

UAV Celestial Navigation with Light Pollution Adaptation

Master of Science Thesis

Cognitive Robotics

Janvi Seth

Student ID: 4645987



UAV Celestial Navigation with Light Pollution Adaptation

by Janvi Seth

to obtain the degree of Master of Science at the Delft University of Technology to be defended publicly on December 22nd, 2025 at 10:30

Student ID: 4645987

Supervisors: Dr. L. Ferranti (TU Delft)
Dr. O. de Groot (Intellic)

Cover Image Credit: Cash Macanaya on Unsplash
<https://unsplash.com/photos/a-view-of-the-earth-from-space-at-night-DF6xU6kXqlI>

Friday 19th December, 2025

Contents

1	Paper	3
1	Introduction	3
i	Contributions	4
ii	Report Outline	5
2	Preliminaries	5
i	Forward Celestial Model	5
3	Problem Formulation	6
i	Observation Model and Star Hierarchy	6
ii	The Inverse Problem: Celestial Localization	6
4	Methodology	8
i	Star Catalog Preprocessing	8
ii	Synthetic Dataset Generation	9
iii	Star Detection	10
iv	Pattern Matching with Adaptive Catalog Selection	10
v	Geographic Position Estimation	10
5	Results	10
i	Dataset Overview and Test Conditions	10
ii	Star Identification and Position Estimation Performance	11
iii	Sensitivity to Environmental Conditions	12
6	Discussion	14
i	The Necessity of Environmental Adaptation	14
ii	Interpreting the Ablation Study: Component Contribution	15
iii	Calibration, Projection Models, and Accuracy Limits	15
iv	Operational Context and Mission Integration	16
v	Limitations and Validity Constraints	16
7	Conclusion	17
i	Future Work	17

Acronyms

CNS Celestial Navigation Systems
Dec Declination
DDR Dynamic Distance-Ratio
ENU East-North-Up
FOV Field of View
GNSS Global Navigation Satellite System
HUD Heads-Up Display
IMU Inertial Measurement Unit
LUT Look-Up Table
PR Polar-star Rejection
RA Right Ascension
RF Radio Frequency
SI Similarity-based Iteration
TRN Terrain-Referenced Navigation
UAV Uncrewed Aerial Vehicle
UTC Coordinated Universal Time
VPR Visual Place Recognition
YBSC Yale Bright Star Catalog

UAV Celestial Navigation with Automatic Light Pollution Adaptation

Janvi Seth 4645987
Technical University of Delft

Abstract— *Uncrewed Aerial Vehicles (UAVs) increasingly require Global Navigation Satellite System (GNSS)-independent positioning for operation in contested or infrastructure-denied environments. This paper presents a vision-based celestial navigation system with automatic adaptation to light pollution through dynamic star catalog selection. The algorithm employs Dynamic Distance-Ratio (DDR) pattern matching with novel polar-star rejection and consensus-driven magnitude refinement to robustly identify observable stars under varying environmental conditions. Evaluation on 200 synthetic night-sky images demonstrates substantially improved star identification robustness compared to fixed-catalog baselines, achieving 71.5% recall at visual magnitude 7 (Bortle 3) and maintaining non-zero performance under severe light pollution (27.6% recall at magnitude 5.0 and 4.5% at magnitude 4.5), where the baseline fails entirely. Across higher limiting magnitudes (6.5–8.0), the adaptive method consistently attains 71.5–82.5% recall. Including misidentifications, the end-to-end system achieves a median geolocation error of 6.80 km, supporting coarse global localization, GNSS integrity monitoring, and long-duration drift bounding in GNSS-denied environments. These results indicate that adaptive catalog selection significantly extends the operational envelope of celestial navigation into light-polluted conditions previously considered infeasible.*

Keywords— Celestial Navigation, UAV, GNSS-denied, Star Identification, Pattern Matching, Light Pollution, Dynamic Distance-Ratio

1. Introduction

UAVs are increasingly deployed in mission-critical applications including defense operations, environmental monitoring, and industrial inspection, where reliable positioning is fundamental to operational success [1]. While GNSS provides worldwide coverage with meter-level accuracy, its dependence on weak Radio Frequency (RF) signals renders it vulnerable to both intentional and unintentional interference. Recent European assessments document a pronounced increase in jamming and spoofing events, with measurable operational impacts on aviation systems and flight efficiency [2]. These developments underscore an operational gap: UAVs require positioning solutions that maintain functionality in GNSS-degraded or denied environments.

Multiple GNSS-independent approaches have emerged to address this challenge, each exploiting different environmental features or infrastructure. Visual Place Recognition (VPR) leverages deep neural networks to match onboard imagery against geo-referenced databases, achieving decimeter-scale accuracy in densely mapped urban environments [3]. However, VPR performance degrades substantially under large viewpoint changes, seasonal variations, or operations in unmapped regions. These limitations constrain its use to well-characterized areas with up-to-date reference data. RF-beacon networks provide decimeter- to centimeter-level positioning through trilateration from known anchor positions, but deployment requires extensive infrastructure installation, calibration, and maintenance of line-of-sight geometry [4]. Terrain-Referenced Navigation (TRN) correlates onboard sensor measurements with digital elevation models to estimate position, with modern vision-based implementations achieving sub-hundred-meter accuracy over textured terrain at appropriate altitudes [5] [6]. Yet TRN exhibits sensitivity to terrain characteristics, requiring sufficient geometric variation and feature density for reliable correlation.



Figure 1: Example of UAV night operations where celestial navigation is applicable. Image illustrates typical low-light conditions under which vision-based Celestial Navigation Systems (CNS) must operate [7].

While these methods demonstrate impressive performance within their operational envelopes, each relies on specific environmental preconditions. VPR requires comprehensive prior mapping, RF-beacons demand deployed infrastructure, and TRN depends on terrain observability and texture. CNS offers a conceptually independent alternative that remains globally available without active emissions, ground infrastructure, or

prior mapping. Unlike terrain- or infrastructure-based techniques, celestial positioning exploits the predictable geometry of stellar positions, a resource that is passive, universally accessible, and invariant to terrestrial conditions. For UAV operations requiring kilometer-scale waypoint navigation, integrity monitoring of GNSS solutions, or coarse reacquisition before GNSS handover, celestial navigation provides a complementary capability that functions independently of infrastructure availability or environmental mapping. Moreover, celestial methods enable night operations (Figure 1) in scenarios where daylight-dependent techniques like VPR and TRN become unavailable or severely degraded, expanding operational flexibility in contested or infrastructure-denied environments.

Recent work has demonstrated the feasibility of strap-down vision-based CNS on low-cost UAV platforms [8], although the reported position errors remain at the kilometer scale under realistic conditions, substantially larger than alternative methods operating within their preferred regimes. The combination of comparatively coarse accuracy and a clear geometric foundation makes CNS a compelling target for algorithmic improvement. While stars are not always visible due to factors such as the day-night cycle and cloud cover, this paper demonstrates that the primary challenge in vision-based celestial navigation lies in the robustness of star identification under varying environmental conditions. Through our findings, we show that the true bottleneck in reliable navigation arises from the difficulty in consistently identifying stars amidst dynamic and degraded environmental factors, rather than from the geometric visibility of stars themselves.

The Star Identification Challenge In small fields of view and under environmental degradation such as light pollution or cloud cover, the stars actually visible in a captured image form a strict and unpredictable subset of those geometrically predicted to lie within the camera footprint. A UAV operating over a dark rural area may observe 2,000+ stars to magnitude 7.5, while the same platform over an urban center detects fewer than 30 bright stars above magnitude 4.0 [9]. Motion blur, atmospheric extinction, and sensor noise further confound pattern formation. Traditional star identification algorithms assume a fixed catalog size, leading to catastrophic failure when the assumed catalog mismatches the actual scene [10]. Sparse catalogs miss faint stars needed for pattern uniqueness, while dense catalogs introduce false match candidates that overwhelm voting schemes. Since identification errors propagate directly to position error [11], robust performance demands adaptive catalog selection that automatically adjusts to observed stellar density.

This thesis addresses the star identification robustness

challenge by developing an adaptive pattern-matching framework that dynamically selects magnitude-limited star catalogs based on environmental observability. The approach builds upon the DDR matching principle [12], which encodes rotation-invariant geometric features, inter-star angular distances and relative bearing angles, around reference stars and retrieves candidates through voting over a precomputed catalog index. While Dai et al. [12] demonstrate DDR’s effectiveness for star identification in controlled scenarios, their work does not address environmental adaptivity or provide a complete navigation pipeline. This thesis implements the full celestial localization system from first principles, including catalog preprocessing, synthetic sky rendering, star detection, DDR-based pattern matching with novel extensions, and geometric position optimization. The implementation introduces three key innovations: (1) consensus-driven magnitude refinement that dynamically matches catalog density to scene observability, (2) Polar-star Rejection (PR) that eliminates geometrically degenerate high-declination matches, and (3) Similarity-based Iteration (SI) reference selection that systematically explores the detected star set when initial matching confidence is insufficient.

i. Contributions

The primary contributions of this thesis are summarized as follows:

- **Adaptive catalog selection mechanism:** A novel consensus-driven magnitude-refinement mechanism is introduced to robustify star matching under light-polluted or low-visibility conditions. By estimating the effective limiting magnitude from the observed stellar density, the algorithm dynamically selects an appropriate magnitude-limited catalog. This allows DDR matching to operate only on stars that are realistically observable in the scene, substantially improving identification robustness in the presence of light pollution and atmospheric attenuation.
- **Polar-star rejection strategy:** Implementation of a geometric filtering technique that identifies and rejects high-declination stars exhibiting degenerate position constraints, reducing systematic localization bias.
- **Similarity-based iterative matching:** Development of an adaptive reference-star selection algorithm that iteratively explores the detected star set to improve identification robustness under sparse visibility conditions.
- **Comprehensive evaluation framework:** Generation of synthetic datasets using both photorealistic rendering (Stellarium) and noise-free geometric projection (Yale catalog), enabling controlled vali-

dation across varying environmental conditions.

ii. Report Outline

The remainder of this report is structured as follows. Section 2 introduces the preliminaries of celestial coordinate systems and the forward celestial model, establishing the mathematical foundation for position estimation. Section 3 formalizes the celestial localization problem, introducing the observation model, star catalog hierarchy, and the inverse mapping from observed star patterns to geographic position. Section 4 details the proposed methodology, including catalog preprocessing with magnitude stratification, synthetic dataset generation using dual-source validation, image-based star detection, DDR pattern matching with adaptive catalog selection, and nonlinear position estimation through Levenberg-Marquardt optimization. Section 5 presents experimental results on synthetic Stellarium imagery, including identification recall metrics, position accuracy statistics, ablation studies quantifying individual component contributions, and sensitivity analyses across varying light pollution conditions. Section 6 discusses the interpretation of results, systematic limitations including computational complexity and rendering biases, and operational considerations for field deployment. Section 7 concludes the report and outlines directions for future work, including real-world validation, computational optimization, and sensor fusion with inertial measurements.

