Information-based State Estimation for Magnet Tracking

Christopher Shallal

Spring 2026

The strategy applied to the state dynamics is to model the magnet movement as a random process. The magnet dynamics are driven by a second-order Markov process which inputs process noise into the system. This second-order Markov process drives the noise in the complementary states; when propagated through the state transition this produces a correlated random-walk in the primary position and orientation states, accurately capturing the stochastic motion of each magnet.

When tracking both the position and orientation (dipole moment) of each magnet, the state vector is extended to incorporate complementary states derived from the second-order Markov model and disturbance states representing the estimated uniform geomagnetic field.

For M magnets indexed by $j = 1, \dots, M$, each magnet's primary state vector $\mathbf{x}_j$ is:

$$\mathbf{x}_j^\mathsf{T} = [x_j \quad y_j \quad z_j \quad \theta_j \quad \phi_j]$$

The vector $\mathbf{x}_j$ contains the *primary* states of the j-th magnet, namely its position coordinates $(x_j,\, y_j,\, z_j)$ and the two angles $(\theta_j,\, \phi_j)$ that specify the direction of its magnetic moment dipole vector.

The auxiliary vector $\tilde{\mathbf{x}}$ holds the *complementary* states introduced by the second-order Markov process. These states are part of the cascaded Markov process that drives the primary variables and where the process noise is injected into, thereby propagating uncertainty from the complementary level into the position and orientation components. Each primary state, including the geomagnetic disturbance, has its own complementary state.

The complementary states derived from the second-order Markov model serve to capture the temporal dynamics and uncertainties associated with the primary states. By modeling the system as a second-order process, the filter accounts for both the current state and its rate of change, enabling more accurate predictions and smoother state estimates. The second-order Markov model facilitates the propagation of process noise through the system, inducing a "random walk" behavior in the primary states. This stochastic behavior enables the filter to adapt to gradual changes and maintain flexibility in tracking the magnet's motion. Incorporating complementary states allows for a more structured covariance matrix, capturing the correlations between positions, orientations, and their derivatives. This leads to improved estimation accuracy and robustness against uncertainties. In addition, this is also a unique way of being able to compute velocity as well rather than methods susceptible to noise, such as a finite-difference derivative.

The vector $\mathbf{x}_{\text{geo}}$ represents the *disturbance* states that are the geomagnetic field and any other slowly varying environmental magnetic biases. Including these terms in the state vector enables the filter to estimate and compensate external disturbances concurrently with the primary and complementary states, improving the accuracy of magnetic dipole state estimates.

Each sensor board is assigned the three states that represent the local, approximately uniform geomagnetic bias measured by that board. Assume B boards, indexed by $b \in \{1, \dots, B\}$, each consisting of an array of magnetic field sensors. Depending on the number of boards, the disturbance vector becomes

$$\mathbf{x}_{\text{geo}}^\mathsf{T} \;=\; [G_{1x} \quad G_{1y} \quad G_{1z} \quad \cdots \quad G_{Bx} \quad G_{By} \quad G_{Bz}], \qquad G_{bx}, G_{by}, G_{bz} \in \mathbb{R},$$

Let $\mathbf{g}_b = [\, G_{bx},\, G_{by},\, G_{bz}\,]^\mathsf{T} \in \mathbb{R}^3$ denote the local, approximately uniform geomagnetic field measured by sensor-board $b \in \{1, \dots, B\}$. Stacking the B board vectors give

$$\mathbf{x}_{\text{geo}}^\mathsf{T} = [\mathbf{g}_1^\mathsf{T} \quad \cdots \quad \mathbf{g}_B^\mathsf{T}] \in \mathbb{R}^{3B}.$$

Disturbance states model the uniform geomagnetic field and any other external magnetic disturbances that affect sensor measurements. The inclusion is critical as by estimating and compensating for external magnetic fields, the filter ensures that the primary state estimates are not biased or distorted by environmental factors. Accounting for this leads to more reliable sensor measurements, as the filter can subtract the estimated disturbances from the raw measurements, isolating the magnetic field contributions from the tracked magnets. Modeling disturbances as part of the state vector allows the filter to adapt to varying environmental conditions, maintaining consistent tracking performance across varying orientations relative to the geomagnetic field. The disturbance states also have their respective complementary states $\tilde{\mathbf{x}}_{\text{geo}}^\mathsf{T}$.

