(c). To address this issue, we introduce two critical pieces of information: the θ_s^m derived from the DDTW-based alignment process and the mathematical model in [24]. The core idea is to convert the problem to a nonlinear optimization problem by analyzing the correlation between the measured and theoretical values.

As shown in Fig. 14(c), s_i and s_j can measure the magnetic field strength created by the magnet using the peak-detection algorithm at time t_1 . The orientation between the magnet and the sensor array is θ_s^m , derived from the alignment process. Hence, the coordinates of s_i' and s_j' at t_1 can be calculated by $(d_s^i \cos \theta, d_s^i \sin \theta, -h)$ and $(-d_s^j \cos \theta, -d_s^j \sin \theta, -h)$, respectively. We denote the measured magnetic field strengths by s_i' and s_j' at t_1 as B_m^i and B_m^j . The theoretical magnetic field at $(d_s^i \cos \theta, d_s^i \sin \theta, -h)$ and $(-d_s^j \cos \theta, -d_s^j \sin \theta, -h)$ are denoted as B_t^i and B_t^j . Combining these results, we can have the following nonlinear-constrained optimization problem:




Figure 15: Real-world tests of tag reconstruction steps.

$$\begin{aligned} \min_{d_s^i, d_s^j} & \left(\frac{B_t^i}{B_t^j} - \frac{B_m^i}{B_m^j} \right)^2 + (d_s^i + d_s^j - d_s)^2 \\ \text{s.t.} & \quad 0 \leq d_s^i, d_s^j \leq d_s. \end{aligned} \quad (3)$$

To solve for the two variables d_s^i and d_s^j with Eq. (3), we can use existing optimization algorithms like Nelder-Mead algorithm [43] and SLSQP [42] thanks to their computational efficiency when handling a limited number of variables. Our evaluations in Sec. 8 also indicate that Polaris can efficiently derive the lateral distance in < 5 ms implemented by C. Based on the estimated d_s^i and d_s^j and inter-sensor distance d_s , we can derive each magnet's lateral position relative to the sensor array's origin point, i.e., s_1 .

Longitudinal position estimation. To compute the longitudinal position of each magnet, we can utilize both the robot's velocity and the sensor array's sampling rate. Specifically, we can denote the longitudinal position of the first magnet detected by the sensor array as zero. Then, we can derive the longitudinal position of each magnet relative to the first magnet by $L_i = \frac{1}{f} \sum_{j=0}^{p_i} v_j$. The variables v , f , and p_i are the robot's velocity, sample rate, and the polarity index of i^{th} magnet in the raw data series.

6.2.4 Reconstruct the tag. Once we determine the lateral and longitudinal position of the magnet relative to the sensor array, the tag layout can be reconstructed. Based on the reconstruction process, Polaris can perform posture estimation and information decoding.

Posture estimation. The posture estimation process first identifies the three vertex magnets from the detected magnets and reconstructs the position-detection pattern using triangulation. For the orientation estimation, it compares the reconstructed pattern with the standard pattern using rotation transformation.

Information decoding. With θ_s^t and θ_s^m , we can derive the polarity orientation of each magnet relative to the tag layout, θ_m^t , based on $\theta_m^t = \theta_s^t + \theta_s^m$. Then, Polaris can decode the tag ID by comparing each magnet's estimated polarity orientation and position with the synthesized tag set.

We performed a benchmark test to illustrate the tag reconstruction process. Specifically, we use the 3-order tag with eight magnets as shown in Fig. 7(b). A robot car equipped with the sensor array (Sec. 7) moves over this tag with a




Figure 16: Fabrication process of a Polaris tag.

heading angle of 0° . The reconstruction process is shown in Fig. 15. It first detects each magnet and derives θ_s^m with the polarity alignment algorithm. Then it localizes each magnet's position relative to the sensor array using θ_s^m . Next, it reconstructs the position-detection pattern to estimate the heading angle θ_s^t . Finally, using θ_s^m and θ_s^t , it determines the polarity orientation (i.e., θ_m^t) of each magnet and decodes the embedded information.

7 HARDWARE CONFIGURATION

7.1 Sensing Array

Sensor setup and circuit design. For the magnetometer, we choose MLX90393 [14], a low-cost (i.e., $< \$2$ each) and energy-efficient triaxial magnetometer with a wide sensing range of $5 - 50$ mT. We use a compact Bluetooth Low Energy (BLE) module MDBT42Q-512KV2 [4] as the microcontroller. This module is integrated with a nRF52832 system-on-chip (SoC) that facilitates power-efficient data processing and communication. The microcontroller communicates with the magnetometers via SPI protocol [1] at a clock frequency of 2 MHz. The measured data can be transmitted with a wired connection or a BLE channel, depending on the specific requirements.

Sensor array construction. We have implemented two sensor array prototypes for two robotic systems: a robot car and a miniature car. The robot car's sensor array measures $16.9\text{cm} \times 1.2\text{cm}$ and comprises 9 magnetometers spaced 2 cm apart. For the miniature platform, we built a compact sensor array consisting of three magnetometers with an inter-sensor distance of only 0.8 cm. Its dimension is $2.5\text{cm} \times 1.2\text{cm}$. The costs (including a microcontroller, magnetometers, and manufacturing costs) of these two sensor arrays, are less than $\$25$ and $\$12$, respectively.

