HKUST PhD Qualifying Examination

Reliability-Aware Active Visual Navigation for Space Robots

A Proposal on Estimator Reliability, Action-Conditioned Prediction, and Task-Constrained Decision Making

Doctoral Research Proposal

November 2026

Abstract

Reliable visual navigation is essential when a space robot must move through terrain or approach a target while its observations and motion model are imperfect. Existing active-navigation methods commonly rank candidate actions by information predicted under a nominal model. In monocular visual–inertial navigation, however, a physically informative maneuver can also amplify persistent camera–IMU timing error, and an information increase need not produce lower trajectory error. This proposal asks how a robot can diagnose the reliability of its observations and constraints, predict the navigation outcome of feasible physical actions, and select actions that balance estimation benefit against task cost.

The proposed method has three connected components. First, an estimator-side reliability module separates weak observability, random measurement degradation, persistent model mismatch, and tracking loss using residual consistency, constraint validity, and the propagation–correction split. Second, an action-conditioned predictor combines a short-horizon physical model with these reliability measurements to forecast navigation error and recovery probability. Split conformal calibration supplies finite-sample prediction bounds under route-level exchangeability. Third, a risk-aware selector chooses from a finite library of physically executable maneuvers, including continued travel, reduced speed, bounded steering excitation, and brief stopping, while accounting for time and distance cost. The shared decision core is paired with separate lunar-rover and spacecraft angles-only adapters.

The project already provides a lunar simulation testbed and seven frozen navigation baselines. Across four benchmark routes, monocular VINS, stereo VINS, and the Stable stereo pipeline meet a 3% terminal-error criterion on all four routes. A controlled timing-offset study over three independent routes shows increasing full-trajectory RMSE from 0 to 40 and 80 ms, while a 15 s straight segment remains comparatively accurate. In one paired turn-entry case at 80 ms, reducing speed by 50% lowers local RMSE from 13.733 cm to 3.140 cm and moves the outgoing scale ratio from 0.701 to 0.954. These observations motivate a systematic study of when nominal information is converted into useful navigation improvement.

The doctoral work will establish the failure mechanism, develop and calibrate an action-conditioned predictor, evaluate closed-loop selection under matched task budgets, and then test transfer to angles-only spacecraft navigation. The central scientific contribution is a measurable link between estimator reliability, predicted information, and the navigation benefit actually obtained from physical motion.

Contents

1 Introduction **1**

1.1 Mission motivation . . . 1

1.2 Research foundation and scope . . . 1

1.3 Research questions . . . 2

2 Related Work and Research Gaps **2**

2.1 Monocular visual–inertial estimation and observability . . . 2

2.2 Planetary and lunar surface navigation . . . 3

2.3 Spacecraft angles-only navigation and informative maneuvers . . . 3

2.4 Active SLAM, information-theoretic planning, and dual control . . . 3

2.5 Estimator integrity, mismatch detection, and calibrated uncertainty . . . 4

2.6 Closest-work comparison and research gaps . . . 4

3 Problem Formulation and First-Order Error Analysis **6**

3.1 A shared estimation model . . . 6

3.2 First-order Gauss–Newton error decomposition . . . 7

3.3 An action helps only when variance reduction exceeds bias growth . . . 7

3.4 Weak observability and motion-dependent mismatch . . . 8

3.5 Nominal information, measured benefit, and decision regret . . . 9

3.6 Testable hypotheses . . . 10

4 Proposed Research Framework **10**

4.1 Overview: reliability, prediction, and decision . . . 10

4.2 Stage 1: observation and constraint reliability . . . 11

4.3 Stage 2: physics-guided prediction with conformal calibration . . . 11

4.4 Stage 3: task-constrained risk-aware selection . . . 12

4.5 Planned ablations . . . 12

5 Experimental Platform and Protocol **13**

5.1 Lunar simulation and sensor configuration . . . 13

5.2 Paired physical-action protocol . . . 14

5.3 Scoring and decision metrics . . . 14

6 Preliminary Results and Baseline Comparisons **15**

6.1 Seven-estimator lunar benchmark . . . 15

6.2 Controlled timing-offset response . . . 16

6.3 Reliability signals and estimator correction . . . 16

6.4 Physical-action evidence 17

7 Doctoral Work Plan and Acceptance Criteria **18**

7.1 Phase I: diagnose reliability under controlled mismatch 18

7.2 Phase II: predict and select under matched budgets 19

7.3 Phase III: transfer the decision interface to angles-only navigation 19

7.4 Statistics, reproducibility, and decision gates 19

8 Expected Contributions and Scope **20**

8.1 Expected contributions 20

8.2 Scope and limitations 21

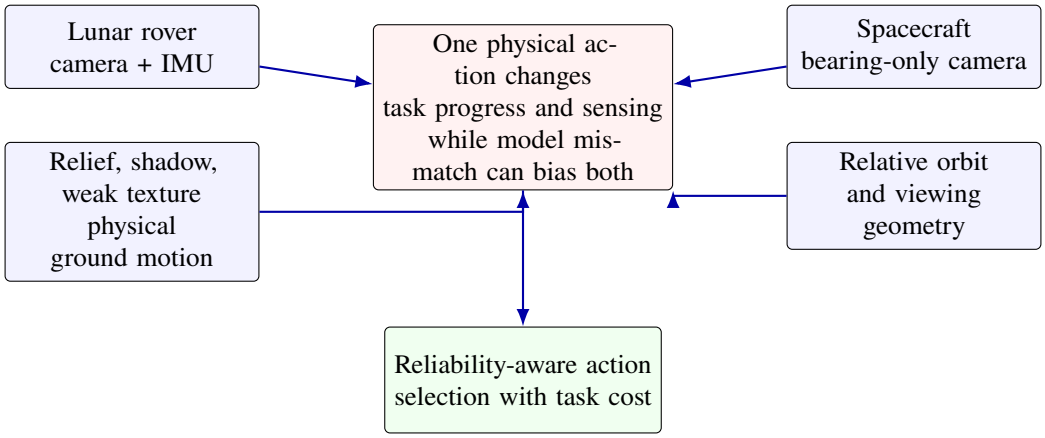


Figure 1: The proposal studies a shared decision problem across a lunar rover and a spacecraft: physical motion changes observation geometry, while persistent model mismatch can separate nominal information from useful navigation improvement.

1 Introduction

1.1 Mission motivation

Space robots must infer motion and target geometry from restricted sensing while operating under tight limits on time, energy, visibility, and safe motion. A lunar rover may encounter long shadows, weak texture, repeated crater patterns, dust, or terrain whose geometry differs from the model used by its estimator. A spacecraft using bearing-only measurements faces a different sensor geometry, but a related difficulty: its range direction may remain weakly observable during nearly coasting motion. In both settings, movement changes both task progress and the observations available for navigation.

This coupling motivates active visual navigation. A rover can slow down, stop, or make a bounded steering maneuver to improve its next observations; a spacecraft can modify its relative trajectory to change viewing geometry. Such motion is valuable only when the predicted benefit survives the sensor and estimator used in operation. A nominal Fisher-information increase measures sensitivity under an assumed measurement model. It does not by itself establish that the estimator will reduce position error when feature associations, calibration, timing, or motion assumptions are imperfect.

The scientific problem is therefore not merely to seek more observations. It is to identify the reliability of the constraints currently shaping the estimate, predict the consequences of each feasible physical maneuver under that reliability state, and choose an action whose expected navigation improvement justifies its task cost. This question matters particularly for visual–inertial systems: camera and inertial measurements are fused over time, so small temporal or geometric mismatches can be absorbed into velocity, scale, bias, or pose corrections before appearing as a large trajectory error.

1.2 Research foundation and scope

The project combines prior work on angles-only spacecraft guidance with a lunar navigation platform. The lunar simulator provides synthetic and Chang'e-5 terrain products, synchronized camera and inertial streams, and offline ground truth for evaluation. The current algorithm set includes stereo visual–inertial baselines and a ground-constrained monocular method. The initial doctoral study will keep one monocular visual–inertial estimator fixed and examine a controlled camera–IMU timing mismatch, because temporal alignment is both physically interpretable and measurable through motion-dependent residuals.

The project treats lunar-rover navigation as the primary experimental domain. Spacecraft angles-only relative

navigation is an extension of the same decision abstraction, not a claim that the current lunar experiments already validate cross-domain transfer. The shared elements are the state estimate, reliability of measurement constraints, feasible action set, predicted navigation outcome, and task cost. Domain-specific dynamics and sensor models remain separate.

1.3 Research questions

The thesis is organized around three questions:

1. **Reliability:** Which estimator-side measurements can distinguish insufficient geometric information, stochastic observation degradation, persistent model mismatch, and tracking loss at the time an action must be chosen?
2. **Prediction:** How can the navigation outcome of a candidate physical action be predicted over a short horizon when the nominal observation model may be biased?
3. **Decision:** How should a robot select among physically feasible actions when navigation benefit, recovery likelihood, and task cost must be considered together?

The working thesis is that action selection should depend on calibrated predictions of estimator-level navigation outcome, not on nominal information alone. The proposal develops this thesis through a reliability measurement, an action-conditioned prediction, and a risk-aware selection procedure. The problem and early evidence are consistent with active SLAM and visual–inertial navigation research, while the novelty claim is deliberately limited to the specific combination and testable gap identified in section 2.

2 Related Work and Research Gaps

2.1 Monocular visual–inertial estimation and observability

Monocular visual odometry recovers relative motion from image correspondences but has an unknown metric scale. An inertial sensor can make scale observable when acceleration and orientation provide sufficient excitation; under near-constant-velocity or weakly excited motion, scale, gravity direction, and inertial bias can remain strongly coupled. Modern visual–inertial systems address this with tightly coupled nonlinear optimization, multi-state constraints, preintegration, and online calibration. VINS-Mono and related estimators demonstrate robust operation on broad ground and aerial datasets, while camera–IMU calibration work shows that spatial and temporal parameters can be estimated jointly only when motion excites the corresponding modes [1–6].

Observability analyses explain why estimator confidence must be interpreted with the motion history. A low residual can arise because a parameter is well estimated, but it can also arise because the current trajectory supplies little sensitivity to that parameter. Recent work on motion excitation and scale observability, and ground-constrained VIO, illustrates both opportunities and limits: extra geometric constraints can improve metric scale, yet their effect depends on how accurately the ground model represents the scene [7–9]. The sensor-specific literature covers initialization, online calibration, robust updates, and tightly coupled estimation across monocular and stereo settings [10–24].

The present proposal does not replace the backend estimator. It asks how an external decision layer can determine when the estimator's current constraints are reliable enough to support a prediction and whether a physical maneuver is likely to improve its output. That distinction separates estimator design from action-conditioned navigation prediction.

2.2 Planetary and lunar surface navigation

Planetary rover navigation combines long traverses, repetitive terrain, limited illumination, and sparse ground truth. Mars rover systems have developed visual odometry, terrain-relative localization, and map-based correction under substantial latency and communication constraints. Lunar navigation adds low-sun illumination and permanently shadowed regions, where crater geometry or active illumination may replace ordinary feature tracking. LunarNav and ShadowNav exemplify crater-based localization, and recent datasets and simulators provide increasingly realistic terrain, image formation, and reference geometry [25–29].

Terrain and illumination datasets broaden evaluation beyond a single visual condition. Polar traversal imagery, lunar surface rendering, and crater-sensitive pipelines enable comparisons of feature support, shadow structure, and terrain-relative localization [30–42]. These studies motivate the testbed used here and provide a path toward external validation. They do not remove the need to distinguish the effect of terrain appearance from the estimator's internal model assumptions.

The lunar experiments will therefore report terrain, lighting, motion, sensor, and scoring conditions explicitly. The primary mismatch study changes only the camera–IMU timing relation while keeping the estimator configuration fixed. Later studies can add geometric mismatch, illumination, or inertial noise one factor at a time.

2.3 Spacecraft angles-only navigation and informative maneuvers

Bearing-only spacecraft navigation estimates relative state from image directions without direct range measurements. The instantaneous bearing Jacobian is insensitive to displacement along the line of sight, producing a range-related weak direction. Relative orbital dynamics can make range observable over time, but the quality of that information depends on the trajectory and viewing geometry. Classical observability analyses and optimal maneuver studies describe how deliberate relative motion can excite otherwise weak modes [43–48].

Recent spacecraft work applies active SLAM, factor-graph information criteria, tree search, and learned guidance to proximity operations. These methods demonstrate how candidate trajectories can be compared by information, fault discrimination, or task reward [49–62]. The research here shares the use of physical maneuvers to alter observability, while focusing on whether the expected estimator-level gain remains valid under persistent sensor or model mismatch and explicit task cost.

2.4 Active SLAM, information-theoretic planning, and dual control

Active SLAM and belief-space planning provide mature tools for selecting motion using expected covariance, entropy, Fisher information, graph connectivity, or goal reward. Generalized belief-space methods propagate joint state and landmark uncertainty through candidate actions; data-association-aware planning adds discrete ambiguity; perception-aware planning and Fisher information fields connect viewpoint selection to expected image geometry [51, 63–73].

Perception competence and active viewpoint selection move closer to estimator performance by learning whether a view is likely to localize successfully. Learned forward models predict physical response or future visual state, while dual control treats actions as both task progress and information acquisition [72, 74–85]. The strongest implication is methodological: nominal information criteria remain useful components, but their value must be tested against the estimator outcome and the cost of the motion that produced it.

