Gimballed Stable-Platform Seeker: Line-of-Sight Rate Estimation via Pixel-Plane Kalman Filtering and Gimbal-Mounted Gyro Fusion

Rasit Evduzen

July 18, 2026

Abstract

This paper presents a line-of-sight (LOS) rate estimation method for a gimballed, stable-platform electro-optical seeker. Unlike a strapdown seeker, where the camera is rigidly fixed to the vehicle body and the body's own rotation must be subtracted out of the measurement, a gimballed seeker carries the camera on an actively steered platform whose gyro senses the platform's own inertial rate directly. This removes the need for a body-rate correction term, at the cost of having to model and compensate for the finite bandwidth of the gimbal servo itself. The method combines a pixel-plane Kalman filter, which tracks the target's normalized image coordinates and their rates, with the gimbal-mounted gyro rate through the transport theorem, producing a three-component inertial LOS rate vector free of az/el singularities. The approach is tested on a single, deliberately severe scenario: the interceptor body executes an independent roll/pitch/yaw disturbance while translating along a sharp-cornered path, and the target simultaneously traces a square path in azimuth and elevation relative to the body, with a finite gimbal servo time constant and realistic pixel-tracking noise. Under this scenario the estimator achieves a steady-state root-mean-square (RMS) LOS-rate error of 1.02 deg/s and an overall RMS error of 2.29 deg/s (max 12.2 deg/s), with the largest errors concentrated at the sharp corners of the target's square path, where the LOS rate itself changes discontinuously.

1 Introduction

Line-of-sight rate is the single most important quantity in a homing guidance loop: proportional navigation and its variants steer the interceptor directly from an estimate of $\boldsymbol{\omega}_{LOS}$, so the accuracy of that estimate largely determines miss distance [2]. Two broad seeker architectures are used to obtain the pixel measurements that feed this estimate. A strapdown seeker mounts the camera rigidly to the vehicle body; the LOS vector is measured in body coordinates and must be rotated into an inertial frame using the body's own attitude and angular rate, both of which are themselves noisy and delayed. A gimballed, or stable-platform, seeker instead mounts the camera on an actively steered two- or three-axis gimbal that keeps the camera boresight pointed at the target largely independent of body motion [4, 3]. The rate gyro in a gimballed seeker sits on the gimbal itself rather than on the body, so it senses the platform's own inertial rotation directly. This is the central practical advantage of the gimballed architecture: the LOS-rate reconstruction no longer needs a separate body-rate term or the geometric composition that a strapdown seeker requires.

The trade-off is that the gimbal is a real servo with finite bandwidth. It cannot track a maneuvering target with zero lag, so the target is never held exactly at the image center; this residual pointing error is what gives the pixel channel a genuine, non-trivial signal to work with. A gimbal that is fast enough to hold the target essentially on boresight produces near-zero pixel-plane rates, and the LOS-rate estimate becomes effectively gyro-only; a gimbal with a visible lag, as modeled here, forces the pixel-derived rate term to do real estimation work alongside the gyro term.

This paper develops the gimbal-frame analogue of a standard strapdown LOS-rate estimator: a Kalman filter tracks the normalized pixel coordinates of the target and their rates, and the transport theorem is applied in the gimbal frame (using the gimbal's own gyro-sensed rate) to reconstruct the full three-component inertial LOS rate vector. The method is demonstrated on one scenario, chosen to be the most demanding case considered during development: simultaneous body rotational disturbance, body translation, a sharp-cornered target maneuver, finite gimbal bandwidth, and pixel measurement noise, all active at once.

2 Problem Formulation

Four reference frames are used: the world (inertial, NED) frame W, the interceptor body frame B, the gimbal frame G, and the camera frame C. The gimbal rotates relative to the body under servo control, and the camera is rigidly mounted to the gimbal through a fixed rotation $R_{cam,gmb}$ and a small lever arm. The rotation from gimbal to world, R_{GW}, is the platform's own attitude and is what the gimbal-mounted gyro measures the rate of.