7.2 Tag Fabrication

For the disc magnet, we use Neodymium (NdFeB) magnets [3] due to their high magnetic strength and low cost. We determined two specific configurations of magnets for our testing platforms. The magnet for the robot car has a diameter of 2 mm and a height of 0.5 mm. For the miniature car, the diameter of the magnet is 1 mm, and the height is 0.5 mm. The tiny magnet is extremely low-cost, costing $< \$0.015$.

To fabricate the tag, we propose a four-step process illustrated in Fig. 16, allowing users to fabricate with minimal




Figure 17: Different failure cases of a Polaris tag.

effort. First, we label the N/S polarity of each magnet with an ultra-fine marker. Note that the magnet manufacturer can perform the labeling process during production. We use an ultra-thin and durable PVC plane with a thickness of 0.2 mm as the substrate. Then, we draw the K -order tag layout with the ultra-fine marker on this substrate. Finally, we place each magnet at the intersection point on the plane using plastic tweezers and fix it with an industrial-grade glue.

7.3 Robustness of Polaris Tag

Real-world disturbances. Real-world disturbances include non-magnetic and ferromagnetic materials. Magnetic fields can penetrate non-magnetic materials, ensuring reliable detection of Polaris tags in everyday scenarios such as homes. In harsh scenarios such as factories, accumulated ferromagnetic debris (e.g., iron filings) on the ground may distort the magnetic field. To address this issue, the peak detection and polarity alignment algorithm (Sec. 6.2) effectively eliminates the disturbance using reliable derivative signals. The experimental results (Sec. 8.3.6) also demonstrate Polaris can accurately interpret a tag under various disturbances.

Fabrication defects. During placement, the polarity orientation of a small magnet might not align well with the desired orientation. For example, Fig. 17(A) shows a desired tag configuration. Fig. 17(B) shows an inner magnet has an incorrect orientation due to fabrication defects. This error may incur errors in decoding the tag ID, e.g., decoding the desired tag ID to a neighboring tag ID, when using a fine-grained polarity orientation for encoding. Thanks to the encoding capability enabled by MOSK, we can mitigate this issue by sustaining the minimum Hamming distance, i.e., the minimum distance between the encoding bits of two tags.

Damaged tags. Passive magnets are made of rigid ferromagnetic metal that can survive high pressure. For example, a neodymium magnet has a compressive and tensile strength of 137,800 psi and 80 MPa [15], respectively. However, in extreme conditions, it is possible that the magnet(s) of a Polaris tag is damaged. We now analyze the robustness of our tag under such extreme cases. There are two types of failure cases based on the position of the magnet in the 2D tag. In Fig. 17(C), one magnet of the position-detection pattern is damaged. This could incur a failure of the tag reconstruction process in Sec. 6.2, which will be reported to the robot. Fig. 17(D) shows a magnet damaged in the remaining intersections. This damage leads to failure detection of




Figure 18: The experiment setup: (A) shows the settings for the robot car; (B) represents the testing scenario of the miniature car.

the magnet’s polarity orientation and position, resulting in missing bits. To tackle this issue, redundant magnets can be placed at the idle intersections of the 2D tag using the existing forward error correction (FEC) schemes [27], e.g., the Reed-Solomon code [64] used in QR codes. In this case, Polaris only employs the polarity orientation of each magnet to encode data. The intersection points are preserved for redundancy requirements. The error correction level depends on the number of intersection points, which is related to the minimal inter-magnet distance (Sec. 5). A small inter-magnet distance allows the Polaris tag to achieve high-level error correction with a compact layout. Our experimental results in Sec. 8.4 indicate that Polaris can decode a tiny tag with an inter-magnet distance of only 1.6 cm, demonstrating the potential for integrating existing FEC schemes into our tag.

8 EVALUATION

8.1 Experimental Settings

we evaluate Polaris with two representative ground robots — a robot car and a miniature car — as shown in Fig. 18. The robot car has a dimension of 24 cm \times 19 cm, a typical size for robots such as vacuum cleaners (e.g., Roomba), household robots (e.g., Amazon Astro), and warehouse robots. We integrate Polaris into the car’s operating system, i.e., Ubuntu 20.04. The default configuration of the testing tag is $\{K=3, M=5, P=8, d=4 \text{ cm}\}$, and each magnet’s default orientation is 0° . We first evaluate the impact of varying headings. Then, we assess Polaris’ robustness and high encoding capability while varying the polarity orientation, tag density, and template sizes. To demonstrate the usability and practicality of Polaris, we evaluate its performance under various real-world noises and build an end-to-end (e2e) system that controls the robot’s movement trajectory. We demonstrate Polaris’ low-power and compact features with the miniature car scenario. First, we show that Polaris’ sensing pipeline can be integrated with low-power nodes such as ESP32 S3 chip. Then, we validate Polaris’ efficacy of detecting a tiny tag. Specifically, we build a thumbnail-sized tag, i.e., $1.6 \times 1.6 \text{ cm}^2$, consisting of three magnets. We also measure the power consumption of Polaris, including the hardware front end and computational components.