The comprehensive state vector $\mathbf{x}$ combines all the previously defined components:

so that the full state vector is

$$\mathbf{x}^\mathsf{T} = \left[\mathbf{x}_1^\mathsf{T} \ \dots \ \mathbf{x}_M^\mathsf{T} \ \ \mathbf{x}_{\text{geo}}^\mathsf{T} \ \tilde{\mathbf{x}}^\mathsf{T} \ \tilde{\mathbf{x}}_{\text{geo}}^\mathsf{T}\right]$$

During filtering, these components absorb slow variations of the geomagnetic field and other quasi-static magnetic biases local to each board, thereby preventing such disturbances from corrupting the magnet state estimates, and allowing the filter to predict magnet estimates at distances away from the board that would cause the strength generated by the dipoles to be less than the geomagnetic field.

The process model governs the evolution of the state vector over time, incorporating both the dynamics of the magnetic dipole and process noise. This serves as the state propagation layer. The discrete-time state transition is modeled as:

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \Gamma \mathbf{w}_{k-1}$$

where $\mathbf{F}$ is the state transition matrix, Γ is the input matrix, and $\mathbf{w}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$ is the process noise, with state covariance matrix $\mathbf{Q}$.

The state transition matrix $\mathbf{F}$ is constructed based on the time constant τ and the timestep dt:

$$\mathbf{F} = e^{\mathbf{A} \cdot dt}$$

where $\mathbf{A}$ is the continuous-time state transition matrix defined as:

$$\mathbf{A} = \begin{bmatrix} -\frac{1}{\tau}\mathbf{I} & \frac{1}{\tau}\mathbf{I} \\ \mathbf{0} & -\frac{1}{\tau}\mathbf{I} \end{bmatrix}$$

$\mathbf{A}$ has size $N \times N$, where $N = 2(5M + dB)$, $d = 3$. The factor $5M$ counts *primary* (x, y, z, θ, ϕ) states for the M magnets, $dB = 3B$ accounts for the geomagnetic vectors $\mathbf{g}_b$, and the doubling adds the corresponding complementary states introduced by the second-order Markov model. Hence $\mathbf{A}$ and $\mathbf{F}$ are $N \times N$ matrices , Γ is $N \times (N/2)$ and $\mathbf{Q}$ is $(N/2) \times (N/2)$.

The continuous-time process is equivalent to two identical first-order low-pass filters cascaded in series (a critically damped second-order system), each with time constant τ. After discretization with sampling period dt the corresponding poles are located at $e^{-dt/\tau}$. Increasing τ moves the poles closer to 1, so the filter averages a longer window of time-correlated measurements: the process noise injected into the complementary state is attenuated twice before it reaches the position state, reducing the estimator's variance but narrowing the bandwidth of response. Conversely, decreasing τ raises the effective cut-off frequency, yielding faster response to rapid motion at the cost of higher estimation variance. Selecting τ to match the dominant time scale of the muscle movement (or other application) therefore provides a principled trade-off between responsiveness and precision. This intuition drives the state propagation model.

The process noise covariance matrix $\mathbf{Q}$ accounts for uncertainties in the state dynamics and is composed of process noise terms for the primary states: position, orientation, and geomagnetic disturbance components.

The discrete-time noise–input matrix Γ is part of the second–order Markov model that injects white process noise $\mathbf{w}_{k-1}$ only into the complementary layer, so that for every scalar pair $(x, \tilde{x})$

$$\begin{bmatrix} x_k \\ \tilde{x}_k \end{bmatrix} = \begin{bmatrix} \mathrm{e}^{-dt/\tau} & \dfrac{dt}{\tau}\mathrm{e}^{-dt/\tau} \\ 0 & \mathrm{e}^{-dt/\tau} \end{bmatrix} \begin{bmatrix} x_{k-1} \\ \tilde{x}_{k-1} \end{bmatrix} + \begin{bmatrix} 0 \\ 1-\mathrm{e}^{-dt/\tau} \end{bmatrix} \mathbf{w}_{k-1}, \qquad \mathbf{w}_{k-1} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$$

In the discretized second-order Markov model the noise–input matrix Γ has a very specific structure: every row that corresponds to a *primary* state is zero, whereas the rows associated with the complementary states contain the non-zero entry $1 - e^{-dt/\tau}$ derived from the continuous-to-discrete conversion. Thus, the white-noise sequence $\mathbf{w}_k$ is injected *only* into the complementary layer; at the next prediction step the state-transition matrix $\mathbf{F}$ integrates that perturbation forward into the position–orientation layer. This arrangement recreates the intended second-order, correlated random walk that suppresses high-frequency jitter in the primary states while still allowing quick response driven by the process noise. This formulation with the complementary state can also be used to recover velocity using τ and x_k denoting the primary state portion of the state vector.

$$\dot{\mathrm{x}}_k = \frac{1}{\tau}(\tilde{x}_k - \mathrm{x}_k)$$

The measurement model relates the state vector to the sensor measurements, which are the magnetic field vectors observed by each sensor. The measurement at timestep k is given by:

$$\mathbf{z}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k$$

where $\mathbf{h}(\mathbf{x}_k)$ is the nonlinear measurement function and $\mathbf{v}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R})$ is the measurement noise, with noise covariance matrix $\mathbf{R}$ of size $3I \times 3I$ for I magnetic field sensors, each providing 3-axis measurements.