A pinhole camera model relates a 3D point in the camera frame to its pixel projection:

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

where (f_x, f_y) are the focal lengths in pixels and (c_x, c_y) is the principal point. The inverse mapping gives normalized image coordinates,

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

so that the unit LOS vector in the camera frame is $\mathbf{r}_c = [o_x,\, o_y,\, 1]^\top / \|[o_x,\, o_y,\, 1]^\top\|$.

3 Proposed Method

3.1 Pixel-plane Kalman filter

The filter state is $\mathbf{x} = [o_x,\, o_y,\, \dot{o}_x,\, \dot{o}_y]^\top$, propagated with a constant-velocity model and corrected by pixel measurements $\mathbf{z} = [o_x,\, o_y]^\top$ obtained from the inverse pinhole mapping applied to the noisy tracker output (u, v):

$$\mathbf{x}_k^- = A\,\mathbf{x}_{k-1}, \qquad P_k^- = A P_{k-1} A^\top + Q \tag{3}$$

$$K_k = P_k^- H^\top \left(H P_k^- H^\top + R \right)^{-1} \tag{4}$$

$$\mathbf{x}_k = \mathbf{x}_k^- + K_k \left(\mathbf{z}_k - H \mathbf{x}_k^- \right), \qquad P_k = (I - K_k H) P_k^- (I - K_k H)^\top + K_k R K_k^\top \tag{5}$$

with A the standard constant-velocity state transition matrix, $H = [I_2 \;\; 0_2]$, Q a discretized constant-velocity process noise matrix, and R the pixel measurement noise covariance. The Joseph-form covariance update is used for numerical stability.

3.2 Gimbal-frame LOS rate reconstruction

From the filtered state, the unit LOS vector and its derivative in the camera frame follow from the derivative of a normalized vector:

$$\mathbf{r}_c = \frac{\mathbf{s}}{\|\mathbf{s}\|}, \qquad \dot{\mathbf{r}_c} = \frac{\dot{\mathbf{s}}}{\|\mathbf{s}\|} - \frac{\mathbf{s}\,(\mathbf{s}\cdot\dot{\mathbf{s}})}{\|\mathbf{s}\|^3}, \qquad \mathbf{s} = [o_x,\, o_y,\, 1]^\top,\;\; \dot{\mathbf{s}} = [\dot{o}_x,\, \dot{o}_y,\, 0]^\top \tag{6}$$

These are rotated into the gimbal frame through the fixed camera mount, $\mathbf{r}_g = R_{cam,gmb}\,\mathbf{r}_c$ and $\dot{\mathbf{r}_g} = R_{cam,gmb}\,\dot{\mathbf{r}_c}$, and then into the world frame using the gimbal's own attitude and gyro-sensed rate:

$$\boxed{\mathbf{r}_w = R_{GW}\,\mathbf{r}_g, \qquad \dot{\mathbf{r}_w} = R_{GW}\left(\boldsymbol{\omega}_g \times \mathbf{r}_g + \dot{\mathbf{r}_g}\right), \qquad \boldsymbol{\omega}_{LOS} = \mathbf{r}_w \times \dot{\mathbf{r}_w}} \tag{7}$$

This is structurally identical to the strapdown formula, with the gimbal's rotation and gyro-sensed rate $\boldsymbol{\omega}_g$ substituted for the body's. The three-component cross-product form is used rather than an az/el rate to avoid the coordinate singularity that az/el parameterizations develop near the poles of the LOS sphere; $\boldsymbol{\omega}_{LOS}$ is guaranteed to lie in the plane perpendicular to $\mathbf{r}_w$, since $\mathbf{r}_w \cdot (\mathbf{r}_w \times \dot{\mathbf{r}_w}) = 0$ for any vector crossed with itself, which is exactly the two rotational degrees of freedom a LOS direction actually has.