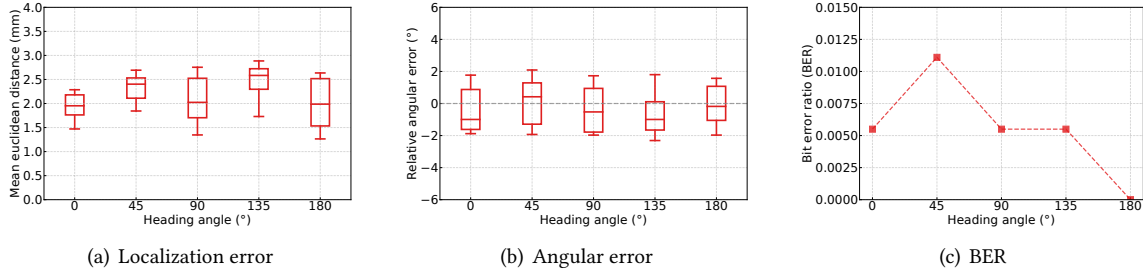


Figure 19: The impact of varying heading angles.

8.2 Evaluation Metrics

Information decoding performance is evaluated by measuring the bit error ratio (BER). To assess the posture estimation performance, we analyze the error of the measured heading angle relative to the ground truth. To evaluate the magnet localization performance, we measure the mean Euclidean distance between the reconstructed coordinate and the ground-truth coordinate of a Polaris tag. For example, for a tag consisting of M magnets, the ground-truth coordinate of each magnet is denoted as (x_i^g, y_i^g) , and the reconstructed coordinate of each magnet is (x_i^m, y_i^m) . The mean Euclidean distance can be derived as:

$$\frac{1}{M} \sum_{i=1}^M \sqrt{(x_i^g - x_i^m)^2 + (y_i^g - y_i^m)^2}. \quad (4)$$

8.3 Robot Car

8.3.1 Implementation. As shown in Fig. 18(A), the testing tag is positioned on a grid paper with a 1 mm granularity. This tag can encode 18 bits using Eq. (1). The sensor array in Sec. 7 is mounted at the front of the robot car with a ground clearance of 10 mm. The sampling rate is 100 Hz. The car moves autonomously with a speed of 10 cm/s, controlled by its onboard system. The sensing pipeline is integrated into the onboard computing module, an NVIDIA Jetson Nano 4GB development board [13]. The sensor readings are transmitted to the computing unit via a flexible printed circuit (FPC) wire. The polarity alignment process uses a template set with $36 \times 5 = 180$ templates, with each template contains 80 data points. With our efficient C architecture, the total time delay of the sensing pipeline is less than 50 ms, ensuring the robot car operates in real time.

8.3.2 Varying headings. Polaris is required to accurately estimate the posture and decode embedded information with arbitrarily changing headings. Thus, we test the detection performance with the default tag under five different heading angles, i.e., 0°, 45°, 90°, 135°, and 180°. For each heading, the robot car moves over the tag in a straight path controlled by the onboard system and repeats this process 10 times.

The results indicate there were no missed detections of each magnet under each heading angle with our peak-detection algorithm. Fig. 19(a) shows the localization performance of each magnet in the tag. The average positioning error is 1.93 mm, 2.33 mm, 2.07 mm, 2.49 mm, and 2.01 mm, respectively. Fig. 19(b) shows the estimation performance of the heading angle. Specifically, the average angular errors are 1.35°, 1.25°, 1.35°, 1.39°, and 1.07°, respectively. The information decoding performance is shown in Fig. 19(c). The maximum BER is 0.011, resulting in a total error of only 2 bits during interpreting the tag across 10 experiments. These results demonstrate the accuracy and robustness of Polaris for interpreting the magnetic tag under various headings.

8.3.3 Varying polarity orientations. We assess the impact of different polarity orientations. Specifically, we set the polarity orientation of each magnet in the tag as 0°, 45°, 90°, 135°, and 180°, respectively. The robot car moves over this tag at 0° and repeats the process 10 times. The [Mean, STD] of the angular error is (1.50°, 0.84°). The [Mean, STD] of the positional error is (1.98 mm, 0.65 mm). The BER of information decoding is 0.006. These results indicate Polaris' performance is agnostic to the polarity orientations of each magnet.

8.3.4 Varying tag density. We now assess the detection performance of Polaris decoding a complex tag layout with a higher magnet density. Specifically, we increased the number of magnets in the default tag to eight, the maximum number of a third-order tag, thus enabling 24-bit encoding. The robot car moves over this tag with a heading angle of 0° and repeats the process ten times. The average angular and positional error is 1.28° and 2.27 mm, respectively. The BER is only 0.0125, remaining consistent with the result of the defaulting tag. The results indicate the configuration of inter-magnet distance ensures the specified sensing accuracy given high magnet density.

8.3.5 Varying template sizes. A fine-grained template can improve the alignment accuracy of θ_s^m and thus distinguishes more polarity orientations in MOSK modulation (Sec. 6.2.2). However, a finer-grained template may incur more time. For




Figure 20: The impact of template size.

