Strapdown Seeker Line-of-Sight Rate Estimation via Pixel-Plane Kalman Filtering and Gyro Fusion

Rasit Evduzen

July 2026

Abstract

Proportional-navigation (PN) guidance requires an accurate estimate of the inertial line-of-sight (LOS) rotation rate between an interceptor and its target. In a strapdown (body-fixed, non-gimbaled or partially stabilized) electro-optical seeker, this rate cannot be measured directly: the seeker only observes the target's position in the image plane, while the vehicle's own rotation is sensed separately by a gyro. We present a complete estimation pipeline that (i) tracks the target's normalized image-plane coordinates with a constant-velocity Kalman filter, (ii) reconstructs the camera-frame LOS rate from the filtered pixel state, and (iii) fuses it with the measured body angular rate through the rigid-body transport theorem to recover the inertial LOS rate. The pipeline is validated across six progressively realistic simulation scenarios: a deterministic zero-noise baseline, a maneuvering target combined with a finite-bandwidth gimbal, a 1000-run Monte Carlo sweep under pixel and gyro noise/bias, a closed-loop three-dimensional true-PN intercept guided entirely by the estimated LOS rate, a Kalman process-noise sensitivity sweep exposing the classic tracking-lag/noise-rejection trade-off, and a pixel-channel latency sweep quantifying the effect of uncompensated tracker delay. The estimator reproduces the ground-truth LOS rate to within 3.13×10^{-4} deg/s RMS in the noise-free baseline, remains accurate under target maneuvers and gimbal lag, achieves a 0.899 ± 0.028 deg/s RMS error under realistic sensor noise, and drives a closed-loop intercept to a 1.71 m miss distance.

Keywords: strapdown seeker, line-of-sight rate estimation, Kalman filter, proportional navigation, sensor fusion, missile guidance.

1 Introduction

True proportional navigation (PN) commands an interceptor's lateral acceleration in proportion to the closing velocity and the inertial rate of rotation of the line of sight (LOS) to the target, $\mathbf{a}_{\text{cmd}} = NV_c\,\boldsymbol{\omega}_{\text{LOS}} \times \hat{\mathbf{r}_w}$ [2]. The quality of the LOS-rate estimate directly determines guidance performance, so its accurate reconstruction from onboard sensors is a central problem in missile and interceptor design [3].

A gimbaled, inertially stabilized seeker can measure the LOS rate almost directly from the gimbal torque-to-balance loop. A *strapdown* seeker, in contrast, has its camera rigidly (or semi-rigidly, through a finite-bandwidth pointing controller) mounted to the vehicle body. It observes only the target's bearing in the image plane; the platform's own rotation must be measured separately by a strapdown gyro and removed from the observed image motion before an inertial LOS rate can be recovered [3]. This means fusing two different sensor channels, a vision-based bearing tracker and an inertial rate gyro, through the kinematics of a rotating, translating camera.

This paper presents a complete pipeline for this fusion problem and tests it under several conditions that simple treatments often skip: a maneuvering, non-cooperative target; a seeker head with finite pointing bandwidth (so the target is not always centered in the image); pixel tracker noise together with gyro noise and bias; closed-loop use of the estimate to actually fly a PN intercept; the Kalman tuning trade-off between tracking lag and noise rejection; and the effect of uncompensated image-processing latency.

The paper is organized as follows. Section 2 formulates the LOS-rate reconstruction problem and fixes notation. Section 3 derives the pixel-plane Kalman filter and the gyro-fusion reconstruction formula. Section 4 presents the six simulation scenarios and their results. Section 5 concludes.

2 Problem Formulation

2.1 Frames and Camera Model

Three right-handed frames are used: the inertial/NED world frame W, the interceptor body frame B, and the camera frame C, related by $R_{WB}(t) \in SO(3)$ (attitude, time-varying) and a constant mount rotation R_{BC} together with a small lever arm $\mathbf{d}_{BC}$ from the body origin to the camera center.

The camera is modeled as an ideal pinhole with focal lengths f_x, f_y and principal point (c_x, c_y). For a target at camera-frame position $\mathbf{d}_c = [X_c, Y_c, Z_c]^\top$ (with $Z_c > 0$ the range along the boresight), the pixel projection is

$$u = f_x \frac{X_c}{Z_c} + c_x, \qquad v = f_y \frac{Y_c}{Z_c} + c_y, \tag{1}$$

and the corresponding *normalized image coordinates* are recovered by the inverse pinhole map

$$o_x = \frac{u - c_x}{f_x}, \qquad o_y = \frac{v - c_y}{f_y}, \tag{2}$$

so that the (unnormalized) target bearing vector in the camera frame is proportional to $\mathbf{s} = [o_x,\ o_y,\ 1]^\top$, independent of the unknown range Z_c.

2.2 Goal

Let $\mathbf{r}_w(t) \in \mathbb{S}^2$ be the unit LOS vector from the camera to the target, expressed in the inertial frame. The PN guidance law requires the inertial LOS rotation rate $\boldsymbol{\omega}_{\mathrm{LOS}}$, defined through $\dot{\mathbf{r}_w} = \boldsymbol{\omega}_{\mathrm{LOS}} \times \mathbf{r}_w$, equivalently

$$\boldsymbol{\omega}_{\mathrm{LOS}} = \mathbf{r}_w \times \dot{\mathbf{r}_w}, \tag{3}$$

which follows from $\mathbf{r}_w \times (\boldsymbol{\omega}_{\mathrm{LOS}} \times \mathbf{r}_w) = \boldsymbol{\omega}_{\mathrm{LOS}}(\mathbf{r}_w \cdot \mathbf{r}_w) - \mathbf{r}_w(\mathbf{r}_w \cdot \boldsymbol{\omega}_{\mathrm{LOS}}) = \boldsymbol{\omega}_{\mathrm{LOS}}$ for a unit vector with $\boldsymbol{\omega}_{\mathrm{LOS}} \perp \mathbf{r}_w$. The seeker has access only to: (a) noisy pixel measurements (u_k, v_k)

of the target at discrete times $t_k = kT_s$, and (b) a body-rate (gyro) measurement $\boldsymbol{\omega}_b(t_k)$. Range, target velocity, and $\dot{\mathbf{r}_w}$ are *not* directly observable. The goal is to reconstruct $\boldsymbol{\omega}_{\text{LOS}}(t_k)$ from these two sensor streams alone.