2. Preliminaries

Before formalizing the mathematical framework, it is instructive to consider the fundamental principle underlying celestial navigation. When an observer stands at a specific location on Earth at a particular time, they see a unique configuration of stars in the sky above them. This configuration is determined by three factors: the observer’s latitude, their longitude, and the time of observation [13]. Conversely, if one can identify which stars are visible and measure their positions in the sky, it becomes possible to work backwards to determine one’s location on Earth. This inverse problem forms the basis of celestial navigation and has guided mariners for centuries [14].

The challenge lies in bridging two fundamentally different coordinate systems. Catalogued stars are defined in the equatorial frame by Right Ascension (RA) and Declination (Dec), which is fixed relative to Earth’s rotation axis and, apart from slow effects such as precession and proper motion, remains effectively constant over human timescales [15]. Observers, however, perceive the sky in the horizontal coordinate system, where positions are described relative to the local horizon and cardinal directions (azimuth and altitude). As Earth rotates, encoded locally by the **lst!** (**lst!**), equatorial co-

ordinates map to the local horizon, so a time-stamped star image provides strong constraints on the observer’s geodetic latitude and longitude. The transformation between these systems directly encodes both geographic position and observation time.

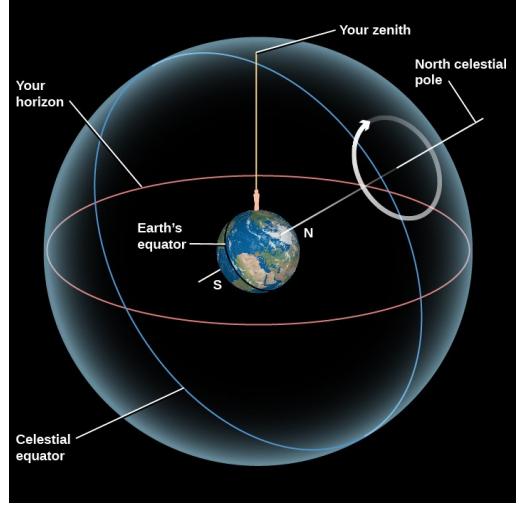


Figure 2: Illustration of the celestial sphere from Fraknoi et al. [16].

Stars are treated as fixed on a sphere surrounding Earth. The apparent rotation of the sky around the pole arises from Earth’s own rotation.

i. Forward Celestial Model

Let $\mathcal{S}_{\text{cat}} = \{s_i\}_{i=1}^{N_{\text{cat}}}$ denote the Star Catalog used for celestial localization. Each element $s_i \in \mathcal{S}_{\text{cat}}$ represents a catalogued star defined by its equatorial coordinates and apparent magnitude,

$$s_i = (\alpha_i, \delta_i, m_i),$$

where $\alpha_i \in [0, 2\pi]$ is the right ascension, $\delta_i \in [-\frac{\pi}{2}, \frac{\pi}{2}]$ is the declination, and $m_i \leq 6.5$ denotes the visual magnitude range covered by the Yale Bright Star Catalog (YBSC).

For an observer located at geodetic coordinates (ϕ, λ) and observation time t , the apparent position of each star s_i in the local horizon system is described by its altitude a_i and azimuth A_i , with

$$(A_i, a_i) \in [0, 2\pi) \times [-\pi/2, \pi/2].$$

The transformation from equatorial to horizontal coordinates is given by the following relations [17]:

$$\sin a_i = \sin \phi \sin \delta_i + \cos \phi \cos \delta_i \cos H_i \quad (2.1)$$

$$\tan A_i = \frac{-\sin H_i \cos \delta_i}{\sin \delta_i \cos \phi - \cos \delta_i \sin \phi \cos H_i} \quad (2.2)$$

where H_i is the hour angle of the star,

$$H_i = \text{LST}(t, \lambda) - \alpha_i, \quad H_i \in (-\pi, \pi] \quad (2.3)$$

and $\text{LST}(t, \lambda)$ denotes the local sidereal time, which depends on the observer's longitude and the Earth's rotation at time t .

Figure 2 illustrates the geometric relationship encoded in these equations. Stars are treated as fixed on the celestial sphere surrounding Earth, with their positions described by equatorial coordinates (RA α , Dec δ). As Earth rotates, the celestial sphere appears to revolve around the celestial pole, causing each star's hour angle H to increase continuously with time (Equation 2.3). The transformation from equatorial to horizontal coordinates (Equations 2.1-2.2) projects this rotating celestial configuration onto the observer's local horizon, where altitude a represents elevation above the horizon and azimuth A indicates direction relative to true north. This apparent motion—the rising and setting of stars—directly encodes both the observer's geographic position and the time of observation.

Equations 2.1–2.3 together define the **forward celestial model** g , which maps the observer's geographic position and observation time to the set of apparent star positions in the local horizon system:

$$g : (\phi, \lambda, t) \mapsto \{(A_i, a_i) \mid s_i \in \mathcal{S}_{\text{cat}}\} \quad (2.4)$$

The conversion from horizontal coordinates (A, a) to a unit line-of-sight vector \hat{v}_i is obtained through the standard spherical-to-Cartesian transformation used in positional astronomy [18]:

$$\hat{v} = \begin{bmatrix} \cos a \cdot \sin A \\ \cos a \cdot \cos A \\ \sin a \end{bmatrix} \quad (2.5)$$

These unit vectors \hat{v}_i represent the same star directions on the unit sphere and are later used for geometric localization.

3. Problem Formulation

In the absence of GNSS signals, the objective of this work is to estimate the geodetic coordinates (latitude ϕ and longitude λ) of an observer UAV from a single image of the night sky acquired at a known Coordinated Universal Time (UTC) time t . Building on the forward celestial model introduced in the preliminaries, this section formalizes the corresponding inverse problem: determining the observer's location from the observed stellar configuration.

i. Observation Model and Star Hierarchy

A captured image $I(x, y)$ undergoes star detection to produce a set of observed centroids in pixel coordinates:

$$\mathcal{P}_{\text{obs}} = \{p_j\}_{j=1}^M = \{(u_j, v_j)\}_{j=1}^M. \quad (3.1)$$

Each detected pixel observation p_j can be projected to a unit direction vector in the camera frame using the known camera intrinsic matrix K [19]:

$$\mathbf{r}_i^{\text{obs}} = \frac{K^{-1}[u_i, v_i, 1]^T}{\|K^{-1}[u_i, v_i, 1]^T\|} \quad (3.2)$$

The fundamental challenge is that \mathcal{P}_{obs} represents only a partial, noisy view of the catalog. The relationship between observed detections and the full stellar catalog is structured by environmental and sensor constraints, formalized as a hierarchy of nested subsets:

$$\mathcal{S}_{\text{vis}}(\phi, \lambda, t; \varepsilon) \subseteq \mathcal{S}_{\text{geom}}(\phi, \lambda, t) \subseteq \mathcal{S}_{\text{cat}} \quad (3.3)$$

Here:

- \mathcal{S}_{cat} is the complete star catalog containing all cataloged stars with their equatorial coordinates and magnitudes.
- $\mathcal{S}_{\text{geom}}(\phi, \lambda, t)$ contains stars that are geometrically visible above the horizon at location (ϕ, λ) at time t : $\mathcal{S}_{\text{geom}}(\phi, \lambda, t) = \{s_i \in \mathcal{S}_{\text{cat}} \mid a_i(\phi, \lambda, t) > 0\}$
- $\mathcal{S}_{\text{vis}}(\phi, \lambda, t; \varepsilon)$ contains the subset of geometrically visible stars that are additionally bright enough to be detected under environmental conditions ε (e.g., limiting magnitude m_{lim} , atmospheric extinction, light pollution): $\mathcal{S}_{\text{vis}}(\phi, \lambda, t; \varepsilon) = \{s_i \in \mathcal{S}_{\text{geom}}(\phi, \lambda, t) \mid m_i < m_{\text{lim}}(\varepsilon)\}$

The detected set \mathcal{P}_{obs} in pixel space corresponds to an unknown subset of \mathcal{S}_{vis} in catalog space. Critically, \mathcal{P}_{obs} may be incomplete due to missed detections (star blending, clouds, detection threshold) and may contain false positives (noise or artifacts). While \mathcal{P}_{obs} is operationally represented in pixel coordinates, the catalog subsets \mathcal{S}_{cat} , $\mathcal{S}_{\text{geom}}$, and \mathcal{S}_{vis} are all defined in celestial coordinates. This distinction is fundamental: the localization problem requires bridging these two representations.

ii. The Inverse Problem: Celestial Localization

The localization task requires establishing which catalog star $\{s_i\}$ each observed detection $\{p_j\}$ corresponds to. This star identification problem can be formatted through a correspondence mapping C :

$$C : \mathcal{P}_{\text{obs}} \rightarrow \mathcal{S}_{\text{cat}} \cup \{\emptyset\} \quad (3.4)$$

where $C(p_j) = s_i$ indicates that pixel detection p_j is identified as catalog star s_i , and $C(p_j) = \emptyset$ indicates a false positive detection with no valid catalog match. A valid correspondence must satisfy geometric consistency: if

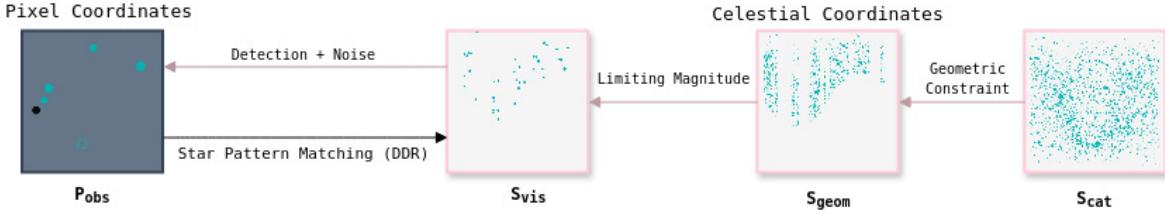


Figure 3: Hierarchy of star catalog subsets in the celestial localization problem. The complete catalog (S_{cat}) is progressively filtered by geometric visibility (S_{geom}) and environmental observability (S_{vis}). The detected set P_{obs} , represented in pixel space, corresponds to a noisy, incomplete subset of S_{vis} as described in Section 3–i.