3.3 Gimbal pointing law and servo lag

The gimbal boresight is commanded to track the instantaneous LOS direction, $\hat{\mathbf{b}}_{Z,cmd} = -\hat{\mathbf{los}}$, but the platform does not reach this command instantaneously. A first-order lag models the finite servo bandwidth:

$$\hat{\mathbf{b}}_Z[k] = \hat{\mathbf{b}}_Z[k-1] + \frac{T_s}{\tau_{gmb}}\left(\hat{\mathbf{b}}_{Z,cmd}[k] - \hat{\mathbf{b}}_Z[k-1]\right), \quad \text{renormalized to unit length} \tag{8}$$

with τ_{gmb} the gimbal time constant and T_s the sample period. The remaining two gimbal axes are completed with a Gram-Schmidt construction against a fixed reference to avoid a singular frame near the poles.

3.4 Algorithm summary

1. Read pixel measurement (u, v); convert to normalized coordinates (o_x, o_y) via the inverse pinhole mapping.
2. Update the Kalman filter state $\mathbf{x} = [o_x, o_y, \dot{o}_x, \dot{o}_y]^\top$.
3. Reconstruct $\mathbf{r}_c, \dot{\mathbf{r}_c}$ from the filter state.
4. Rotate into the gimbal frame via the fixed camera mount, then into the world frame via R_{GW} and the gimbal-mounted gyro rate $\boldsymbol{\omega}_g$.
5. Form $\boldsymbol{\omega}_{LOS} = \mathbf{r}_w \times \dot{\mathbf{r}_w}$.
6. Repeat at the next sample.

4 Simulation Results

4.1 Scenario parameters

A single scenario is used throughout, chosen as the most demanding combination of effects exercised during development: the interceptor body flies its own square path (lateral and vertical motion only, no forward drift) while independently wobbling in roll, pitch, and yaw; the target holds a constant 50 m range and traces its own square path in azimuth and elevation, starting dead-center on boresight; the gimbal servo has a finite, visible time constant; and the pixel tracker carries realistic 1-pixel measurement noise. No Monte Carlo averaging is performed; all results are from this single deterministic run (seeded noise). Table 1 lists the parameters.

Table 1: Scenario parameters

Parameter	Value
Sample period T_s	0.02 s (50 Hz)
Run duration	8 s
Interceptor square path side	10 m (lateral/vertical)
Body roll / pitch / yaw amplitude	15° / 8° / 5°
Body roll / pitch / yaw frequency	0.8 / 0.5 / 0.3 Hz
Target range	50 m (constant)
Target az/el square amplitude	20° / 20°
Gimbal servo time constant τ_{gmb}	0.25 s
Pixel measurement noise σ_{px}	1 px
Camera horizontal FOV	40°
Image size	640 × 480 px

4.2 Kinematics

Figure 1 shows the 3D scene at the end of the run: the target's square path in the world frame, the interceptor body and gimbal frames, and the true and estimated LOS/LOS-rate arrows overlaid at the camera origin. The gimbal frame (G) visibly counters the body's roll/pitch/yaw wobble to keep the camera boresight tracking the target despite the disturbance.