The magnetic field generated by a dipole at a sensor location is given by the fundamental dipole field equation:

$$\mathbf{B}_j^i = \frac{\mu_0}{4\pi}\left(\frac{3\mathbf{r}_{ij}(\mathbf{m}_j \cdot \mathbf{r}_{ij})}{|\mathbf{r}_{ij}|^5} - \frac{\mathbf{m}_j}{|\mathbf{r}_{ij}|^3}\right)$$

where μ_0 is the permeability of free space, $\mathbf{r}_{ij} = [s_x^i - x_j,\ s_y^i - y_j,\ s_z^i - z_j]^\mathsf{T}$ is the distance vector from dipole j to sensor i, and $\mathbf{m}_j = [m_{xj},\ m_{yj},\ m_{zj}]^\mathsf{T}$ is the dipole-moment vector of magnet j.

The corresponding unit vector $\hat{\mathbf{m}}_j$ is obtained by the standard spherical-to-Cartesian mapping:

$$\hat{\mathbf{m}}_j = \begin{bmatrix} \sin\theta_j \cos\phi_j \\ \sin\theta_j \sin\phi_j \\ \cos\theta_j \end{bmatrix} = \begin{bmatrix} \hat{m}_{xj} \\ \hat{m}_{yj} \\ \hat{m}_{zj} \end{bmatrix}$$

These equations calculate the estimated magnetic field vector at sensor i given the position and magnetic moment vector of dipole j, based on the distance vector between the magnet and sensor locations in the defined reference frame. Throughout this work the magnitude of every magnet's dipole moment is treated as a *known* scalar rather than an estimated state. The permeability of free space μ_0 and any other scalar terms are implicitly included into this lumped value. In all field equations we therefore write: $\mathbf{m}_j = m_0\, \hat{\mathbf{m}}_j$ where $\hat{\mathbf{m}}_j$ is the unit vector encoded by the orientation states (θ_j, ϕ_j). Process noise is not applied to m_0, reflecting the assumption that the magnetization of each bead remains constant. The total magnetic field at sensor i is the sum of contributions from all dipoles and any disturbance:

$$\mathbf{z}^i = \sum_{j=1}^{M} \mathbf{B}_j^i + \mathbf{G}^i + \mathbf{v}^i$$

where $\mathbf{G}^i$ represents the external disturbance (such as the geomagnetic field) at sensor i. Accurate computation of the Jacobian matrix $\mathbf{H}_k$ is critical for the tracking performance. The Jacobian is calculated due to the nonlinearities in the measurement model $\mathbf{h}(\mathbf{x}_k)$ as a way to linearize the function. The nominal operation point that the Jacobian is linearized about is the *a priori* (predicted) state:

$$\mathbf{x}_{k|k-1} = \mathbf{F}\,\mathbf{x}_{k-1|k-1}$$

The resulting Jacobian is

$$\mathbf{H}_k = \left.\frac{\partial \mathbf{h}}{\partial \mathbf{x}}\right|_{\mathbf{x}=\mathbf{x}_{k|k-1}} \in \mathbb{R}^{m \times n}$$

where n and m are the dimensions of the state and measurement vectors, respectively. Accurate, analytical evaluation of $\mathbf{H}_k$ is essential. It supplies the first-order sensitivity of the measurement to the state and determines the information gain used to map the measurement residual back into state space. Any error in $\mathbf{H}_k$ therefore propagates directly into the state update and degrades tracking accuracy. Therefore, we analytically compute the Jacobian rather than relying on numerical methods such as finite difference which reduce accuracy and increase latency.