$p_j \leftrightarrow s_i$, then the observed direction $\mathbf{r}_j^{\text{obs}}$ should align with the predicted direction $\hat{\mathbf{v}}_i(\varphi, \lambda, t)$ from the true observer location. This is achieved with the following transform constraint:

$$\mathbf{r}_j^{\text{obs}} \approx \mathbf{R}_{\text{cam}} \cdot \hat{\mathbf{v}}_i(\varphi, \lambda, t) \quad (3.5)$$

where \mathbf{R}_{cam} represents the camera's orientation relative to the local East-North-Up (ENU) frame. The correspondence problem is challenging because: (1) the sets S_{vis} and S_{geom} depend on the unknown observer coordinates (φ, λ) , creating a circular dependency; (2) only $M = |\mathcal{P}_{\text{obs}}|$ stars are detected, where $M \ll |\mathcal{S}_{\text{cat}}$, and M may be smaller than $|S_{\text{vis}}|$ due to missed detections; (3) without absolute orientation or position information, multiple catalog configurations may appear geometrically similar; and (4) detection errors, atmospheric effects, and false positives degrade correspondence quality.

The correspondence problem can be addressed by exploiting geometric invariants rather than absolute celestial positions. The approach recognizes that certain geometric relationships between stars are preserved regardless of the observer's unknown location or camera orientation [20] [21]. Specifically, for any pair of stars s_i and s_k in the catalog, their angular separation depends only on their intrinsic equatorial coordinates, not on (φ, λ) . Similarly, the relative bearing angles between stars form rotation-invariant patterns. Given a reference detection $p_r \in \mathcal{P}_{\text{obs}}$ and its neighboring detections $\{p_k\}$ within the field of view, one can compute observable geometric features. Consistent with the neighbourhood-based pattern-matching paradigm described in Spratling and Mortari [11], we formalise the observed and catalog star patterns using local geometric feature sets:

$$\mathcal{F}_{\text{obs}}(p_r) = \{(\theta_{rk}, \beta_{rk}) \mid p_k \in \text{neighbors}(p_r)\} \quad (3.6)$$

where θ_{rk} is the angular distance from p_r to p_k and β_{rk} is the relative bearing angle. For each catalog star

$s_i \in \mathcal{S}_{\text{cat}}$, an analogous pattern can be precomputed from catalog coordinates:

$$\mathcal{F}_{\text{cat}}(s_i) = \{(\theta'_{ik}, \beta'_{ik}) \mid s_k \in \text{neighbors}(s_i)\} \quad (3.7)$$

The star identification problem then reduces to finding the catalog star s^* whose geometric pattern best matches the observed pattern:

$$s^* = \arg \max_{s_i \in \mathcal{S}_{\text{cat}}} \text{similarity}(\mathcal{F}_{\text{obs}}(p_r), \mathcal{F}_{\text{cat}}(s_i)) \quad (3.8)$$

where the similarity metric quantifies geometric agreement between patterns, typically through voting schemes that count matching feature pairs within tolerance thresholds. This pattern-based approach enables correspondence establishment without prior knowledge of (φ, λ) , though it assumes sufficient stellar density and distinct geometric configurations to avoid ambiguities.

Once a set of correspondences $\{(p_j, s_i)\}$ is established via pattern matching, the observer location can be estimated by finding (φ, λ) that minimizes the angular discrepancy between observed and predicted star directions [5,17]. The overall celestial localization task can be summarized as the mapping:

$$f : (I(x, y), t) \mapsto (\hat{\varphi}, \hat{\lambda}) \quad (3.9)$$

where $f(\cdot)$ represents the full computational pipeline that includes: (1) star detection from the image, (2) identification of catalog correspondences, and (3) geometric optimization to estimate position. For a hypothesized observer position (φ, λ) and known time t , the predicted star directions $\hat{\mathbf{r}}_i(\varphi, \lambda, t)$ are obtained using the forward model. The goal is to find the observer coordinates that minimize the mean angular error between the observed and predicted directions:

$$(\hat{\varphi}, \hat{\lambda}) = \arg \min_{\varphi, \lambda} \frac{1}{N} \sum_{i=1}^N \rho \left(\arccos \left(\hat{\mathbf{v}}_i(\varphi, \lambda, t) \cdot \mathbf{r}_i^{\text{obs}} \right) \right) \quad (3.10)$$

where N is the number of matched star correspondences used in the optimization, where $\rho(\cdot)$ is a robust loss function that reduces the effect of mismatched or noisy detections. Longitude λ primarily governs the rotational alignment of the observed sky through the sidereal-time term, while latitude φ determines the altitude distribution of the celestial sphere. Together, these parameters uniquely define the observer's position on Earth. The optimization is inherently nonlinear due to the trigonometric relationships in Equations (2.1)–(2.3), and non-convex with potential local minima, particularly at high latitudes where the celestial pole geometry becomes degenerate.

The celestial localization problem can be formally stated as follows. Given a single image $I(x, y)$ of the night sky, observation time t (UTC), camera intrinsic parameters \mathbf{K} , camera orientation \mathbf{R}_{cam} (known or approximately zenith-pointing), and star catalog $\mathcal{S}_{\text{cat}} = \{(\alpha_i, \delta_i, m_i)\}$, the objective is to determine observer geodetic coordinates $(\hat{\varphi}, \hat{\lambda})$ such that: (1) a geometrically consistent correspondence $C : \mathcal{P}_{\text{obs}} \rightarrow \mathcal{S}_{\text{cat}}$ is established via star pattern matching, (2) the predicted star directions from $(\hat{\varphi}, \hat{\lambda})$ align with observed directions according to the optimization criterion in Equation 3.10, and (3) the solution is robust to detection noise, correspondence errors, and environmental uncertainty. This problem decomposes into two coupled subproblems: the correspondence problem (identifying which catalog stars are observed via geometric pattern matching) and the estimation problem (computing the location that explains those observations through nonlinear optimization). Known quantities include observation time t (from UAV clock), camera calibration matrix \mathbf{K} , and detected star centroids \mathcal{P}_{obs} . Unknown quantities are observer coordinates (φ, λ) [primary objective], environmental conditions ε , and true correspondences C . The methodology in Section 4 details the algorithms used to jointly solve both components of this coupled inverse problem.

4. Methodology

This section describes the complete computational pipeline for estimating the observer's geographic position from a single nighttime image. The approach addresses the fundamental challenge that the number and brightness of visible stars varies unpredictably with environmental conditions. A city sky may reveal only ~ 365 bright stars while a suburban location shows ~ 2110 or more [9]. Traditional star identification algorithms assume a fixed catalog, leading to catastrophic failure when the assumed catalog mismatches the actual scene [10]. The proposed adaptive methodology dynamically selects the appropriate magnitude-limited catalog based on observed stellar density, then employs geometric pattern matching to identify stars without requiring absolute orientation or position knowledge.

The pipeline consists of catalog preprocessing, synthetic data generation, image-based star detection, DDR-based pattern matching, and geometric position estimation. Figure 4 illustrates the overall data flow in the implemented system. Importantly, the entire pipeline was implemented entirely from scratch, as no existing codebase or datasets were available for this task.

i. Star Catalog Preprocessing

The YBSC serves as the celestial reference, containing 9,110 stars with equatorial coordinates (RA, Dec) and visual magnitudes [22]. To enable efficient pattern matching under varying observability conditions, the catalog undergoes magnitude-stratified preprocessing.

To emulate different sky-visibility regimes, the catalog is stratified by visual magnitude thresholds that reflect the environmental observability under varying light pollution levels. Twenty-six discrete catalog subsets are generated corresponding to visual magnitude thresholds $m_{\text{vis}} \in \{2.0, 2.1, 2.2, \dots, 4.4, 4.5\}$ in increments of 0.1 magnitude.

Pattern Database Construction: For each catalog subset, a pattern database is precomputed by selecting each star as a potential reference and encoding its neighborhood geometry. Given a reference star s_r and radius Field of View (FOV), neighboring stars within this angular distance are identified using great-circle separation [23]:

$$\theta_{rk} = \arccos(\sin \delta_r \sin \delta_k + \cos \delta_r \cos \delta_k \cos(\alpha_r - \alpha_k)) \quad (4.1)$$

For each reference-neighbor pair, two features are computed: (1) angular distance θ_{rk} and (2) distance ratio $\rho_k = \theta_{rk}/\theta_{\text{max}}$, where θ_{max} is the maximum neighbor distance. Additionally, relative bearing angles between neighbors are encoded to form a complete geometric descriptor.

Look-Up Table (LUT) Construction: To enable rapid candidate retrieval during online matching, the pattern database is indexed into a discretized lookup table. The continuous feature space (θ, ρ) is quantized using bin sizes $\Delta_\theta = 0.5^\circ$ and $\Delta_\rho = 0.02$, derived from empirical analysis of detection noise characteristics. Each LUT entry maps to a list of catalog star identifiers whose patterns contain features within that bin, enabling voting-based matching with $O(1)$ retrieval complexity per feature.

Density Statistics: For each magnitude-limited catalog, the mean neighbor count within r_{FOV} is computed and stored. These statistics enable automatic LUT selection during matching by comparing detected star density to precomputed catalog densities (Section III-iv).

