InfoMagNav: Real-Time Uncertainty-Aware Motion Planning for Magnetic-Based Navigation

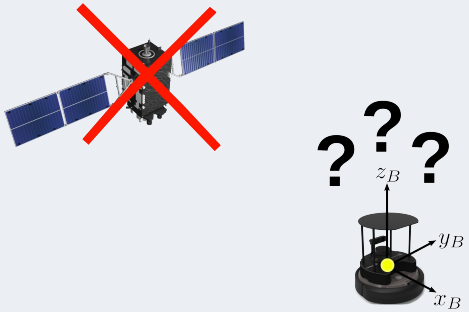
By Aditya Penumarti

April 3rd, 2026

Motivation and Background

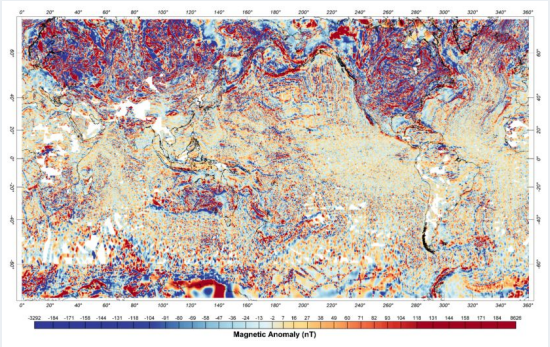
Motivation: Magnetic Anomaly-based Navigation (MagNav)

Long-duration operation without GPS



- GPS not available (jammed)
- Weather-dependent
- Not robust in repeated features

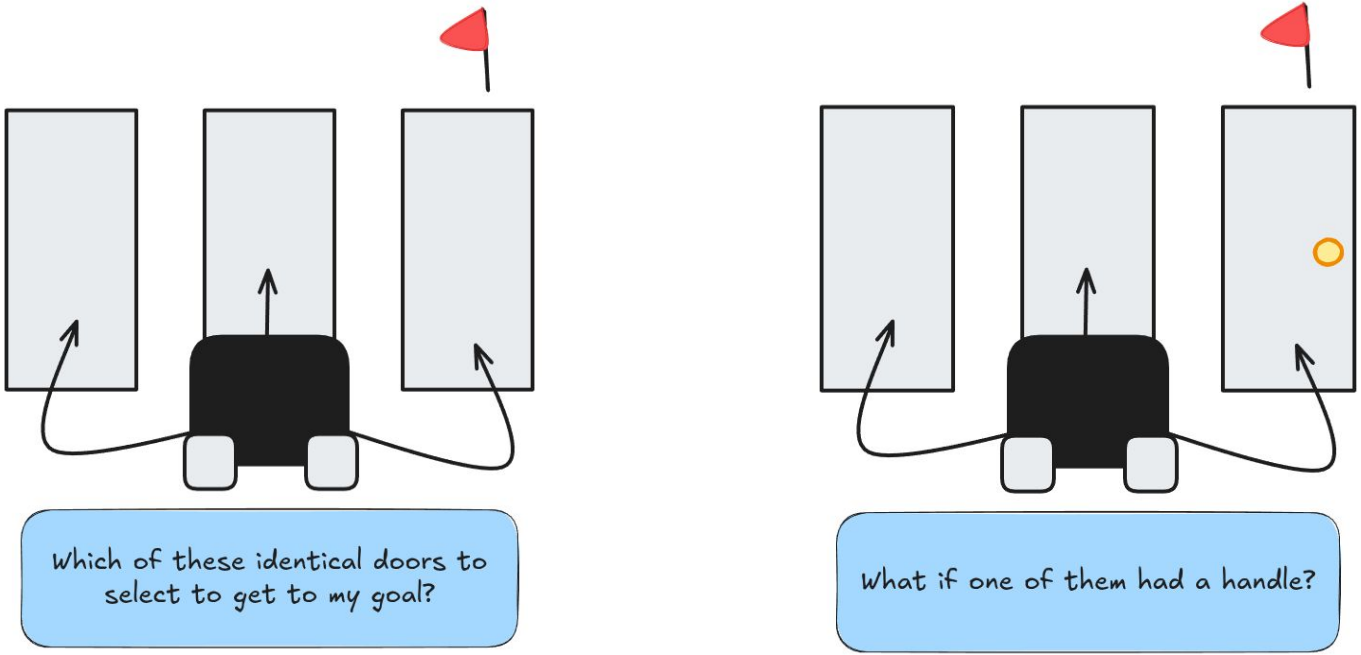
Magnetic Anomaly map



- Available globally
- Robust to dynamic changes unlike magnetic fields

How can we utilize magnetic anomalies for GPS-denied navigation?

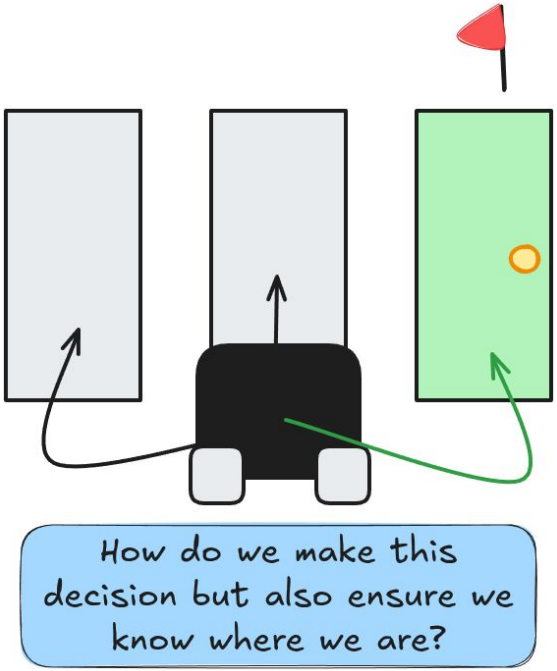
Motivation: Uncertainty Reduction



How can we use information encoded in our environment to make better navigation decisions for our localization?

The Problem

Using information can we reduce our uncertainty about our state while completing our task?



- Use the **information** available to make **decisions**
- These **decisions** as a consequence will help us know **where we are**

Background – Particle Filter

- Particle Filtering [1] is used to localize in magnetic fields
- The filter has state estimate guesses initialized as:
  - a standard normal distribution with uniformly distributed weights
- The filter is propagated using a *unicycle* motion model

$$x_{k+1} = x_k + V\cos\theta\Delta t + \epsilon_x$$
$$y_{k+1} = y_k + V\sin\theta\Delta t + \epsilon_y$$
$$\theta_{k+1} = \theta_k + \omega\Delta t + \epsilon_\theta$$

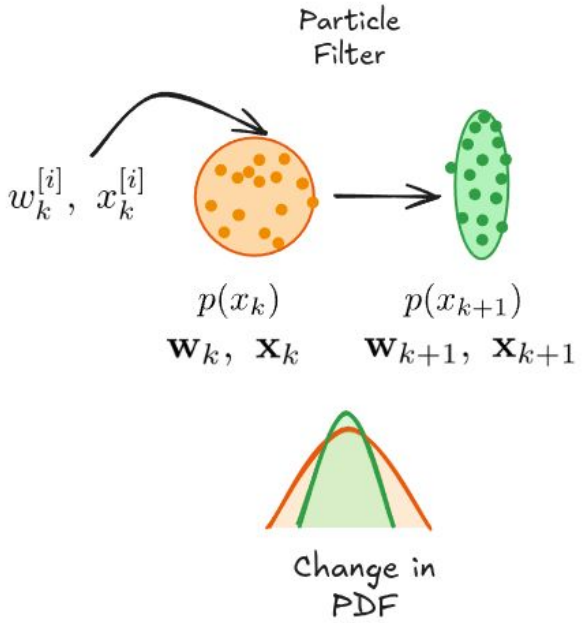


Fig. 2: Particle Filter Architecture

[1] F. Dellaert, D. Fox, W. Burgard, and S. Thrun, "Monte Carlo localization for mobile robots," in *Proceedings 1999 IEEE International Conference on Robotics and Automation (Cat. No.99CH36288C)*, May 1999, pp. 1322–1328 vol.2. doi: 10.1109/ROBOT.1999.772544.