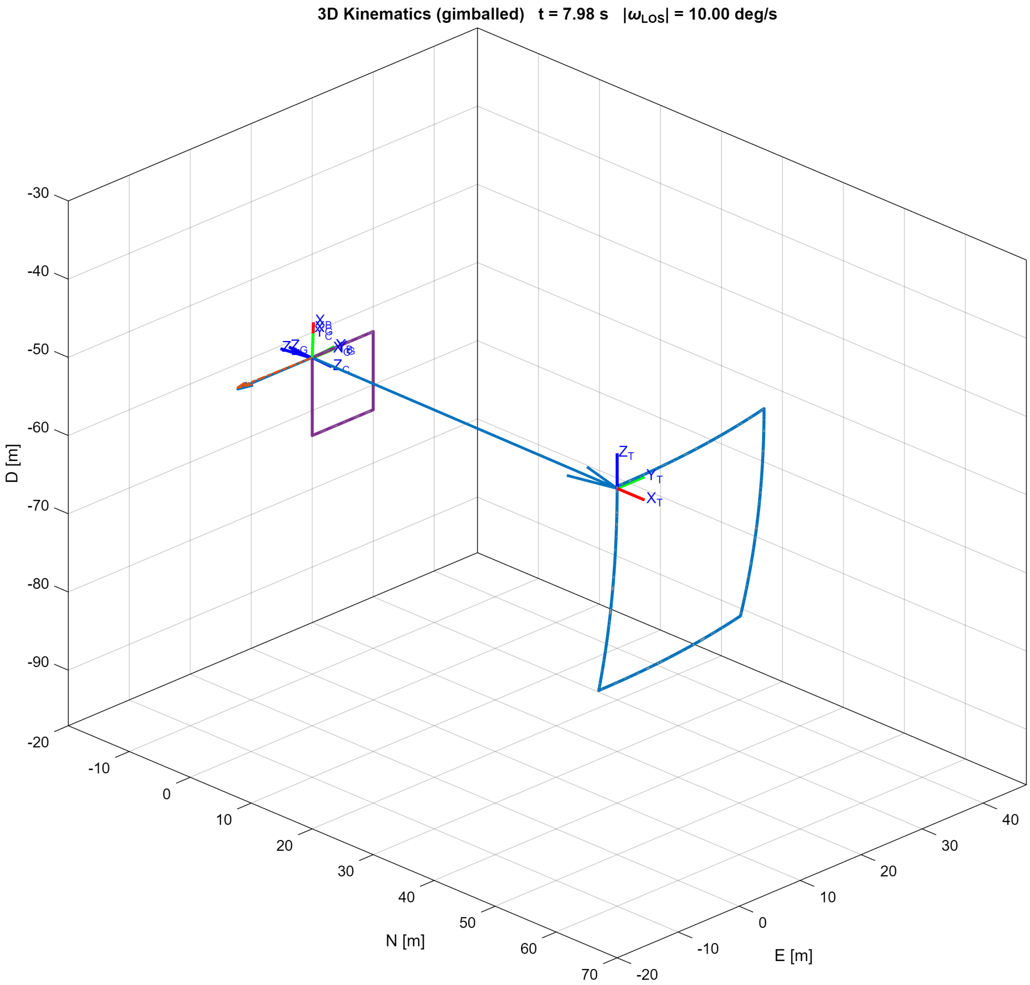

Figure 1: 3D kinematics: interceptor body, gimbal, camera, and target trajectory, with true (solid) and estimated (dashed) LOS and LOS-rate vectors at the final sample.

4.3 LOS-rate estimation

Figure 2 compares the truth and estimated LOS rate on each NED axis, along with the error norm over time. The target's square path in az/el produces a piecewise-constant true LOS rate with sharp step transitions at each corner; the estimator tracks the constant-rate segments closely, with the largest transient errors occurring exactly at these corners, where the true LOS rate itself is discontinuous and no causal estimator can respond instantaneously.