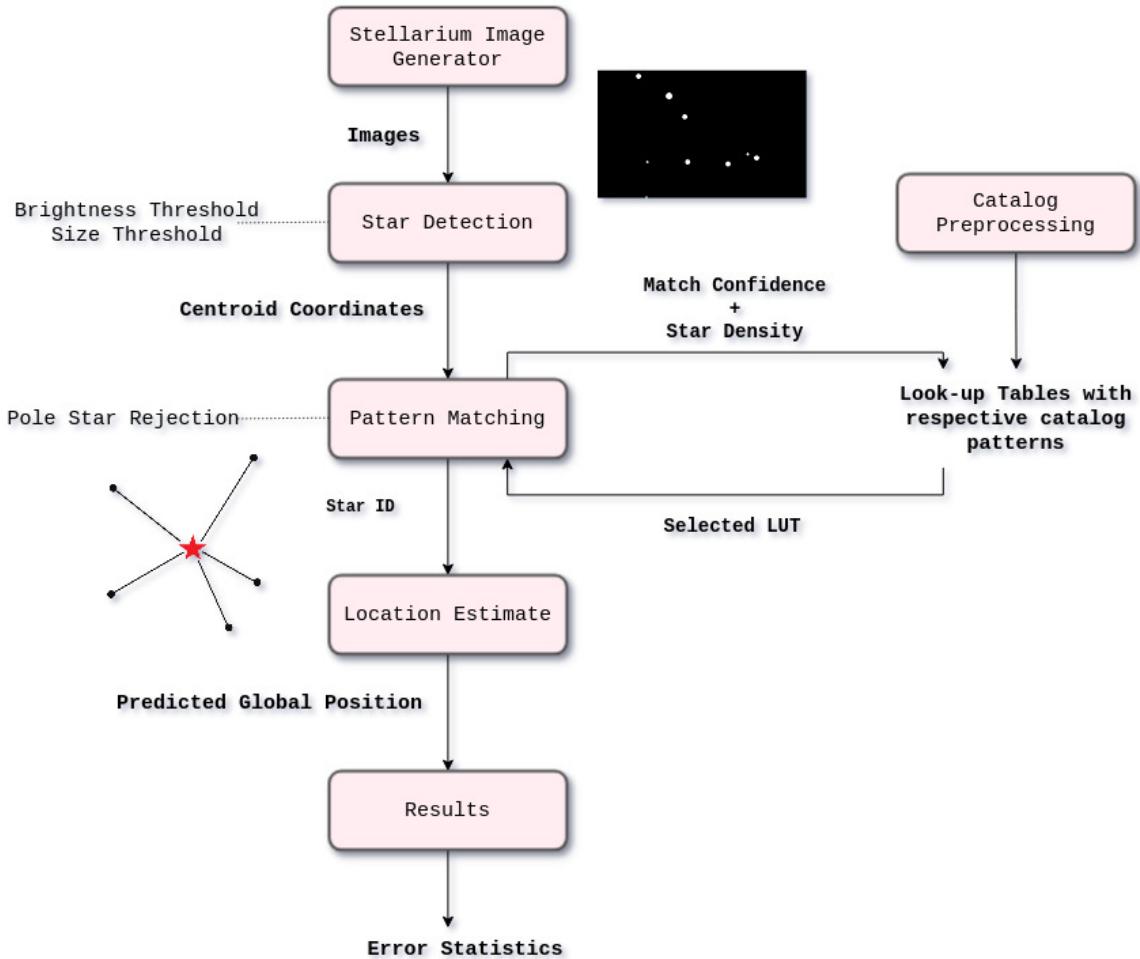


Figure 4: Overview of the data flow and processing pipeline.

ii. Synthetic Dataset Generation

Two complementary datasets were constructed programmatically: a photorealistic Stellarium dataset and a noise-free Yale-projection dataset.

1. Stellarium Dataset

Using the Stellarium planetarium engine, sky images are rendered for randomly sampled global locations and UTC timestamps. The renderer incorporates atmospheric extinction, sky brightness, and image artifacts, producing highly realistic nighttime conditions. Image generation is automated via Stellarium's HTTP remote control API.

Each rendered image is paired with precise metadata, including:

- latitude and longitude,
- timestamp (UTC),
- camera orientation,
- optical parameters (FOV and image resolution).

This dataset simulates real-world UAV imaging conditions, including imperfections introduced by projection, atmospheric modeling, and rendering. Notably, subsequent analysis revealed systematic biases in Stellarium's camera projection model (Section 6-iii), motivating the parallel Yale dataset for controlled validation.

2. Yale-Projection Dataset

To isolate algorithmic performance from rendering noise, YBSC stars are projected directly into pixel coordinates using:

- equatorial → horizontal coordinate conversion,
- ENU transformation based on the metadata,
- the pinhole camera projection model.

Only stars above the horizon are retained. The resulting skymaps provide exact, noise-free ground truth for star detection and DDR matching, enabling controlled evaluation and fine-grained analysis of each pipeline component.

iii. Star Detection

Given an image $I(x, y)$, star detection extracts pixel centroids: $\mathcal{P}_{\text{obs}} = \{(u_j, v_j)\}_{j=1}^M$. Images are converted to grayscale, Heads-Up Display (HUD) overlays are masked, and a flood-fill algorithm identifies connected bright components above a brightness and area threshold. Each connected component is reduced to a centroid, representing a detected star.

iv. Pattern Matching with Adaptive Catalog Selection

Traditional CNS pipelines assume a fixed magnitude-limited catalog [24]. However, the number of visible stars in UAV imagery varies sharply with light pollution, atmospheric transparency, exposure time, and noise [25] [26]. To handle this, the present work introduces a three-stage adaptive LUT selection mechanism:

1. Density-based LUT estimation,
2. Similarity-based probing across adjacent LUTs,
3. Consensus-driven refinement.

This forms one of the thesis's primary technical contributions.

1. Density-Based LUT Estimation

The observed stellar density is estimated by counting neighbors within the field-of-view radius in pixel space. This density is compared with each catalog subset's precomputed density. The LUT whose density most closely matches the observed density is chosen as the initial hypothesis.

2. Similarity-Based LUT Probing

Next, DDR matching is performed against:

- the density-suggested LUT,
- its brighter and dimmer neighboring LUTs,
- and the current consensus-preferred LUT, if any.

For each LUT, the DDR matching routine computes a normalized *similarity score*, quantifying the geometric agreement between the observed pattern and catalog patterns. The LUT yielding the highest similarity score becomes the provisional selection.

3. Consensus-Driven Refinement

A running history stores magnitude choices and similarity scores from successful matches. Once enough high-similarity matches accumulate, a stable consensus magnitude limit emerges. If the consensus LUT achieves similarity comparable to the provisional best, it overrides the selection. When the consensus becomes highly stable (exceeding 80%), the algorithm enters *fast-consensus mode*, skipping LUT probing entirely.

This process yields a robust, visibility-adaptive catalog selection mechanism tailored to the observed star density in each frame.

v. Geographic Position Estimation

Once a catalog correspondence has been established, the identified star provides a geometric constraint linking the observer's geodetic coordinates (ϕ, λ) to the measured image-ray direction. The observed pixel location is converted to a unit line-of-sight vector in the camera frame using the calibrated intrinsic matrix (3.2). This vector is then rotated into the local ENU frame (3.5) using the camera attitude provided by Stellarium metadata.

For the matched catalog star, its catalog coordinates (α, δ) are propagated from the catalog epoch to the time of observation using the full chain of astrometric corrections implemented in the solver: proper motion, precession, nutation, and annual aberration [27]. The corrected coordinates (α', δ') are then mapped to apparent azimuth–elevation through the forward celestial model introduced in Section 2.

The location estimate is obtained by solving for the observer coordinates (ϕ, λ) that minimise the angular discrepancy between the predicted and observed star directions as equation 3.10 states. A two-stage optimisation is employed. A coarse grid search over latitude and longitude provides an initial estimate, followed by iterative refinement using a Levenberg–Marquardt [28] update. The solver jointly minimises errors in azimuth and elevation and incorporates atmospheric refraction (Bennett model [29]), ensuring consistency between the predicted apparent elevation and the measured ray direction. The final output is the geodetic position $(\hat{\phi}, \hat{\lambda})$ that yields maximal agreement with the observed stellar geometry.

5. Results

i. Dataset Overview and Test Conditions

The proposed adaptive celestial navigation algorithm was evaluated on synthetic datasets of 200 star field images generated using Stellarium planetarium software (Version 24.4) [30]. The dataset was designed to simulate diverse observing conditions encountered in real-world UAV operations, with systematic variation across geographic, temporal, and environmental parameters.

- **Geographic Coverage:** Test locations were randomly sampled from a uniform distribution spanning latitudes $\phi \in [-90^\circ, +90^\circ]$ and longitudes $\lambda \in [-180^\circ, +180^\circ]$, ensuring coverage of Earth's surface including challenging high-latitude regions near the celestial poles.

- **Temporal Coverage:** Observation timestamps were sampled uniformly between 1 January 2000 and 31 December 2025 and restricted to times when the solar altitude was below $< -6^\circ$ (nautical twilight or darker [31]), ensuring stars are observable while providing coverage across all seasons and local times.
- **Camera Configuration:** All images were generated with a 120° horizontal field of view at 1080x720 pixel resolution, simulating wide-angle optics typical of UAV navigation cameras. The camera was configured as zenith-pointing with a north-aligned azimuth, representing an idealized strap-down mounting configuration.
- **Environmental Conditions:** Multiple synthetic datasets were generated at discrete limiting visual magnitude levels (8.0, 7.5, 7.0, 6.5, ..., 4.5, 4.0). For each magnitude threshold, a separate set of 200 Stellarium images was generated, enabling controlled analysis across systematically varying stellar densities.
- **Ground Truth:** For each synthetic observation, Stellarium recorded precise metadata including geodetic coordinates (latitude, longitude), UTC timestamp, and camera orientation parameters, enabling quantitative evaluation of position estimation accuracy via geodesic distance comparison.

ii. Star Identification and Position Estimation Performance

Although the algorithm was tested across all limiting magnitude levels, detailed performance analysis in this section focuses on the representative limiting visual magnitude of 7. This choice is justified by three factors:

1. **Alignment with prior research:** Teague and Chahl [8] report that low-cost UAV cameras reliably detect stars up to magnitude 6–7. Dai et al. [12] similarly use catalogs filtered to ≤ 6 mag for DDR-based identification. Focusing on a limiting visual magnitude 7 thus aligns with the typical sensitivity range used in prior celestial navigation studies.
2. **Realistic star-density regime:** Magnitude 7 provides a balanced scene density representative of practical UAV night-sky imaging. Brighter limits ($m \leq 6$) yield too few stars for stable pattern formation, while fainter limits ($m \geq 7.5$) create unrealistically dense fields given sensor noise, exposure constraints, and atmospheric extinction.
3. **Most informative for evaluating adaptivity:** At this intermediate density, enough stars are present to exercise the adaptive reference selection and polar-star rejection logic, without the trivial (dense) or degenerate (sparse) cases seen at the extremes.

Definition — Recall. Since ground-truth star identities are unavailable in the Stellarium dataset, recall is evaluated indirectly through localization accuracy. A detection is counted as correctly identified when the resulting position error is below 10 km, consistent with the confidence–error relationship shown in this section.