Background – Particle Filter

- Sensor model is a queried map sensor value with stochastic noise:

$$z_k = g(\mathbf{x}_k, m)$$

$$z_k \sim \mathcal{N}\left(m(\mathbf{x}_k), \sigma^2\right)$$

- This sensor model is then used to update the state estimate based on the prior distribution.
- The localization is highly dependent on where the anomalies are

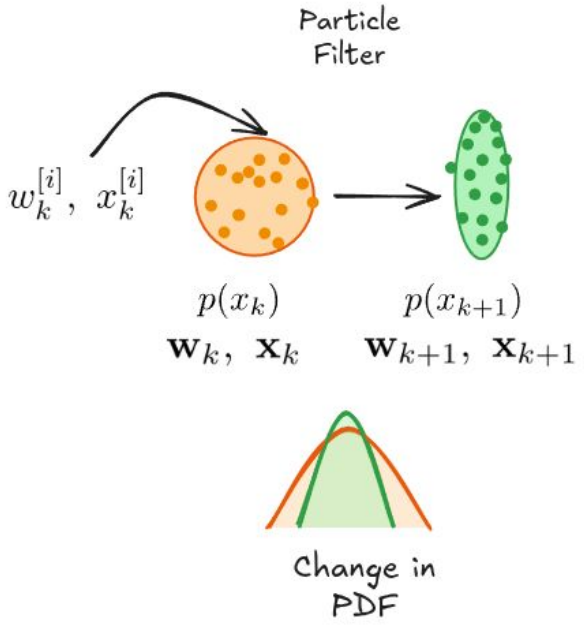


Fig. 2: Particle Filter Architecture

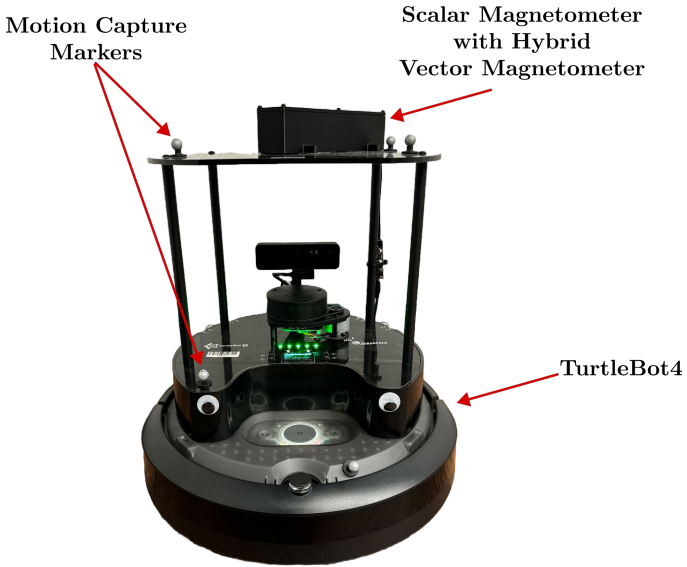
Background – Map Collection

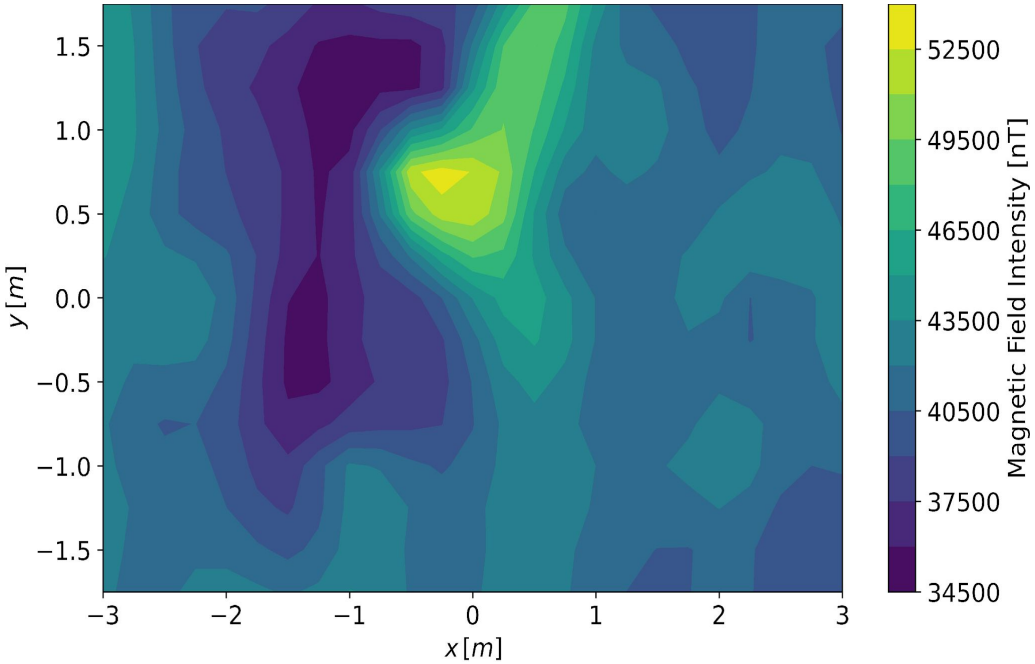
- MagNav requires highly sensitive magnetometers to be used
  - An optically pumped quantum magnetometer is used
- The magnetometer produces a scalar intensity output
- The sensitivity of the magnetometer used is $3\ pT/\sqrt{Hz}$



QTFM Magnetometer
(https://quspin.com/qtfm-gen-2/)