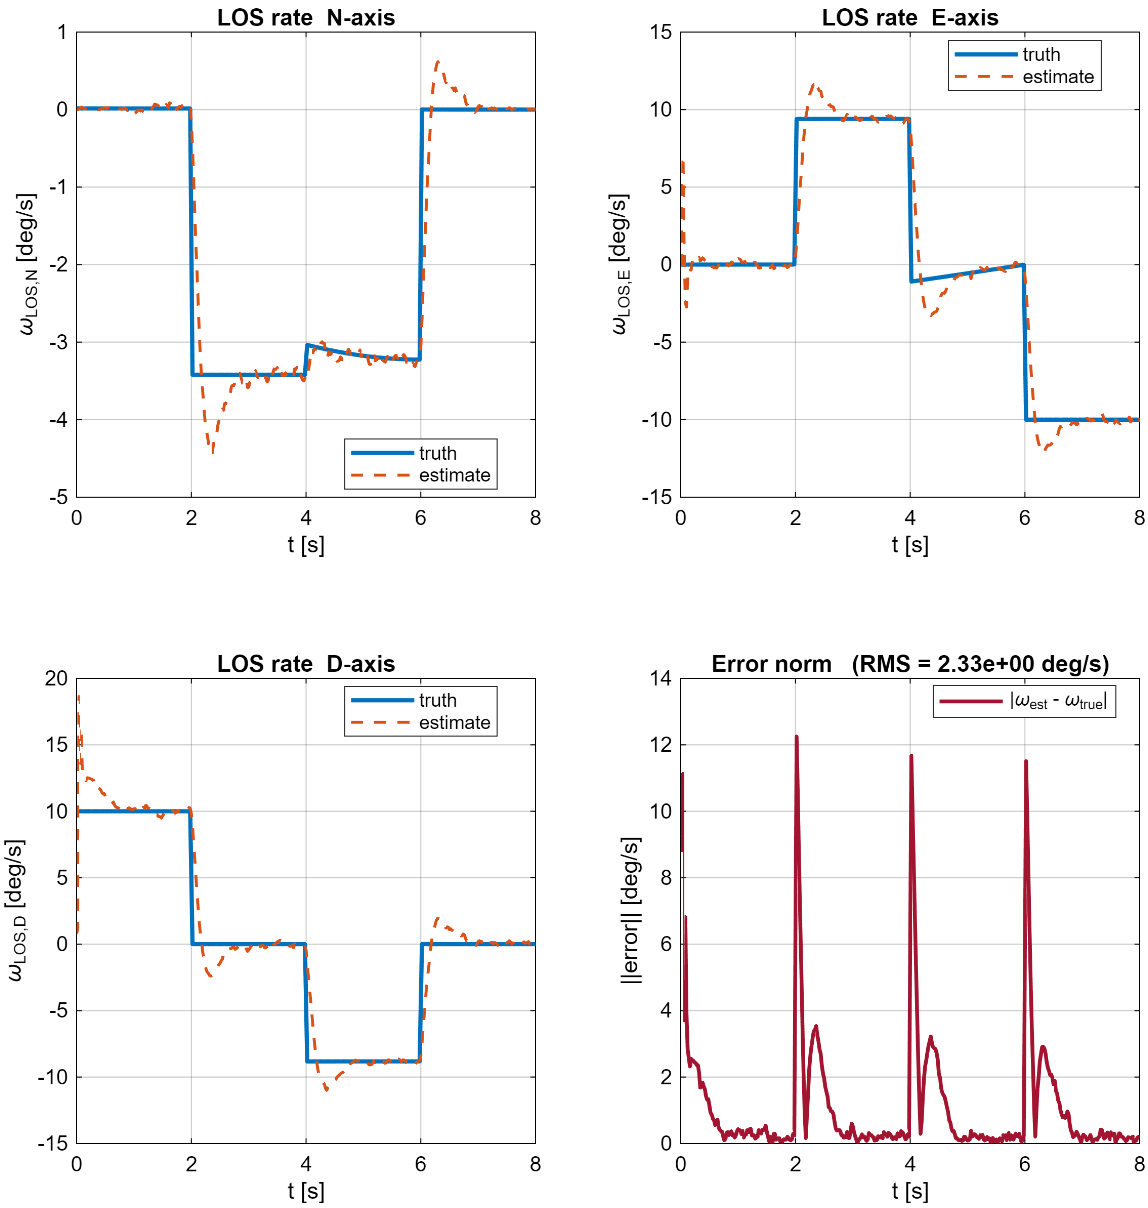

Figure 2: LOS rate estimation: truth vs. estimate on the N, E, and D axes, and the error norm over time (steady-state RMS = 1.02 deg/s, overall RMS = 2.29 deg/s).

4.4 Body disturbance and gimbal response

Figure 3 shows the body's independent roll/pitch/yaw disturbance alongside the gimbal's own roll/pitch/yaw (relative to the body), which combines the disturbance-rejection motion with the target-tracking motion. The gimbal angles do not simply mirror the body disturbance, since the gimbal is simultaneously also steering to follow the target's own az/el square path; the two motions are superimposed in the gimbal's angular history.

