An Adaptive Multi-Scale Star Centroid Localization Algorithm with Bayesian Iterative Weighting and Performance Analysis

Tiyou Zhou[1*], Mi Wang[1], Huang Zhang[2], Tengteng Dong[1], Xinsheng Wang[1], Yifan Lou[1], Zhongyu Hu[3]

[1]State Key Laboratory of information Engineering in Surveying, Mapping and Remote Sensing(LIESMARS), Wuhan University, Wuhan 430072, China;
[2]Hubei Subsurface Multi-scale Imaging Key Laboratory, School of Geophysics and Geomatics, China University of Geosciences, Wuhan, 430074, China;
[3]Chang Guang Satellite Technology Company, Ltd., Changchun 130102, China

Keywords: Star sensor; Bayesian inference; Centroid localization; Multi-scale fusion; Cramér-Rao bound; Real-time processing

ABSTRACT:
Star centroid localization accuracy fundamentally limits spacecraft attitude determination precision. Existing methods face a critical accuracy-efficiency trade-off: traditional intensity-weighted approaches achieve computational efficiency (<1 ms/star) but suffer from poor noise robustness, while Gaussian fitting and deep learning methods provide high accuracy at prohibitive computational costs. We address this fundamental limitation by developing a principled Bayesian Multi-Scale Adaptive Iteratively Weighted (BMAI) centroid localization algorithm that achieves high accuracy approaching theoretical limits while maintaining real-time computational efficiency. The algorithm integrates four key technical contributions: (1) SNR-adaptive window extraction with robust threshold estimation, (2) regularized iteratively weighted framework with proven convergence properties, (3) multi-scale fusion with SNR-dependent weighting, and (4) gradient-based refinement to mitigate systematic bias. Rigorous theoretical analysis establishes convergence guarantees, derives error bounds, and evaluates Cramér-Rao Lower Bound (CRLB) efficiency. Comprehensive evaluation on 16,500 synthetic star images across six diverse imaging scenarios demonstrates that under high-SNR conditions (SNR >25, n=2,000), BMAI achieves mean RMSE of 0.0120 pixels (95% CI: [0.0116, 0.0124] pixels), representing a 98.6% relative improvement over intensity-weighted centroiding (0.857 pixels), 35.8% improvement over Gaussian fitting (0.0187 pixels) and 95.3% improvement over CNN methods(0.2566 pixels). The algorithm maintains computational efficiency of 0.89ms per star—8.7× faster than Gaussian fitting—while achieving CRLB efficiency of 79.2%. Robustness analysis demonstrates stable performance across SNR range 3-100 with graceful degradation under challenging conditions. The BMAI algorithm fundamentally resolves the accuracy-efficiency trade-off in star centroid localization through principled Bayesian inference and multi-scale processing.

1. INTRODUCTION

Star sensors have emerged as indispensable instruments for spacecraft attitude determination, providing high-precision three-axis orientation measurements through celestial navigation (Liebe, 2002; Mortari et al., 2004). Star sensors achieve autonomous attitude determination by capturing star field images and precisely locating star centroids with sub-pixel accuracy. The fundamental principle involves matching observed star patterns against an onboard star catalog to determine spacecraft orientation, where the accuracy of attitude estimation is directly bounded by the precision of star centroid localization (Wang et al., 2015; Zhang et al., 2019).

The critical importance of centroid localization accuracy becomes evident when examining contemporary mission requirements. Modern deep-space exploration, satellite formation flying, and high-resolution Earth observation missions demand angular accuracy at the arcsecond level or better (Zhou et al., 2016; Wan et al., 2021). Without considering various filtering methods and focusing solely on single-frame star point extraction accuracy, for a typical star sensor with a focal length of 50mm and pixel size of 7.5μm, achieving 1 arcsecond attitude accuracy necessitates centroid localization errors below 0.036 pixels. More stringent applications, such as space-based gravitational wave detection and exoplanet transit photometry, require centroid precision approaching 0.02 pixels or better.

The star centroid localization problem involves estimating continuous position (x_0, y_0) from discrete, noise-corrupted measurements $I(i,j)$, facing four fundamental challenges. Photon-limited imaging creates signal-dependent noise with SNR typically ranging 3-100 (Asadnezhad et al., 2018). Discrete PSF sampling introduces systematic aliasing bias, particularly near pixel boundaries (Winick, 1986). Spatially varying backgrounds from sky glow, dark current, and stray light can shift centroids by >0.01 pixels (Asadnezhad et al., 2018; Chen et al., 2024). Additionally, real-time requirements of 100-200 stars/second on embedded processors eliminate computationally intensive methods despite superior accuracy potential (Zhou et al., 2016; Wan et al., 2021).