Table 1: Methodological taxonomy across the five literature themes reviewed in this proposal.

Theme	Main quantity	Typical method role	Connection addressed in this proposal
Monocular VIO	Metric scale, pose, velocity, and sensor bias	Tightly coupled estimation and online calibration [1–3]	Determine when motion provides enough excitation for reliable state estimation.
Planetary surface navigation	Terrain-relative pose and map alignment	Visual odometry, crater localization, and terrain-relative navigation [8, 25, 26]	Relate terrain and illumination evidence to estimator-level reliability.
Spacecraft angles-only navigation	Relative state and weakly observed range	Dynamic observability analysis and trajectory design [43, 45, 49]	Connect physical maneuvers to navigation error under bearing-only sensing.
Active SLAM and belief-space planning	Expected covariance, information, or task reward	Rank candidate views or trajectories under an assumed observation model [64, 67, 68]	Test whether nominal information predicts the error produced by the estimator.
Integrity and calibrated prediction	Constraint consistency, localization competence, and predictive coverage	Detect unreliable estimates or calibrate outcome intervals [86–88]	Combine online reliability evidence with calibrated, task-cost-aware action choice.

2.5 Estimator integrity, mismatch detection, and calibrated uncertainty

Robust estimation uses robust losses, switchable constraints, data-association hypotheses, innovation tests, or fault-detection logic to reduce the influence of inconsistent measurements. Integrity monitoring asks whether an estimate can be trusted for a specified application and risk bound. Perception-competence methods estimate the probability that a scene or viewpoint will support localization. Conformal prediction adds finite-sample marginal coverage guarantees under exchangeability without requiring a fully specified residual distribution [86–105].

These tools address complementary questions. A robust loss limits the effect of outliers; an integrity monitor quantifies the status of an estimate; a competence model forecasts whether a view is usable; conformal calibration controls coverage of a prediction set. The proposed work combines estimator-side reliability evidence with an action-conditioned prediction of navigation error, then uses the prediction in a task-constrained selector. It does not treat a robust loss or a covariance criterion alone as a new contribution.

2.6 Closest-work comparison and research gaps

The five literature themes address different parts of the navigation problem. Table 1 summarizes their measurement and action assumptions and identifies the connection left open in the reviewed studies. The following comparison then focuses on fourteen locally reviewed studies selected for proximity to the proposed problem. It summarizes what each method predicts, what it assumes about measurement validity, and which part of the proposed chain remains open. This comparison is a scoped reading set rather than a census of the full field.

Table 2: Four-axis comparison of closely related studies.

Study	Demonstrated capability	Assumption or open question	Relation to this proposal
Placed et al. [51]	Surveys active SLAM objectives and action-selection methods.	The review identifies evaluation diversity; it does not supply an estimator reliability interface.	Establishes the broader active-SLAM design space.
Lauri et al. [63]	Unifies robotics planning under partially observed state and reward.	Model and observation distributions are specified for planning.	Supports belief-state formulation, not mismatch diagnosis.
Indelman et al. [64]	Propagates trajectory and landmark belief through candidate actions.	Predicted measurements and model structure are nominal.	Provides the belief-space basis for action-conditioned analysis.
Pathak et al. [65]	Incorporates data-association uncertainty into robust belief-space planning.	Association ambiguity is represented; persistent calibration mismatch is distinct.	Motivates separating discrete ambiguity from systematic bias.
Rahman et al. [66]	Constrains uncertainty in belief-space trajectory optimization.	Covariance propagation follows the assumed observation model.	Provides a task-constrained planning comparator.
Zhang et al. [67]	Builds a differentiable Fisher-information field for perception-aware planning.	Landmark geometry and matchability are available to the planner.	Motivates nominal information as a comparator, not a final objective.
Perception-aware receding-horizon planning [68]	Couples visual–inertial uncertainty with short-horizon motion planning.	The predictor is based on the estimator's assumed measurement model.	Closest planning interface for a rover-like platform.
Li et al. [74]	Learns localization success as a function of viewpoint.	Success probability is not a calibrated estimate of physical action benefit under mismatch.	Closest competence-aware viewpoint method.
Ragan et al. [50]	Selects spacecraft maneuvers to distinguish fault hypotheses.	The hypotheses are discrete hardware faults and task cost is limited.	Shows active diagnosis by motion in a neighboring domain.
Ticozzi and Tsiotras [49]	Uses factor-graph information to select spacecraft pointing actions.	Candidate information is evaluated under a nominal model.	Closest angles-only active-SLAM comparator.
Zhou et al. [8]	Adds ground geometry and online geometric calibration to monocular VIO.	Ground constraints may be systematically wrong if terrain departs from the assumed surface.	Directly motivates constraint-validity measurement.
Cauligi et al. [25]	Uses crater geometry for lunar localization in dark regions.	Localization assumes a usable crater map and known observation model.	Grounds the lunar application and environmental challenge.
Roth et al. [75]	Predicts action-conditioned physical response for safe navigation.	The prediction target is physical motion and risk, not estimator-level gain.	Informs the physics backbone for short-horizon prediction.
Yang et al. [76]	Predicts future visual representations for action selection.	Latent similarity is not a calibrated metric navigation error.	Motivates direct evaluation in estimator units.

Three gaps follow from this comparison. **Gap 1: reliability is not yet a decision variable.** Robust losses and competence estimates can reject or down-weight information, but the estimator-side evidence that distinguishes weak geometry, random corruption, persistent mismatch, and tracking loss is not consistently passed to a navigation-outcome predictor. **Gap 2: predicted information is not the same as estimator-level benefit.**

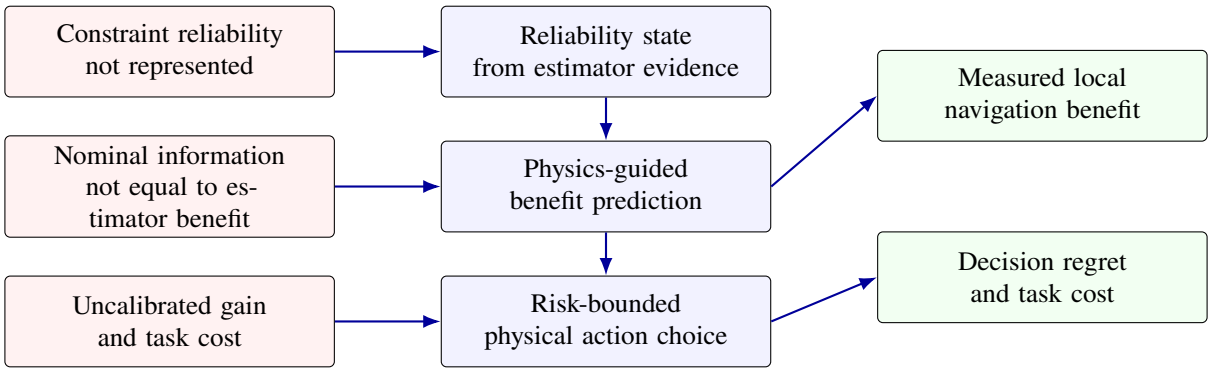


Figure 2: Mapping the reviewed research gaps to the proposal's three connected contributions and their evaluation outputs.

Fisher information and covariance predict performance under a model; they do not guarantee a reduction in trajectory error when the model is biased. **Gap 3: calibration and task cost must accompany prediction.** A prediction that ranks actions correctly on average may still fail when its uncertainty is miscalibrated or when an informative maneuver consumes excessive time or distance.

The proposed innovations map one-to-one to these gaps: (i) a reliability representation grounded in residual and correction behavior; (ii) a physics-guided predictor whose target is short-horizon navigation benefit with calibrated uncertainty; and (iii) a risk-aware action selector evaluated against fixed task-only, fixed-recovery, and nominal-information baselines. The finite maneuver library and the shared lunar/spacecraft interface are experimental structure, not standalone novelty claims.

3 Problem Formulation and First-Order Error Analysis

3.1 A shared estimation model

Let x_k denote the navigation state at time k, θ the sensor and environment parameters, and z_i a camera or inertial observation. The MAP estimate over a sliding window is

$$\hat{\mathcal{X}} = \arg\min_{\mathcal{X}} \frac{1}{2} \sum_{i \in \mathcal{F}} r_i(\mathcal{X}; \theta)^\top W_i r_i(\mathcal{X}; \theta), \tag{1}$$

where $\mathcal{F}$ is the set of factors and W_i is the inverse covariance assigned to factor i. For monocular VIO, a state contains camera/IMU pose, velocity, accelerometer and gyroscope bias, and feature depth; visual reprojection factors couple poses to landmarks, and preintegrated inertial factors couple consecutive motion states. For a spacecraft observed only by bearing, the state contains relative position and velocity, and each camera factor constrains the line of sight to the target. The same MAP form applies, while the dynamics, measurement functions, and feasible actions differ by domain [1, 43, 49, 64].

The important distinction is between random error and persistent model mismatch. Let the estimator use nominal parameters θ_0, while the actual sensor or environment follows $\theta_0 + \delta\theta$, with $\delta\theta$ approximately constant over a short window. A camera–IMU time offset, an incorrect camera–IMU rotation, or an inaccurate ground-plane model can produce such a persistent perturbation. The estimator can reduce its residual by adjusting the state, even when that adjustment moves the trajectory away from the physical state.

The state, position, and error quantities are expressed in SI units in the equations; result tables convert position errors to centimetres. Image residual units are pixels, timing is in seconds, and information increments use natural logarithms.

Table 3: Recurring symbols and dimensions used in the estimation and action analysis.

Symbol	Dimension or unit	Meaning
x, δx	$\mathbb{R}^n$	Gauge-fixed state and its first-order estimation error.
θ, $\delta\theta$	$\mathbb{R}^q$	Sensor/environment parameter and persistent mismatch.
r, ϵ	$\mathbb{R}^m$	Stacked residual and zero-mean measurement noise.
J_x, J_θ	$m \times n$, $m \times q$	Residual Jacobians with respect to state and mismatch.
W, R	$m \times m$	Residual information weight and noise covariance, with $W = R^{-1}$.
Λ	$n \times n$	Gauss–Newton information matrix on the observable subspace.
P	$3 \times n$	Selector that extracts position error from the state.
δt, H	s	Camera–IMU timing offset and local evaluation horizon.
E_H, ΔE	m	Local position RMSE and action-induced error reduction.
ΔI_{nom}	nats	Nominal log-determinant information increment.

3.2 First-order Gauss–Newton error decomposition

Let $r \in \mathbb{R}^m$ stack the residuals in one optimization window and let $x \in \mathbb{R}^n$ be the observable state coordinates after fixing gauge freedoms. Linearizing the residual about the nominal optimum x_0 gives

$$r(x_0 + \delta x; \theta_0 + \delta\theta) \simeq r_0 + J_x \delta x + J_\theta \delta\theta + \epsilon, \qquad J_x \in \mathbb{R}^{m\times n}, \quad J_\theta \in \mathbb{R}^{m\times q}, \tag{2}$$

where $\epsilon \in \mathbb{R}^m$ is zero-mean random measurement noise with covariance $R \in \mathbb{R}^{m\times m}$. Let $W = R^{-1}$ and $\Lambda = J_x^\top W J_x \in \mathbb{R}^{n\times n}$. Since x_0 is a nominal optimum, $J_x^\top W r_0 = 0$. The first-order normal equations therefore yield

$$\delta x = -\Lambda^{-1} J_x^\top W (J_\theta \delta\theta + \epsilon). \tag{3}$$

Define $C = \Lambda^{-1} J_x^\top W \in \mathbb{R}^{n\times m}$ and let $P \in \mathbb{R}^{3\times n}$ select position error. The position error has a random component $-PC\epsilon$ and a systematic component $-PCJ_\theta\delta\theta$.

Proposition. First-order position MSE under persistent mismatch *Assume the residual model is differentiable in a neighborhood of x_0, Λ is positive definite on the gauge-fixed observable subspace, $\mathbb{E}[\epsilon] = 0$, and $\mathbb{E}[\epsilon\epsilon^\top] = R$. For a fixed mismatch $\delta\theta$, the first-order position-error mean-square magnitude is*

$$\mathbb{E}\left[\|P\delta x\|^2\right] = \operatorname{tr}\left(PCRC^\top P^\top\right) + \|PCJ_\theta\delta\theta\|^2 . \tag{4}$$

The first term is propagated random uncertainty; the second is squared bias induced by the persistent mismatch.

Proof. Substituting equation (3) gives $P\delta x = -PCJ_\theta\delta\theta - PC\epsilon$. The first term is deterministic for fixed $\delta\theta$. The second has zero mean. Expanding the squared norm, the expected cross term is zero, while $\mathbb{E}[\|PC\epsilon\|^2] = \operatorname{tr}(PC\mathbb{E}[\epsilon\epsilon^\top]C^\top P^\top)$. Replacing the expectation by R gives equation (4). *Q.E.D.*

The decomposition explains why a smaller nominal covariance does not guarantee smaller navigation error. A candidate action may add rows to J_x and increase Λ, reducing the covariance term, while also increasing the mismatch sensitivity PCJ_θ. A D-optimal criterion sees the change in $\det(\Lambda)$; it does not include the bias term unless mismatch is represented in the state or the objective. The diagnostic in figure 3 illustrates this distinction: an estimate can become more concentrated and less accurate at the same time.