Each dipole contributes the five *primary* states $(x_j, y_j, z_j, \theta_j, \phi_j)$ and the corresponding five *complementary* states. With I sensors, M dipoles and B sensor boards the measurement Jacobian has

$$\mathbf{H}_k \in \mathbb{R}^{3I \times (5M + 5M + 3B + 3B)}$$

Using the sensor index $i = 1, \ldots, I$ and dipole index $j = 1, \ldots, M$ we write

$$\mathbf{H}_k = \begin{bmatrix} \mathbf{H}_{11} & \ldots & \mathbf{H}_{1M} & \mathbf{G}_1 & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{H}_{21} & \ldots & \mathbf{H}_{2M} & \mathbf{G}_2 & \mathbf{0} & \ldots & \mathbf{0} \\ \vdots & \ddots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{H}_{I1} & \ldots & \mathbf{H}_{IM} & \mathbf{G}_I & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}$$

where $\mathbf{H}_{ij} \in \mathbb{R}^{3\times 5}$ carries the derivatives of the field measured by sensor i with respect to the five primary states of dipole j; the $3I \times (5M + 3B)$ zero block corresponds to the complementary states, for which $\partial \mathbf{h}/\,\partial \tilde{\mathbf{x}} = \mathbf{0}$; $\mathbf{G}_i \in \mathbb{R}^{3\times 3B}$ inserts the 3×3 identity in the three columns associated with the board that contains sensor i (geomagnetic disturbance states) and zeros elsewhere. The assumption is that the disturbance field is seen uniformly across all the sensors in the same way, resulting in the Jacobian being simplified to an identity matrix. For a particular sensor–dipole pair the 3×5 block is

$$\mathbf{H}_{ij} = \begin{bmatrix} \frac{\partial \mathbf{B}_j^i}{\partial x_j} & \frac{\partial \mathbf{B}_j^i}{\partial y_j} & \frac{\partial \mathbf{B}_j^i}{\partial z_j} & \frac{\partial \mathbf{B}_j^i}{\partial \theta_j} & \frac{\partial \mathbf{B}_j^i}{\partial \phi_j} \end{bmatrix}$$

We write the fixed-magnitude dipole moment as $\mathbf{m}_j = m_0 \widehat{\mathbf{m}}_j$, where $\widehat{\mathbf{m}}_j$ is the unit vector encoded by the orientation states (θ_j, ϕ_j).

The magnetic-field measurement function $\mathbf{h}(\mathbf{x})$ depends only on the current positions, orientation angles, and disturbance field estimation; it is independent of the complementary states introduced for the second-order Markov process. Therefore $\partial \mathbf{h} / \partial \tilde{\mathbf{x}} = \mathbf{0}$, and the corresponding columns in $\mathbf{H}_k$ are identically zero. The magnetic-field measurement $\mathbf{h}(\mathbf{x})$ depends only on the *instantaneous* positions, orientations, and board-level geomagnetic biases; it contains no explicit dependence on the complementary states that were introduced solely to model the process dynamics. Consequently, the partial derivatives with respect to the complementary states are zero: $\frac{\partial \mathbf{h}}{\partial \tilde{\mathbf{x}}} = \mathbf{0}$, and the corresponding columns in $\mathbf{H}_k$ are identically zero. Physically this is consistent with the fact that a magnetic field sensor responds to the magnetic field produced by the magnets' present configuration, not to how fast that field is changing (which would be a derivative of the flux).

Implementation

The extended information filter (INFO) operates with the information matrix $\mathbf{Y}$ and the information vector $\mathbf{y}$, which are the inverse of the covariance matrix $\mathbf{P}$ and the product of $\mathbf{P}^{-1}$ with the state mean $\hat{\mathbf{x}}$, respectively. Large diagonal values within the covariance matrix intuitively imply a lack of precision or a larger degree of uncertainty regarding those state estimate components. Conversely, large diagonal elements imply a high degree of precision, or a smaller degree of uncertainty, of those state estimate components. Think of the covariance $\mathbf{P}$ as describing state uncertainty, the goal is to minimize the covariance matrix. When viewing the problem from the information lens, the goal is to maximize information gain $\mathbf{Y}$.

First, to initialize the state vector $\mathbf{x}_0$, the covariance matrix $\mathbf{P}_0$, and compute the initial information matrix and vector $\mathbf{Y}_0$ and $\mathbf{y}_0$, the system can be driven into steady state by

$$\mathbf{P}_{k+1} = \mathbf{F}\,\mathbf{P}_k\,\mathbf{F}^{\mathsf{T}} + \Gamma\,\mathbf{Q}\,\Gamma^{\mathsf{T}}, \qquad k = 0,1,\dots$$

starting from $\mathbf{P}_0 = \mathbf{I}$ and continuing until $\mathbf{P}$ has reached steady-state (large number of k iterations). And, for the full state vector of dimension $N = 2(5M + 3B)$,

$$\mathbf{P}_0 = \mathbf{P}^{\star} \in \mathbb{R}^{N \times N}$$

Finally, we convert the steady-state covariance to information form:

$$\boxed{\mathbf{Y}_0 = \mathbf{P}_0^{-1}, \qquad \mathbf{y}_0 = \mathbf{Y}_0\,\mathbf{x}_0}$$

so that the filter starts with information consistent with the assumed process model and noise.