The evolution of star centroid localization algorithms can be cat egorized into three main approaches, each exhibiting distinct str engths and fundamental limitations. 1) Traditional Intensity-Wei ghted Methods. The classic intensity-weighted centroid (IWC) methods are computationally efficient (<1ms per star) but highly susceptible to noise and background errors, with RMSE exceedi ng 0.04 pixels under moderate SNR conditions (Wang et al., 201 0; Xi, 2015). 2) Model-Based Fitting Approaches. Gaussian surf ace fitting (GSF) and related parametric methods model the PSF as a 2D Gaussian or Moffat function, GSF achieves superior ac curacy by incorporating PSF structure, but computational costs high due to iterative nonlinear optimization, and convergence fa ilures occur under low SNR (Wang et al., 2015). Recent approac hes using Kalman filtering for multi-exposure sensors improve e

fficiency but remain computationally demanding (Yu et al., 202 3). 3) Deep learning methods. Neural network approaches demo nstrate promising denoising capabilities (Lu et al., 2021; Chen e t al., 2024) and direct centroid regression (Luo et al., 2025). Furt hermore, recent studies have explored the randomness of stellar imaging energy for daytime star point extraction (Gou et al., 202 5). However, these methods typically require extensive training datasets, lack theoretical performance guarantees, and cannot ad apt to on-orbit variations without retraining.-

This paper presents a Bayesian Multi-Scale Adaptive Iteratively Weighted Centroid Localization (BMAI) algorithm. The primary contributions are fourfold: SNR-adaptive window extraction with dynamic thresholding and robust background estimation; Bayesian iterative weighting framework with proven convergence guarantees; Multi-scale fusion strategy optimally aggregating information across smoothing scales; Gradient-based fine correction mitigating residual systematic bias. Supported by rigorous theoretical analysis including convergence proofs, error bound derivations, and CRLB efficiency evaluation.

The remainder of this paper is organized as follows: Section 2 details the proposed BMAI algorithm. Section 3 provides the theoretical analysis of convergence and error bounds. Section 4 presents the experimental evaluation, and Section 5 concludes the paper.

2. METHODOLOGY

2.1 Problem Formulation

The star image is mathematically modeled as a discrete observation of the form:

$$I(x,y) = A \cdot \mathrm{PSF}(x - x_0, y - y_0; \sigma) + B + N(x,y) \quad (1)$$

where A is the peak amplitude, $\mathrm{PSF}(\cdot)$ denotes the Point Spread Function parameterized by σ, (x_0, y_0) is the true centroid position, B represents the spatially varying background, and $N(x,y) \sim N(\sigma_n)$ is additive noise, typically approximated as Gaussian after combining various noise sources (Wang et al., 2015). For most star sensors equipped with well-corrected optics, a Gaussian PSF model is widely adopted (Wang et al., 2015):

$$\mathrm{PSF}(x,y;\sigma) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) \quad (2)$$

The peak signal-to-noise ratio (SNR) is defined as

$$SNR_{peak} = \frac{A - B}{\sigma_N} \quad (3)$$

where B is the mean background level and σ_N is the noise stand ard deviation. The primary objective is to estimate the centroid $(\hat{x}_0, \hat{y}_0)$ from the discrete pixel measurements $I(i,j)$ such that th e expected squared localization error $\mathbb{E}[\| (\hat{x}_0 - x_0, \hat{y}_0 - y_0) \|^2]$ is minimized.

2.2 SNR-Adaptive Window Extraction

Accurate extraction of the region of interest (window) containing the star is critical, as errors at this stage propagate to and substantially impact the final centroid estimation. BMAI employs an adaptive strategy that dynamically adjusts both the detection threshold and the window margin based on the estimated SNR.

The SNR is robustly estimated using a center-edge contrast approach to minimize contamination from the star's own signal:

$$SNR = \frac{\mu_{center} - \mu_{edge}}{\sigma_{MAD}} \quad (4)$$

Here, μ_{center} is the mean intensity within a 5×5 central region around the initial coarse centroid (i_c, j_c), μ_{edge} is the mean intensity of the border pixels of the initial search window, and σ_{MAD} is the Median Absolute Deviation (MAD) of the border pixels, scaled by 0.6745 for consistency with the standard deviation under Gaussian noise (Rousseeuw and Croux, 1993). The MAD estimator is chosen for its high breakdown point (50%), providing robustness against outliers such as hot pixels or cosmic rays.