An important consideration in evaluating baseline performance is the influence of polar star detections. The baseline (non-adaptive) algorithm occasionally matched high-declination stars (near the celestial pole, $|\delta| > 75^\circ$) which led to degenerate position solutions. If one were to exclude these polar-star cases from the baseline’s results, the baseline’s recall on the remaining images would increase (for instance, from 39.0% on all 200 images to 43.5% on the 177 non-polar images). At first glance this suggests an improved baseline performance; however, such post-hoc filtering introduces two biases that make the comparison unfair:

1. **Sample size asymmetry:** Removing images with polar-star matches reduces the number of test cases for the baseline (from 200 to 177 in this example), whereas the adaptive algorithm (which internally rejects polar matches) was evaluated on the full set of 200 images. This creates unequal denominators and artificially inflates the baseline’s success rate.
2. **Masking of algorithmic deficiency:** The fundamental issue is not whether polar stars are present in the image, but whether the algorithm identifies them as matches. The baseline’s 23 polar identifications represent algorithmic failures, cases where the system selected geometrically degenerate matches that should have been rejected. Removing these cases from evaluation obscures this failure mode rather than accounting for it.

For these reasons, all performance metrics are reported on the complete 200-image set for each method, without excluding difficult cases. This ensures a fair comparison between the adaptive system (which inherently avoids certain failure cases) and the baseline (which would otherwise require result filtering).

Under the magnitude 7 test scenario (Table 1.1), the adaptive algorithm achieved a 71.5% star identification recall across all 200 images, whereas the baseline (fixed-catalog) algorithm managed only 39.0% recall. In other words, the adaptive approach successfully localized in nearly 72% of the trials, versus 39% for the conventional baseline. This 32.5 percentage point improvement in recall demonstrates the effectiveness of incorporating dynamic catalog selection and geometric filtering into the star identification process. Among the images that were successfully localized by each algorithm, the adaptive approach also attained higher accuracy: the median

Table 1.1: Ablation study results showing recall rate, computation time, and median localization error across algorithmic configurations at magnitude 7. SI: Similarity-based Iteration; PR: Polar-star Rejection.

Configuration	Recall (%)	Computation Time (s)	Median Error (km)
Baseline Algorithm	39.0	5.64	13.80
Baseline + SI	42.5	47.21	11.96
Baseline + SI + PR	41.5	52.62	10.47
Adaptive Algorithm	71.5	5.23	6.80

position error of its solutions was 6.80 km, which is a 50.7% reduction compared to the baseline’s median error of 13.80 km (see Table 1.1). These results indicate that adaptivity not only increases the success rate but also improves the solution precision for this dataset.

Table 1.2 quantifies the relationship between match confidence and position accuracy. Frames with confidence scores between 0.40–0.60 exhibited median errors of 6.87 km with a 33.8% failure rate (error ≥ 10 km). Performance improved in the 0.60–0.80 range (median error 7.11 km, 24.0% failure rate), though surprisingly the >0.80 range showed slightly higher median error (7.31 km) with a 29.2% failure rate. This counter-intuitive result at very high confidence may reflect a few incorrectly identified stars. When a star is misidentified, its position is essentially arbitrary, leading to large errors in the position estimate despite high confidence. The single low-confidence match (<0.40) resulted in catastrophic failure (5,875.8 km error). This strong correlation between confidence and accuracy validates the normalized vote metric as a reliable quality indicator for autonomous mission planning.

Table 1.2: Position error and failure rate stratified by DDR match confidence score. Failure is defined as localization error ≥ 10 km.

Confidence Range	Count	Median Error (km)	Failure Rate
< 0.40	1	5,875.8	100.0%
0.40–0.60	71	6.87	33.8%
0.60–0.80	104	7.11	24.0%
> 0.80	24	7.31	29.2%

To isolate the contribution of each algorithmic innovation, we examine the ablation results in more detail. All experiments used identical test data, catalog preprocessing, and optimization parameters, varying only the presence of adaptive components.

Algorithmic Configurations (Ablation Studies):

- **Baseline Algorithm:** Implements the core star identification via DDR pattern matching against a fixed star catalog ($m_{\text{lim}} = 4.0$). It selects reference stars in a single pass (based purely on image centrality) and performs no iterative refinement or special filtering. This represents a traditional star-ID approach assuming known environmental conditions.

- **Baseline + SI:** Adds Similarity-based Iteration to the baseline. When initial matching confidence falls below threshold $\tau_c = 0.45$, the algorithm iteratively attempts alternative reference candidates from the detected star set until sufficient identification reliability is achieved. This configuration retains the fixed catalog assumption but addresses low confidence through adaptive reference selection.

- **Baseline + SI + PR:** Incorporates Polar-star Rejection on top of SI. High-declination catalog matches ($|\delta| > 75^\circ$) are filtered because near-pole stars exhibit vanishing altitude variation with observer latitude, rendering position estimation geometrically degenerate. Such matches trigger an automatic reference reselection (to avoid the near-polar degeneracy where changes in latitude produce minimal changes in observed angles for stars near the celestial pole). This configuration still uses the fixed catalog but now incorporates both of the new strategies to improve reliability.

- **Adaptive Algorithm:** This is the full proposed system, integrating all components: (1) *Dynamic Catalog Selection*, the star catalog’s magnitude limit is adjusted on the fly based on the observed star density, using a consensus-driven refinement to find the appropriate limit, (2) *Similarity-based Iteration* for robust reference star selection under ambiguous or low-density conditions, and (3) *Polar-star Rejection* for filtering out geometrically degenerate matches. The adaptive algorithm automatically tunes itself to the environmental conditions in each image while maintaining efficiency through consensus convergence (rather than brute-force testing of all possibilities).

iii. Sensitivity to Environmental Conditions

To further evaluate robustness, we analyzed the algorithm’s performance across the range of sky brightness conditions represented in our datasets. In practical terms, this is a sensitivity analysis with respect to light pollution and observability. We systematically varied the limiting visual magnitude of the star field (which directly controls how many stars are detectable in the

image). Astronomically, the limiting magnitude defines the faintest stars visible. Higher limits (e.g. 7.5–8.0) correspond to dark sky conditions where even very dim stars can be seen, whereas lower limits (e.g. 5.0 or 4.0) correspond to bright, light-polluted urban skies where only the brightest stars remain detectable [32]. For interpretability, we map each limiting-magnitude setting to its approximate Bortle sky class, which describes overall night-sky quality rather than a strict magnitude threshold. For example, a limiting magnitude of \sim 7.5–8.0 corresponds roughly to Bortle Classes 1–3 (excellent dark-sky conditions), while a limiting magnitude of \sim 5.0 corresponds to about Bortle 7–8 (heavy urban light pollution).

Figure 5 summarizes the star identification recall achieved by the Adaptive algorithm (with dynamic catalog, SI, and PR) versus the Baseline algorithm (fixed mag 4.0 catalog) across all these environmental conditions. The adaptive approach maintains high recall over the full range of sky qualities, whereas the baseline’s performance drops off sharply as light pollution increases. In particular:

- **Dark sky conditions** (bortle 1–2): Under very dark skies with abundant visible stars, both algorithms achieve high recall rates above 70%. The adaptive approach reaches 82.5% recall at Bortle 1 and 81% at Bortle 2, while the baseline achieves 74.5% and 71% respectively. In these pristine conditions, the baseline’s fixed catalog proves adequate, as suffi-

cient bright stars remain visible for reliable pattern matching. The adaptive algorithm demonstrates modest improvements of 8–10 percentage points by effectively applying SI and PR while using a similar selected catalog as the baseline

- **Rural/Suburban conditions** (bortle 3–5): A dramatic divergence emerges as conditions degrade. At Bortle 3 (limiting magnitude \approx 7.0), the adaptive algorithm maintains 71.5% recall while the baseline drops precipitously to 39%. This crossover marks the critical threshold where dynamic catalog selection becomes essential. As light pollution intensifies to Bortle 4 (magnitude \approx 6.5), the baseline collapses to just 3% while the adaptive sustains 71.5% recall. By Bortle 5 (magnitude \approx 6.0), the adaptive algorithm achieves 61% recall compared to the baseline’s 2%. This regime represents the critical operational envelope for UAV navigation in populated areas, where catalog adaptation transitions from advantageous to absolutely necessary.
- **Urban conditions** (bortle 6–8): Under heavy light pollution, where only the brightest stars remain detectable, the adaptive approach maintains operational capability while the baseline fails completely. At Bortle 6 (limiting magnitude \approx 5.5), the adaptive algorithm achieves 49% recall compared to the baseline’s 1%. At Bortle 7 (magnitude \approx 5.0), the adaptive maintains 27.6% recall while the baseline reaches 0%, the baseline becomes entirely non-

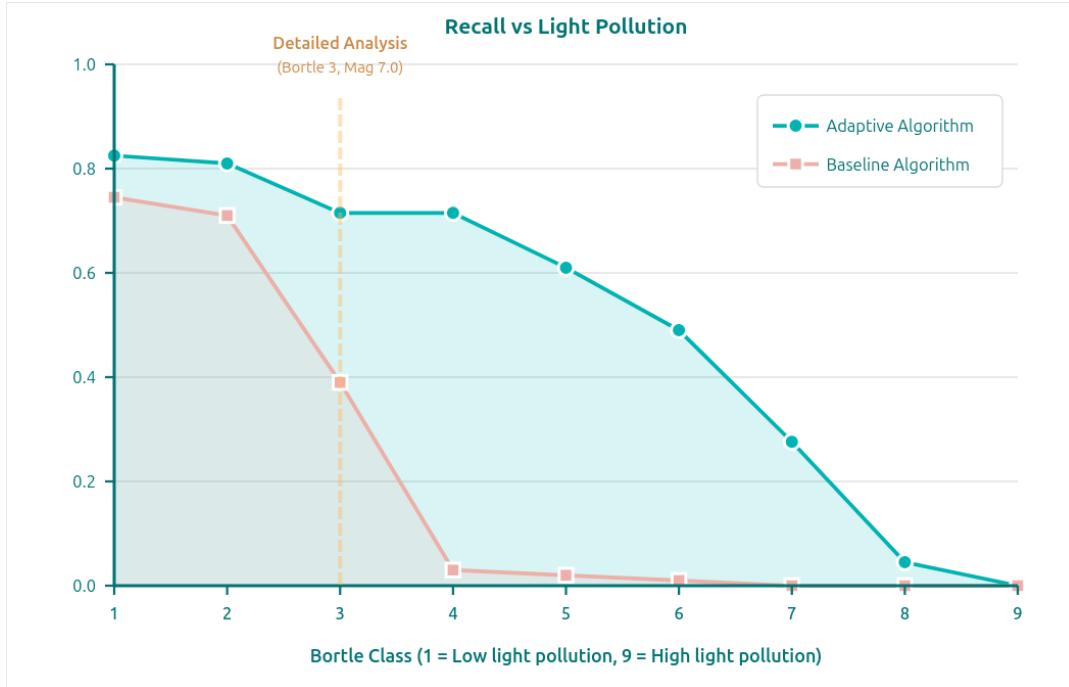


Figure 5: Recall comparison of the Adaptive Algorithm and Baseline Algorithm across varying Bortle classes, with detailed analysis in Section 5– ii focusing on Bortle 3 (Magnitude 7) conditions.

functional. Even at Bortle 8 (magnitude ≈ 4.5), the adaptive algorithm sustains 4.5% recall, whereas the baseline remains at 0%. At Bortle 9 (magnitude ≈ 4.0), both algorithms fail as too few stars remain visible for any reliable pattern matching.