3.3 An action helps only when variance reduction exceeds bias growth

Let Λ_0 be the information matrix at a decision point and let a candidate physical action contribute measurements with Jacobian $J_a \in \mathbb{R}^{m_a\times n}$ and weight $W_a \in \mathbb{R}^{m_a\times m_a}$. Its nominal information matrix is $\Lambda_a = \Lambda_0 + J_a^\top W_a J_a$.

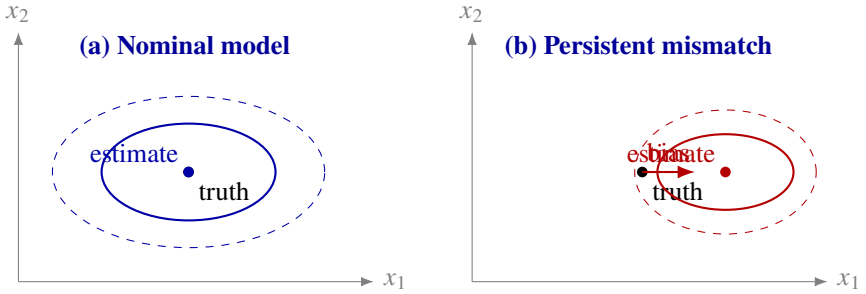


Figure 3: A persistent mismatch can shift the estimate away from truth even when additional information narrows its random uncertainty. Solid and dashed ellipses represent two covariance levels.

On the gauge-fixed observable subspace, the corresponding first-order covariance approximation is Λ_a^{-1}. Let $e_0 \in \mathbb{R}^n$ denote the state bias accumulated before the action, $b_0 = Pe_0$ its position component, and $g_a = J_a^\top W_a G_a \delta\theta$ the action's first-order mismatch gradient, where $G_a \in \mathbb{R}^{m_a \times q}$ is the residual Jacobian with respect to the persistent mismatch $\delta\theta \in \mathbb{R}^q$. The post-action bias is

$$b_a = P\Lambda_a^{-1}(\Lambda_0 e_0 - g_a), \qquad \mathrm{MSE}(a) = \mathrm{tr}(P\Lambda_a^{-1}P^\top) + \|b_a\|^2 . \tag{5}$$

The projected position selector P has dimension $3 \times n$; $g_a \in \mathbb{R}^n$ and $b_a \in \mathbb{R}^3$. This expression uses the same first-order, fixed-mismatch, locally Gaussian assumptions as Proposition 1. It separates the uncertainty reduction produced by added geometry from the state shift produced by mismatch-sensitive measurements.

Define $\Delta V(a) = \mathrm{tr}\big(P(\Lambda_0^{-1} - \Lambda_a^{-1})P^\top\big)$ and $\Delta B(a) = \|b_a\|^2 - \|b_0\|^2$. The action's expected squared position error changes by

$$\mathrm{MSE}(0) - \mathrm{MSE}(a) = \Delta V(a) - \Delta B(a). \tag{6}$$

Thus, an action improves this local error criterion exactly when its variance reduction exceeds its increase in squared bias. A nominal information score can rank the action highly while the inequality reverses, because $\log\det\Lambda_a$ describes information volume and contains no term for $\delta\theta$. The proposed method tests whether online residual and correction signals can predict the sign and magnitude of $\Delta V - \Delta B$ before the maneuver is executed.

3.4 Weak observability and motion-dependent mismatch

In monocular VIO, global position and global yaw are gauge freedoms. Metric scale is observable through inertial measurements only when motion supplies sufficient excitation; nearly straight constant-speed travel can provide weak scale and bias information. For a bearing measurement $u = \rho/\|\rho\|$, where ρ is the relative position vector, the first-order bearing Jacobian is

$$H_\rho = \frac{1}{\|\rho\|}\left(I - uu^\top\right), \qquad H_\rho u = 0. \tag{7}$$

The instantaneous measurement is insensitive to range change along the line of sight. Relative motion can rotate the line of sight and accumulate range information, but long coasting segments may preserve a weak radial mode. These are different physical systems with a common mathematical signature: a nearly singular information matrix and a maneuver-dependent measurement Jacobian.

For a camera–IMU timing perturbation δt, the pose discrepancy over a short interval is approximately

$$\delta p \simeq v\,\delta t, \qquad \delta\phi \simeq \omega\,\delta t, \tag{8}$$

where v is translational velocity and ω is angular velocity. The image-level sensitivity follows from the camera projection. Let $p_c = R_{wc}^{\top}(P_w - p_w)$ be a static scene point in camera coordinates, with camera-to-world rotation R_{wc}, camera position p_w, and scene point P_w. Under camera-frame angular velocity ω_c and world-frame translational velocity v_w, its instantaneous camera-frame velocity is

$$\dot{p}_c = -[\omega_c]_\times p_c - R_{wc}^{\top} v_w, \qquad J_{\text{feature}} = J_\pi(p_c)\dot{p}_c, \tag{9}$$

where $[\omega_c]_\times$ is the skew-symmetric cross-product matrix and $J_\pi = \partial\pi(p_c)/\partial p_c$ is the 2×3 projection Jacobian. For a pinhole camera with $p_c = (X, Y, Z)^{\top}$, J_π has rows $(f_x/Z, 0, -f_xX/Z^2)$ and $(0, f_y/Z, -f_yY/Z^2)$. A small timestamp error changes the projected feature coordinate by $\delta z \simeq J_{\text{feature}}\delta t$; the corresponding residual receives this term or its negative according to the residual sign convention. Its magnitude depends on image depth and both translational and angular motion; the angular term grows with $\|\omega_c\|\ |\delta t|$.

Straight motion with low angular rate may make this mismatch difficult to distinguish from weak geometric excitation. A turn increases the rotational timing sensitivity and can make the error more visible while also increasing its effect on pose. A maneuver can therefore improve nominal observability and worsen the estimate under a biased model. The same derivative links physical motion, image residuals, and navigation error without assuming that every informative maneuver is beneficial.

3.5 Nominal information, measured benefit, and decision regret

For a candidate action a, let J_a and R_a denote the predicted measurement Jacobian and covariance under the nominal model. Over the observable subspace, define

$$\Lambda_{\text{nom}}(a) = \Lambda_0 + J_a^{\top} R_a^{-1} J_a, \qquad \Delta I_{\text{nom}}(a) = \log\det\Lambda_{\text{nom}}(a) - \log\det\Lambda_0. \tag{10}$$

If the matrix has unobservable modes, the determinant is computed after projection to the observable subspace or after applying the same fixed prior to every candidate. This score is an information proxy, measured in nats; it is not navigation error.

Let $E_H(a)$ be position RMSE over a fixed short horizon H, or over an equal-progress interval when action duration differs. The measured benefit relative to continuing the task is

$$\Delta E_{\text{real}}(a) = E_H(a_0) - E_H(a), \qquad \rho_E(a) = \frac{\Delta E_{\text{real}}(a)}{E_H(a_0) + \varepsilon}. \tag{11}$$

Positive ΔE_{real} indicates reduced error; negative values indicate that the action worsened navigation. ρ_E is the fractional error reduction used for observed action comparisons. The proposed predictor will learn a calibrated mapping from online state and ΔI_{nom} to ΔE_{real}; the conversion is estimated from route-level calibration data rather than assumed from the units of information gain.

For a decision point with feasible action set $\mathcal{A}$, the oracle regret of a selected action is

$$\mathcal{R} = C(a_{\text{selected}}) - \min_{a\in\mathcal{A}} C(a), \tag{12}$$

where C is frozen before evaluation and contains navigation error and separately reported task costs. Ground truth is used only to compute this retrospective benchmark and action label; it is unavailable to the online estimator, predictor, and selector.

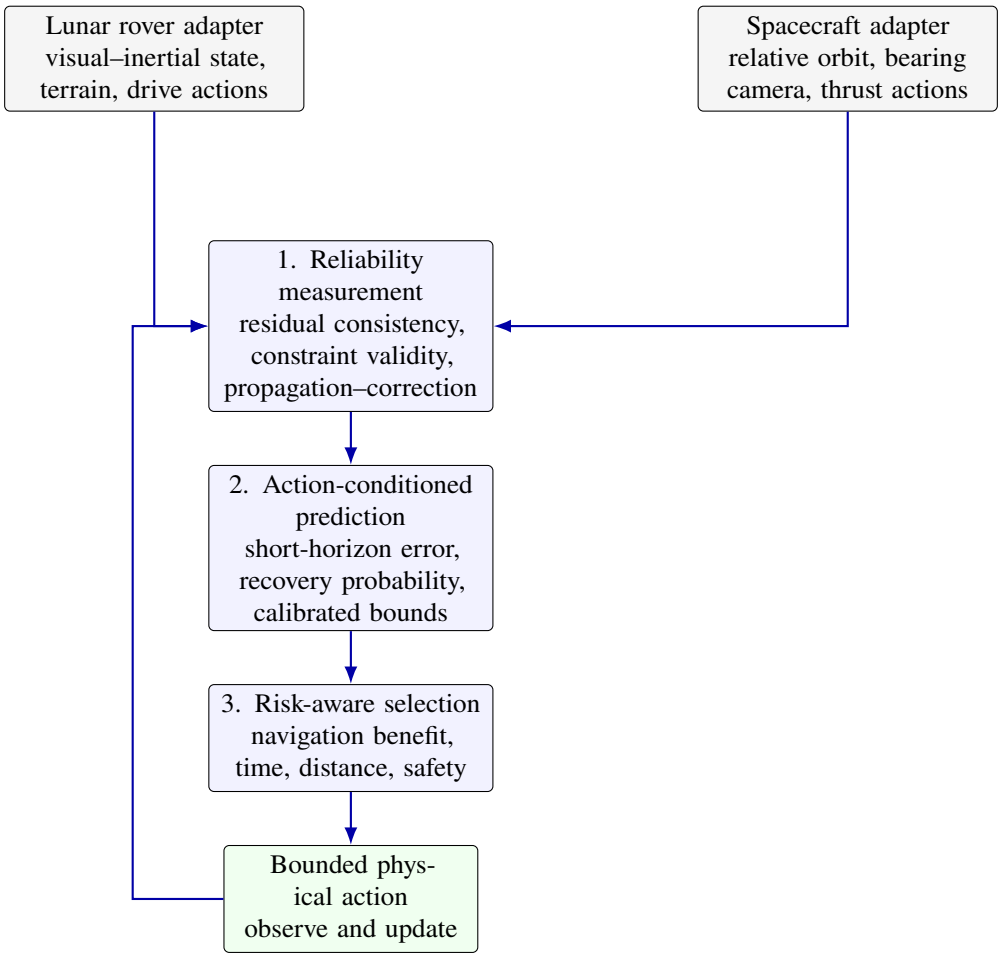


Figure 4: A shared decision core receives domain-specific state and action models through separate adapters. It evaluates reliability, predicts short-horizon navigation outcomes, selects a bounded physical maneuver, and repeats after new observations.

3.6 Testable hypotheses

H1: Estimator residual and correction measurements will distinguish persistent timing mismatch from weak visual observability and random image degradation better than a residual-magnitude threshold alone.
H2: A physics-guided, reliability-conditioned predictor will forecast short-horizon navigation benefit more accurately and with better calibration than nominal information, frequency, and fixed-rule baselines.
H3: Under matched task budgets, risk-aware selection will reduce route-level decision regret relative to task-only, fixed-recovery, and nominal-information policies without lowering task completion beyond a pre-registered non-inferiority margin.

4 Proposed Research Framework

4.1 Overview: reliability, prediction, and decision

The method separates three operations that are often collapsed into one information score. A reliability module summarizes whether the measurements and geometric constraints currently shape the state in a trustworthy way. An action-conditioned predictor estimates the navigation consequences of each feasible physical maneuver. A risk-aware selector trades the lower-bound estimate improvement against task cost and motion constraints. The estimator remains a separate component and can be replaced without changing the decision interface.

Table 4 gives the shared interface. Only the adapter supplies domain-dependent dynamics and feasible actions.

Table 4: Shared decision variables and domain adapters.

Interface	Lunar rover	Spacecraft relative navigation
State	Pose, velocity, inertial bias, landmarks, local terrain geometry	Relative position and velocity, target attitude or shape state
Observation	Monocular/stereo images and inertial measurements	Camera bearing or image-plane target features
Weak mode	Scale, gravity, or motion bias under weak excitation	Range along the line of sight during coasting
Physical action	Continue, reduce speed, bounded steering excitation, brief stop	Continue, bounded relative maneuver, hold/coast, viewpoint-changing thrust
Primary task cost	Time, route distance, turning and control effort	Time, propellant/fuel, collision and plume constraints
Evaluation	Trajectory RMSE, recovery, terminal error, action regret	Relative-state error, range observability, fuel, terminal safety

This design supports a rover-first study while retaining a principled route to angles-only spacecraft navigation.

4.2 Stage 1: observation and constraint reliability

The first component updates an online reliability vector from estimator quantities already available during operation. The initial feature set includes (i) pre- and post-optimization visual residual cost, (ii) whitened inertial residual magnitude, (iii) grouped normalized innovation squared, and (iv) the split between propagated motion and backend correction. A residual cost may be written

$$V_{\text{ratio}} = \frac{\text{median}(V_{\text{pre}})}{\text{median}(V_{\text{post}}) + \varepsilon}, \quad W_k = \frac{\left\| \Sigma_{I,k}^{-1/2} r_{I,k} \right\|}{\sqrt{d_I}}, \quad \text{NIS}_g = r_g^\top S_g^{-1} r_g. \tag{13}$$

Here d_I is the residual dimension for the inertial group, S_g is the innovation covariance for group g, and ε prevents division by zero. The grouped NIS can be compared with a χ^2 interval when its distributional assumptions are reasonable; it is not treated as a standalone correctness certificate.