3 Proposed Method

3.1 Pixel-Plane Kalman Filter

The filter state is the normalized image coordinates and their rates, $\mathbf{x} = [o_x,\ o_y,\ \dot{o}_x,\ \dot{o}_y]^\top \in \mathbb{R}^4$, propagated by a discrete constant-velocity (CV) model with sample time T_s:

$$A = \begin{bmatrix} 1 & 0 & T_s & 0 \\ 0 & 1 & 0 & T_s \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \qquad H = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix}. \tag{4}$$

Process noise follows the standard discrete white-noise-acceleration model,

$$Q = q \cdot \begin{bmatrix} T_s^4/4 & 0 & T_s^3/2 & 0 \\ 0 & T_s^4/4 & 0 & T_s^3/2 \\ T_s^3/2 & 0 & T_s^2 & 0 \\ 0 & T_s^3/2 & 0 & T_s^2 \end{bmatrix}, \qquad R = \sigma_{px}^2/f_x^2 \cdot I_2, \tag{5}$$

where q (`q_scale`) is a tunable process-noise scale and σ_{px} is the tracker's pixel-noise standard deviation. The filter uses the standard predict/update equations with a Joseph-form covariance update for numerical robustness:

$$\mathbf{x}_{k|k-1} = A\mathbf{x}_{k-1|k-1}, \qquad P_{k|k-1} = AP_{k-1|k-1}A^\top + Q, \tag{6}$$

$$K_k = P_{k|k-1}H^\top (HP_{k|k-1}H^\top + R)^{-1}, \tag{7}$$

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + K_k(\mathbf{z}_k - H\mathbf{x}_{k|k-1}), \tag{8}$$

$$P_{k|k} = (I - K_kH)P_{k|k-1}(I - K_kH)^\top + K_kRK_k^\top. \tag{9}$$

Measurement noise is floored at $\sigma_{px} = 10^{-6}$px-equivalent to avoid a singular gain in the deterministic (zero-noise) test scenario.

3.2 LOS-Rate Reconstruction: Gyro-Pixel Fusion

At every filter step the camera-frame bearing and its rate are recovered from the filter state via

$$\mathbf{s} = [\hat{o}_x,\ \hat{o}_y,\ 1]^\top, \quad \mathbf{r}_c = \frac{\mathbf{s}}{\|\mathbf{s}\|}, \quad \dot{\mathbf{r}_c} = \frac{\dot{\mathbf{s}}}{\|\mathbf{s}\|} - \frac{\mathbf{s}(\mathbf{s}^\top\dot{\mathbf{s}})}{\|\mathbf{s}\|^3}, \qquad \dot{\mathbf{s}} = [\hat{\dot{o}}_x,\ \hat{\dot{o}}_y,\ 0]^\top, \tag{10}$$

the standard derivative of a normalized vector. Because the camera mount is rigid, $\mathbf{r}_b = R_{BC}\mathbf{r}_c$ and $\dot{\mathbf{r}_b} = R_{BC}\dot{\mathbf{r}_c}$. The rigid-body transport theorem gives the inertial-frame derivative of a body-fixed vector as

$$\dot{\mathbf{r}_w} = R_{WB}\big(\boldsymbol{\omega}_b \times \mathbf{r}_b + \dot{\mathbf{r}_b}\big), \tag{11}$$

i.e., the rate seen in the body frame ($\dot{\mathbf{r}}_b$, from the pixel channel) plus the rate induced by the body's own rotation ($\boldsymbol{\omega}_b \times \mathbf{r}_b$, from the gyro channel), rotated into the world frame. Substituting into (3) with $\mathbf{r}_w = R_{WB}\mathbf{r}_b$ gives the final reconstruction

$$\boxed{\boldsymbol{\omega}_{\text{LOS}} = \mathbf{r}_w \times \dot{\mathbf{r}}_w}, \tag{12}$$

which is the quantity used throughout the results in Section 4. Equation (11) shows explicitly that $\boldsymbol{\omega}_{\text{LOS}}$ requires *both* sensor channels: the pixel-derived term $\dot{\mathbf{r}}_b$ alone omits the platform's own rotation, and the gyro-derived term $\boldsymbol{\omega}_b \times \mathbf{r}_b$ alone omits the target's apparent motion in the image. When the seeker's pointing controller keeps the target near boresight ($\mathbf{r}_b \approx \text{const}$), the gyro term dominates; when the target drifts off-boresight, the pixel-rate term becomes significant (Section 4.3).

3.3 Estimation Algorithm

The full per-sample procedure is:

1. Convert the raw pixel measurement (u_k, v_k) to normalized image coordinates via the inverse pinhole map (2).

2. Update the pixel-plane Kalman filter, Eqs. (4)–(9), to obtain $\hat{\mathbf{x}}_{k|k} = [\hat{o}_x, \hat{o}_y, \hat{\dot{o}}_x, \hat{\dot{o}}_y]$.

3. Form the camera-frame bearing and its rate, Eq. (10).

4. Rotate into the body frame via the (constant) mount rotation R_{BC}.

5. Fuse with the gyro measurement $\boldsymbol{\omega}_b(t_k)$ through the transport theorem, Eq. (11), and rotate into the world frame with the current attitude $R_{WB}(t_k)$.

6. Recover $\boldsymbol{\omega}_{\text{LOS}}(t_k)$ via Eq. (12).

This is a single forward pass with $\mathcal{O}(1)$ cost per sample; no iteration or batch processing is required.

4 Simulation Results

4.1 Common Setup

An interceptor and a target move in 3D (NED). The interceptor body attitude follows a pursuit-pointing law (body $-Z$ toward the instantaneous LOS), either ideally (ideal boresight-locked gimbal) or through a first-order lag of time constant τ_{gimbal} on the boresight axis, to emulate a finite-bandwidth stabilization loop. Body angular rate (the gyro "truth") is extracted from consecutive attitude matrices via the matrix logarithm, $\boldsymbol{\omega}_b = \text{vee}\big(\log(R_{k-1}^{\top} R_k)\big)/T_s$, where $\log(\cdot)$ maps a rotation matrix to its skew-symmetric matrix logarithm and $\text{vee}(\cdot)$ reads the three independent entries of that skew-symmetric matrix off as a vector. Camera intrinsics and the shared simulation parameters are listed in Table 1.

Table 1: Common camera and simulation parameters.