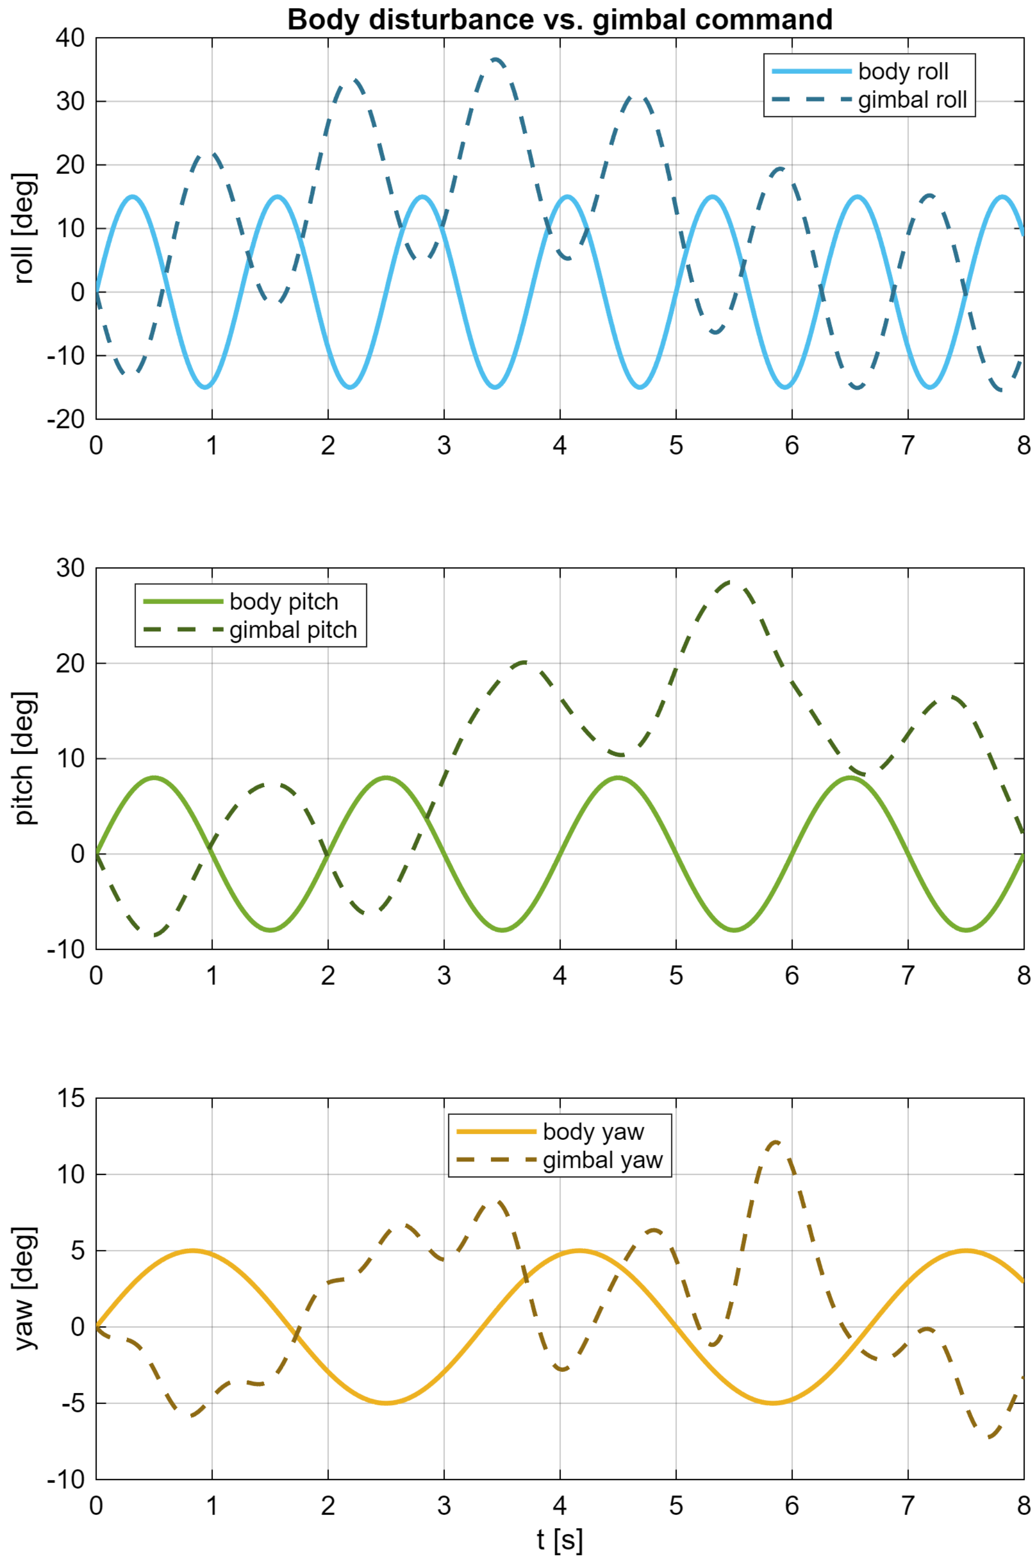

Figure 3: Body attitude disturbance (roll/pitch/yaw) versus the resulting gimbal angles relative to the body.

4.5 Pixel-plane track

Figure 4 shows the Kalman-estimated pixel-plane position and rate, plotted parametrically (o_x vs. o_y, and $\dot{o}_x$ vs. $\dot{o}_y$) rather than against time. The finite gimbal bandwidth keeps the target measurably off-center throughout the run, tracing the same square shape as the commanded az/el path; this off-center excursion is