In the prediction step, the filter propagates the state and information matrices forward based on the previously described process model. The state prediction incorporates the state transition matrix $\mathbf{F}$, the process noise covariance $\mathbf{Q}$ and process input Γ.

$$
\begin{aligned}
\mathbf{x}_{k|k-1} &= \mathbf{F}\mathbf{x}_{k-1|k-1} \\
\mathbf{P}_{k|k-1} &= \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^{\top} + \Gamma\mathbf{Q}\Gamma^{\top} \\
\mathbf{Y}_{k|k-1} &= \mathbf{P}_{k|k-1}^{-1} \\
\mathbf{y}_{k|k-1} &= \mathbf{Y}_{k|k-1}\mathbf{x}_{k|k-1}
\end{aligned}
$$

The inverse of the predicted covariance matrix $\mathbf{Y}_{k|k-1}$ is computed by propagating the previous covariance through the state transition matrix $\mathbf{F}$ and adding the process noise covariance $\mathbf{Q}$ through process input matrix Γ. The predicted information vector $\mathbf{y}_{k|k-1}$ is obtained by multiplying the predicted information matrix $\mathbf{Y}_{k|k-1}$ with the propagated state vector $\mathbf{F}\mathbf{x}_{k-1|k-1}$. This step prepares the information vector for incorporation of new measurements. Defining the measurement residual $\boldsymbol{\nu}_k = \mathbf{z}_k - \mathbf{h}(\hat{\mathbf{x}}_{k|k-1})$ and the *augmented innovation*

$$
\tilde{\mathbf{z}}_k = \boldsymbol{\nu}_k + \mathbf{H}_k\,\hat{\mathbf{x}}_{k|k-1}, \qquad \mathbf{H}_k = \left.\frac{\partial \mathbf{h}}{\partial \mathbf{x}}\right|_{\hat{\mathbf{x}}_{k|k-1}}
$$

The information-domain update is then

$$
\begin{aligned}
\mathbf{S}_k &= \mathbf{H}_k^{\top}\mathbf{R}^{-1}\mathbf{H}_k, \\
\mathbf{Y}_k &= \mathbf{Y}_{k|k-1} + \mathbf{S}_k, \\
\mathbf{y}_k &= \mathbf{y}_{k|k-1} + \mathbf{H}_k^{\top}\mathbf{R}^{-1}\tilde{\mathbf{z}}_k,
\end{aligned}
$$

Here $\mathbf{S}_k$ is the innovation information contributed by the new measurement, while $\mathbf{H}_k^{\top}\mathbf{R}^{-1}\tilde{\mathbf{z}}_k$ is the information *increment* derived from the residual. When $\mathbf{h}(\mathbf{x})$ is linear the term $\mathbf{H}_k\hat{\mathbf{x}}_{k|k-1}$ cancels, recovering the standard linear information filter equations. This additional term is an artifact of the first-order Taylor series used to convert the nonlinear measurement function into the equivalent *linear* form required by the information space matrix operations.

This concise formulation shows that all covariance and residual effects are accumulated in information space; the updated mean and covariance are recovered only once per timestep *k* through the inversion of $\mathbf{Y}_k$. The state estimate update step finalizes the incorporation of the new measurements by computing the updated state vector and covariance matrix from the information matrix and information vector.

$$
\begin{aligned}
\mathbf{x}_{k|k} &= \mathbf{Y}_k^{-1}\mathbf{y}_k \\
\mathbf{P}_{k|k} &= \mathbf{Y}_k^{-1}
\end{aligned}
$$

The updated state estimate $\mathbf{x}_{k|k}$ is obtained by multiplying the inverse of the updated information matrix $\mathbf{Y}_k^{-1}$ with the updated information vector $\mathbf{y}_k$. This operation yields the most probable state estimate given the prior predictions and the new measurements. The updated covariance matrix $\mathbf{P}_{k|k}$ is simply the inverse of the updated information matrix $\mathbf{Y}_k^{-1}$. This matrix represents the uncertainty associated with the updated state estimate. The process then loops again for the next state estimate.