The propagation–correction interface records the position increment before and after optimization:

$$\mathbf{p}_k^+ - \mathbf{p}_{k-1}^+ = \Delta\mathbf{p}_{\text{prop}} + \Delta\mathbf{p}_{\text{corr}}. \tag{14}$$

A large correction relative to the propagated increment indicates that visual optimization is materially changing the inertial prediction. The diagnostic module combines this signal with residual consistency and geometric support, rather than interpreting a small post-optimization residual as proof of a correct trajectory.

The reliability representation is a four-class belief $b_k \in \Delta^3$: insufficient geometric information, stochastic observation degradation, persistent model mismatch, and tracking loss. Labels are defined at the window or scene level when feature-level ground truth is unavailable. The output is a calibrated belief and feature summary; robust weights remain an estimator concern and are evaluated separately.

4.3 Stage 2: physics-guided prediction with conformal calibration

For each action a, a short-horizon physics model predicts the motion and observation geometry over $H = 15$ s, or over a 1.0–1.5 m progress interval when action durations differ. The timing-mismatch term couples the commanded translational and rotational motion to camera–IMU alignment through $v\delta t$ and $\omega\delta t$. Terrain geometry, estimated feature support, and the current reliability belief modify the nominal observation contribution. The predictor outputs a distribution or interval for local position RMSE, recovery probability,

route-progress error, and action duration.

The initial model comparison is deliberately hierarchical: a fixed-frequency baseline, a transparent rule based on track quality and residual growth, a nominal geometric/Fisher predictor, a geometry-plus-reliability predictor, and a lightweight learned predictor only if the first four leave structured residual error. Every input must be available before action execution. Future ground truth is used only to create labels and is excluded from deployment features.

Uncertainty calibration uses route-grouped split conformal prediction with simultaneous coverage over the frozen action set. Let $\mathcal{D}_r$ be the decision points on calibration route r, $\mathcal{A}_{r,d}$ the candidate actions at point d, $\widehat{\Delta E}_{r,d,a}$ the predicted error reduction, $\Delta E_{r,d,a}$ the measured reduction, and $\hat{\sigma}_{r,d,a} > 0$ a scale predicted by the frozen model. The route-level nonconformity score is the largest one-sided overprediction across all decision points and candidate actions:

$$S_r = \max_{d \in \mathcal{D}_r,\, a \in \mathcal{A}_{r,d}} \left[\frac{\widehat{\Delta E}_{r,d,a} - \Delta E_{r,d,a}}{\hat{\sigma}_{r,d,a}} \right]_+, \qquad q_{1-\alpha} = S_{(\lceil (n+1)(1-\alpha) \rceil)}, \tag{15}$$

where $[x]_+ = \max(0, x)$ and the order statistic is taken over n independent calibration routes; if its index exceeds n, set $q_{1-\alpha} = +\infty$. The lower bound for a new route is

$$\underline{\Delta E}_{d,a} = \widehat{\Delta E}_{d,a} - q_{1-\alpha} \hat{\sigma}_{d,a}, \qquad \Pr\Big(\Delta E_{d,a} \geq \underline{\Delta E}_{d,a} \;\; \forall d, a \Big) \geq 1 - \alpha, \tag{16}$$

under exchangeability of route groups and with the predictor, decision-point rule, and candidate-action family frozen before calibration. Because the score takes a maximum over all candidate actions on a route, the coverage statement remains valid when the deployed selector chooses among them using its predictions. It is a route-marginal statement, not conditional coverage for each terrain or a guarantee under distribution shift. At $\alpha = 0.10$, at least nine independent calibration routes are needed for a finite split-conformal quantile.

4.4 Stage 3: task-constrained risk-aware selection

At each decision time, the planner first removes actions that violate vehicle, visibility, collision, or task constraints. Among the remaining actions, it considers the lower confidence bound on navigation benefit and upper confidence bounds on task costs:

$$a^* = \arg \max_{a \in \mathcal{A}_{\text{safe}}} \left[\underline{\Delta E}(a) - \lambda_T \overline{\Delta T}(a) - \lambda_D \overline{\Delta D}(a) - \lambda_C C_{\text{control}}(a) \right] \quad \text{subject to } R_U(a) \leq \alpha. \tag{17}$$

Here $R_U(a)$ is a route-marginal upper bound on the probability that an admitted action fails to reduce error. Set $R_U(a) = \alpha$ when the simultaneous conformal lower bound on $\Delta E(a)$ is positive and $R_U(a) = 1$ otherwise. The joint calibration event covers every candidate before selection, so the bound also applies to the action chosen by the selector. If no active maneuver passes the risk gate, the system follows the task-only policy. Weights and safety constraints are frozen on development routes before held-out testing.

The objective keeps navigation benefit, time, distance, and control effort separately interpretable. Energy is not used as a primary cost until a validated torque–power model is available. The action library contains real changes to vehicle speed or heading; changing feature identities or blending alternative software trajectories is outside the action space.

4.5 Planned ablations

The experiments isolate one component at a time. First, compare the reliability module against residual thresholding, robust weighting, and calibration-only baselines at matched observation retention. Second,

Algorithm 1 Online reliability-aware active navigation

Initialize the estimator, domain adapter, reliability state, and frozen action library.
while the navigation task is active **do**
    Update the state estimate from currently available sensor data.
    Compute residual consistency, geometric support, and propagation–correction features.
    Update the calibrated reliability belief b_k.
    for all feasible physical actions $a \in \mathcal{A}_{\text{safe}}$ **do**
        Predict short-horizon motion, navigation error, recovery probability, and task cost.
        Apply route-level conformal calibration to obtain benefit and cost bounds.
    end for
    Select the action maximizing the risk-bounded objective in equation (17).
    Execute a bounded segment, obtain new observations, and repeat.
end while

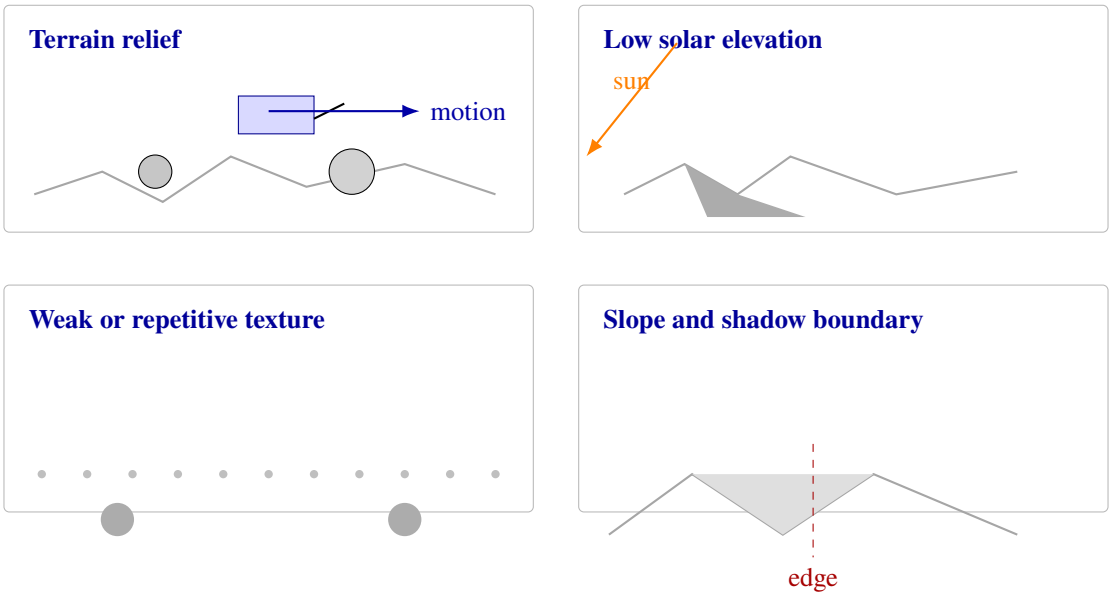


Figure 5: Schematic of the visual conditions represented in the simulator. Panels illustrate relief, low sun, repetitive texture, and slope/shadow boundaries; they are not simulator screenshots.

compare action-outcome prediction against fixed frequencies, simple rules, and nominal geometry on route-held-out data. Third, compare closed-loop policies under matched time and distance budgets: task-only motion, fixed recovery, nominal-information selection, predictor without reliability features, and the full method. These groups determine whether reliability diagnosis improves prediction and whether prediction changes the outcome of the physical action selector.

5 Experimental Platform and Protocol

5.1 Lunar simulation and sensor configuration

The primary testbed is a lunar surface simulation built in Isaac Sim with PhysX dynamics, synthetic and Chang'e-5 elevation products, and lunar visual textures. The environment can vary terrain relief, texture, light direction, and solar elevation while maintaining synchronized motion and sensor streams. The benchmark conditions include a synthetic-Dem polyline, a Chang'e-5 multi-turn route, a synthetic low-sun route, and a Chang'e-5 serpentine traverse.

The image stream is sampled at 2 Hz with 1600×900 pixels, and the inertial stream at 120 Hz. Baseline route

Table 5: Platform and information-access protocol.

Item	Frozen setup and role
Physics and environment	Isaac Sim 5.1 / PhysX; lunar gravity 1.62 m/s^2; synthetic and Chang'e-5 elevation/texture sources; lighting and route conditions recorded per session.
Camera and IMU	2 Hz, 1600×900 imagery; 120 Hz IMU; timing and calibration files remain fixed within each comparison.
Benchmark routes	Polyline (9.393 m), multi-turn (9.396 m), low-sun (7.599 m), and serpentine (21.099 m); each route is shared across the seven estimator baselines.
Timing-mismatch routes	Three independent composite routes with straight travel and a waypoint turn; an identical noisy IMU realization and online bias estimation across the 0/40/80 ms conditions.
Estimator inputs	Images, IMU, calibration, and the action's physical sensor stream only. Ground truth is not delivered to the estimator or online selector.
Ground-truth role	Route design, offline quality checks, scoring, and retrospective action labels. It is excluded from estimator, prediction, and online action-selection inputs.
Pose scoring	First valid pose is aligned with a single rigid transform; scale is fixed at 1. No whole-trajectory scale fit is applied to the primary RMSE or endpoint metric.
Action evaluation	Actions physically change speed, heading, or temporary motion; outcomes are scored on actual estimator output frames and equal-progress windows when duration differs.

speed is 0.02 m/s. Lunar gravity is 1.62 m/s^2. The controlled mismatch experiment uses the same fixed visual–inertial backend while the camera timestamps receive a 0, 40, or 80 ms offset that is ramped linearly from 15 to 25 s and held thereafter. All three conditions begin from the same initialization interval; online inertial-bias estimation remains enabled.

The benchmark's inertial stream is generated from the recorded motion model, and the controlled timing study adds a reproducible noisy sensor model with online bias estimation. These protocols answer different questions: the first compares estimator behavior on the frozen four-route suite; the second isolates a motion-dependent timing mismatch. Neither is presented as a substitute for a hardware IMU characterization.

5.2 Paired physical-action protocol

For action studies, each route and decision context has a common pre-decision state. The candidate maneuvers are continued travel, 50% speed reduction, a bounded S-shaped steering maneuver, and a 5 s stop. They alter the physical sensor trajectory and therefore the subsequent observations. The evaluation replays each accepted sensor log under the same timing-offset ramp and uses the same fixed backend. Before the decision point, branch input streams and physical states are checked for agreement; the reported prefix discrepancy is at most 0.166 mm in position and 0.00622° in orientation. Six replay repeats show a maximum local RMSE variation of 0.000940 cm.

The information score is computed from a nominal, causal candidate-action model and a fixed landmark map; it does not consume future ground truth. Ground truth is used after each run to score local error, route progress, scale ratio, and task cost. The experimental unit is an independent route or route-level paired branch, not an individual video frame. This choice preserves the dependency among frames and prevents artificial sample-size inflation.

5.3 Scoring and decision metrics

Let $\mathcal{I}$ be the set of camera timestamps with a valid estimator output, and let $\tilde{p}_i$ and p_i be estimated and ground-truth positions after the single rigid alignment defined in table 5. Scale remains fixed. Full-trajectory

Table 6: Full-trajectory position RMSE (cm) on four frozen lunar scenarios. The endpoint pass criterion is below 3%; DBA-Fusion outputs are incomplete.

Estimator	Polyline	Multi-turn	Low sun	Serpentine	Pass / status
VINS monocular	1.658	2.371	1.245	3.699	4/4
VINS stereo	0.358	0.402	0.389	0.684	4/4
Stable stereo	0.568	0.531	1.166	2.738	4/4
Mono MSCKF	1.91×10^4	344.38	1.40×10^6	1.80×10^6	0/4
PM-SLAM	313.958	290.731	1190.0	346.683	0/4
Selective-VIO	198.529	248.985	181.558	624.919	0/4
DBA-Fusion	227.409†	$1.11 \times 10^{6\dagger}$	$1.12 \times 10^{6\dagger}$	$1.64 \times 10^{6\dagger}$	0/4 complete

†Partial estimator output; the corresponding terminal trajectory was not complete. RMSE values describe the available output interval only.

position error and endpoint error are

$$E_{\text{full}} = \sqrt{\frac{1}{|\mathcal{I}|}\sum_{i\in\mathcal{I}} \|\tilde{p}_i - p_i\|^2}, \qquad e_{\text{end}} = 100\frac{\|\tilde{p}_N - p_N\|}{L_{\text{GT}}}\%, \tag{18}$$

where L_{GT} is the cumulative ground-truth path length evaluated at camera timestamps. For a decision window $\mathcal{I}_H$, the same root-mean-square definition gives E_H; timestamps are mapped back through any imposed camera-time ramp before pairing, and only actual estimator output frames are scored. No pose is synthesized between outputs. Valid-output coverage is $|\mathcal{I}|/N_{\text{camera}}$, reported separately so a low error over a short surviving segment cannot appear equivalent to complete tracking.