Background – Map Collection





Background – MagNav Localization

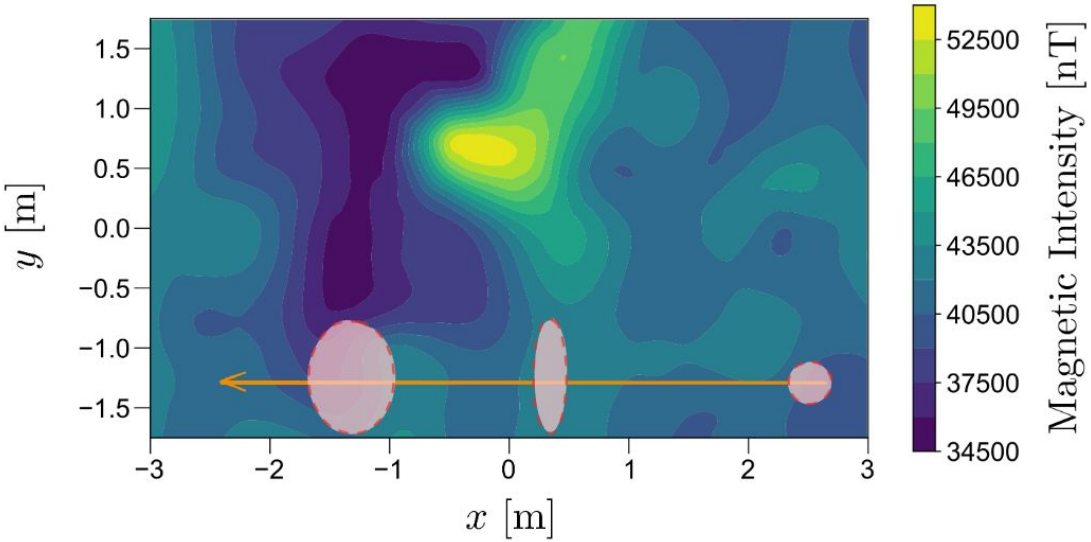
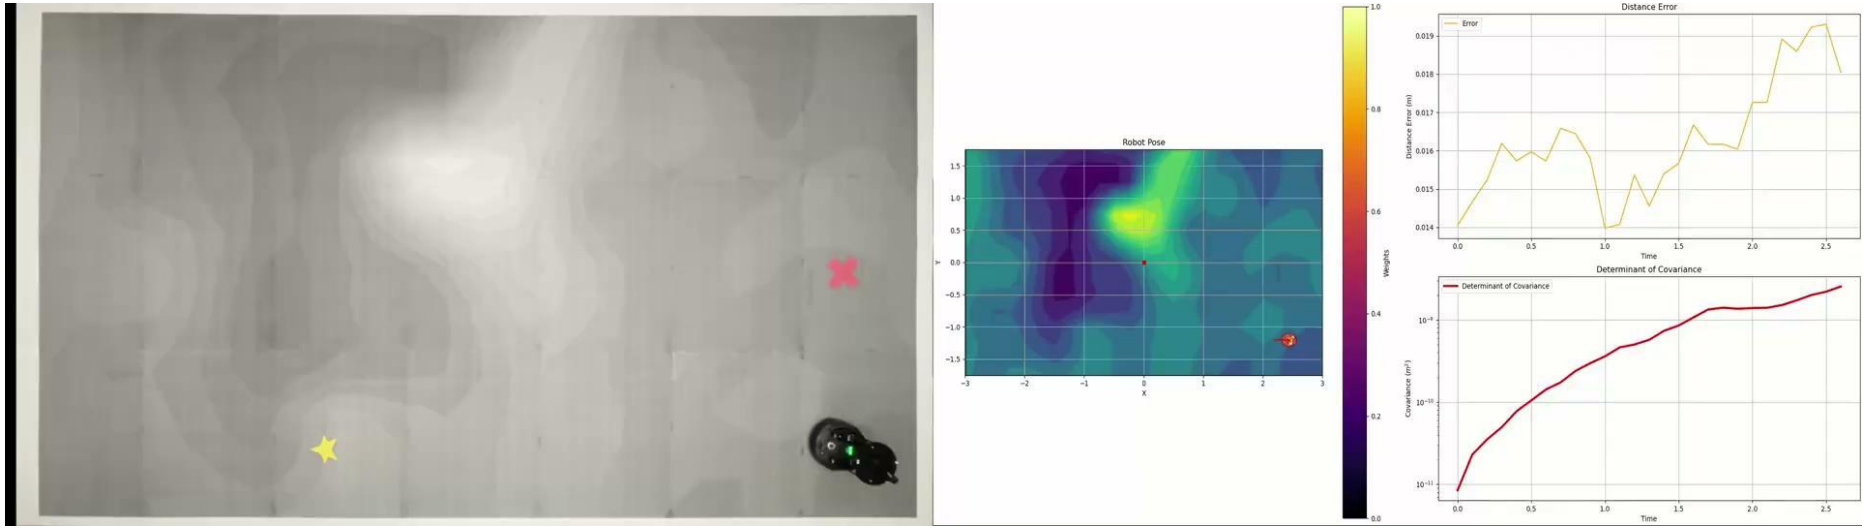


Fig. 3: Navigation through uniform/flat gradients

Navigating through low gradients the uncertainty increases as there is no uniqueness.

Background – MagNav Localization



Navigating through low gradients the uncertainty is increased.
Localization is completed using a particle filter.

Background – MagNav Localization

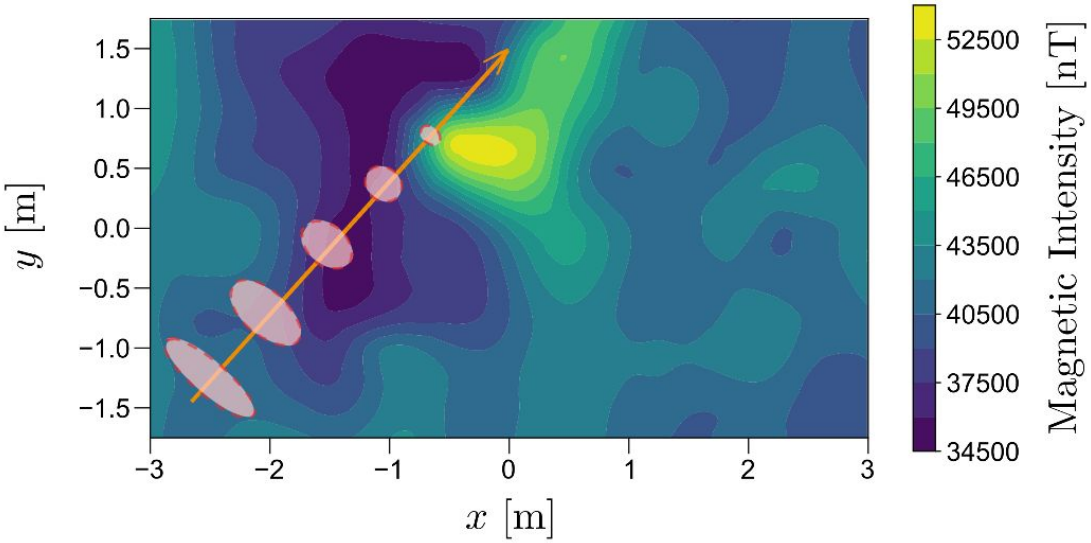
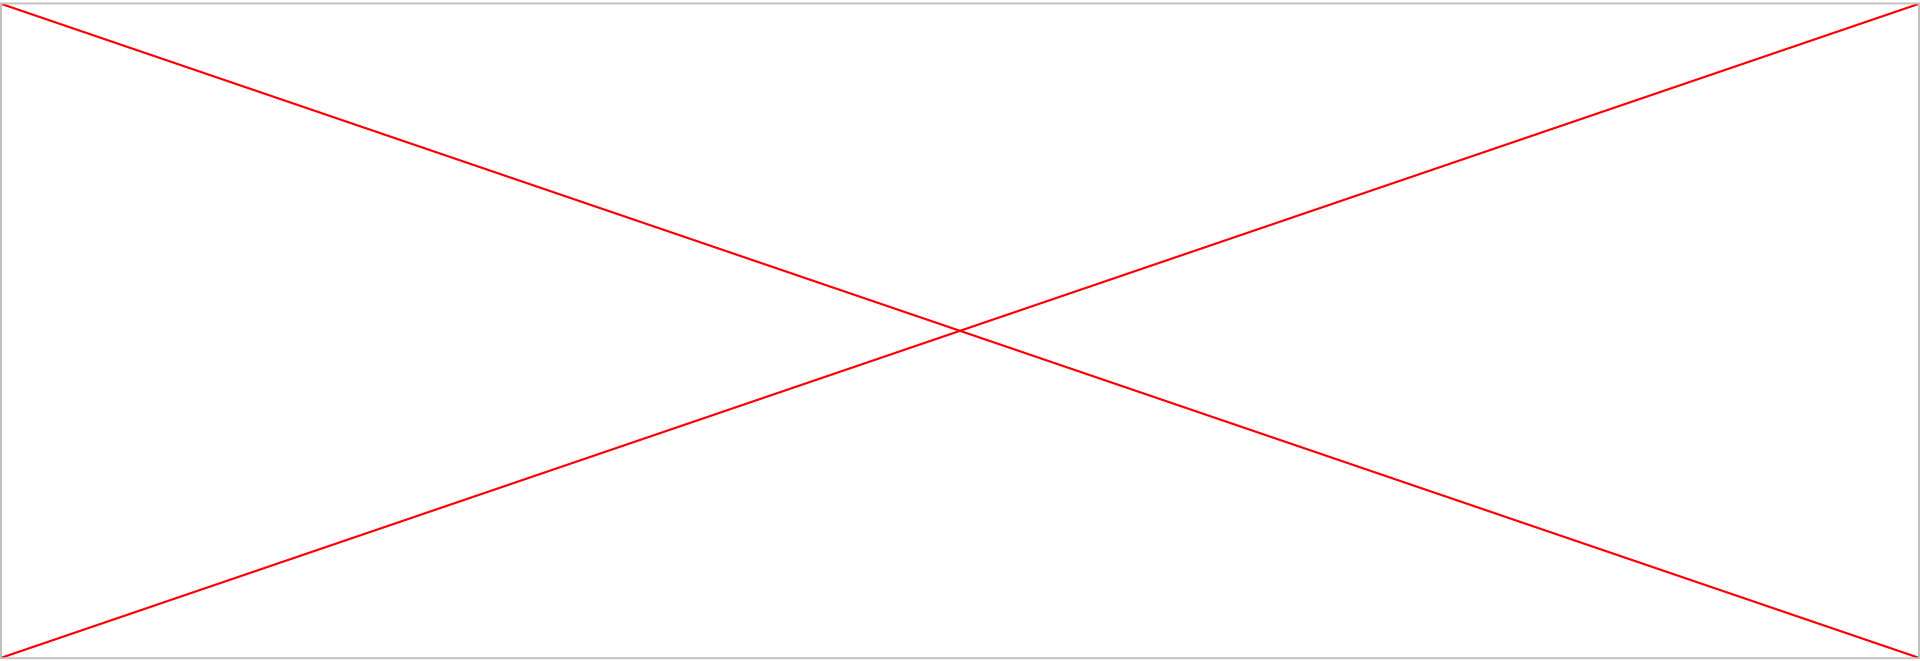