Parameter	Symbol	Value
Image size	$\mathrm{img}_W \times \mathrm{img}_H$	640×480 px
Horizontal FOV	fov_h	$40°$
Focal length	$f_x = f_y$	$(\mathrm{img}_W/2)/\tan(\mathrm{fov}_h/2)$ px
Sample period	T_s	$0.02\,\mathrm{s}$ (50 Hz)
Camera lever arm	$\mathbf{d}_{BC}$	$[0.072,\ 0,\ 0]\,\mathrm{m}$
Camera mount rotation	R_{BC}	$\begin{bmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}$
KF process-noise scale (default)	q	1.0
KF initial covariance	P_0	I_4

4.2 Deterministic Baseline

The first scenario validates the reconstruction pipeline itself, isolated from any sensor imperfection: zero pixel noise, an ideal (zero-lag) boresight-locked gimbal, and an ideal gyro. The interceptor flies a constant-velocity leg while a laterally-offset constant-velocity target generates a smoothly time-varying LOS rate. Figure 1 shows the engagement geometry, and Figure 2 compares the estimated LOS rate against ground truth (obtained by numerically differentiating the true LOS unit vector) on all three NED axes. The estimator reproduces the truth curve visually exactly, with an RMS error of 3.13×10^{-4} deg/s; the small residual is dominated by a one-sample Kalman-filter initialization transient (the filter has no velocity information until the second measurement arrives) and by the truncation error of the centered-difference ground truth itself, not by any defect in the reconstruction formula.

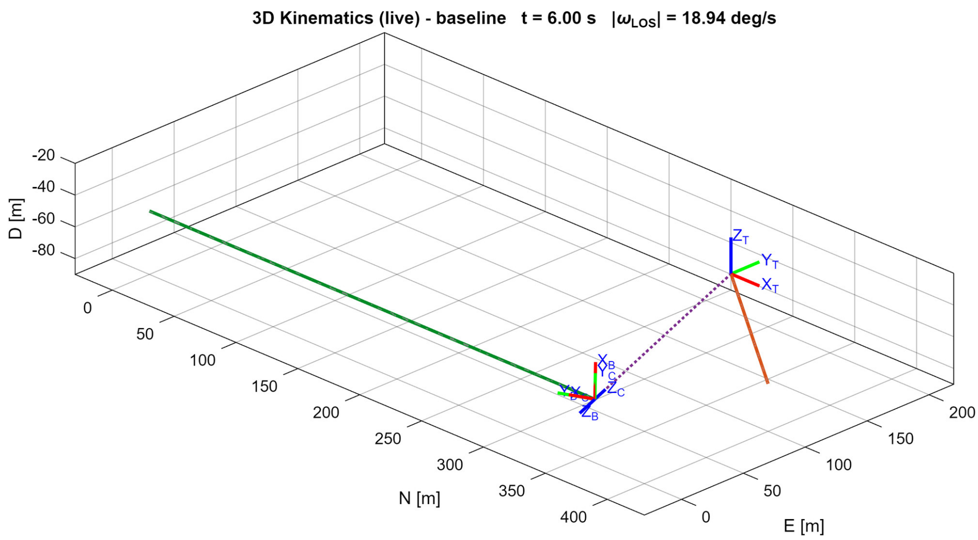

Figure 1: Engagement geometry, deterministic baseline: interceptor (green) and target (orange) trajectories with body (B), camera (C), and target (T) frames, and the instantaneous line of sight (purple).

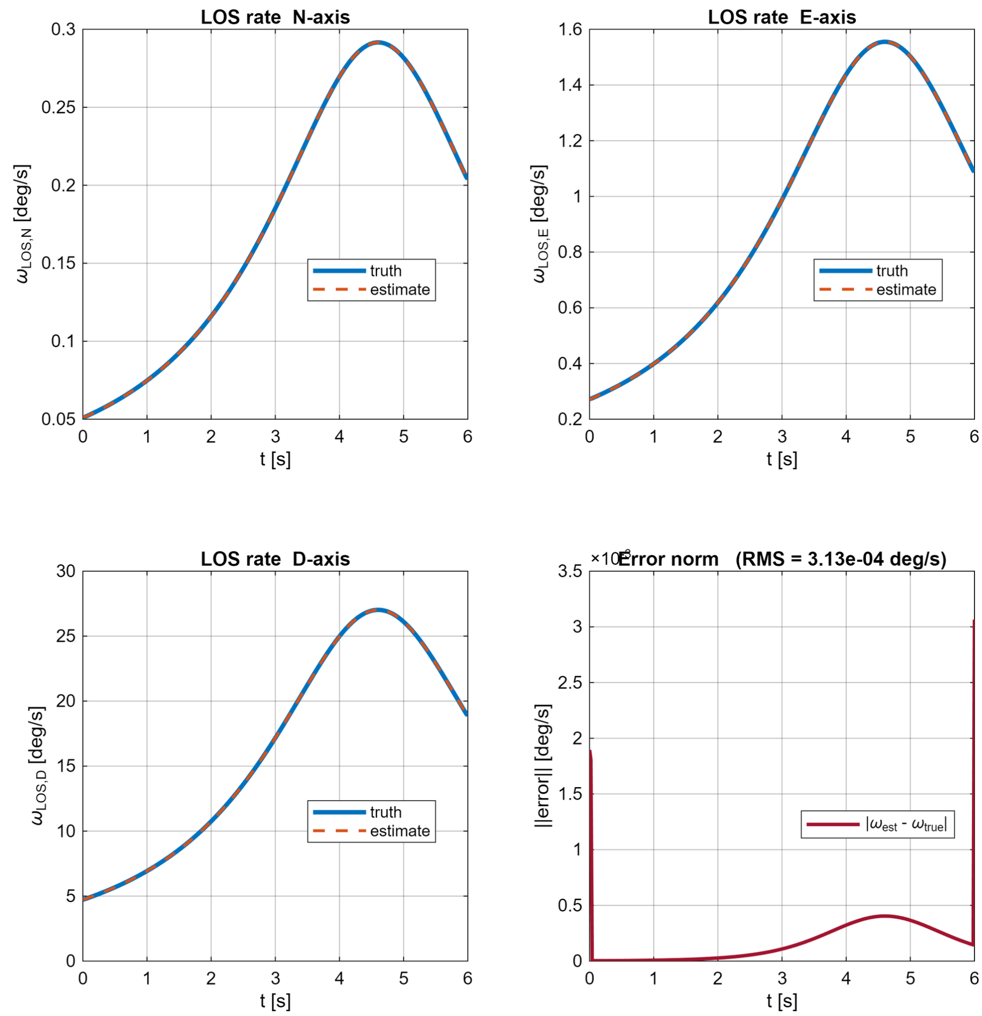

Figure 2: LOS rate estimation, deterministic baseline (zero measurement noise, ideal boresight-locked gimbal): per-axis (N/E/D) truth vs. estimate and the error norm (RMS $= 3.13 \times 10^{-4}$ deg/s).