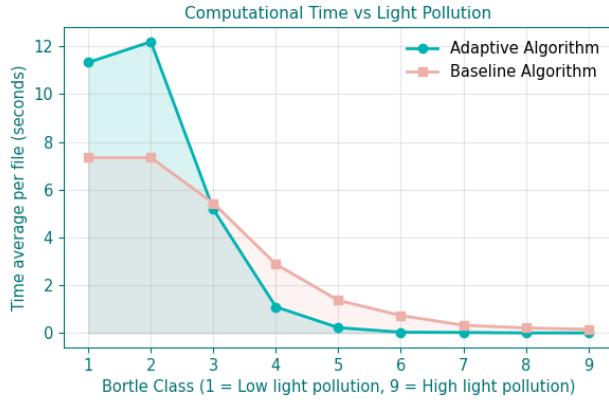


Figure 6: Average computation time of the Adaptive and Baseline algorithms across varying Bortle classes. Dark skies increase runtime for the Adaptive method due to higher star density, whereas light-polluted conditions reduce the workload, making it comparable to or faster than the Baseline.

The performance trends reveal a fundamental operational trade-off. In pristine conditions (Bortle 1–2), the baseline’s simpler fixed-catalog approach proves competitive, achieving 88–90% of the adaptive algorithm’s recall. However, this narrow performance gap masks a critical vulnerability: the baseline exhibits catastrophic degradation once conditions exceed Bortle 2. The transition from Bortle 2 to Bortle 3 sees the baseline’s recall plummet from 71% to 39%, while the adaptive algorithm maintains stable 71.5% performance. Beyond Bortle 4, the baseline becomes operationally unusable, sustaining less than 3% recall. In contrast, the adaptive algorithm demonstrates graceful degradation across the full environmental spectrum. While its recall does decline from 82.5% at Bortle 1 to 27.6% at Bortle 7, it maintains functional capability throughout. This extended operational envelope is critical for real-world UAV deployment, where environmental conditions cannot be guaranteed and may vary substantially during a single mission. The adaptive algorithm’s ability to sustain 27.6% recall at Bortle 7, where the baseline achieves zero, represents the difference between operational capability and complete system failure in urban environments.

In addition to identification success, we examined the computational efficiency of each algorithm under varying conditions. Figure 6 illustrates the average processing time per image for the adaptive and baseline methods as a function of sky brightness. A clear trade-off emerges. Under star-rich dark sky conditions, the

adaptive algorithm incurs significantly higher computation time than the baseline, because it must process a much larger catalog and often perform iterative matching. In our tests, the adaptive system analyzes several times more stars than the baseline in dark conditions, leading to longer runtimes. However, as light pollution increases (fewer stars visible), the adaptive algorithm’s workload decreases. It automatically restricts the star catalog to the sparse visible set, which cuts down the pattern-matching complexity. By the time we reach the most light-polluted scenarios, the adaptive algorithm’s average runtime per image drops below that of the baseline. In those cases, the baseline doesn’t gain any speed advantage from the simpler sky, since it still performs its fixed set of operations, whereas the adaptive method is effectively doing less work on the small set of stars. In summary, the adaptive system pays a computational cost in pristine conditions (where it processes more data than necessary for the baseline), but becomes comparably efficient in poor conditions, all while maintaining far superior recall. This adaptivity in computation is an advantageous side-effect. The algorithm expends effort proportional to the complexity of the scene, and in the very scenarios where the baseline is fastest (because there are few stars, but also few navigation cues) the adaptive method is actually both fast and accurate.

6. Discussion

The experimental results validate the central premise that adapting to environmental conditions is essential for robust UAV celestial navigation across realistic operational scenarios. While the results section documented substantial performance improvements. 71.5% recall versus 39.0% for the baseline at magnitude 7, with an even more dramatic divergence in light-polluted conditions. This discussion examines the deeper implications of these findings for system design, deployment strategy, and the broader role of celestial navigation in GNSS-denied operations.

i. The Necessity of Environmental Adaptation

The performance divergence between adaptive and baseline algorithms reveals that environmental adaptation is not an incremental enhancement but a fundamental architectural requirement. Three observations support this conclusion.

First, the existence of a sharp performance threshold near Bortle 3–4 (limiting magnitude 6.0–6.5) indicates that fixed-catalog approaches operate within a narrow margin of environmental tolerance. A system tested under pristine dark-sky conditions may exhibit acceptable performance during validation, yet fail catastrophically when deployed in even moderately light-polluted environments. This brittleness poses significant operational

risk. Missions planned assuming dark-sky performance may encounter complete navigation failure when atmospheric conditions degrade or when route planning necessitates transit through populated areas. The adaptive algorithm eliminates this vulnerability by treating environmental conditions as an observable variable rather than a design constant.

Second, the graceful degradation exhibited by the adaptive system, maintaining 27.6% recall at Bortle 7 where the baseline achieves zero, demonstrates that catalog adaptation extracts meaningful navigation capability from severely constrained observability. This is not merely a performance optimization but an expansion of the operational envelope into conditions previously considered infeasible for celestial navigation. For missions where alternatives are unavailable (no GNSS, no terrain features, no visual landmarks), even 27.6% success rate provides intermittent position updates that can bound Inertial Measurement Unit (IMU) drift and enable mission continuation.

Third, the computational efficiency characteristics suggest that adaptation provides compounding benefits. In the precise conditions where environmental challenges are greatest (sparse star visibility, high uncertainty), the adaptive algorithm becomes computationally efficient by restricting its search space. The adaptive system expends maximum effort when conditions permit and automatically reduces complexity when resources are constrained. A fixed-catalog system, by contrast, maintains constant computational load regardless of whether that effort produces useful results.

ii. Interpreting the Ablation Study: Component Contribution

The ablation results illuminate which algorithmic innovations contribute to robustness and under what conditions. Dynamic catalog selection emerged as the dominant factor, contributing approximately 30 percentage points of recall improvement. This finding was initially surprising, intuitively, one might expect SI (which explicitly addresses detection incompleteness) or PR (which eliminates geometric degeneracy) to be more impactful. However, the result makes sense upon reflection: if the algorithm attempts to match observed stars against an inappropriate catalog (too dense or too sparse), no amount of iterative refinement can recover correct correspondences. The catalog selection problem is foundational, other innovations address second-order effects.

SI contributed modestly (+3.5 percentage points) at substantial computational cost (8.4 \times runtime increase). This suggests SI is most valuable as a fallback mechanism rather than a primary strategy. In the majority of cases where initial reference selection succeeds, SI

provides no benefit and merely delays the result. However, in the subset of cases where initial matching fails due to unfortunate reference star selection (e.g., the brightest detected star happens to be near an edge or in a sparse region), SI enables recovery. This pattern, low average benefit but high value in edge cases, suggests SI should be implemented with early termination. If initial matching achieves sufficient confidence, skip iterative exploration entirely.

PR demonstrated a subtle but critical contribution. While it did not dramatically increase overall recall, it eliminated a specific failure mode that produced systematic rather than random errors. The baseline's 11.5% incidence of polar identifications means that roughly one in nine images would produce a solution with arbitrarily large latitude uncertainty. For an autonomous system, such systematic failures are more problematic than random noise because they appear geometrically consistent (high confidence) while being fundamentally wrong. The value of PR lies not in average-case performance but in eliminating a tail risk that could compromise mission safety.

iii. Calibration, Projection Models, and Accuracy Limits

Beyond the performance metrics, the experiments also revealed a subtle but important point regarding system calibration and accuracy. We observed an unexpected phenomenon. The highest-confidence star matches did not always yield the lowest errors. In fact, the dataset showed that images with an extremely high identification confidence (> 0.80) sometimes produced slightly larger median position errors than those with more moderate confidence (0.6–0.8). This counter-intuitive result prompted a closer examination of the imaging geometry and revealed a bias in the synthetic data generation. Stellarium was configured for a 120° field-of-view to simulate a wide-angle lens, however, upon comparison with the Yale catalog's projections, it became clear that Stellarium's effective FOV was not exactly as specified. The rendered images were more "zoomed out" than a true 120° projection. This mis-calibration meant that stars near the edges of the Stellarium images had greater positional errors than anticipated, because the assumed camera model in the algorithm did not perfectly match Stellarium's projection. In essence, there was a slight discrepancy between how Stellarium mapped celestial coordinates to pixel coordinates and how our algorithm thought the mapping should be. As a result, some high-confidence solutions were actually using subtly incorrect geometry, especially if they involved stars toward the periphery of the image.

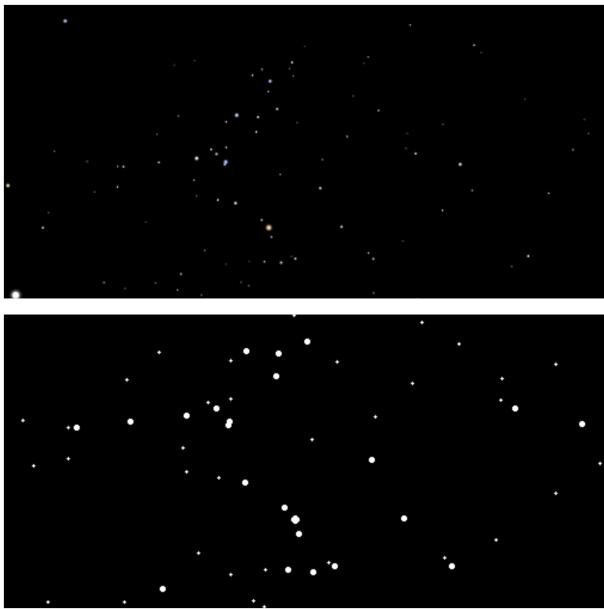


Figure 7: Comparison between a raw Stellarium render (top) and its corresponding YBSC projection (bottom). The mismatch between Stellarium’s projection and the algorithm’s camera model introduces peripheral geometric distortion, explaining why some high-confidence matches yield higher localisation error.