Fig. 4: Navigation through spatial varied gradients

Navigating through high gradients the uncertainty is reduced.

Background – MagNav Localization



- Navigating through high gradients the uncertainty is reduced.
- Using a particle filter the localization algorithm is completed.

Methodology

Methodology – Uncertainty-Aware MagNav

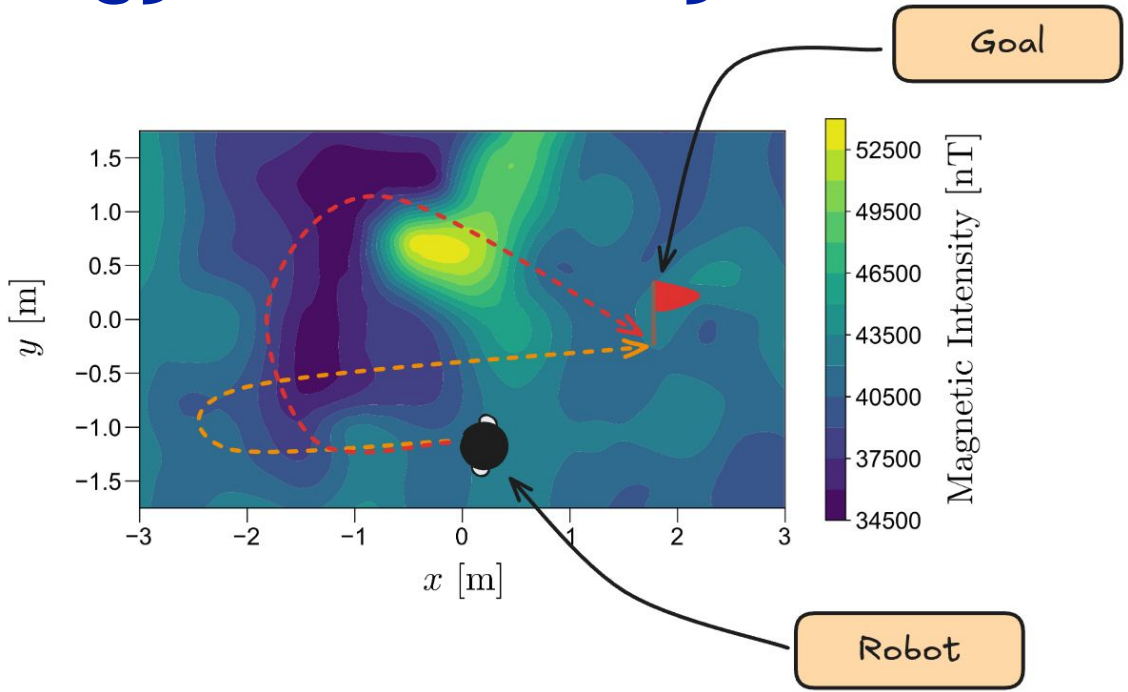


Fig. 6: Leveraging the gradients in the magnetic anomaly map, the robot will be able to make intelligent actions to reach the goal.

Using the information-rich magnetic anomaly map, can we navigate through the magnetic fields intelligently to improve localization?

Methodology – Uncertainty-Aware MagNav

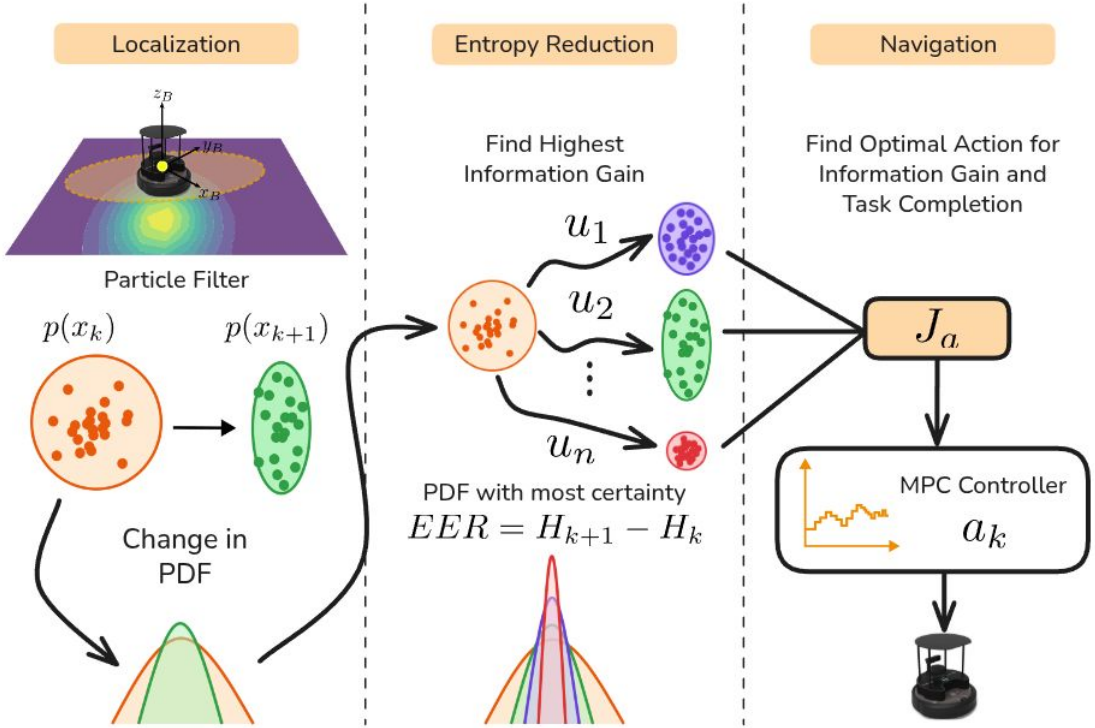


Fig. 7: Architectural overview of InfoMagNav for the uncertainty-aware navigation