4.3 Maneuvering Target and Finite-Bandwidth Gimbal

The ideal boresight-locked baseline keeps the target essentially at the image center at all times, so the pixel-rate term of Eq. (11) is negligible and the reconstruction is effectively gyro-only. This is a favorable but unrealistic case. To exercise the pixel channel, the target is given a weaving lateral acceleration of $25\,\mathrm{m/s^2}$ (onset at $t = 2\,\mathrm{s}$) and the gimbal is given a first-order pointing lag of $\tau_{\mathrm{gimbal}} = 0.25\,\mathrm{s}$, which prevents perfect boresight tracking. Figure 3 confirms this: the target drifts up to $\sim 5.9^\circ$ off boresight, giving the pixel channel a genuine, non-trivial signal. Figures 4 and 5 show that, despite the constant-velocity Kalman process model now being mismatched with the true (maneuvering) image-plane motion, the estimator still tracks the LOS rate closely, with an RMS error of 3.49×10^{-1} deg/s against a peak true LOS rate of $\sim$27 deg/s (a relative error of roughly 1.3%).

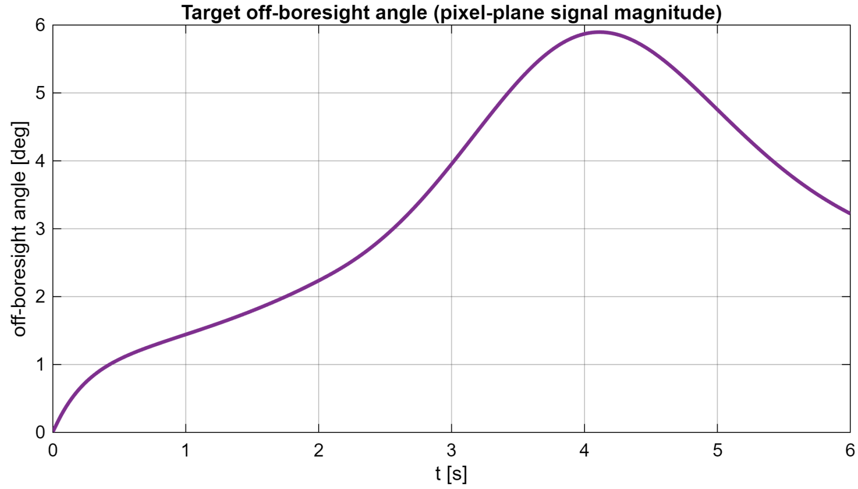

Figure 3: Target off-boresight angle over time under the weaving maneuver and $\tau_{\text{gimbal}} = 0.25\,\text{s}$ pointing lag, confirming a non-trivial pixel-plane signal (peak $\approx 5.9^\circ$).

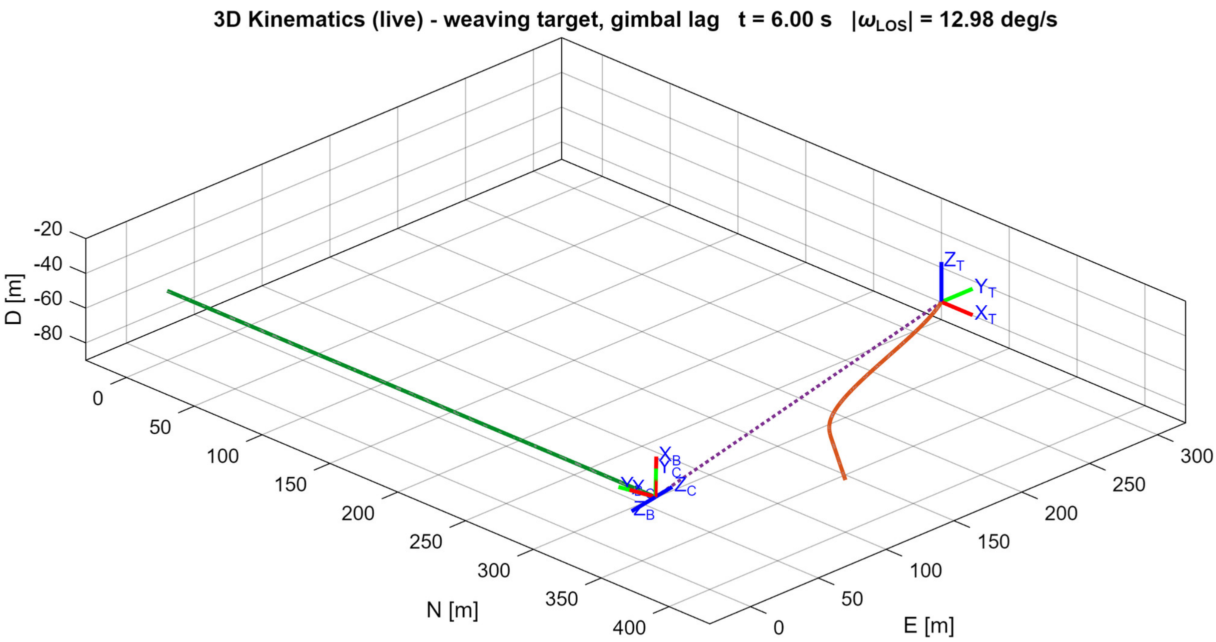

Figure 4: Engagement geometry for the weaving-target, finite-bandwidth gimbal case. The curved target track (orange) reflects the lateral weaving maneuver.