Real-world interference	Smartphone	Metal chair	Metal umbrella stand	Iron filings	Iron oxide (Fe ₂ O ₃)	Dirt	Muddy water
Experimental settings							
Angular Err. (Mean, STD)	(1.10, 0.28)	(1.29, 0.40)	(1.65, 0.47)	(1.33, 0.57)	(1.93, 0.30)	(1.12, 0.23)	(1.24, 0.66)
Localization Err. (Mean, STD)	(1.91, 0.32)	(2.15, 0.65)	(2.21, 0.51)	(2.63, 0.47)	(2.89, 0.64)	(1.98, 0.34)	(2.10, 0.46)

Figure 21: The performance of Polaris under real-world interference. Units of angular and localization errors are $^{\circ}$ and mm , respectively.

template construction, we consider three angular resolutions: 1° , 10° , and 20° . The corresponding template sizes are 1,800, 180, and 90, respectively. At each resolution, the robot car detects the default tag at 0° and repeats the process 10 times.

As shown in Fig. 20, posture estimation and information decoding performance decrease with the angular resolution. Specifically, Fig. 20(a) shows Polaris achieves a $< 2mm$ localization error. Even at a resolution of 20° , the positioning error remains below 3 mm. As shown in Fig. 20(b), at 20° -degree resolution, Polaris achieves the maximum angular error within 2° . Fig. 20(c) shows that the BER increases with the angular granularity. Polaris can reliably decode the tag with a < 0.03 BER at a 20° -resolution. At 1° resolution, the BER decreases to zero. The alignment process at 1° , 10° , and 20° requires 350.2 ms, 29.4 ms, and 16.7 ms, respectively.

8.3.6 Real-world magnetic interference and debris. We now evaluate Polaris’ resilience against real-world magnetic interference and magnetic debris that will introduce signal noise into the systems. As shown in Fig. 21, we emulate seven exemplary signal noise sources in three practical scenarios: household objects (electric device, metal chair and umbrella stand), heavy-duty factories (iron oxide and iron debris), and harsh environments (dirt and muddy water). The robot car moved over the tag in a straight path and repeated the process 5 times. The average BER during the experiments is 0.013. The angular and localization errors are shown in Fig. 21, demonstrating no significant performance degradation with these real-world marker interference sources.

8.3.7 Orientation calibration by Polaris. We developed an end-to-end (e2e) system of Polaris and installed it in the robot car as shown in Fig. 22(A). Two Polaris tags with consecutive IDs are placed 1 m apart, creating a straight-line path. Upon detecting tag “1”, the e2e system decodes the




Figure 22: End-to-end experiment. (A) the experimental setup and (B) a trajectory calibrated by Polaris with iron filings.




Figure 23: Results of trajectory calibration. (a) w/o interference (iron filings); (b) w/ interference.

embedded information “Start” and estimates the robot car’s orientation in real time. Then, it performs orientation calibration, i.e., to control the car to return to the straight path relative to tag “1”. Finally, the car will stop when detecting tag “2” embedded with “Stop”. Fig. 22(B) shows an example of the above orientation calibration process.

We assessed our e2e system using three different heading angles: 30° , 45° , and 60° . To collect the ground truth, we affixed a camera to a tripod for each heading to capture video at UHD resolution and 30 FPS. We recorded a video to track the robot car’s position on the graph paper. We first tested it in an ideal scenario, i.e., no disturbances/occlusion to the two tags. Polaris successfully interpreted the tag and calibrated the car’s posture during all experiments. The reconstructed trajectories are shown in Fig. 23(a). Then, we evaluated the system’s efficiency in challenging scenarios by covering both tags with iron filings (Fig. 22(B)). Fig. 23(b) show that Polaris can interpret the tag and accurately calibrate the trajectory in all experiments, even with the iron filings, indicating that Polaris can serve as a reliable complement to existing visual fiducials.

8.4 Miniature Car

8.4.1 Implementation. we implemented a $1.6 \times 1.6\text{cm}^2$ tag with three disc magnets (Fig. 18(B)). This tiny tag can encode 9 bits with 8 polarity configurations according to Eq. (1). The compact sensor array ($1.2 \times 2.5\text{cm}^2$) is installed at the front of the miniature car and operates at a sample rate of 17 Hz.

8.4.2 Practicability of Polaris. We integrated the sensing pipeline into an ESP32-S3-WROOM-1-N16R8 [5], a low-power MCU-based SoC, with ESP-IDF programming guide [11]. This MCU has an SRAM of 8 MB and a clock frequency of 240 MHz. The real-time sensing data is transmitted to the MCU via BLE transmission. For the miniature platform, the template size is $18 \times 1 = 18$. Each template has 15 data points. The template alignment and magnet localization process took 0.15 s and 0.60 s on average. That is, detecting each magnet only takes 0.75 s in total. Note that the miniature robots usually operate at low speed when performing critical tasks in confined spaces. For example, a miniature in-pipe inspection robot has a speed of <1 mm/s to ensure the damage detection accuracy [53, 57].