For an action a compared with continuing motion a_0, the local error reduction is $\Delta E(a) = E_H(a_0) - E_H(a)$. Positive values indicate a smaller RMSE under the maneuver. When action duration differs, the primary comparison uses a common progress interval and reports elapsed time and traveled distance alongside error. A policy's retrospective decision regret uses the frozen task cost C and feasible action set $\mathcal{A}$:

$$\mathcal{R} = C(a_{\text{selected}}) - \min_{a\in\mathcal{A}} C(a). \tag{19}$$

Ground truth contributes to these offline scores and regret labels only. Estimation, prediction, and action selection receive the live sensor stream and causal state history described in section 4.

6 Preliminary Results and Baseline Comparisons

6.1 Seven-estimator lunar benchmark

Seven frozen estimators were evaluated on four identical camera–IMU route logs. Terminal error is the final 3D position error divided by the camera-time ground-truth path length; the engineering threshold is 3%. The monocular VINS, stereo VINS, and Stable stereo pipeline each pass all four routes. The four-route mean RMSE is 2.243 cm for monocular VINS, 0.458 cm for stereo VINS, and 1.250 cm for Stable stereo. The baseline therefore establishes a strong traditional stereo reference and a strong monocular reference for the proposed diagnosis and action study.

On 21 additional straight-cruise runs, monocular VINS terminal error ranges from 3.27% to 8.77% of route length, and no run meets the 3% criterion. This contrast with the four mixed-route cases motivates a motion-dependent analysis: an estimator that performs well when turns and visual geometry supply excitation may accumulate substantially more drift during long, nearly straight travel.

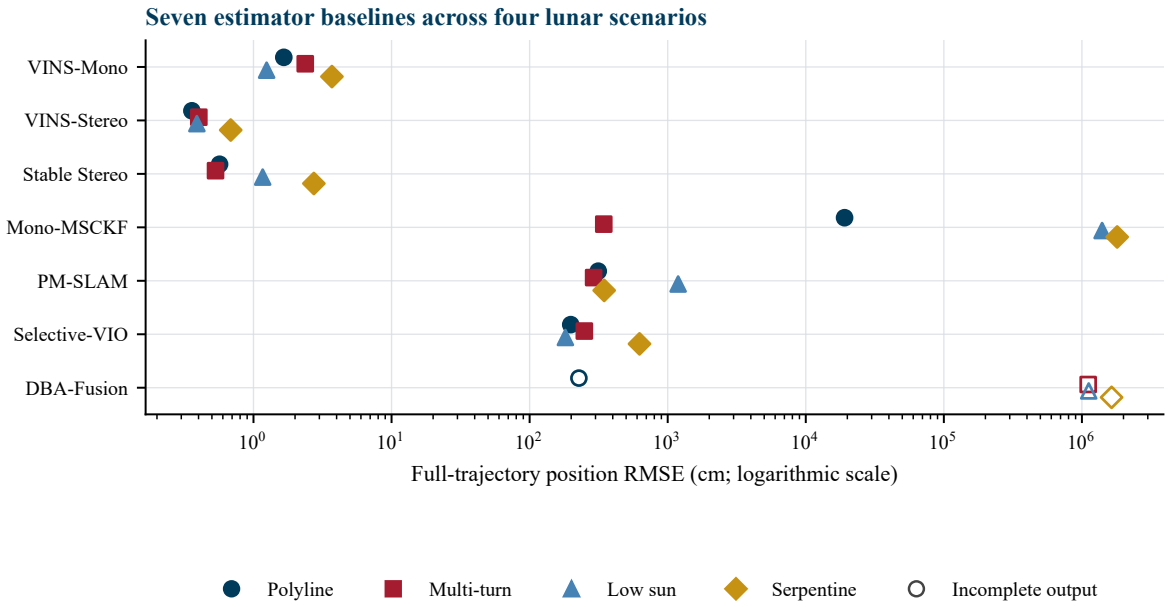


Figure 6: Seven-estimator position RMSE over the four lunar scenarios. The logarithmic axis keeps complete and divergent outputs visible in a single panel; open symbols denote incomplete output.

Table 7: Controlled camera–IMU timing offsets: full-route RMSE (cm) and terminal error (% of route length).

Offset	Route 1	Route 2	Route 3	Mean residual indicators
0 ms	10.23 / 1.32%	6.89 / 1.10%	3.23 / 1.37%	$V_{\text{ratio}} = 1.38$, $W = 0.36$
40 ms	33.06 / 4.04%	24.13 / 5.35%	12.33 / 4.19%	$V_{\text{ratio}} = 6.38$, $W = 2.05$
80 ms	70.37 / 8.89%	31.55 / 7.15%	27.46 / 10.35%	$V_{\text{ratio}} = 10.78$, $W = 3.48$

6.2 Controlled timing-offset response

The fixed estimator initialized on all three controlled routes, produced no backend restarts, and covered approximately 95% of camera frames. The camera timestamp was held unshifted through 15 s, ramped to its target offset by 25 s, and then held. With online bias estimation enabled and the same noisy IMU bottom stream, full-route RMSE increased across the 0, 40, and 80 ms conditions on every route.

The 40 ms condition raises full-route RMSE by 3.23–3.82 times over the matched condition; 80 ms raises it by 4.58–8.50 times. The 15 s straight window from 25 to 40 s remains comparatively accurate, with RMSE between 0.45 and 2.98 cm. The timing penalty concentrates around motion with stronger angular change: turn-window RMSE increases monotonically at both offset levels on all three routes. The result agrees with the first-order coupling in equation (8): translation contributes through $v\delta t$, while turning adds orientation error through $\omega\delta t$.

6.3 Reliability signals and estimator correction

The online signals separate the three offset levels. The three-route mean visual cost before optimization rises from 18.11 to 128.70 and 283.17; post-optimization visual cost rises more slowly, from 13.06 to 20.68 and 26.45. The median visual-cost absorption ratio V_{ratio} therefore increases from 1.38 to 6.38 and 10.78. The whitened inertial residual W increases from 0.36 to 2.05 and 3.48. These signals provide a measurable reliability input for the proposed predictor, even when backend optimization partially absorbs the visual inconsistency.

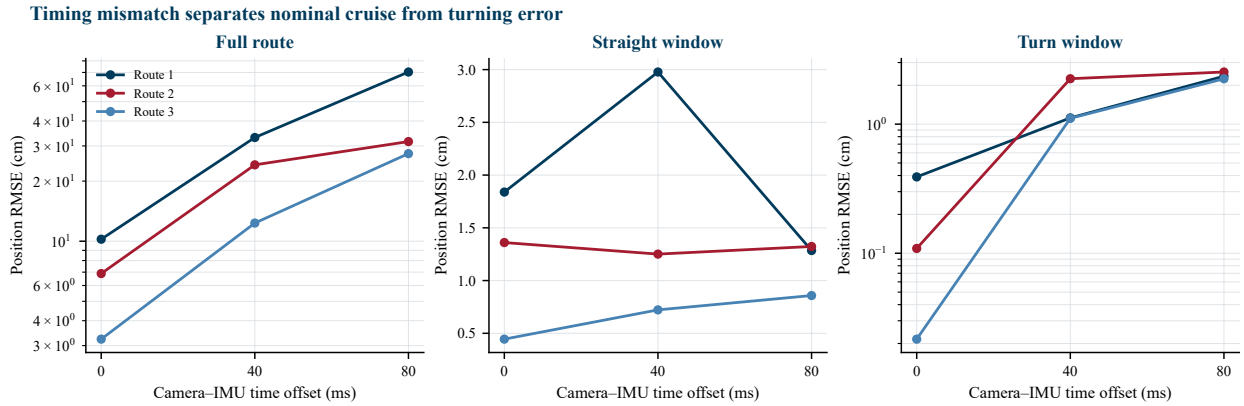


Figure 7: Full-route, straight-window, and turn-window position RMSE under controlled camera–IMU time offsets. Route colors are held constant across panels; the straight segment stays near its matched-condition error while turn and full-route error rise.

Table 8: Available straight-cruise action comparisons: equal-progress local RMSE (cm). A dash denotes an action branch not present in the accepted data.

Route	Offset	Continue	Reduce speed 50%	S-curve	Pause 5 s
1	0 ms	1.947	0.867	2.839	0.837
1	40 ms	4.942	0.899	5.675	5.132
1	80 ms	1.857	14.324	6.081	8.598
2	0 ms	2.547	1.202	3.103	—
2	40 ms	2.903	0.665	8.653	—
2	80 ms	2.517	19.259	9.836	—

A separate short-window propagation–correction decomposition attributes 89.3–94.9% of the measured displacement deficit to backend position correction, while inertial velocity-integration drift contributes 0.2–0.6%. This result cautions against interpreting a small post-optimization residual as proof that the propagated and optimized trajectories are physically correct. The proposed reliability module uses the pre/post residual relation together with motion and constraint evidence.

6.4 Physical-action evidence

The paired action study shows that an action's value depends on both motion context and mismatch strength. Table 8 gives the available equal-progress straight-cruise comparisons. On Route 1 at 80 ms, reducing speed worsens RMSE from 1.857 to 14.324 cm; the S-shaped maneuver and pause also worsen it. On Route 2, the same 80 ms slowdown increases error from 2.517 to 19.259 cm. At 40 ms on Route 2, the S-shaped maneuver increases RMSE from 2.903 to 8.653 cm despite a positive nominal information score.

Turn-entry outcomes are more favorable for reducing speed under the strongest tested offset. At 80 ms, the local RMSE falls from 13.733 cm under continued travel to 3.140 cm with a 50% speed reduction, a 77.1% decrease. The outgoing scale ratio changes from 0.701 to 0.954. At 40 ms, however, the lowest RMSE is obtained by the S-shaped maneuver, not by slowing down; at 0 ms, a brief pause has the lowest local RMSE. Thus, the same physical action does not dominate every condition.

The action data were checked against shared pre-decision state and repeated replay. They support a local mechanism claim: physically reducing angular-rate demand can mitigate timing-induced turn error, while the benefit depends on context and mismatch level. The action predictor and full closed-loop policy remain the next research objectives.

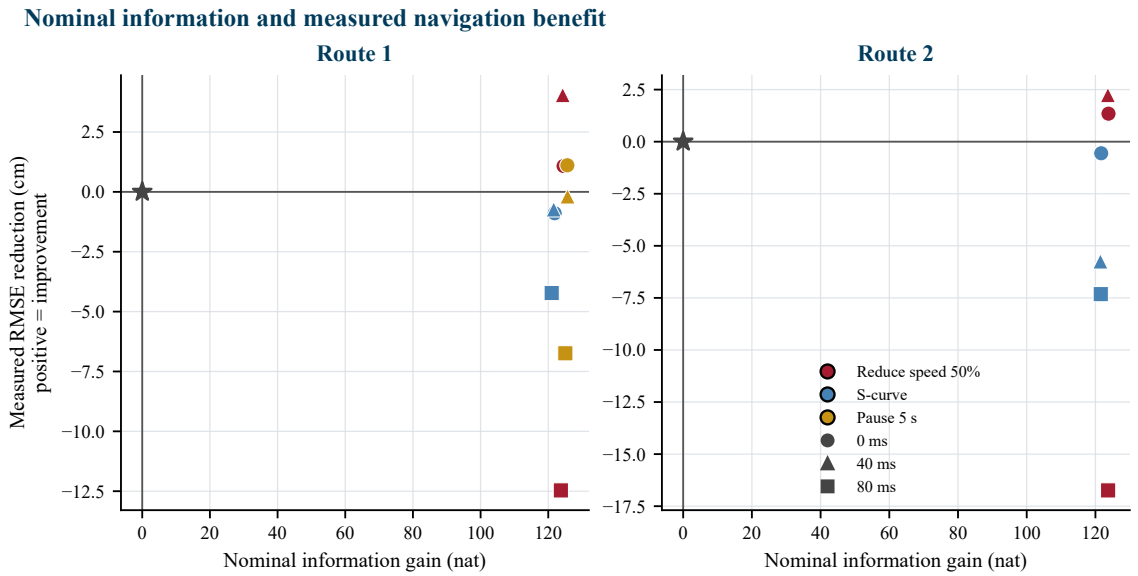


Figure 8: Nominal information gain and measured equal-progress RMSE reduction for the available straight-cruise action branches. Points below zero indicate that the maneuver worsened navigation.

Table 9: Available turn-entry comparisons on one route: local RMSE (cm) and outgoing scale ratio.