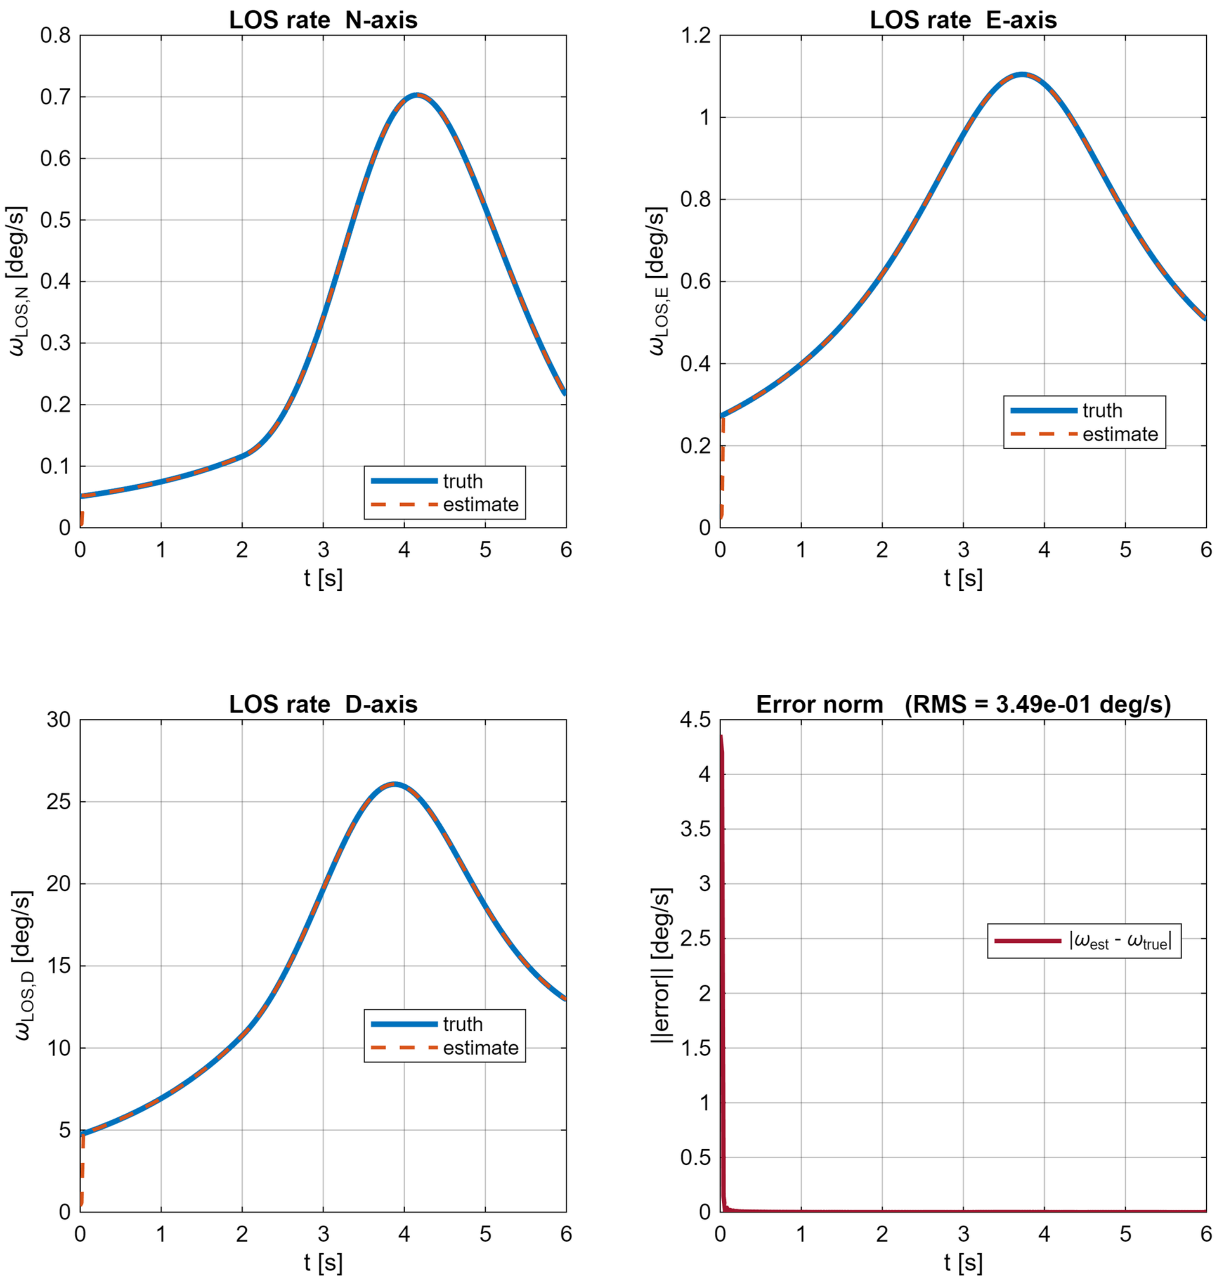

Figure 5: LOS rate estimation under target maneuver and gimbal lag: truth vs. estimate per axis and error norm (RMS $= 3.49 \times 10^{-1}$ deg/s). Both the pixel-rate and the gyro terms of Eq. (11) are now exercised.

4.4 Sensor Noise and Bias: Monte Carlo

Realistic tracker pixel noise ($\sigma_{px} = 0.3$ px, 1σ) and gyro noise/bias ($\sigma_{\text{gyro}} = 0.05\,°/\text{s}$ angle random walk plus a $0.02\,°/\text{s}$ constant bias drawn independently per run) are injected on the constant-velocity baseline trajectory (ideal gimbal, isolating sensor-noise robustness from the maneuver-induced model mismatch of Section 4.3). Over $N = 1000$ independent Monte Carlo runs, Figure 6 shows the resulting RMS-error distribution (mean 0.899, std. 0.028 deg/s), an example run's per-axis error time history and LOS-rate magnitude tracking, and the RMS error achieved on each individual run. The tight spread of the distribution (std/mean $\approx 3\%$) indicates that performance is dominated by the fixed noise levels rather than by particular noise realizations, i.e. the filter behaves consistently run to run.