what gives the pixel-derived rate term real information content rather than tracking pure noise near zero.

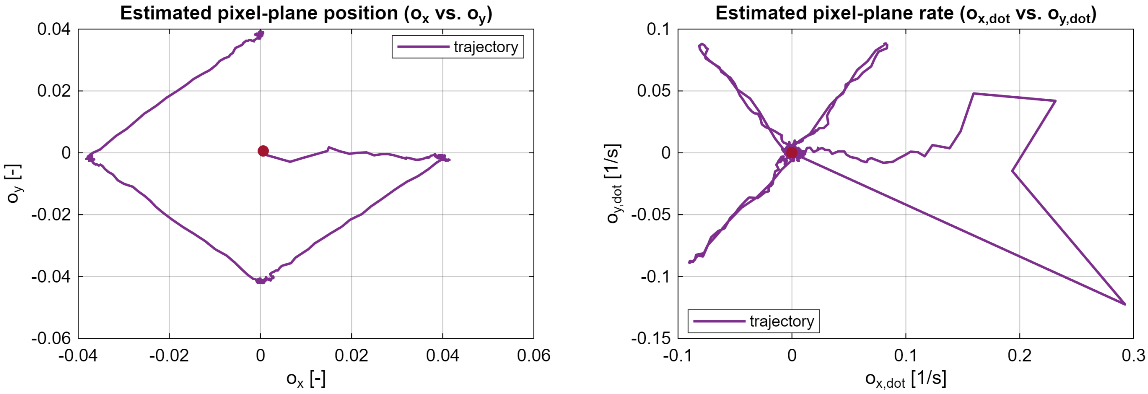

Figure 4: Kalman-estimated normalized pixel-plane position (left) and rate (right), plotted parametrically.

4.6 Azimuth/elevation tracking

Figure 5 compares the commanded target az/el square path against the az/el implied by the estimated world-frame LOS vector. The estimate follows the square path closely, with visible overshoot only at the sharp corners.