For experimentation purposes, we manually moved the car on a straight-line by aligning the car with a ruler, with a heading angle of 0° and repeated the process 10 times. The average velocity is controlled to be about 5 mm/s. This reconstruction process was consistent with the robot car's, and we used the ESP32 chip as the computing module.

The BER result is 0.033, with the number of polarity orientations set to 8. The (mean, STD) pair of the Euclidean distance between the measured and ground truth coordinates is (0.58 mm, 0.08 mm). That is, Polaris can achieve a sub-millimeter level localization accuracy for the miniature robot. For the angle estimation, the (mean, STD) pair of the estimation result is (0.997° , 0.125°), again showing Polaris' practicability as a fiducial system for miniature robots.

8.5 Overhead analysis

The energy consumption of Polaris is critical for the battery-powered robots. To evaluate Polaris' energy consumption, we analyzed each module, including the sensing hardware and the computing unit. We profile the energy consumption for the sensor array by measuring the battery supply voltage and current. Specifically, the sensor array on the robot car consumed 77.22 mW. For the miniature car, the sensor array with three magnetometers consumed only 25.08 mW. We also measured a compact solar panel of 30 mm by 25 mm in an outdoor environment under 100,000 lux illuminance. The panel produces 85 mW of power, demonstrating the feasibility of Polaris on a solar-powered miniature platform.

For the computing unit, we profile it on the ESP32 chip using a power meter. We calculated the overhead of the computing unit by subtracting the idle power consumption and

repeating this process 5 times. The average power consumption is only 2.8 mW, corroborating the energy efficiency of our template matching scheme.

9 DISCUSSION

Tackling miss detection. For small robots, detecting only part of a Polaris tag can miss bits and yield inaccurate posture estimation. To address this problem, additional redundant tags can be used as they are compact and cost-effective (Sec. 5). The experimental results in Sec. 8.4 show that Polaris can accurately decode tags with an inter-magnet distance of only 1.6 cm, demonstrating its efficacy in detecting compact tags. Note that Polaris can achieve high encoding capacity with small tags by employing a higher orientation-shifting modulation, e.g., 20-MOSK. The improved order of modulation does not require extra costs of hardware or magnets. Specifically, it requires a template with more synthesized trajectories to distinguish more polarity orientations (Sec. 6.2.2). With our efficient C-based architecture, this can be implemented effectively, e.g., taking 350 ms for 20-MOSK as elaborated in Sec. 8.3.5.

Improving Polaris. Polaris complements existing ground-mounted visual fiducials in challenging scenarios where cameras fail to detect them even at close ranges (Sec. 2.1). To handle higher speeds, one can increase the magnetometer's sampling rate to capture the magnetic field created by the passive magnet. To extend the sensing range, stronger passive magnets and/or more sensitive magnetometers can be used. To minimize the sensing latency, Polaris can reduce tag sizes and employ higher detecting speeds with an increased sampling rate. For tag fabrication, the PVC substrate is durable and resistant to deformation. One can further enhance the durability by placing the magnet on rigid, non-magnetic substrates.

10 CONCLUSION

We have presented the first vision-free fiducial marker system, Polaris, which enables posture calibration and simple message communication with a new magnetic sensing pipeline and tag design methods. Compared to the existing fiducial systems that rely on imagery tags and cameras, Polaris offers several features, such as robustness, energy efficiency, and privacy protection. Therefore, Polaris demonstrates unique advantages in preparing future robotic systems for challenging tasks.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers and shepherd for their suggestions. This work was supported by the National Natural Science Foundation of China under Grants No. 61960206002, T2421002, 62102256, and 62472283. Dongyao Chen is the corresponding author.