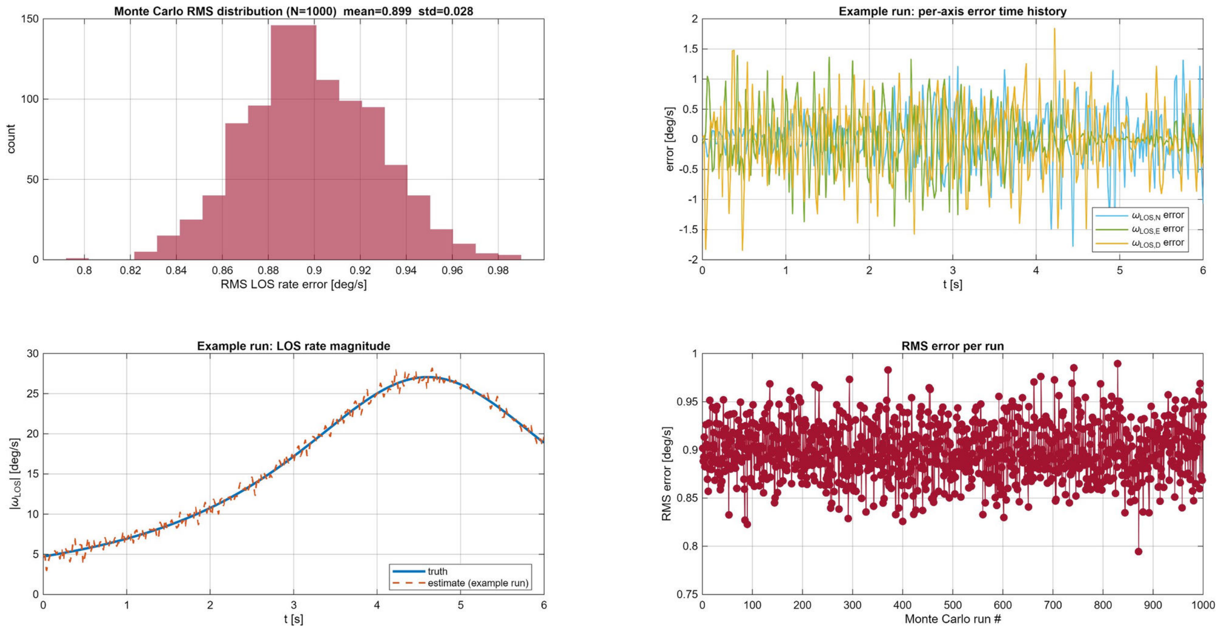

Figure 6: Monte Carlo robustness under pixel tracker noise ($\sigma_{px} = 0.3\,\text{px}$) and gyro noise/bias ($\sigma_{\text{gyro}} = 0.05°/\text{s}$, bias $0.02°/\text{s}$, $N = 1000$ runs): RMS error distribution (mean $= 0.899$, std. $= 0.028\,\text{deg/s}$), an example run's per-axis error and LOS-rate magnitude, and RMS error per run.

4.5 Closed-Loop Intercept Driven by the Estimator

The estimator is finally closed in the loop: a true-PN guidance law, $\mathbf{a}_{\text{cmd}} = NV_c\,(\boldsymbol{\omega}^{\text{est}}_{\text{LOS}} \times \mathbf{r}_w)$ with navigation ratio $N = 3$, drives a constant-speed ($250\,\text{m/s}$) interceptor toward a non-maneuvering but crossing-geometry target from an initial range of $\sim 6070\,\text{m}$. Here $\boldsymbol{\omega}^{\text{est}}_{\text{LOS}}$ is entirely the output of the estimation pipeline of Section 3; the guidance loop never has access to the true LOS rate. Figure 7 shows a near-terminal snapshot of the resulting geometry ($t = 16.00\,\text{s}$, range $\approx 23.5\,\text{m}$, one animation frame before intercept). Figure 8 shows the full engagement: the LOS rate is progressively nulled by PN guidance, from $\sim 1\,\text{deg/s}$ down to a fraction of a degree per second through most of the flight, with a sharp terminal spike as range $\to 0$. This is the well-known LOS-rate singularity of a non-zero, shrinking miss distance, and it is visible in both the truth and the (correctly lagging, bounded) estimate. The commanded acceleration decays from $\sim 2\,g$ to well under $1\,g$ as the geometry converges toward a collision triangle, again spiking only in the terminal frame. The engagement ends at $t_{\text{go}} = 16.06\,\text{s}$ with a miss distance of $1.71\,\text{m}$, a hit-to-kill-quality intercept flown entirely on the estimated LOS rate.

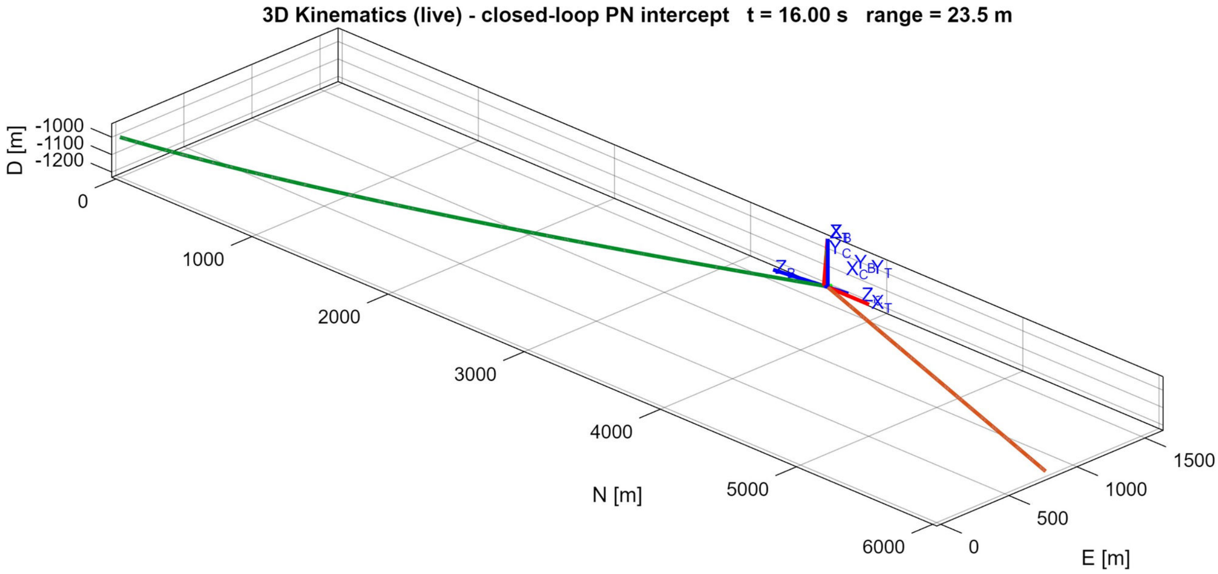

Figure 7: Closed-loop true-PN intercept, near-terminal animation snapshot ($t = 16.00\,\mathrm{s}$, range $\approx 23.5\,\mathrm{m}$; final miss distance at $t_{\mathrm{go}} = 16.06\,\mathrm{s}$ is $1.71\,\mathrm{m}$, see Fig. 8), guided entirely by the estimated LOS rate (navigation ratio $N = 3$).