Methodology - Information Gain

The method uses the entropy from the previous action and entropy from the current action to calculate **Information Gain:**

$$I(z_{k+1}, u_k) = H(bel(\mathbf{x}_k \mid z_k)) - H(bel(\mathbf{x}_{k+1} \mid \mathbf{x}_k, u_k, z_{k+1}))$$

As we want to predict the next best action to taken, we need the expectation over all the future sensor measurements with respect to information gain to calculate the **Expected Entropy Reduction :**

$$EER(u_k) = \mathbb{E}_{z_{k+1}}[I(u_k)] = \int_{z_{k+1}} p(z_{k+1} \mid u_k)\ I(z_{k+1}, u_k) dz_{k+1}$$

For efficient calculation, the particles from the filter can be used to calculate the **EER:**

$$\sum_{j=1}^{N_{\text{meas}}} p(z_{k+1}^{[j]}) I^{[j]}(z_{k+1}, u_k)$$

Methodology – Entropy Approximation

- The approximation for the entropy is a computationally intractable problem, **largely due to the estimation of the evidence term**.
- To combat this the particle filter is leveraged to estimate the entropy from the state-estimation distribution
- Often *naively* the following formula is used:

$$H(p(\mathbf{x}_k \mid z_k)) \approx -\sum_{i=1}^{N} w_k^{[i]} \log w_k^{[i]}$$

- Often in a low-dimensional feature space like MagNav, simply accounting for the weights is misleading

Methodology – Entropy Approximation

- The resampling step, as it **discards** and **clusters highly weighted** particles can destroy the **notion of entropy**, if it was solely dependent on weights.
- Instead a position, sensor and weight based entropy is used [2]

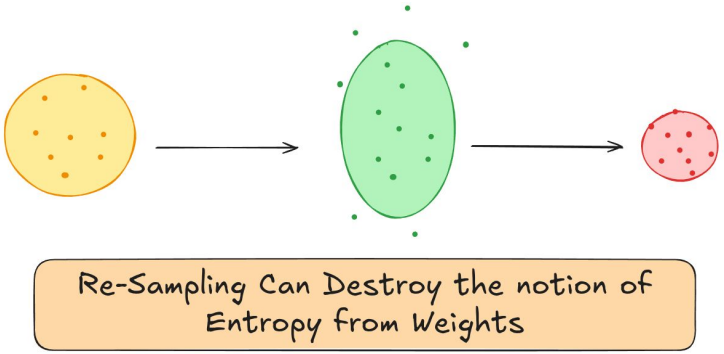


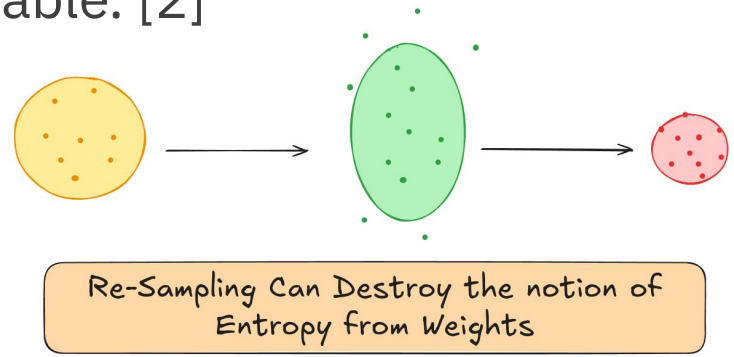
Fig. 8: Resampling step destruction of information

[2] Y. Boers, H. Driessen, A. Bagchi, and P. Mandal, "Particle filter based entropy," in *2010 13th International Conference on Information Fusion*, Jul. 2010, pp. 1–8. doi: 10.1109/ICIF.2010.5712013.

Methodology - Entropy Approximation

$$H\left(p\left(\mathbf{x}_k \mid z_k\right)\right) \approx \log\left(\sum_{i=1}^{N} p\left(z_k \mid \mathbf{x}_k^{[i]}\right) w_{k-1}^{[i]}\right) - \sum_{i=1}^{N} \log\left(p(\left(z_k \mid \mathbf{x}_k^{[i]}\right) \cdot \left(\sum_{j=1}^{N} p\left(\mathbf{x}_k^{[i]} \mid \mathbf{x}_{k-1}^{[j]}\right) w_{k-1}^{[j]}\right) w_k^{[i]}\right)$$

- Particle estimation gives us **good convergence** to the **true distribution** and maintains the notion of spatial variance affecting information available. [2]



[2] Y. Boers, H. Driessen, A. Bagchi, and P. Mandal, "Particle filter based entropy," in *2010 13th International Conference on Information Fusion*, Jul. 2010, pp. 1–8. doi: 10.1109/ICIF.2010.5712013.

Methodology – Action Selection

- To navigate to the goal a **receding-horizon** based optimization is deployed.
  - The system optimizes between **exploitation (goal-reachability)** and **exploration (uncertainty-reduction)**.
- Optimizing over multiple objectives gives us a intelligent navigation algorithm that maintains a **certain localization belief.**