REFERENCES

- [1] 2003. SPI Protocol. <https://web.archive.org/web/20150413003534/http://www.ee.nmt.edu/~teare/ee3081/datasheets/S12SPIV3.pdf>.
- [2] 2007. *Dynamic Time Warping*. Springer Berlin Heidelberg, Berlin, Heidelberg, 69–84. https://doi.org/10.1007/978-3-540-74048-3_4
- [3] 2017. NdFeB magnet technical datasheet. <https://www.arnoldmagnetics.com/wp-content/uploads/2017/10/Catalog-151021.pdf>.
- [4] 2018. RAYTAC MDBT42Q-512KV2. https://www.raytac.com/product/ins.php?index_id=31.
- [5] 2021. ESP32-S3 Series Datasheet. https://www.espressif.com/sites/default/files/documentation/esp32-s3_datasheet_en.pdf.
- [6] 2022. Amazon's tiny robot drives do the heavy lifting. <https://www.amazon.science/latest-news/amazon-robotics-autonomous-drive-units-hercules-pegasus-xanthus-xbot>.
- [7] 2023. Data Matrix. https://en.wikipedia.org/wiki/Data_Matrix.
- [8] 2024. AprilTag fiducial system. <https://github.com/AprilRobotics/apriltag>.
- [9] 2024. Diodes: AH8502 Linear Hall-Effect Sensor. <https://www.diodes.com/assets/Datasheets/AH8502.pdf>.
- [10] 2024. Docking and undocking Spot. <https://support.bostondynamics.com/s/article/Docking-and-undocking-Spot>.
- [11] 2024. ESP-IDF Programming Guide. <https://docs.espressif.com/projects/esp-idf/en/stable/esp32/index.html>.
- [12] 2024. HM01B0 Ultralow Power CIS. <https://www.himax.com.tw/products/cmos-image-sensor/always-on-vision-sensors/hm01b0/>.
- [13] 2024. Jetson Nano Developer Kit. <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>.
- [14] 2024. Melexis: MLX90393 Triaxis Micropower Magnetometer. <https://www.melexis.com/-/media/files/documents/datasheets/mlx90393-datasheet-melexis.pdf>.
- [15] 2024. Neodymium Magnet Specifications. <https://www.kjmagnetics.com/neodymium-magnet-specifications.asp>.
- [16] 2024. Orbbec Astra Pro Plus camera. <https://shop.orbbec3d.com/Astra-Pro-Plus>.
- [17] 2024. Repository for Polaris. <https://github.com/wjk5117/Polaris>.
- [18] 2024. Roomba Combo j7+ Robot Vacuum and Mop. https://www.irobot.com/en_US/roomba-combo-j7plus-robot-vacuum-and-mop/C755020.html.
- [19] 2024. Roomba j9+ Self-Emptying Robot Vacuum. https://www.irobot.com/en_US/roomba-j9plus-self-emptying-robot-vacuum/J955020.html.
- [20] 2024. Spot-The Agile Mobile Robot. <https://bostondynamics.com/products/spot/>.
- [21] Muhammad Aizat, Ahmad Azmin, and Wan Rahiman. 2023. A Survey on Navigation Approaches for Automated Guided Vehicle Robots in Dynamic Surrounding. *IEEE Access* 11 (2023), 33934–33955. <https://doi.org/10.1109/ACCESS.2023.3263734>
- [22] H. Aoyama, A. Himoto, O. Fuchiwaki, D. Misaki, and T. Sumrall. 2005. Micro hopping robot with IR sensor for disaster survivor detection. In *IEEE International Safety, Security and Rescue Robotics, Workshop, 2005*. 189–194. <https://doi.org/10.1109/SSRR.2005.1501249>
- [23] Yang Bai, Nakul Garg, and Nirupam Roy. 2022. SPiDR: ultra-low-power acoustic spatial sensing for micro-robot navigation. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services (Portland, Oregon) (MobiSys '22)*. Association for Computing Machinery, New York, NY, USA, 99–113. <https://doi.org/10.1145/3498361.3539775>
- [24] Alessio Caciagli, Roel J. Baars, Albert P. Philipse, and Bonny W.M. Kuipers. 2018. Exact expression for the magnetic field of a finite cylinder with arbitrary uniform magnetization. *Journal of Magnetism and Magnetic Materials* 456 (2018), 423–432. <https://doi.org/10.1016/j.jmmm.2018.02.003>
- [25] Dongyao Chen, Mingke Wang, Chenxi He, Qing Luo, Yasha Iravantchi, Alanson Sample, Kang G. Shin, and Xinning Wang. 2021. MagX: Wearable, Untethered Hands Tracking with Passive Magnets. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (New Orleans, Louisiana) (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 269–282. <https://doi.org/10.1145/3447993.3483260>
- [26] Ke-Yu Chen, Shwetak N. Patel, and Sean Keller. 2016. Finexus: Tracking Precise Motions of Multiple Fingertips Using Magnetic Sensing. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (San Jose, California, USA) (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1504–1514. <https://doi.org/10.1145/2858036.2858125>
- [27] George C Clark Jr and J Bibb Cain. 2013. *Error-correction coding for digital communications*. Springer Science & Business Media.
- [28] Houde Dai, Pengfei Guo, Hongyu Chen, Silin Zhao, Penghua Liu, and Guijuan Lin. 2021. An Improved Magnetic Spot Navigation for Replacing the Barcode Navigation in Automated Guided Vehicles. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. 786–791. <https://doi.org/10.1109/ICRA48506.2021.9561316>
- [29] Raffaello D'Andrea. 2012. Guest Editorial: A Revolution in the Warehouse: A Retrospective on Kiva Systems and the Grand Challenges Ahead. *IEEE Transactions on Automation Science and Engineering* 9, 4 (2012), 638–639. <https://doi.org/10.1109/TASE.2012.2214676>
- [30] C. De Wagter, S. Tijmons, B. D. W. Remes, and G. C. H. E. de Croon. 2014. Autonomous flight of a 20-gram Flapping Wing MAV with a 4-gram onboard stereo vision system. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 4982–4987. <https://doi.org/10.1109/ICRA.2014.6907589>
- [31] Joseph DeGol, Timothy Bretl, and Derek Hoiem. 2017. ChromaTag: A Colored Marker and Fast Detection Algorithm. In *2017 IEEE International Conference on Computer Vision (ICCV)*. 1481–1490. <https://doi.org/10.1109/ICCV.2017.164>
- [32] Daniel A Duecker, Nathalie Bauschmann, Tim Hansen, Edwin Kreuzer, and Robert Seifried. 2020. Towards Micro Robot Hydrobatcs: Vision-based Guidance, Navigation, and Control for Agile Underwater Vehicles in Confined Environments. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1819–1826. <https://doi.org/10.1109/IROS45743.2020.9341051>
- [33] M. Fiala. 2005. ARTag, a fiducial marker system using digital techniques. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 2. 590–596 vol. 2. <https://doi.org/10.1109/CVPR.2005.74>
- [34] Matthew Gadd and Paul Newman. 2015. A framework for infrastructure-free warehouse navigation. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 3271–3278. <https://doi.org/10.1109/ICRA.2015.7139650>
- [35] Benjamin Goldberg, Raphael Zufferey, Neel Doshi, Elizabeth Farrell Helbling, Griffin Whittredge, Mirko Kovac, and Robert J. Wood. 2018. Power and Control Autonomy for High-Speed Locomotion With an Insect-Scale Legged Robot. *IEEE Robotics and Automation Letters* 3, 2 (2018), 987–993. <https://doi.org/10.1109/LRA.2018.2793355>
- [36] R. W. Hamming. 1950. Error detecting and error correcting codes. *The Bell System Technical Journal* 29, 2 (1950), 147–160. <https://doi.org/10.1002/j.1538-7305.1950.tb00463.x>
- [37] Richard Hartley and Andrew Zisserman. 2003. *Multiple view geometry in computer vision*. Cambridge university press.
- [38] Vikram Iyer, Ali Najafi, Johannes James, Sawyer Fuller, and Shyamnath Gollakota. 2020. Wireless steerable vision for live insects and insect-scale robots. *Science Robotics* 5, 44 (2020), eabb0839. <https://doi.org/10.1126/scirobotics.abb0839>