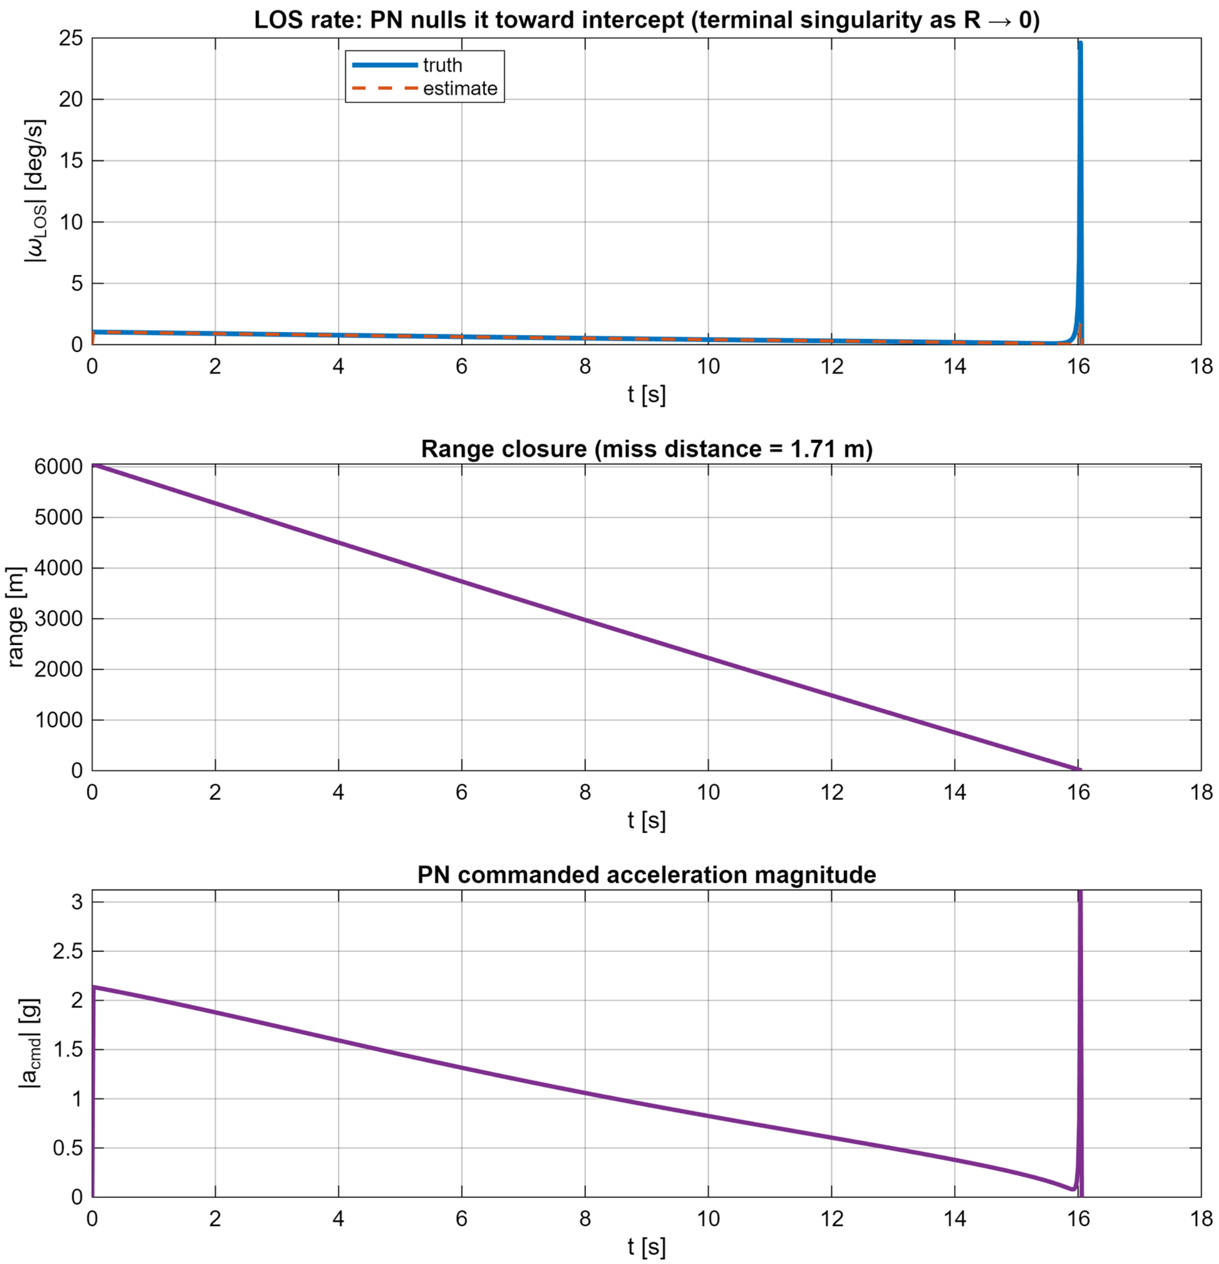

Figure 8: Closed-loop intercept results: LOS-rate magnitude (truth vs. estimate, showing the terminal singularity as range $\to 0$), range closure (miss distance $= 1.71\,\mathrm{m}$), and commanded acceleration magnitude.

4.6 Kalman Process-Noise Sensitivity

The process-noise scale q in Eq. (5) trades off tracking lag against noise sensitivity: a small q over-trusts the constant-velocity model (smooth but lagging under maneuvers), while a large q over-trusts the raw measurement (agile but noisy). This trade-off is only visible when both a maneuver *and* measurement noise are present simultaneously; Figure 9 sweeps $q \in [10^{-3}, 10^{3}]$ on the weaving-target/gimbal-lag trajectory of Section 4.3 with $\sigma_{px} = 0.5\,\mathrm{px}$ pixel noise (25 Monte Carlo runs per point). The mean RMS error is minimized near $q \approx 0.01$ and rises steadily on both sides, from $\sim 0.3\,\mathrm{deg/s}$ at the optimum to $\sim 10\,\mathrm{deg/s}$ at $q = 1000$, more than an order of magnitude worse from poor tuning alone.

This shows that q cannot be chosen arbitrarily; it should be tuned against the expected target maneuver level and sensor noise floor.

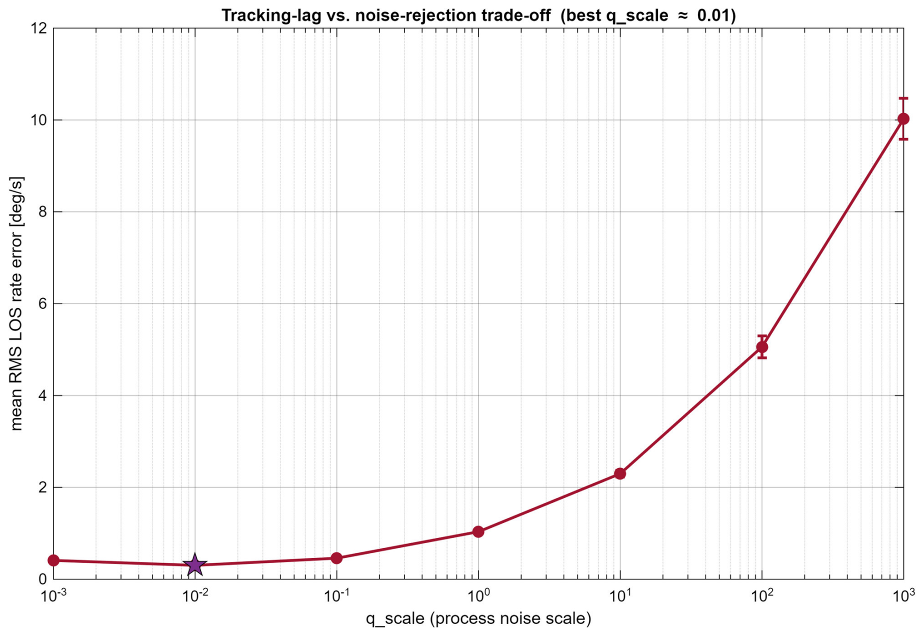

Figure 9: Process-noise scale (q) sensitivity sweep on the maneuvering/gimbal-lag trajectory with pixel noise ($\sigma_{px} = 0.5\,\text{px}$): mean RMS LOS-rate error vs. q, showing the classic tracking-lag/noise-rejection trade-off (minimum at $q \approx 0.01$, star marker).