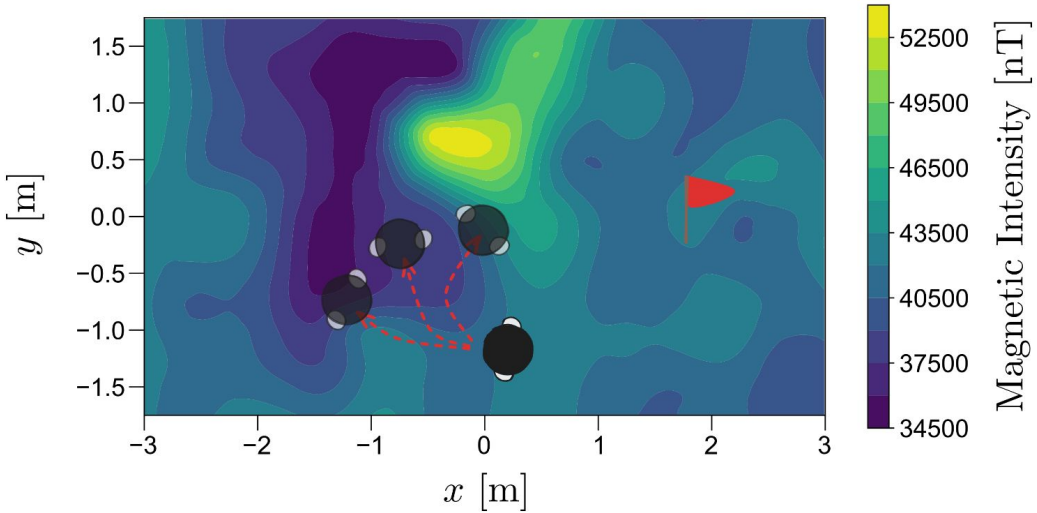
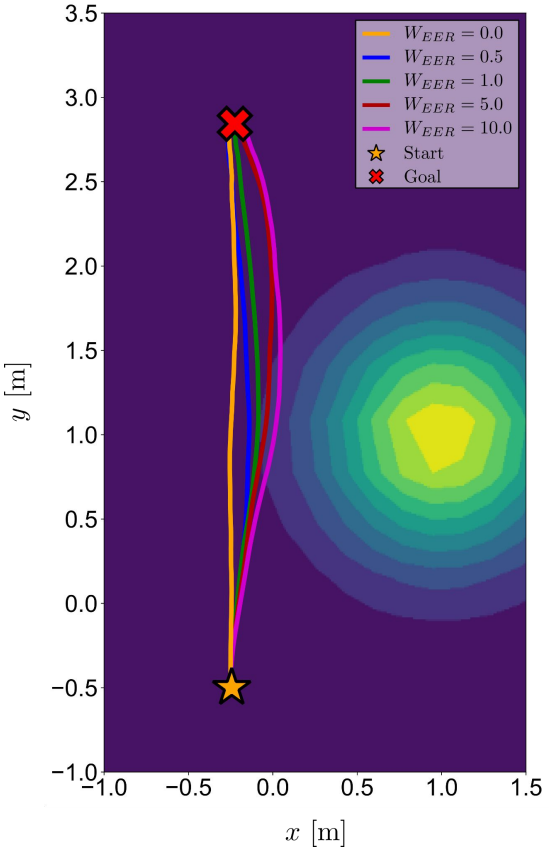


Fig. 9: Receding-horizon based optimization for navigation with exploitation and exploration

Experiments

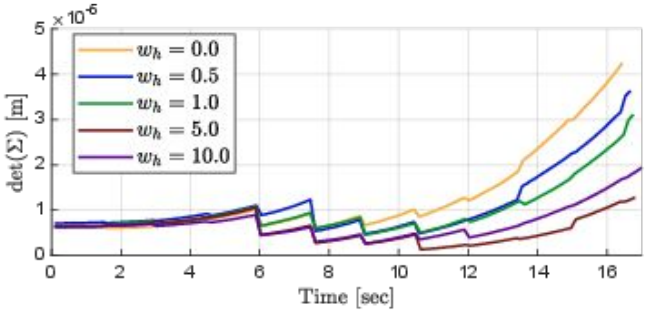
Experiments – Case Study

- Planning algorithm considers **uncertainty measurements more** with a **higher weight**
- When information in **unweighted** the robot trajectory is almost a **straight line**
- When the localization uncertainty is **weighted more heavily** in the objective function, the estimation covariance is **significantly reduced**, leading to **lower covariance over time**



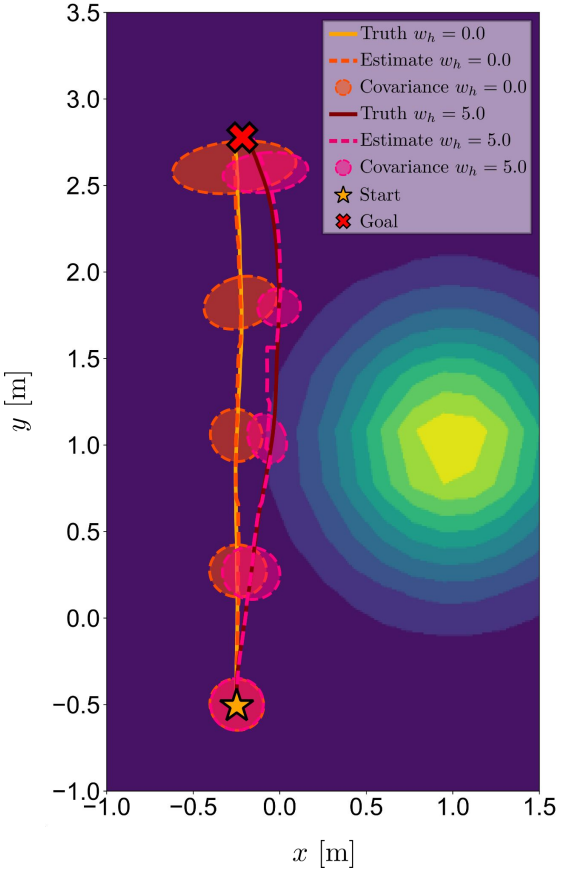
Experiments – Case Study

- Planning algorithm considers **uncertainty measurements more** with a **higher weight**
- When information in **unweighted** the robot trajectory is almost a **straight line**
- When the localization uncertainty is **weighted more heavily** in the objective function, the estimation covariance is **significantly reduced**, leading to **lower covariance over time**



Experiments – Case Study

- A second study with no information weight and high information weight indicate **uncertainty movement**
- The trajectory generated with **no weight** has a **larger covariance** due to the **uniform magnetic field** across the area through which the straight trajectory passes
- With **high weight** the covariance initially **increases**; however, as the robot navigates the edge of the gradient, the **covariance significantly improves**



Experiments – Simulation

The method is validated inside a simulation environment that has:

- a map with distinct anomalies
- randomly selected start and end positions
- and robot and sensor modeled on hardware we have available

The simulated experiments consist of *Monte Carlo* runs to statistically validate localization performance improvements.

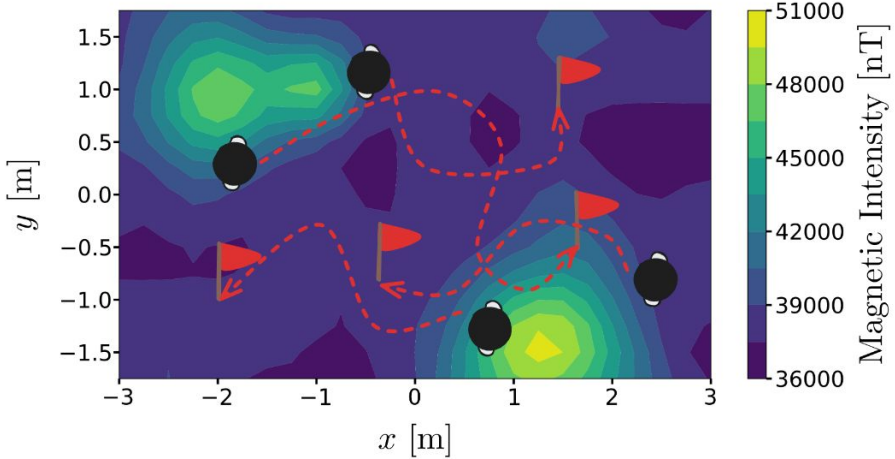


Fig. 10: Monte-Carlo simulation for the navigation algorithm for completion of statistical validation.