Offset	Continue	Reduce speed 50%	S-curve	Pause 5 s	Scale ratio, continue / reduce speed
0 ms	1.583	2.710	1.840	0.778	0.990 / 0.951
40 ms	6.914	5.383	2.886	3.134	0.860 / 0.925
80 ms	13.733	3.140	8.288	7.826	0.701 / 0.954

7 Doctoral Work Plan and Acceptance Criteria

The doctoral study proceeds from a controlled mechanism to a calibrated predictor, then to closed-loop rover evaluation and a spacecraft transfer test. The phases are sequential in their acceptance gates but partially overlap in implementation. Each data split is grouped by independent route and terrain identity; all branches from a common pre-decision prefix remain in the same split. Parameters are selected on development routes, frozen before the held-out evaluation, and never updated using the final test outcomes.

7.1 Phase I: diagnose reliability under controlled mismatch

Months 1–8 will establish the first controlled mismatch and characterize when current estimator constraints remain informative. The study will retain one monocular visual–inertial backend and compare nominal timing, moderate offset, and stronger offset under a fixed 15 s horizon. At least 18 independently seeded physical routes will be assembled before fitting a high-dimensional predictor. Route-level repetition will be selected after measuring action variability and mismatch repeatability; 18 is a planning floor, not a claim of statistical sufficiency.

The reliability module will first be compared with residual thresholds, robust weighting, and local geometric correction. The main quantities are detection delay, missed mismatch events, false acceptance, valid-observation retention, and the downstream change in local navigation error. The analysis will distinguish geometry-level observability from measurement inconsistency and tracking loss. If the reliability output does not improve a fixed prediction or selection rule, its downstream value will not be claimed.

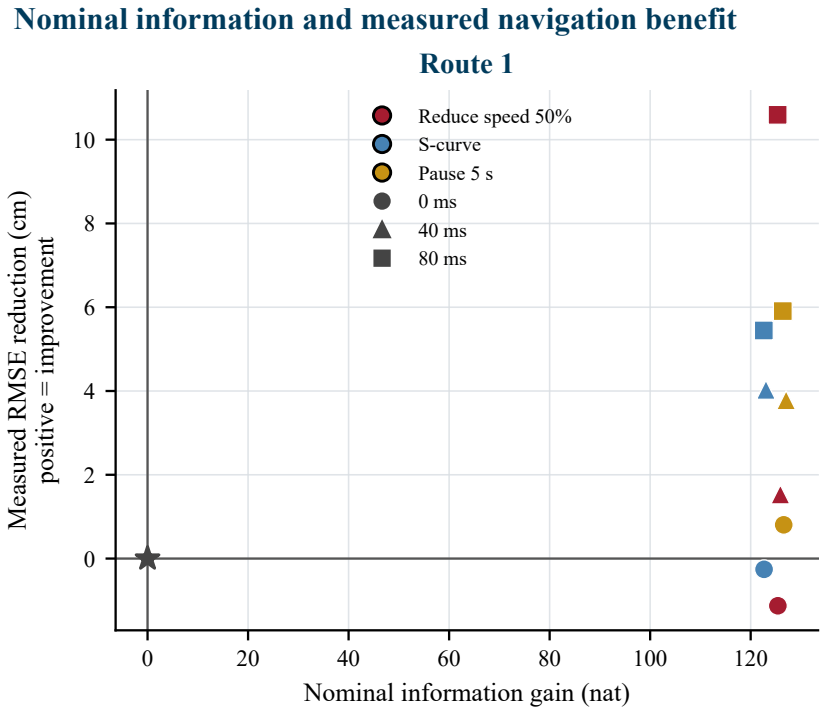


Figure 9: Nominal information gain and measured RMSE reduction for available turn-entry action branches. The lower-left and upper-right quadrants distinguish actions that worsen navigation from those that improve it.

7.2 Phase II: predict and select under matched budgets

Months 7–17 will develop the short-horizon action predictor, calibrate its uncertainty by route group, and test closed-loop rover motion. Candidate predictor classes are fixed-frequency, transparent rule-based, nominal geometric/Fisher, geometry plus reliability, and a lightweight learned component only if a structured residual remains. The learned component will use only decision-time features and will be compared against simpler models on held-out routes.

The core closed-loop groups are task-only motion, fixed slowdown or recovery, nominal-information selection, the proposed predictor without reliability features, and the complete reliability-aware predictor. Each group uses the same estimator, initial-state distribution, sensor configuration, target, safety constraints, and maximum task budget. Navigation error, endpoint error, recovery, completion rate, time, route distance, turning effort, and constraint violations are reported separately. Energy is deferred until a validated actuator power model is available.

7.3 Phase III: transfer the decision interface to angles-only navigation

Months 18–24 will instantiate the shared interface with spacecraft relative dynamics and bearing-only observations. The spacecraft adapter will supply a relative-state model, feasible bounded thrust maneuvers, fuel/time cost, and safety constraints. Initial tests will focus on range-related observability under coasting and view-changing motion; the same reliability and calibrated-benefit interface will be reused where the observation process permits.

7.4 Statistics, reproducibility, and decision gates

The independent unit is a route or route-level paired decision point. Video frames from the same route are not counted as independent samples. Main comparisons use paired route differences with route-cluster bootstrap

Table 10: Planned evaluation matrix and quantitative acceptance rules. Thresholds that depend on task tolerance are frozen on development routes before the held-out test.

Claim	Comparison	Primary metrics	Acceptance rule
Reliability diagnosis	Residual threshold, robust weighting, local geometric correction	Missed mismatch rate, false acceptance, detection delay, observation retention	Compare at matched observation-retention or false-acceptance rate; route-clustered confidence intervals must show improved detection or lower downstream navigation error.
Action-outcome prediction	Frequency, rule, nominal geometry/Fisher, geometry plus reliability, optional learning	RMSE, MAE, rank correlation, Brier score, route-simultaneous 90% coverage	Held-out route MAE and ranking loss must improve over the strongest fixed/rule baseline; report simultaneous candidate-set coverage and bound width with a 95% binomial confidence interval.
Closed-loop selection	Task-only, fixed recovery, nominal-information, predictor without reliability, full method	Route-level decision regret, local/full RMSE, completion, time, distance, constraints	Paired route-level regret must improve over the strongest fixed policy; completion non-inferiority margin is fixed from task tolerance before testing.
Spacecraft transfer	Coasting, fixed excitation, nominal-information, calibrated active selection	Relative-state and range error, risk calibration, propellant/time, safety events	A separate held-out orbit set must show valid calibration and lower regret under matched fuel/time constraints.

intervals or an appropriate paired permutation test. Probabilistic outputs are evaluated with Brier scores and reliability curves; prediction intervals are assessed for empirical coverage and width. All failures, initialization failures, timeouts, action rejections, and non-recovery events remain in the denominator for completion and reliability metrics.

The second-phase gate requires the frozen predictor to outperform fixed or simple rule baselines on held-out routes and to maintain its stated calibration. The closed-loop gate requires lower route-level regret under matched budgets while preserving task completion within the pre-registered tolerance. If a simpler fixed maneuver performs equally well after costs are matched, the simpler policy is retained and the proposed selector's claim is narrowed. The spacecraft phase proceeds only after the rover decision interface and calibration protocol are stable.

8 Expected Contributions and Scope

8.1 Expected contributions

The thesis is expected to make three methodological contributions.

1. **Estimator-grounded reliability representation.** A compact state will describe whether current visual–inertial constraints are geometrically weak, stochastically degraded, persistently mismatched, or unavailable because tracking has failed. Its value will be measured not only by classification accuracy but by whether it improves the downstream prediction and selection task.
2. **Physics-guided prediction of navigation benefit.** Candidate physical actions will be evaluated in estimator-level quantities such as local trajectory error, recovery probability, and route-progress error. A

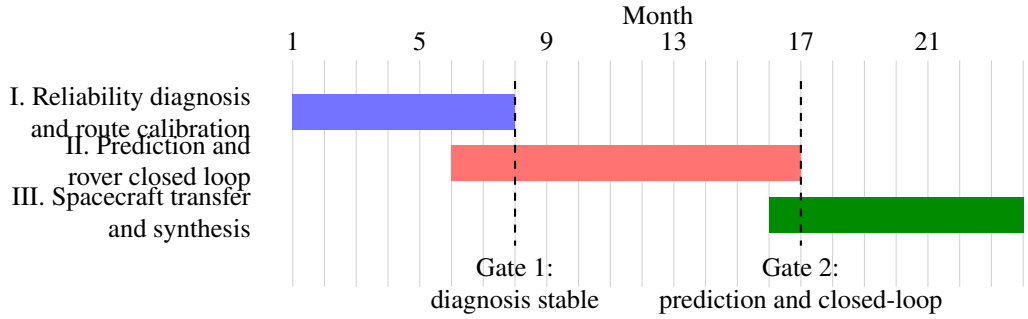


Figure 10: Twenty-four-month plan with route-grouped acceptance gates. The second phase begins after the first reliability and data protocol is frozen; spacecraft transfer follows the rover closed-loop test.

calibrated prediction interval will separate nominal geometric information from the navigation benefit supported by held-out route data.

3. **Risk-aware active navigation under task cost.** A bounded selector will use predicted benefit and uncertainty alongside time, distance, turning or propellant cost, and safety constraints. The evaluation will report decision regret against a finite action oracle and compare against frozen task-only, fixed-recovery, and nominal-information baselines.

The scientific contribution lies in connecting reliability diagnosis to calibrated prediction of the estimator outcome and then to a physically executed action. Information gain, robust weighting, and task-cost optimization are established tools; this proposal tests how the decision changes when their inputs are informed by the reliability of the observations and constraints.

8.2 Scope and limitations

The present evidence establishes a controlled timing-offset response and several route-specific physical-action outcomes in a lunar simulator. The timing study includes three routes; the accepted action branches cover only a subset of route and context combinations. No calibrated predictor or full closed-loop policy has yet passed held-out route evaluation. The action observations therefore motivate and constrain the method rather than serve as a final performance claim.

The current camera and inertial streams are simulation products. The deterministic motion-derived inertial baseline and the controlled noisy sensor model answer different questions; neither substitutes for characterization of the intended hardware IMU, clock, camera, or thermal behavior. The four-route baseline set also does not span the full range of regolith deformation, dust, hardware vibration, or shadow dynamics encountered in field operation. Data with single images can validate appearance and feature processing but cannot establish trajectory accuracy or action benefit without temporal sequences, calibration, reference motion, and an action execution model.

The shared spacecraft interface is a transfer hypothesis. Angles-only relative navigation has a different state, dynamics, and safety envelope from rover VIO, and requires its own representative orbit dataset and resource model. Finite-sample calibration is marginal under exchangeability of route groups; it does not guarantee conditional coverage on a shifted terrain or untested sensor regime. These boundaries define the next validation steps rather than weaken the central question.

References

[1] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 1004–1020, 2018. [Online]. Available: https://doi.org/10.1109/tro.2018.2853729

[2] L. von Stumberg and D. Cremers, "Dm-vio: Delayed marginalization visual-inertial odometry," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 1408–1415, 2022. [Online]. Available: https://doi.org/10.1109/lra.2021.3140129

[3] M. Li and A. I. Mourikis, "Online temporal calibration for camera–imu systems: Theory and algorithms," *The International Journal of Robotics Research*, vol. 33, no. 7, pp. 947–964, 2014. [Online]. Available: https://doi.org/10.1177/0278364913515286

[4] T. Qin and S. Shen, "Online temporal calibration for monocular visual-inertial systems," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3662–3669. [Online]. Available: https://doi.org/10.1109/iros.2018.8593603

[5] Y. Niu, J. Liu, X. Wang, W. Hao, W. Li, and L. Chen, "Accurate and robust odometry by fusing monocular visual, inertial, and wheel encoder," *CCF Transactions on Pervasive Computing and Interaction*, vol. 2, no. 4, pp. 275–287, 2020. [Online]. Available: https://doi.org/10.1007/s42486-020-00040-4

[6] W. Huang, H. Liu, and W. Wan, "Online initialization and extrinsic spatial-temporal calibration for monocular visual-inertial odometry," *arXiv*, 2020. [Online]. Available: https://doi.org/10.48550/arxiv.2004.05534

[7] J. Delaune, D. S. Bayard, and R. Brockers, "Range-visual-inertial odometry: Scale observability without excitation," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2421–2428, 2021. [Online]. Available: https://doi.org/10.1109/lra.2021.3058918

[8] Y. Zhou, X. Li, S. Li, X. Wang, and Z. Shen, "Ground-vio: Monocular visual-inertial odometry with online calibration of camera-ground geometric parameters," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 10, pp. 14 328–14 343, 2024. [Online]. Available: https://doi.org/10.1109/tits.2024.3393125

[9] H. Hailu, B. Gebregziabher, S. Gudipudi, and M. Bhatta, "Quantifying motion excitation for metric scale observability in monocular visual-inertial odometry," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2603.26740

[10] S. Li, J. Nie, C. Guo, Y. Yang, and L. Mei, "A novel online time calibration framework via double-stage ekf for visual-inertial odometry," *Robotica*, vol. 42, no. 11, pp. 3675–3691, 2024. [Online]. Available: https://doi.org/10.1017/s0263574724001589

[11] W. Zhou, Y. Pan, J. Liu, T. Wang, and H. Zha, "Visual-inertial-wheel odometry with wheel-aided maximum-a-posteriori initialization for ground robots," *IEEE Robotics and Automation Letters*, vol. 9, no. 5, pp. 4814–4821, 2024. [Online]. Available: https://doi.org/10.1109/lra.2024.3386448