4.7 Pixel-Channel Latency Sensitivity

Real vision-based trackers introduce processing latency between image capture and the availability of a bearing measurement. To quantify its effect, a naive, uncompensated zero-order-hold delay of 0–400 ms (0–20 samples at 50 Hz) is applied to the pixel channel only (the gyro is treated as effectively delay-free, consistent with typical inertial sensor bandwidths), again on the maneuvering/gimbal-lag trajectory. Figure 10 shows that RMS error grows monotonically and substantially with latency, from effectively zero at 0 ms to 1.34 deg/s at 400 ms, with the maximum instantaneous error growing even faster (from near-zero to 7.3 deg/s). The bottom panel shows the mechanism directly: increasing delay introduces an increasing phase lag between the true and estimated LOS-rate time histories, exactly as expected for an uncompensated transport delay. This result quantifies why real strapdown seekers generally require explicit delay compensation (e.g. out-of-sequence-measurement handling or extrapolation) rather than treating tracker output as instantaneous.

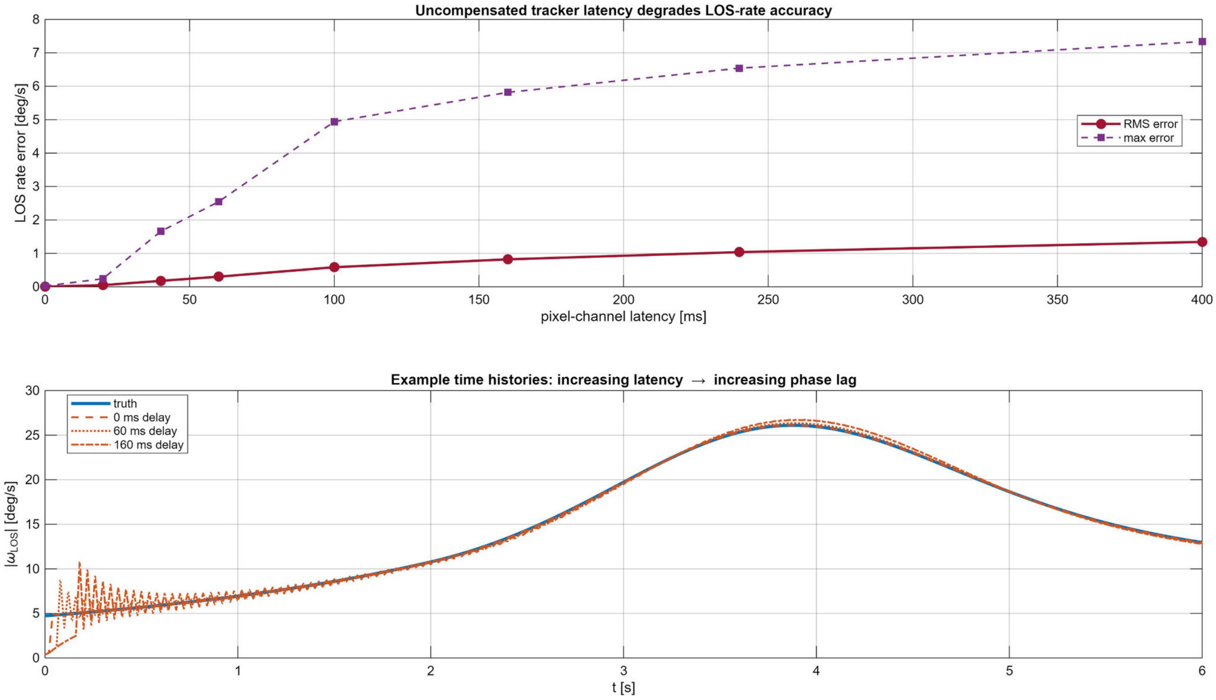

Figure 10: Effect of uncompensated pixel-channel latency on LOS-rate estimation accuracy: RMS/max error vs. latency (top), and example LOS-rate time histories showing the growing phase lag (bottom).

5 Conclusions

This paper presented a complete strapdown-seeker LOS-rate estimation pipeline that fuses a pixel-plane constant-velocity Kalman filter with a strapdown gyro through the rigid-body transport theorem, and validated it across six scenarios of increasing realism. The main findings are:

(a) The reconstruction formula $\boldsymbol{\omega}_{\text{LOS}} = \mathbf{r}_w \times \dot{\mathbf{r}_w}$, built from gyro and filtered pixel-rate terms, reproduces numerically differentiated ground truth to 3.13×10^{-4} deg/s RMS in a deterministic baseline, and remains accurate (3.49×10^{-1} deg/s RMS against a 27 deg/s peak rate) under a maneuvering target and a finite-bandwidth gimbal.

(b) Under realistic pixel and gyro noise/bias, the estimator achieves 0.899 ± 0.028 deg/s RMS error over 1000 Monte Carlo runs, with low run-to-run variance.

(c) The estimate is accurate enough to close the guidance loop: a true-PN law driven entirely by the estimated (never the true) LOS rate achieves a 1.71 m miss distance.

(d) The Kalman process-noise scale exhibits a genuine tracking-lag/noise-rejection trade-off, with more than an order-of-magnitude RMS degradation possible from poor tuning alone. It must be tuned against the expected maneuver level and sensor noise.

(e) Uncompensated pixel-channel latency degrades LOS-rate accuracy substantially and monotonically (RMS error growing from ≈ 0 to 1.34 deg/s over 0–400 ms), motivating explicit delay compensation in a fielded system.

Future work includes extending the pixel-channel delay compensation to an out-of-sequence-measurement Kalman update, evaluating performance against an evading/maneuvering intercept target, and characterizing sensitivity to range-dependent camera lever-arm effects at short final-approach ranges.

References

[1] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, pp. 35–45, 1960.

[2] P. Zarchan, *Tactical and Strategic Missile Guidance*, 6th ed. AIAA, Reston, VA, 2012.

[3] G. M. Siouris, *Missile Guidance and Control Systems*. Springer, New York, 2004.

[4] D. Simon, *Optimal State Estimation: Kalman, H_∞, and Nonlinear Approaches*. Wiley, 2006.