- arXiv:<https://www.science.org/doi/pdf/10.1126/scirobotics.abb0839>
- [39] Michail Kalaitzakis, Brennan Cain, Sabrina Carroll, Anand Ambrosi, Camden Whitehead, and Nikolaos Vitzilaios. 2021. Fiducial markers for pose estimation: Overview, applications and experimental comparison of the artag, apriltag, aruco and stag markers. *Journal of Intelligent & Robotic Systems* 101 (2021), 1–26.
- [40] Saber Kazeminasab, Neda Sadeghi, Vahid Janfaza, Moein Razavi, Samira Ziyadidegan, and M. Katherine Banks. 2021. Localization, Mapping, Navigation, and Inspection Methods in In-Pipe Robots: A Review. *IEEE Access* 9 (2021), 162035–162058. <https://doi.org/10.1109/ACCESS.2021.3130233>
- [41] Eamonn J. Keogh and Michael J. Pazzani. [n.d.]. *Derivative Dynamic Time Warping*. 1–11. <https://doi.org/10.1137/1.9781611972719.1> arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9781611972719.1>
- [42] Dieter Kraft. 1994. Algorithm 733: TOMP–Fortran Modules for Optimal Control Calculations. *ACM Trans. Math. Softw.* 20, 3 (sep 1994), 262–281. <https://doi.org/10.1145/192115.192124>
- [43] Jeffrey C. Lagarias, James A. Reeds, Margaret H. Wright, and Paul E. Wright. 1998. Convergence Properties of the Nelder–Mead Simplex Method in Low Dimensions. *SIAM Journal on Optimization* 9, 1 (1998), 112–147. <https://doi.org/10.1137/S1052623496303470> arXiv:<https://doi.org/10.1137/S1052623496303470>
- [44] Zewen Li, Fan Liu, Wenjie Yang, Shouheng Peng, and Jun Zhou. 2022. A Survey of Convolutional Neural Networks: Analysis, Applications, and Prospects. *IEEE Transactions on Neural Networks and Learning Systems* 33, 12 (2022), 6999–7019. <https://doi.org/10.1109/TNNLS.2021.3084827>
- [45] Yibo Liu, Hunter Schofield, and Jinjun Shan. 2022. Intensity Image-Based LiDAR Fiducial Marker System. *IEEE Robotics and Automation Letters* 7, 3 (2022), 6542–6549. <https://doi.org/10.1109/LRA.2022.3174971>
- [46] Atanu Maity, S. Majumder, and Sukamal Ghosh. 2010. An experimental hyper redundant serpentine robot. In *2010 IEEE International Conference on Systems, Man and Cybernetics*. 3180–3185. <https://doi.org/10.1109/ICSMC.2010.5642271>
- [47] Edwin Olson. 2011. AprilTag: A robust and flexible visual fiducial system. In *2011 IEEE International Conference on Robotics and Automation*. 3400–3407. <https://doi.org/10.1109/ICRA.2011.5979561>
- [48] Michael Ortner and Lucas Gabriel Coliado Bandeira. 2020. Magpylib: A free Python package for magnetic field computation. *SoftwareX* 11 (2020), 100466.
- [49] Edward Ramsden. 2011. *Hall-effect sensors: theory and application*. Elsevier.
- [50] Michael Rubenstein, Christian Ahler, and Radhika Nagpal. 2012. KiloBot: A low cost scalable robot system for collective behaviors. In *2012 IEEE International Conference on Robotics and Automation*. 3293–3298. <https://doi.org/10.1109/ICRA.2012.6224638>
- [51] Ksenia Shabalina, Artur Sagitov, Leysan Sabirova, Hongbing Li, and Evgeni Magid. 2020. ARTag, AprilTag and CALTag Fiducial Systems Comparison in a Presence of Partial Rotation: Manual and Automated Approaches. In *Informatics in Control, Automation and Robotics*, Oleg Gusikhin and Kurosh Madani (Eds.). Springer International Publishing, Cham, 536–558.
- [52] Jerome Svigals. 2012. The long life and imminent death of the mag-stripe card. *IEEE Spectrum* 49, 6 (2012), 72–76. <https://doi.org/10.1109/MSPEC.2012.6203975>
- [53] Minoru Takagi, Kazunari Yoshida, Hiroyuki Hoshino, Riichiro Tadakuma, Yoshiyuki Suzuri, and Hidemitsu Furukawa. 2019. Sliding Walk With Friction Control of Double-Network Gel on Feet of Inchworm Robot. *Frontiers in Mechanical Engineering* 5 (2019). <https://doi.org/10.3389/fmech.2019.00044>