To mitigate this, we performed an offline camera model tuning for the Stellarium dataset. We adjusted the preprocessing pipeline (specifically, the image scaling and normalization steps) to better align the Stellarium-generated star positions with the expected 120° projection. These tuning adjustments substantially improved the consistency of the star maps. After calibration, the algorithm’s estimated FOV more closely matched Stellarium’s, and the projection errors for stars at the image edges were reduced. This correction was crucial for ensuring that the localization error truly reflected algorithmic performance rather than a simulator artifact.

iv. Operational Context and Mission Integration

The results indicate that the demonstrated system is best regarded as an auxiliary, infrastructure-independent sensor within a multi-sensor navigation suite rather than a drop-in replacement for GNSS. With a median geolocation error of approximately 6.8 km, celestial fixes are well suited to coarse waypoint-level routing, periodic bounding of IMU drift during long-range transits, and continuity or integrity monitoring for TRN, vision-, or GNSS-based solutions, particularly during night operations over unstructured terrain.

A critical operational prerequisite is an upward-looking camera with unobstructed sky view. While the algorithm assumes a zenith-pointing configuration, practical UAV deployments must account for platform at-

titude variations, which can be addressed through IMU-based attitude knowledge or gimbal stabilization. Forward- or downward-looking cameras used for visual odometry or terrain navigation cannot provide celestial observations, necessitating either a dedicated zenith camera or a multi-axis gimbal system that can reorient during celestial update cycles. This hardware requirement distinguishes celestial navigation from terrain- or feature-based methods and must be factored into system integration decisions.

At the same time, residual camera-calibration and projection-model errors, together with the finite astrometric accuracy of the star catalog, impose a hard accuracy floor. At the wide fields of view typical for UAVs, a one-pixel centring error can translate into kilometre-scale position offsets, such that the method is not appropriate for terminal guidance, obstacle-proximity manoeuvres, or other applications requiring sub-hundred-metre precision. The value proposition of celestial navigation therefore lies in diversification rather than displacement, as a low-rate absolute positioning input that exploits independent observables and failure modes, it enhances the resilience of integrated architectures that fuse GNSS, IMU, visual odometry, and TRN.

v. Limitations and Validity Constraints

Several limitations constrain the generalizability of these findings. First, evaluation used synthetic imagery with known ground truth but imperfect realism. Stellarium rendering incorporates atmospheric extinction and sky brightness, but does not model (thin) clouds, optical aberrations, or sensor-specific noise characteristics. Real-world performance may degrade relative to synthetic results if these factors prove significant. Second, the assumption of known camera orientation simplifies the problem substantially. Operational systems must either integrate an IMU for attitude knowledge or implement lost-in-space pattern recognition that jointly solves for position and orientation. The latter is computationally expensive and may not be feasible for real-time UAV applications. IMU integration introduces drift in attitude knowledge, which propagates to position error through the coordinate transformations. Quantifying this effect requires hardware-in-the-loop testing. Third, motion blur during image exposure was not addressed. UAV platform motion during typical 0.5–2 second exposures will streak star images, degrading centroid accuracy and potentially causing missed detections. The magnitude of this effect depends on platform stability and exposure duration, factors that trade against each other, since shorter exposures reduce blur but also reduce signal-to-noise ratio for faint stars. Fourth, the PR threshold ($|\delta| > 75^\circ$) was chosen heuristically. While this eliminates geometric degeneracy

in mid-latitude regions, it may be overly conservative (rejecting useful stars) or insufficiently conservative (permitting marginal matches) depending on latitude and required accuracy. A more principled approach would compute condition numbers for the position estimation Jacobian and reject matches that exceed a specified threshold, adapting the rejection criterion to the geometric configuration rather than using a fixed declination cutoff.

7. Conclusion

This thesis has demonstrated that robust vision-based celestial navigation in variable environments requires dynamic adaptation to observability conditions. The developed system achieves this through three integrated innovations: consensus-driven magnitude refinement that automatically adjusts catalog density to match visible stellar populations, PR that eliminates geometrically degenerate solutions, and SI matching that maintains identification robustness under sparse detection conditions.

Evaluation on 200 synthetic images spanning pristine to heavily light-polluted skies validates the approach. Under dark rural conditions (Bortle 3, limiting magnitude 7), the adaptive algorithm achieves 71.5% identification recall compared to 39.0% for conventional fixed-catalog matching, while simultaneously reducing median localization error by 50.7% (6.80 km versus 13.80 km). More critically, the system sustains operational capability throughout the environmental spectrum: maintaining 27.6% recall at Bortle 7 (urban/suburban) where baseline methods achieve zero, and extending marginally functional operation to Bortle 8 (4.5% recall) despite only the brightest stars remaining visible.

These results address a central challenge identified throughout prior work on autonomous celestial navigation: that robust star identification, rather than the subsequent geometric estimation, constitutes the primary bottleneck for reliable operation in lost-in-space scenarios. The literature consistently treats star identification as a distinct and unresolved problem, devoting entire algorithmic studies to correspondence establishment under limited fields of view, sparse observations, and environmental degradation, while assuming that attitude or position estimation becomes well-conditioned once correct correspondences are available. By explicitly adapting to environmental uncertainty instead of assuming fixed observability conditions, the proposed system demonstrates resilience across regimes in which traditional fixed-catalog approaches exhibit catastrophic failure rather than graceful degradation.

The achieved localization accuracy remains substantially coarser than GNSS or visual SLAM under favorable conditions and is therefore not suitable for

precision waypoint following or terminal guidance. Instead, the demonstrated performance supports coarse global localization, GNSS integrity monitoring, and periodic bounding of inertial drift in GNSS-denied environments, where absolute position estimates on the order of several kilometers can still provide operational value. The limitations revealed through synthetic validation, particularly sensitivity to camera calibration and wide-angle projection errors, underscore the necessity of real-world flight testing with calibrated hardware. The reported 6.80 km median error should be interpreted as an upper bound on algorithmic performance under idealized detection and known camera parameters. Operational accuracy will depend critically on sensor quality, optical characterization, and environmental effects not fully captured in synthetic rendering.

i. Future Work

Hardware Integration and Flight Testing Real-world validation on physical UAV platforms remains the critical next step. This requires integration with flight control systems, power management, and thermal considerations. Flight testing under varied weather conditions, atmospheric turbulence, and actual urban light pollution will reveal performance characteristics not captured in synthetic simulation. Particular attention must be paid to vibration isolation of the imaging sensor, as it directly impact star detection reliability [33].

Attitude Estimation and Motion Compensation The current implementation assumes known camera orientation from Stellarium metadata. Operational systems must either integrate with an IMU for attitude knowledge, or perform autonomous attitude estimation by identifying stars and solving Wahba's problem to recover the camera orientation from matched star direction vectors [34]. IMU fusion introduces additional complexity, as gyroscope drift must be bounded through periodic celestial attitude updates, while accelerometer measurements provide complementary information for altitude estimation. Motion blur during image exposure presents a related challenge: UAV platform motion during typical 0.5–2 second exposures can streak star images, shifting measured centroids and thereby degrading attitude accuracy. Mitigation strategies include active gimbal stabilization, shorter exposures with increased sensor gain, or computational de-blurring using IMU-measured motion during exposure.

Multi-Frame Temporal Fusion Current performance is evaluated on single-frame solutions. Sequential estimation over multiple images would enable Kalman filtering or particle filtering approaches that explicitly model process noise (platform dynamics) and measurement noise (identification uncertainty). Temporal fusion offers two advantages: statistical averaging re-

duces random error, and consistency checking across time rejects transient misidentifications that appear geometrically valid in single frames but exhibit impossible motion [35].

Sensor Fusion Architectures Integration with complementary positioning modalities, such as visual odometry, TRN, and magnetometer-based heading, would enable a resilient hybrid navigation system suitable for GNSS-degraded or denied UAV operations. Conventional visual odometry provides high-precision relative motion estimates but degrades in low-light or night-time conditions and suffers from unbounded drift without external corrections [36]. Recent work investigates infrared and thermal cameras for odometry under darkness, showing improved robustness in low-visibility environments but continued sensitivity to scene structure and thermal contrast [37]. CNS, provides an absolute global reference with coarse accuracy, while TRN offers position estimates when reliable geospatial databases are available. An unified probabilistic fusion framework that dynamically weights each modality based on environmental conditions and estimated uncertainty is therefore essential to maximize availability and robustness across operational UAV scenarios [38].

Computational Optimization Beyond algorithmic robustness, the current prototype leaves substantial headroom for computational optimisation. The adaptive matcher in its Python implementation requires on the order of 5–50 s per frame depending on configura-

tion (Table 1.1), with runtime dominated by high-level Python loops in DDR feature extraction and repeated JSON-backed LUT queries. Re-implementing star detection and pattern construction in C++ with vectorised CPU instructions (e.g., NEON/AVX), and loading catalog data into contiguous in-memory arrays at start-up, would substantially reduce interpreter and parsing overhead [39], [40]. On platforms with embedded GPUs, star detection can be expressed as convolution and reduction operations that map naturally to CUDA or OpenCL [41]. Since SI evaluates candidate reference subsets independently, it is amenable to coarse-grained parallelisation across CPU cores or GPU threads [42]. Combined with confidence-based early termination, this trades a modest increase in mismatch probability for a significant reduction in worst-case runtime, consistent with real-time embedded design principles [43].

Acknowledgements

I would like to express my sincere gratitude to my daily supervisor, Oscar de Groot, for his continuous guidance, critical feedback, and practical insight throughout this work. I also thank my supervisor at TU Delft, Laura Ferranti, for her academic support and oversight. I am grateful to my parents for their unwavering support and encouragement, and to my friends for their understanding and motivation. Finally, I would like to thank the people I met at Intelic for providing an inspiring environment and valuable discussions that contributed to this research.