[12] W. Li, G. Wang, Q. Zhang, and J. Liu, "Lg-viwo: Visual-inertial-wheel odometry leveraging the depth-aided local ground constraints for mobile robots," *IEEE Transactions on Instrumentation and Measurement*, vol. 74, pp. 1–11, 2025. [Online]. Available: https://doi.org/10.1109/tim.2025.3565096

[13] Y. Fan, T. Zhao, L. Guo, C. Chen, X. Wang, and F. Zhou, "Universal online temporal calibration for optimization-based visual-inertial navigation systems," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2025, pp. 1–7. [Online]. Available: https://doi.org/10.1109/icra55743.2025.11127433

[14] Y. Yoon, S. Yang, and J. Lim, "Lirs: Linear relative scale constraints for compact parameterization of monocular visual-inertial initialization," *IEEE Robotics and Automation Letters*, vol. 11, no. 9, pp. 10 409–10 416, 2026. [Online]. Available: https://doi.org/10.1109/lra.2026.3710045

[15] T. Hua, T. Li, and L. Pei, "Piekf-viwo: Visual-inertial-wheel odometry using partial invariant extended kalman filter," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2023, pp. 2083–2090. [Online]. Available: https://doi.org/10.1109/icra48891.2023.10160380

[16] R. Tian, Y. Zhang, D. Zhu, S. Liang, S. Coleman, and D. Kerr, "Accurate and robust scale recovery for monocular visual odometry based on plane geometry," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2021, pp. 5296–5302. [Online]. Available: https://doi.org/10.1109/icra48506.2021.9561215

[17] W. Huang, W. Wan, and H. Liu, "Optimization-based online initialization and calibration of monocular visual-inertial odometry considering spatial-temporal constraints," *Sensors*, vol. 21, no. 8, p. 2673, 2021. [Online]. Available: https://doi.org/10.3390/s21082673

[18] S. A. Cerezo, S. H. Lee, and J. Civera, "An efficient closed-form solution to full visual-inertial state initialization," *IEEE Robotics and Automation Letters*, vol. 11, no. 6, pp. 6919–6926, 2026. [Online]. Available: https://doi.org/10.1109/lra.2026.3682536

[19] Z. Feng, J. Li, L. Zhang, and C. Chen, "Online spatial and temporal calibration for monocular direct visual-inertial odometry," *Sensors*, vol. 19, no. 10, p. 2273, 2019. [Online]. Available: https://doi.org/10.3390/s19102273

[20] C. Mu, D. Feng, Q. Zheng, and Y. Zhuang, "A robust and efficient visual-inertial initialization with probabilistic normal epipolar constraint," *IEEE Robotics and Automation Letters*, vol. 10, no. 4, pp. 3590–3597, 2025. [Online]. Available: https://doi.org/10.1109/lra.2025.3544522

[21] A. Abi, M. Boughellaba, M. Wang, and A. Tayebi, "Reduced order observers for monocular visual inertial odometry," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2609.26612

[22] G. Peng, Z. Lu, C. Bocheng, S. Chen, and D. He, "Robust tightly coupled pose estimation based on monocular vision, inertia, and wheel speed," *arXiv*, 2020. [Online]. Available: https://doi.org/10.48550/arxiv.2003.01496

[23] J. Song, A. Richard, and M. A. Olivares-Méndez, "Unleashing the power of discrete-time state representation: Ultrafast target-based imu-camera spatial-temporal calibration," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2026, pp. 12 755–12 762. [Online]. Available: https://doi.org/10.1109/icra57385.2026.11696581

[24] ——, "Improving monocular visual-inertial initialization with structureless visual-inertial bundle adjustment," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2025, pp. 1039–1045. [Online]. Available: https://doi.org/10.1109/icra55743.2025.11127304

[25] A. Cauligi, R. M. Swan, H. Ono, S. Daftry, J. Elliott, L. H. Matthies, and D. Atha, "Shadownav: Crater-based localization for nighttime and permanently shadowed region lunar navigation," in *IEEE Aerospace Conference (AERO)*, 2023, pp. 1–12. [Online]. Available: https://doi.org/10.1109/aero55745.2023.10115745

[26] S. Daftry, Z. Chen, C. Yang, S. Tepsuporn, S. Khattak, L. H. Matthies, B. Coltin, U. Naal, L. M. Ma, and M. C. Deans, "Lunarnav: Crater-based localization for long-range autonomous lunar rover navigation," in *IEEE Aerospace Conference (AERO)*, 2023, pp. 1–15. [Online]. Available: https://doi.org/10.1109/aero55745.2023.10115640

[27] D. Atha, R. M. Swan, A. Cauligi, A. Bettens, E. Y. Goh, D. Kogan, L. H. Matthies, and M. Ono, "Shadownav: Autonomous global localization for lunar navigation in darkness," *IEEE transactions on field robotics.*, vol. 1, pp. 213–230, 2024. [Online]. Available: https://doi.org/10.1109/tfr.2024.3462391

[28] L. H. Matthies, S. Daftry, S. Tepsuporn, Y. Cheng, D. Atha, R. M. Swan, S. Ravichandar, and M. Ono, "Lunar rover localization using craters as landmarks," in *IEEE Aerospace Conference (AERO)*, 2022, pp. 1–17. [Online]. Available: https://doi.org/10.1109/aero53065.2022.9843714

[29] J. Strader, K. Otsu, and A. Agha–mohammadi, "Perception-aware autonomous mast motion planning for planetary exploration rovers," *Journal of Field Robotics*, vol. 37, no. 5, pp. 812–829, 2019. [Online]. Available: https://doi.org/10.1002/rob.21925

[30] B.-H. Chen, P. Negrut, T. Liang, N. Batagoda, H. Zhang, and D. Negruţ, "Polar-sim: Augmenting nasa's polar dataset for data-driven lunar perception and rover simulation," *IEEE Aerospace and Electronic Systems Magazine*, vol. 41, no. 1, pp. 4–18, 2025. [Online]. Available: https://doi.org/10.1109/maes.2025.3617426

[31] R. Castilla-Arquillo and M. A. Olivares-Méndez, "Mare: An open-source lunar rover simulator with shadow ground truth for perception and autonomy," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2610.28465

[32] J. Liu, X. Luo, Q. Zhang, X. Wan, C. Ye, S. Zhang, Y. Tian, H. Han, Y. Zhao, B. Liu, Z. Zhao, and Z. Kang, "Lusnar: A lunar segmentation, navigation, and reconstruction dataset based on multisensor for autonomous exploration," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 64, p. 4706723, 2026. [Online]. Available: https://doi.org/10.1109/tgrs.2026.3720373

[33] D. Rodríguez-Martínez, D. van der Meer, J. Song, A. Bera, C. J. P. del Pulgar, and M. A. Olivares-Méndez, "Spice-hl3: Single-photon, inertial, and stereo camera dataset for exploration of high-latitude lunar landscapes," *Scientific Data*, vol. 13, no. 1, p. 374, 2026. [Online]. Available: https://doi.org/10.1038/s41597-026-06668-8

[34] C. Ye, Z. Kang, X. Wan, S. Shao, J. Yang, and Z. Cao, "Map-guided coarse-to-fine lunar rover localization via multiview crater matching," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 19 884–19 896, 2025. [Online]. Available: https://doi.org/10.1109/jstars.2025.3594111

[35] P. Mithun, A. S, and A. Z. F, "Autonomous lunar crater detection and matching for precision terrain-relative navigation," in *International Conference on Electronics and Sustainable Communication Systems (ICESC)*, 2026, pp. 438–443. [Online]. Available: https://doi.org/10.1109/icesc70182.2026.11687089

[36] B. Candan and S. Servadio, "Deep learning-based lunar crater terrain relative navigation," in *International Conference on Information Fusion (FUSION)*, 2026, pp. 1–7. [Online]. Available: https://doi.org/10.23919/fusion67272.2026.11726646

[37] A. Sekmen, F. E. Gunes, F. Horoz, H. U. Işık, M. A. Ozaydin, O. A. Topaloglu, Şahin Umutcan Üstündaş, Y. A. Yeni, H. E. Söken, E. Sahin, R. G. Cinbiş, and S. Kalkan, "Lumon: A comprehensive benchmark and development suite with novel datasets for lunar monocular depth estimation," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2604.09352

[38] A. Dai, G. C. Vila, and G. Gao, "Visual slam with dem anchoring for lunar surface navigation," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2603.17229

[39] G. P. Parracino, M. Ceresoli, S. Silvestrini, and M. R. Lavagna, "Integrated optical terrain relative navigation for autonomous lunar landing," *IAF Astrodynamics Symposium*, pp. 355–369, 2024. [Online]. Available: https://doi.org/10.52202/078368-0033

[40] C. Grethen, Y. Shi, S. Gasparini, and G. Morin, "Moonanything: A vision benchmark with large-scale lunar supervised data," in *ACM Multimedia Systems Conference (MMSys)*, 2026, pp. 457–463. [Online]. Available: https://doi.org/10.1145/3793853.3799814

[41] X. Lin, T. Li, B. Hua, L. Li, and C. Zhao, "Cratersense: Efficient and robust crater-based lunar localization system with structure-aware sparse transformer," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 19, pp. 49–69, 2025. [Online]. Available: https://doi.org/10.1109/jstars.2025.3633611

[42] M. Hansen, U. Y. Wong, and T. Fong, "The polar traverse dataset: A dataset of stereo camera images simulating traverses across lunar polar terrain under extreme lighting conditions," *arXiv*, 2024. [Online]. Available: https://doi.org/10.48550/arxiv.2403.12194

[43] D. Woffinden, "Angles-only navigation for autonomous orbital rendezvous," *Digital Commons - USU (Utah State University)*, 2008. [Online]. Available: https://doi.org/10.26076/15dc-267c

[44] D. Woffinden and D. K. Geller, "Observability criteria for angles-only navigation," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 45, no. 3, pp. 1194–1208, 2009. [Online]. Available: https://doi.org/10.1109/taes.2009.5259193

[45] ——, "Optimal orbital rendezvous maneuvering for angles-only navigation," *Journal of Guidance Control and Dynamics*, vol. 32, no. 4, pp. 1382–1387, 2009. [Online]. Available: https://doi.org/10.2514/1.45006

[46] J. Grzymisch and W. Fichter, "Analytic optimal observability maneuvers for in-orbit bearings-only rendezvous," *Journal of Guidance Control and Dynamics*, vol. 37, no. 5, pp. 1658–1664, 2014. [Online]. Available: https://doi.org/10.2514/1.g000612

[47] F. J. Franquiz, J. Muñoz, B. Udrea, and M. J. Balas, "Optimal range observability maneuvers of a spacecraft formation using angles-only navigation," *Acta Astronautica*, vol. 153, pp. 337–348, 2018. [Online]. Available: https://doi.org/10.1016/j.actaastro.2018.02.003

[48] R. Du, W. Liao, and X. Zhang, "Optimal maneuver strategy to improve the observability of angles-only rendezvous," *Journal of Systems Engineering and Electronics*, vol. 34, no. 4, pp. 1020–1032, 2023. [Online]. Available: https://doi.org/10.23919/jsee.2023.000091

[49] L. Ticozzi and P. Tsiotras, "Factor graph-based active slam for spacecraft proximity operations," *arXiv*, 2025. [Online]. Available: https://doi.org/10.48550/arxiv.2501.10950

[50] J. Ragan, B. Rivière, F. Y. Hadaegh, and S. Chung, "Online tree-based planning for active spacecraft fault estimation and collision avoidance," *Science Robotics*, vol. 9, no. 93, p. eadn4722, 2024. [Online]. Available: https://doi.org/10.1126/scirobotics.adn4722

[51] J. A. Placed, J. Strader, H. Carrillo, N. Atanasov, V. Indelman, L. Carlone, and J. A. Castellanos, "A survey on active simultaneous localization and mapping: State of the art and new frontiers," *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1686–1705, 2023. [Online]. Available: https://doi.org/10.1109/tro.2023.3248510

[52] B. Gong, S. Wang, S. Li, and X. Li, "Review of space relative navigation based on angles-only measurements," *Astrodynamics*, vol. 7, no. 2, pp. 131–152, 2022. [Online]. Available: https://doi.org/10.1007/s42064-022-0152-2

[53] M. Dor, T. Driver, K. M. Getzandanner, and P. Tsiotras, "Astroslam: Autonomous monocular navigation in the vicinity of a celestial small body—theory and experiments," *The International Journal of Robotics Research*, 2024. [Online]. Available: https://doi.org/10.1177/02783649241234367

[54] A. Kitanov and V. Indelman, "Topological belief space planning for active slam with pairwise gaussian potentials and performance guarantees," *The International Journal of Robotics Research*, vol. 43, no. 1, pp. 69–97, 2023. [Online]. Available: https://doi.org/10.1177/02783649231204898

[55] I. Gusija, F. Alajaji, and S. Yüksel, "Slam as a stochastic control problem with partial information: Optimal solutions and rigorous approximations," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2604.21693

[56] N. Atanasov, J. L. Ny, K. Daniilidis, and G. J. Pappas, "Decentralized active information acquisition: Theory and application to multi-robot slam," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2015, pp. 4775–4782. [Online]. Available: https://doi.org/10.1109/icra.2015.7139863

[57] M. Wijayatunga, R. Armellin, and H. Holt, "Robust trajectory design and guidance for far-range rendezvous using reinforcement learning with safety and observability considerations," *Aerospace Science and Technology*, vol. 159, p. 109996, 2025. [Online]. Available: https://doi.org/10.1016/j.ast.2025.109996

[58] A. T. Espinoza, A. T. Espinoza, J. B. Folkesson, C. Deutsch, N. Rolleberg, P. Sigray, and J. Kuttenkeuler, "Stern: Simultaneous trajectory estimation and relative navigation for autonomous underwater proximity operations," *IEEE Journal of Oceanic Engineering*, vol. 51, no. 1, pp. 293–316, 2025. [Online]. Available: https://doi.org/10.1109/joe.2025.3624470

[59] J. A. Placed, J. J. G. Rodríguez, J. D. Tardós, and J. A. Castellanos, "Explorb-slam: Active visual slam exploiting the pose-graph topology," *Lecture notes in networks and systems*, pp. 199–210, 2022. [Online]. Available: https://doi.org/10.1007/978-3-031-21065-5_17

[60] F. J. Franquiz, "Spacecraft trajectory planning for optimal observability using angles-only navigation," *Scholarly Commons (Embry–Riddle Aeronautical University)*, 2019.