Experiments – Simulation Results

Algorithm	RMSE X (m) ↓	RMSE Y (m) ↓	Tr. Stat Cov (m^2) ↓	Completion Rate (%) ↑	Computation Time (ms) ↓
Baseline Methods					
Graph	0.0803	0.0410	0.0051	**100**	537.559
MPC Goal Con.	0.0747	0.0419	0.0055	**100**	3.4382
Information Driven Methods					
Info-Gain Greedy	0.0898	0.1118	0.0251	56	33.0905
E-RRT	0.0799	0.0849	0.0137	94	96.1154
DPSolver	**0.0588**	0.0492	0.0042	64	537.621
InfoMagNav	0.0610	0.0451	**0.0040**	98	**3.2816**

Experiments – Simulation Results

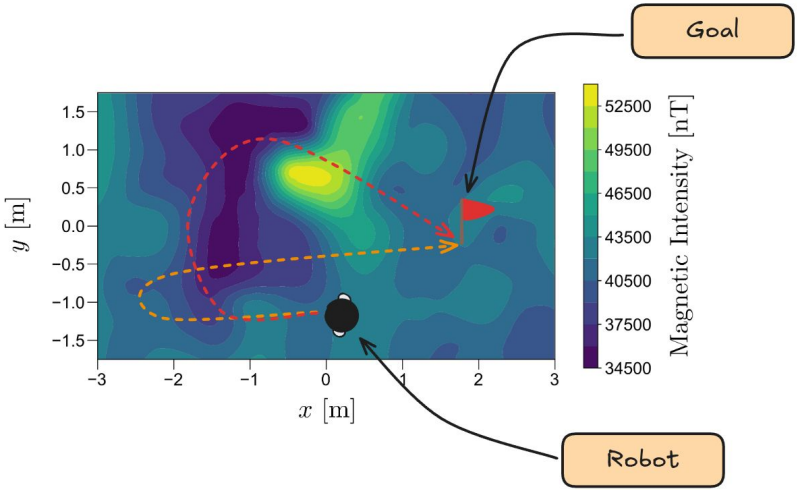
Algorithm	RMSE X (m) ↓	RMSE Y (m) ↓	MAE X (m) ↓	MAE Y (m) ↓	Tr. PF Cov (m²) ↓	Tr. Stat Cov (m²) ↓	Completion Rate (%) ↑
Baseline Methods							
Graph	0.0803	0.0410	0.0600	0.0300	0.0019	0.0051	100
MPC Goal Con.	0.0747	0.0419	0.0601	0.0305	0.0016	0.0055	100
Information Driven Methods							
Info-Gain Greedy	0.0898	0.1118	0.0683	0.0826	0.0017	0.0251	56
E-RRT	0.0799	0.0849	0.0595	0.0616	**0.0015**	0.0137	94
DPSolver	**0.0588**	0.0492	**0.0485**	0.0367	0.0019	0.0042	64
InfoMagNav	0.0610	0.0451	0.0498	0.0323	0.0016	0.0040	98

Experiments – Simulation Results

Algorithm	Path Length (m) ↓	Computation Time (ms) ↓
Baseline Methods		
Graph	2.8147	537.559
MPC Goal Con.	3.0779	3.4382
Information Driven Methods		
Info-Gain Greedy	3.3252	33.0905
E-RRT	3.5616	96.1154
DPSolver	3.1576	537.621
InfoMagNav	3.0639	3.2816

Experiments – Hardware

- The method is run in an indoor setting for real-time hardware validation
- The robot is equipped with a quantum magnetometer for magnetic anomaly detection
- The robot is tracked using a motion capture system for ground truth



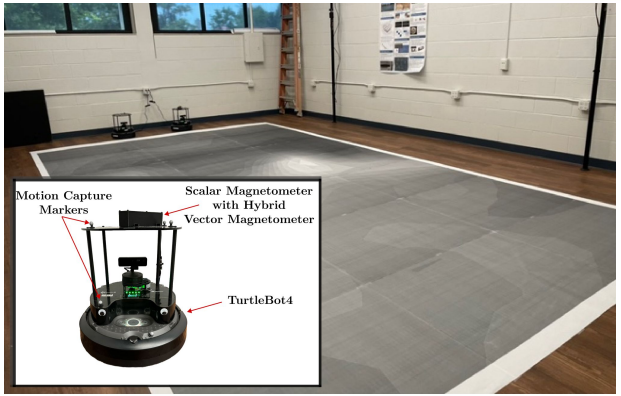
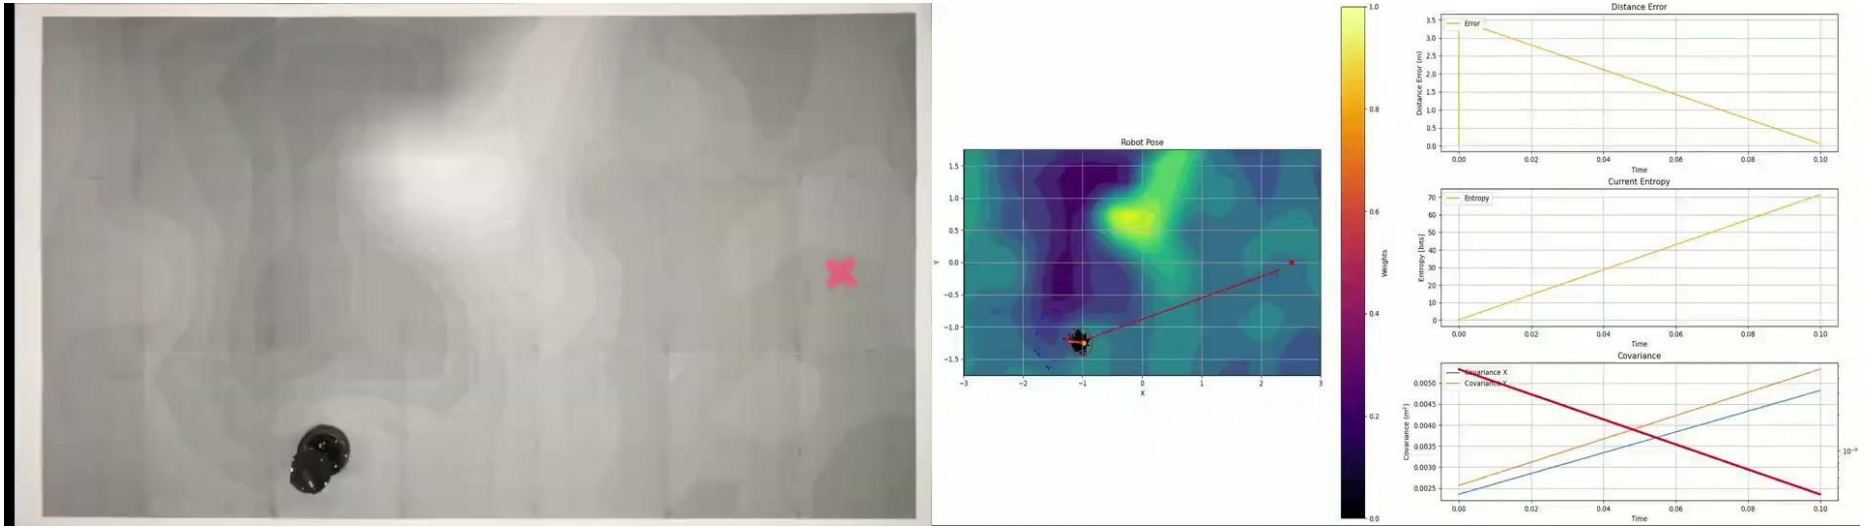