Bibliography

- [1] A. Adel, N. H. S. Alani, S. T. Whiteside, and T. Jan, "Who is watching whom? military and civilian drone: Vision intelligence investigation and recommendations," *IEEE Access*, 2024, Received 20 Oct. 2024, accepted 20 Nov. 2024, published 22 Nov. 2024, current version 5 Dec. 2024. doi: [10.1109/ACCESS.2024.3505034](https://doi.org/10.1109/ACCESS.2024.3505034).
- [2] European Organisation for the Safety of Air Navigation (EUROCONTROL), "Performance review report (prr) 2024: An assessment of air traffic management in europe," Performance Review Commission, EUROCONTROL, Tech. Rep., 2025, Accessed: [insert date you accessed the document]. [Online]. Available: <https://www.eurocontrol.int/sites/default/files/2025-03/eurocontrol-performance-review-report-2024.pdf>.
- [3] S. Garg, T. Fischer, and M. Milford, "Where is your place, visual place recognition?" In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence (IJCAI-21), Survey Track*, Aug. 2021, pp. 4416–4425. doi: [10.24963/ijcai.2021/603](https://doi.org/10.24963/ijcai.2021/603). [Online]. Available: <https://doi.org/10.24963/ijcai.2021/603>.
- [4] J. Qiao, F. Yang, J. Liu, G. Huang, W. Zhang, and M. Li, "Advancements in indoor precision positioning: A comprehensive survey of uwb and wi-fi rtt positioning technologies," *Network*, vol. 4, no. 4, pp. 545–566, 2024. doi: [10.3390/network4040027](https://doi.org/10.3390/network4040027).
- [5] J. P. Golden, "Terrain referenced navigation: History, trends and the unused potential," in *2012 IEEE/ION Position, Location and Navigation Symposium*, IEEE, 2012, pp. 44–52. doi: [10.1109/PLANS.2012.6236949](https://doi.org/10.1109/PLANS.2012.6236949).
- [6] T. Kim, S. Nam, H. Lee, and J. Oh, "Verification of vision-based terrain-referenced navigation using the iterative closest point algorithm through flight testing," *Sensors*, vol. 25, no. 18, p. 5813, 2025. doi: [10.3390/s25185813](https://doi.org/10.3390/s25185813).
- [7] COPTRZ. "Flying drones at night: What is legal and what is not." Accessed: 2025-11-14. [Online]. Available: <https://shop.coptrz.com/blogs/news/flying-drones-at-night-what-is-legal-what-is-not>.
- [8] S. Teague and J. Chahl, "An algorithm for affordable vision-based gnss-denied strapdown celestial navigation," *Drones*, vol. 8, no. 11, p. 652, 2024. doi: [10.3390/drones8110652](https://doi.org/10.3390/drones8110652).
- [9] C. to Protect Rural England (CPRE) and B. A. Association, *Star count 2019: Up to 6 stars (class 9: Inner-city sky) ... up to 33 stars (class 4: Rural/suburban transition)*, PDF document, Available at: https://bathastronomers.org.uk/wp-content/uploads/2019/01/CPRE_Star_Count_2019_20190119.pdf, 2019. [Online]. Available: https://bathastronomers.org.uk/wp-content/uploads/2019/01/CPRE_Star_Count_2019_20190119.pdf.
- [10] D. Rijlaarsdam, H. Yous, J. Byrne, D. Oddenino, G. Furano, and D. Moloney, "A survey of lost-in-space star identification algorithms since 2009," *Sensors*, vol. 20, no. 9, p. 2579, 2020. doi: [10.3390/s20092579](https://doi.org/10.3390/s20092579). [Online]. Available: <https://doi.org/10.3390/s20092579>.
- [11] B. B. Spratling and D. Mortari, "A survey on star identification algorithms," *Algorithms*, vol. 2, no. 1, pp. 93–107, 2009. doi: [10.3390/a2010093](https://doi.org/10.3390/a2010093). [Online]. Available: <https://doi.org/10.3390/a2010093>.
- [12] Y. Dai et al., "Star identification algorithm based on dynamic distance ratio matching," *Remote Sensing*, vol. 17, no. 1, p. 62, 2025.
- [13] N. Bowditch, *The american practical navigator*, Online Course Material, "If the Earth did not rotate ... it is necessary to include time measurements in the linking of the celestial and geographical coordinate systems.", 2005. [Online]. Available: <https://web.pa.msu.edu/courses/2005fall/AST101/coursepk.html>.
- [14] K.-C. Tsai, W.-K. Tseng, C.-L. Chen, and Y.-J. Sun, "A novel analytical solution method for celestial positioning," *Journal of Marine Science and Engineering*, vol. 10, no. 6, p. 771, 2022. doi: [10.3390/jmse10060771](https://doi.org/10.3390/jmse10060771). [Online]. Available: <https://doi.org/10.3390/jmse10060771>.

[15] R. H. M. Bosanquet, "On the relation between precession and proper motion," *Monthly Notices of the Royal Astronomical Society*, vol. 55, no. 8, pp. 481–486, 1895. doi: 10.1093/mnras/55.8.481. [Online]. Available: <https://academic.oup.com/mnras/article/55/8/481/1013800>.

[16] A. Fraknoi, D. Morrison, and S. C. Wolff, "2.1 the sky above," in *Astronomy*, Houston, TX: OpenStax, 2017. Accessed: Nov. 25, 2025. [Online]. Available: <https://pressbooks.online.ucf.edu/ast2002tjb/chapter/2-1-the-sky-above/>.

[17] U.S. Naval Observatory, Astronomical Applications Department, *Computing altitude and azimuth from greenwich apparent sidereal time*, Accessed: 2025-10-20, n.d. [Online]. Available: https://aa.usno.navy.mil/faq/alt_az.

[18] J. Meeus, *Astronomical Algorithms*, 2nd. Richmond, VA: Willmann-Bell, 1998.

[19] R. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd. Cambridge, UK: Cambridge University Press, 2004, Accessed: 2025-11-14. [Online]. Available: http://www.r-5.org/files/books/computers/algo-list/image-processing/vision/Richard_Hartley_Alexander_Zisserman-Multiple_View_Geometry_in_Computer_Vision-EN.pdf.

[20] D. Mortari, "A fast and robust star identification algorithm," *Journal of the Astronautical Sciences*, vol. 45, no. 2, pp. 179–194, 1997.

[21] B. Spratling and D. Mortari, "Star identification and attitude determination with projective cameras," in *AIAA Scitech 2021 Forum*, 2021. doi: 10.2514/6.2021-0375.

[22] D. Hoffleit and W. H. Warren Jr. "Yale bright star catalogue, 5th revised edition (bsc5)," Accessed: Nov. 10, 2025. [Online]. Available: <http://tdc-www.harvard.edu/catalogs/bsc5.html>.

[23] E. W. Weisstein. "Great circle." Accessed: 2025-11-10, Wolfram Research. [Online]. Available: <https://mathworld.wolfram.com/GreatCircle.html>.

[24] G. A. Finney, S. Fox, B. Nemati, and P. J. Reardon, "Extremely accurate star tracker for celestial navigation," in *Proceedings of the 2023 Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS)*, Poster, AMOS, 2023. [Online]. Available: <https://amostech.com/TechnicalPapers/2023/Poster/Finney.pdf>.

[25] Celestron, *Exposure time in celestial photography*, <https://www.celestron.com/blogs/knowledgebase/exposure-time-in-celestial-photography>, Accessed: 2025-11-27, Jul. 2023.

[26] R. S. J. Wright, *Seeing vs. transparency: What's the difference?* <https://skyandtelescope.org/astronomy-blogs/imaging-foundations-richard-wright/seeing-vs-transparency-difference/>, Accessed: 2025-11-27, Dec. 2017.

[27] R. G. Stuart, "Applications of complex analysis to precession, nutation and aberration," *Monthly Notices of the Royal Astronomical Society*, vol. 400, no. 3, pp. 1366–1372, 2009. doi: 10.1111/j.1365-2966.2009.15529.x. [Online]. Available: <https://academic.oup.com/mnras/article/400/3/1366/960867>.

[28] H. P. Gavin, "The levenberg–marquardt algorithm for nonlinear least squares curve-fitting problems," Duke University, Department of Civil and Environmental Engineering, Durham, NC, Tech. Rep., 2024, May 5, 2024. [Online]. Available: <https://people.duke.edu/~hpgavin/lm.pdf>.

[29] T. A. Wilson, "Evaluating the effectiveness of current atmospheric refraction models in predicting sunrise and sunset times," Ph.D. thesis, submitted September 2018, Ph.D. dissertation, Michigan Technological University, 2018. [Online]. Available: https://www.researchgate.net/publication/330425692_Evaluating_the_Effectiveness_of_Current_Atmospheric_Refraction_Models_in_Predicting_Sunrise_and_Sunset_Times.

[30] F. Chéreau et al., *Stellarium: Open-source planetarium software*, version 25.3, 2025. doi: 10.5281/zenodo.17226780. [Online]. Available: <https://stellarium.org/>.

[31] National Weather Service, *Twilight types*, <https://www.weather.gov/lmk/twilight-types>, Accessed: 2025-11-25.

[32] Vito Technology, Inc., *Magnitude in astronomy: How is brightness in astronomy measured*, <https://starwalk.space/en/news/what-is-magnitude-in-astronomy>, Accessed: 2025-11-27, 2025.

[33] S. Bagchi, P. Anastasiou, M. Tetlow, T.-J. Chin, and Y. Latif, "Event-based star tracking under spacecraft jitter," *arXiv:2505.12588*, 2025.

[34] M. L. Psiaki, "Generalized wahba problems for spinning spacecraft attitude and rate determination," *Journal of the Astronautical Sciences*, vol. 57, no. 1-2, pp. 73–92, 2010. [Online]. Available: https://gps.mae.cornell.edu/Psiaki_JAS_1337_finalproofs_datecorrected.pdf.

[35] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960. doi: 10.1115/1.3662552.

- [36] D. Scaramuzza and F. Fraundorfer, "Visual odometry [tutorial]," *IEEE Robotics & Automation Magazine*, vol. 18, no. 4, pp. 80–92, 2011. doi: [10.1109/MRA.2011.943233](https://doi.org/10.1109/MRA.2011.943233).
- [37] S. Khattak et al., "On the benefits of thermal cameras for navigation in challenging conditions," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3540–3547, 2019. doi: [10.1109/LRA.2019.2927124](https://doi.org/10.1109/LRA.2019.2927124).
- [38] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Artech House, 2013.
- [39] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, 2017.
- [40] M. Kleppmann, *Designing Data-Intensive Applications*. O'Reilly Media, 2017.
- [41] J. Sanders and E. Kandrot, *CUDA by Example: An Introduction to General-Purpose GPU Programming*. Addison-Wesley, 2010.
- [42] M. J. Quinn, *Parallel Programming in C with MPI and OpenMP*. McGraw-Hill, 2004.
- [43] G. C. Buttazzo, *Hard Real-Time Computing Systems*. Springer, 2011.