[61] R. Du, W. Liao, and X. Zhang, "Optimal closed-loop robust rendezvous guidance with the observability constraint of angles-only navigation," *Transactions of the Institute of Measurement and Control*, vol. 45, no. 14, pp. 2748–2759, 2023. [Online]. Available: https://doi.org/10.1177/01423312231166044

[62] Y. Zhao, Z. Xiong, J. Wang, L. Zhang, and P. Campoy, "Visual active slam method considering measurement and state uncertainty for space exploration," *Aerospace*, vol. 12, no. 7, p. 642, 2025. [Online]. Available: https://doi.org/10.3390/aerospace12070642

[63] M. Lauri, D. Hsu, and J. Pajarinen, "Partially observable markov decision processes in robotics: A survey," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 21–40, 2022. [Online]. Available: https://doi.org/10.1109/tro.2022.3200138

[64] V. Indelman, L. Carlone, and F. Dellaert, "Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 849–882, 2015. [Online]. Available: https://doi.org/10.1177/0278364914561102

[65] S. Pathak, A. Thomas, and V. Indelman, "A unified framework for data association aware robust belief space planning and perception," *The International Journal of Robotics Research*, vol. 37, no. 2-3, pp. 287–315, 2018. [Online]. Available: https://doi.org/10.1177/0278364918759606

[66] S. Rahman and S. L. Waslander, "Uncertainty-constrained differential dynamic programming in belief space for vision based robots," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3112–3119, 2021. [Online]. Available: https://doi.org/10.1109/lra.2021.3062338

[67] Z. Zhang and D. Scaramuzza, "Fisher information field: an efficient and differentiable map for perception-aware planning," *arXiv*, 2020. [Online]. Available: https://doi.org/10.48550/arxiv.2008.03324

[68] X. Wu, S. Chen, K. Sreenath, and M. W. Mueller, "Perception-aware receding horizon trajectory planning for multicopters with visual-inertial odometry," *IEEE Access*, vol. 10, pp. 87 911–87 922, 2022. [Online]. Available: https://doi.org/10.1109/access.2022.3200342

[69] A. Taylor, T. A. Berrueta, and T. D. Murphey, "Active learning in robotics: A review of control principles," *Mechatronics*, vol. 77, p. 102576, 2021. [Online]. Available: https://doi.org/10.1016/j.mechatronics.2021.102576

[70] Z. Wang, H. Chen, S. Zhang, and Y. Lou, "Active view planning for visual slam in outdoor environments based on continuous information modeling," *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 1, pp. 237–248, 2023. [Online]. Available: https://doi.org/10.1109/tmech.2023.3272910

[71] H. Deguchi, S. Taguchi, and N. Hirose, "Active visual localization using predicted images," in *Annual Conference of the Society of Instrument and Control Engineers of Japan (SICE)*, 2020, pp. 288–293. [Online]. Available: https://doi.org/10.23919/sice48898.2020.9240354

[72] A. Parsi, A. Iannelli, and R. S. Smith, "An explicit dual control approach for constrained reference tracking of uncertain linear systems," *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2652–2666, 2022. [Online]. Available: https://doi.org/10.1109/tac.2022.3176800

[73] J. Chesser, T. Sathyan, and D. C. Ranasinghe, "Mexgen: An effective and efficient information gain approximation for information gathering path planning," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6472–6479, 2024. [Online]. Available: https://doi.org/10.1109/lra.2024.3405388

[74] J. Li, B. Sun, L. D. Giammarino, H. Blum, and M. Pollefeys, "Actloc: Learning to localize on the move via active viewpoint selection," in *Conference on Robot Learning*, 2025. [Online]. Available: https://doi.org/10.48550/arxiv.2508.20981

[75] P. Roth, J. Frey, C. Cadena, and M. Hutter, "Learned perceptive forward dynamics model for safe and platform-aware robotic navigation," *Robotics: Science and Systems*, 2025. [Online]. Available: https://doi.org/10.15607/rss.2025.xxi.001

[76] N. Yang, Y. Huang, K. Peng, Z. He, K. Wang, C. Miao, K. Lyu, G. Li, X. Wang, Z. Zhu, J. Liu, and N. Liu, "Wam-nav: Asymmetric latent world-action modeling for unified visual navigation," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2606.04907

[77] E. Arcari, L. Hewing, M. Schlichting, and M. N. Zeilinger, "Dual stochastic mpc for systems with parametric and structural uncertainty," *Learning for Dynamics and Control*, pp. 894–903, 2019. [Online]. Available: https://doi.org/10.3929/ethz-b-000454951

[78] S. He, Y. Tao, I. Spasojevic, V. Kumar, and P. Chaudhari, "An active perception game for robust information gathering," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2025, pp. 14 168–14 174. [Online]. Available: https://doi.org/10.1109/icra55743.2025.11128798

[79] K. B. Naveed, M. Singh, D. R. Agrawal, and D. Panagou, "Trajectory planning for safe dual control with active exploration," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2604.15507

[80] A. M. Khass, A. Cosse, V. Pandey, and N. Motee, "Conflict-aware active perception and control in 3d gaussian splatting fields via control barrier functions," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2605.20566

[81] I. G. Velentzas and K. Tomita, "Lidia: Localizing in the dark with illumination-awareness toward perception-aware planning," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2026, pp. 13 938–13 945. [Online]. Available: https://doi.org/10.1109/icra57385.2026.11696975

[82] T. M. J. T. Baltussen, W. Heemels, and A. Katriniok, "Dual mpc for active learning of nonparametric uncertainties," *European Journal of Control*, p. 101577, 2026. [Online]. Available: https://doi.org/10.1016/j.ejcon.2026.101577

[83] K. B. Naveed, D. R. Agrawal, and D. Panagou, "A formal gatekeeper framework for safe dual control with active exploration," *arXiv*, 2025. [Online]. Available: https://doi.org/10.48550/arxiv.2510.06351

[84] E. S. Hu, J. Wang, X. Yuan, F. Luo, M. Li, G. Lambrechts, O. Rybkin, and D. Jayaraman, "Real-world reinforcement learning of active perception behaviors," *neural information processing systems*, pp. 129 994–130 030, 2025. [Online]. Available: https://doi.org/10.52202/085713-3910

[85] A. Zhitnikov and V. Indelman, "Simplified continuous high-dimensional belief space planning with adaptive probabilistic belief-dependent constraints," *IEEE Transactions on Robotics*, vol. 40, pp. 1684–1705, 2023. [Online]. Available: https://doi.org/10.1109/tro.2023.3341625

[86] C. Zhu, M. Meurer, and C. G. Günther, "Integrity of visual navigation—developments, challenges, and prospects," *NAVIGATION Journal of the Institute of Navigation*, vol. 69, no. 2, p. navi.518, 2022. [Online]. Available: https://doi.org/10.33012/navi.518

[87] S. Pohland and C. J. Tomlin, "Competency-aware planning for probabilistically safe navigation under perception uncertainty," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 15 291–15 298. [Online]. Available: https://doi.org/10.1109/iros60139.2025.11246073

[88] J. Lekeufack, A. N. Angelopoulos, A. V. Bajcsy, M. I. Jordan, and J. Malik, "Conformal decision theory: Safe autonomous decisions from imperfect predictions," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2024, pp. 11 668–11 675. [Online]. Available: https://doi.org/10.1109/icra57147.2024.10610041

[89] L. Lindemann, M. Cleaveland, G. Shim, and G. J. Pappas, "Safe planning in dynamic environments using conformal prediction," *IEEE Robotics and Automation Letters*, vol. 8, no. 8, pp. 5116–5123, 2023. [Online]. Available: https://doi.org/10.1109/lra.2023.3292071

[90] B. Hang, W. Li, X. Li, R. Fan, H. Sun, and F. Shen, "Qa-msckf: A statistically adaptive measurement update method for visual–inertial odometry," *IEEE Internet of Things Journal*, vol. 13, no. 17, pp. 38 934–38 956, 2026. [Online]. Available: https://doi.org/10.1109/jiot.2026.3704077

[91] C. Tian, N. Hao, and F. He, "T-eskf: Transformed error-state kalman filter for consistent visual-inertial navigation," *IEEE Robotics and Automation Letters*, vol. 10, no. 2, pp. 1808–1815, 2025. [Online]. Available: https://doi.org/10.1109/lra.2024.3524905

[92] H. Wu, C. Qi, and Y. Guo, "Calibrating uncertainty induced by visual perceptual aliasing in embodied navigation," in *International Conference on Artificial Intelligence and Big Data (ICAIBD)*, 2026, pp. 855–860. [Online]. Available: https://doi.org/10.1109/icaibd69640.2026.11637152

[93] T. Faraci, M. Camurri, D. Fontanelli, and L. Palopoli, "Unseen: Uncertainty-aware navigation via sparse estimation in unknown environments," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2606.20755

[94] J. Shin, J. Lee, and I. Yang, "Egocentric conformal prediction for safe and efficient navigation in dynamic cluttered environments," in *IEEE Conference on Decision and Control (CDC)*, 2025, pp. 4910–4917. [Online]. Available: https://doi.org/10.1109/cdc57313.2025.11312645

[95] C. Yang, A. Vadlamani, A. Soloviev, M. J. Veth, and C. N. Taylor, "Feature matching error analysis and modeling for consistent estimation in vision-aided navigation," *NAVIGATION Journal of the Institute of Navigation*, vol. 65, no. 4, pp. 609–628, 2018. [Online]. Available: https://doi.org/10.1002/navi.265

[96] M. Li and A. I. Mourikis, "High-precision, consistent ekf-based visual-inertial odometry," *The International Journal of Robotics Research*, vol. 32, no. 6, pp. 690–711, 2013. [Online]. Available: https://doi.org/10.1177/0278364913481251

[97] A. N. Ikhsan, R. Mardiyanto, K.-Y. Wong, and M. N. Cahyadi, "Perceptual aliasing and uncertainty in repetitive environments for monocular uav navigation using perception-aware reinforcement learning," *IEEE Access*, vol. 14, pp. 159 943–159 978, 2026. [Online]. Available: https://doi.org/10.1109/access.2026.3743906

[98] J. Dong, S. Wu, and Y. Fang, "Differentiable optimization layered safety-critical control for risk-aware navigation via conformal prediction," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2605.16327

[99] B. Du, J. Kim, and Y. Lyu, "Gaussian mixture-based inverse perception contract for uncertainty-aware robot navigation," *arXiv*, 2026. [Online]. Available: https://doi.org/10.48550/arxiv.2603.04329

[100] G. P. Huang, M. Kaess, and J. J. Leonard, "Towards consistent visual-inertial navigation," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2014, pp. 4926–4933. [Online]. Available: https://doi.org/10.1109/icra.2014.6907581

[101] E. Stracca, O. Napolitano, L. Pallottino, and P. Salaris, "A unified local risk map for uncertainty-aware mobile robot navigation in cluttered and dynamic environments," *Sensors*, vol. 26, no. 12, p. 3900, 2026. [Online]. Available: https://doi.org/10.3390/s26123900

[102] D. Kumar, S. Tayebati, F. Migliarba, R. Krishnan, and A. R. Trivedi, "Learnable conformal prediction for safe and efficient robotics under perception and planning uncertainties," in *Proceedings - IEEE International Conference on Robotics and Automation/Proceedings*, 2026, pp. 14 129–14 136. [Online]. Available: https://doi.org/10.1109/icra57385.2026.11697249

[103] V. Peretroukhin, L. Clement, M. Giamou, and J. Kelly, "Probe: Predictive robust estimation for visual-inertial navigation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 3668–3675. [Online]. Available: https://doi.org/10.1109/iros.2015.7353890

[104] C. Tian, F. He, and N. Hao, "Unobservable subspace evolution and alignment for consistent visual-inertial navigation," *arXiv*, 2025. [Online]. Available: https://doi.org/10.48550/arxiv.2511.17992

[105] G. Puthumanaillam, P. Padrão, J. Fuentes, L. Bobadilla, and M. Ornik, "Guided agents: Enhancing navigation policies through task-specific uncertainty abstraction in localization-limited environments," *IEEE Robotics and Automation Letters*, vol. 11, no. 3, pp. 2999–3006, 2026. [Online]. Available: https://doi.org/10.1109/lra.2026.3655198