- [54] Cameron R. Taylor, Haley G. Abramson, and Hugh M. Herr. 2019. Low-Latency Tracking of Multiple Permanent Magnets. *IEEE Sensors Journal* 19, 23 (2019), 11458–11468. <https://doi.org/10.1109/JSEN.2019.2936766>
- [55] Sumit Tiwari. 2016. An Introduction to QR Code Technology. In *2016 International Conference on Information Technology (ICIT)*. 39–44. <https://doi.org/10.1109/ICIT.2016.021>
- [56] Stefan M. Trenkwalder. 2019. Computational Resources of Miniature Robots: Classification and Implications. *IEEE Robotics and Automation Letters* 4, 3 (2019), 2722–2729. <https://doi.org/10.1109/LRA.2019.2917395>
- [57] Ankush Verma, Ayush Kaiwart, Nikhil Dhar Dubey, Farman Naseer, and Swastik Pradhan. 2022. A review on various types of in-pipe inspection robot. *Materials Today: Proceedings* 50 (2022), 1425–1434. <https://doi.org/10.1016/j.matpr.2021.08.335> 2nd International Conference on Functional Material, Manufacturing and Performances (ICFMMP-2021).
- [58] G. Vivó-Truyols, J.R. Torres-Lapasió, A.M. van Nederkassel, Y. Vander Heyden, and D.L. Massart. 2005. Automatic program for peak detection and deconvolution of multi-overlapped chromatographic signals: Part II: Peak model and deconvolution algorithms. *Journal of Chromatography A* 1096, 1 (2005), 146–155. <https://doi.org/10.1016/j.chroma.2005.03.072> Chemical Separations and Chemometrics.
- [59] Chia-Cheng Wang, Jyh-Cheng Chen, Yi Chen, Rui-Heng Tu, Jia-Jiun Lee, Yu-Xin Xiao, and Shan-Yu Cai. 2021. MVP: magnetic vehicular positioning system for GNSS-denied environments. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (New Orleans, Louisiana) (MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 531–544. <https://doi.org/10.1145/3447993.3483264>
- [60] John Wang and Edwin Olson. 2016. AprilTag 2: Efficient and robust fiducial detection. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [61] Jike Wang, Shanmu Wang, Yasha Iravantchi, Mingke Wang, Alanson Sample, Kang G. Shin, Xinbing Wang, Chenghu Zhou, and Dongyao Chen. 2024. METRO: Magnetic Road Markings for All-weather, Smart Roads. In *Proceedings of the 21st ACM Conference on Embedded Networked Sensor Systems (, Istanbul, Turkiye.) (SenSys '23)*. Association for Computing Machinery, New York, NY, USA, 280–293. <https://doi.org/10.1145/3625687.3625809>
- [62] Mingke Wang, Qing Luo, Yasha Iravantchi, Xiaomeng Chen, Alanson Sample, Kang G. Shin, Xiaohua Tian, Xinbing Wang, and Dongyao Chen. 2022. Automatic calibration of magnetic tracking. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking (Sydney, NSW, Australia) (MobiCom '22)*. Association for Computing Machinery, New York, NY, USA, 391–404. <https://doi.org/10.1145/3495243.3558760>
- [63] Philip B. Weerakody, Kok Wai Wong, Guanjin Wang, and Wendell Ela. 2021. A review of irregular time series data handling with gated recurrent neural networks. *Neurocomputing* 441 (2021), 161–178. <https://doi.org/10.1016/j.neucom.2021.02.046>
- [64] Stephen B Wicker and Vijay K Bhargava. 1999. *Reed-Solomon codes and their applications*. John Wiley & Sons.
- [65] Tzu-Hsu Yu and Hsin-Mu Tsai. 2023. *ReMark: Privacy-preserving Fiducial Marker System via Single-pixel Imaging*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3570361.3613289>
- [66] Xu Zhong, Yu Zhou, and Hanyu Liu. 2017. Design and recognition of artificial landmarks for reliable indoor self-localization of mobile robots. *International Journal of Advanced Robotic Systems* 14, 1 (2017), 1729881417693489. <https://doi.org/10.1177/1729881417693489>