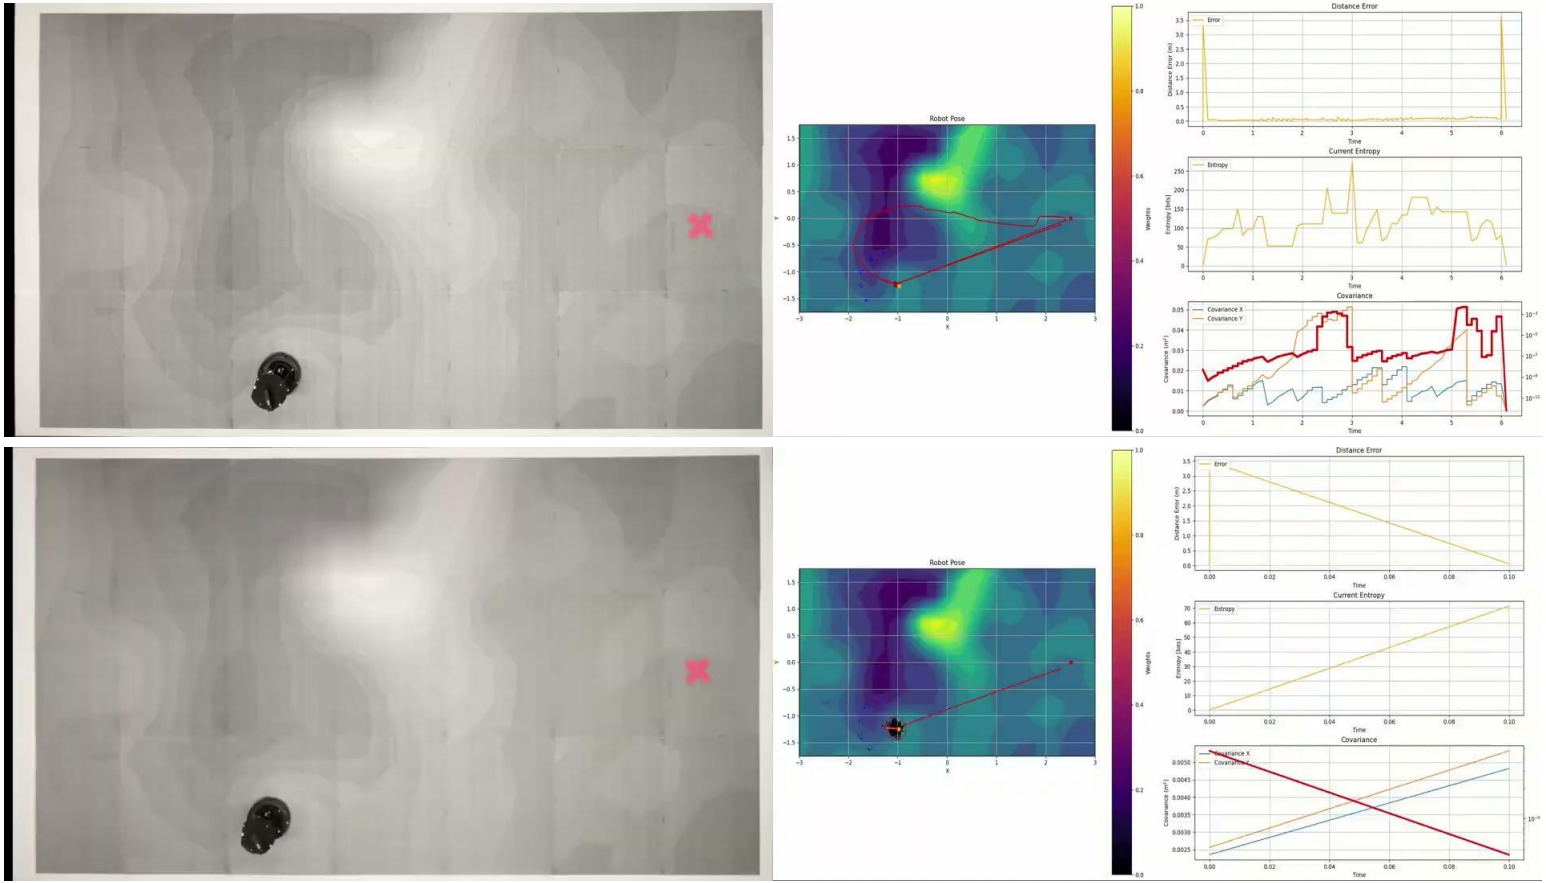
Fig. 10: Hardware experimental setup

Experiments – Hardware Demonstration



- When we consider no information in the navigation, the robot navigates to the goal as fast as it can (with kinematic constraints).
- The localization uncertainty grows when the robot takes this path as it is sub-optimal

Experiments – Hardware Demonstration



Experiments – Hardware Results

- The hardware experiments inform us about the methods robustness to changing environments
- The *Graph* method outperforms all others, but *InfoMagNav* is not significantly worse compared to the SoTA *DPSolver.*

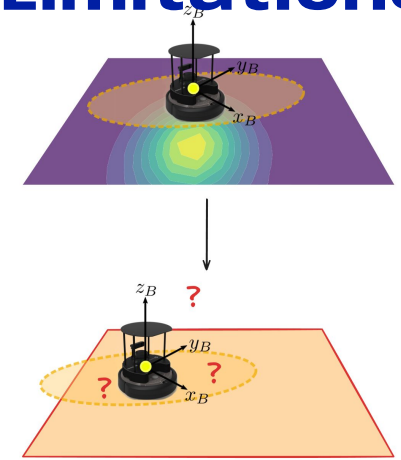
Algorithm	Tr. Statistical Cov (m) ↓	Completion Rate (%) ↑
Baseline Methods		
Graph	**0.0188**	100
MPC Goal Con.	0.0262	100
Information Driven Methods		
DPSolver	0.0546	70
InfoMagNav	0.0217	100

Conclusions and Limitations

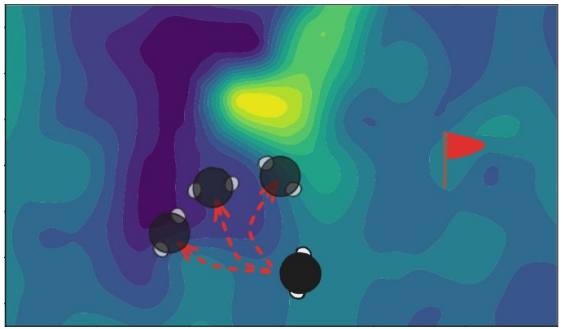
Conclusions

- Presented a **information theoretic** approach to **navigation** in magnetic anomalies
  - The method **applies an estimation** to probabilistic state uncertainty to produce a **calibrated uncertainty propagation** with **reduction**.
- The method presented is significantly more **robust** and **computationally efficient** than the state of the art.
- **Statistically** and **hardware** validated the method to show **robustness** in different maps.

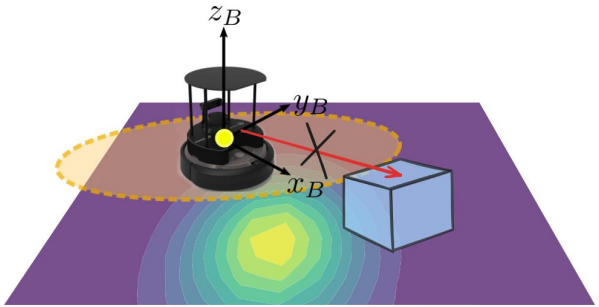
Limitations



Map Dependence



Short Horizon Navigation



Obstacle Detection

Fig. 11: Limitations of the presented method

- The system is completely reliant on the **map** for **localization** and **navigation**
- The system is bound to **short horizon actions** to optimize over
- MagNav cannot inherently **detect any objects** that were not mapped