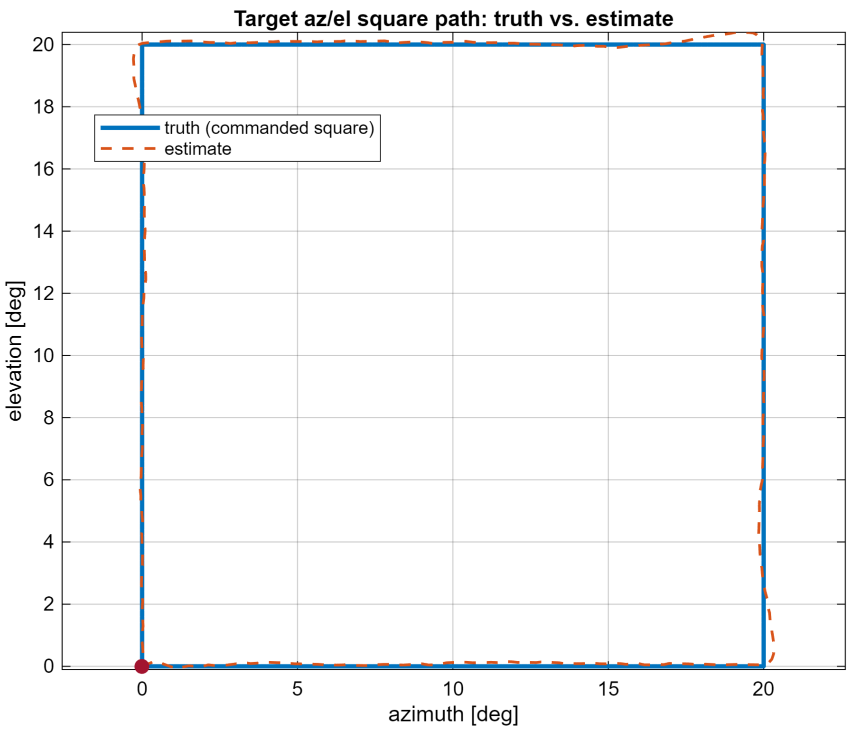

Figure 5: Target azimuth/elevation square path: truth (commanded) vs. estimate derived from the estimated LOS vector.

4.7 Discussion: idealized gimbal pointing

One property of this simulation deserves an explicit note. The gimbal pointing command, $\hat{\mathbf{b}}_{Z,cmd} = -\hat{\mathbf{los}}$, is computed directly from the true target and interceptor positions at each sample, not from a closed-loop control law driven by the body-mounted sensors a real gimbal would use. Because of this, moving or rotating the interceptor body has no effect on the numerical LOS-rate RMS reported here; it only changes what Figures 1 and 3 visually show, since the target's position and az/el path are defined relative to the body regardless of how the body itself moves. A real closed-loop gimbal rejects body disturbance through a rate loop closed on a body-mounted gyro, which introduces its own bandwidth-limited tracking error; that additional error source is not present in this study and would be expected to raise the reported RMS figures somewhat in a more complete model.

5 Conclusions

(a) A gimbal-frame LOS-rate reconstruction, structurally identical to the strapdown formula but using the gimbal's own attitude and gyro-sensed rate in place of the body's, correctly recovers the inertial LOS rate under a combined body-disturbance, target-maneuver, and finite-servo-bandwidth scenario.

(b) Under the single hardest-case scenario tested, the estimator achieves a steady-state RMS LOS-rate error of 1.02 deg/s and an overall RMS of 2.29 deg/s, with a maximum error of 12.2 deg/s occurring at the sharp corners of the target's square path.

(c) A finite gimbal servo bandwidth is what gives the pixel-plane channel genuine information content; an idealized, instantaneous gimbal holds the target essentially exactly on boresight and reduces the pixel term to near-zero-signal noise tracking.

(d) Because the gimbal pointing law in this study is computed directly from ground truth rather than from a body-mounted-gyro rate loop, body motion and disturbance do not affect the reported LOS-rate accuracy numerically, only the visual kinematics; a fully closed-loop gimbal control law would be expected to add its own bandwidth-limited error contribution.

(e) The three-component vector form of $\boldsymbol{\omega}_{LOS}$, obtained from $\mathbf{r}_w \times \dot{\mathbf{r}}_w$, avoids the az/el coordinate singularity while still representing exactly the two physical rotational degrees of freedom a LOS direction possesses.

References

[1] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

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

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

[4] J. M. Hilkert, "Inertially Stabilized Platform Technology," *IEEE Control Systems Magazine*, vol. 28, no. 1, pp. 26–46, 2008.

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