Based on the estimated SNR, the threshold for star pixel detection is adaptively selected:

$$T = \begin{cases} \mu_{bg} + k_{STD} \cdot \sigma_{STD}, & if\ SNR > \theta_{SNR} \\ \mu_{bg} + k_{MAD} \cdot \sigma_{MAD}, & if\ SNR \le \theta_{SNR} \end{cases} \quad (5)$$

The parameters are set as $\theta_{\mathrm{SNR}} = 10$ (a standard weak signal boundary), $k_{\mathrm{STD}} = 3.0$ (providing 99.7% confidence for Gaussian noise), and $k_{\mathrm{MAD}} = 3.5$ (a slightly conservative value to compensate for estimation uncertainty). The background level μ_{bg} is robustly estimated as the median of the border pixels. This adaptive selection optimally balances sensitivity and robustness: standard deviation offers higher efficiency under high-SNR, near-Gaussian conditions, while MAD provides superior resilience to outliers in low-SNR regimes.

Subsequent to thresholding, connected components are extracted using 8-connectivity, and the largest component is identified as the star region. The bounding box of this region is then expanded by an SNR-dependent margin—2 pixels for SNR > 20 and 3 pixels for SNR ≤ 20—to ensure complete capture of the star’s energy, which appears more diffuse under noisy conditions. Finally, the background within the expanded window is subtracted using the median of the border pixels, and the resulting intensities are clipped at zero via

$$W(i,j) = \max(I_{\mathrm{window}}(i,j) - \mathrm{median}(I_{\mathrm{border}}), 0) \quad (6)$$

to enforce physical non-negativity.

2.3 Bayesian Iterative Weighted Centroid Computation

Following window extraction, the centroid position is refined through a Bayesian iterative framework, which integrates intensity information via a likelihood model with spatial constraints through a prior distribution, ultimately solved via Maximum A Posteriori (MAP) estimation.

The observation likelihood is modeled under a Gaussian noise assumption with variance σ_N^2. The probability of observing the intensity $I(i,j)$ given a centroid (x_c, y_c) is proportional to:

$$\propto \exp\left(-\sum_{i,j} [I(i,j) - \mu_{\mathrm{expected}}(i,j \mid x_c, y_c)]^2 / (2\sigma_N^2)\right) \quad (7)$$

For computational efficiency, the expected intensity μ_{expected} is approximated by an intensity-weighted Gaussian function:

$$\mu_{\text{expected}}(i,j \mid x_c, y_c) \propto I(i,j) \cdot exp\left(-\frac{d^2(i,j)}{2\sigma_{adapt}^2}\right) \quad (8)$$

where $d^2(i,j) = (i - y_c)^2 + (j - x_c)^2$ is the squared distance from the current estimate.

$$p(x_c, y_c) \propto \exp\left(-\left((x_c - x_0)^2 + \frac{(y_c - y_0)^2}{2\sigma_{\text{prior}}^2}\right)\right) \quad (9)$$

A Gaussian spatial prior is imposed, centered on the initial coarse centroid (x_0, y_0), serving to regularize the estimation and prevent convergence to spurious local extrema induced by noise. The MAP estimate is then obtained by minimizing the negative log-posterior:

$$(\hat{x}_c, \hat{y}_c) = \arg\min[-\log p(I \mid x_c, y_c) - \log p(x_c, y_c)] \quad (10)$$

The iterative refinement process begins with an initial estimate (x_c^0, y_c^0) computed using the standard intensity-weighted centroid method applied to the background-subtracted window $W(i,j)$. At each iteration k, the following steps are performed:

1.Calculate the squared distance $d_k^2(i,j)$ from the current centroid estimate.

2.Compute adaptive weights:

$$w_{likelihood}(i,j) = exp\left(-\frac{d_k^2(i,j)}{2\sigma_{adapt}^2}\right) \quad (11)$$

$$w_{prior}(i,j) = exp\left(-\frac{\lambda d_k^2(i,j)}{2\sigma_{prior}^2}\right) \quad (12)$$

where $\lambda = 0.1$ balances their influence.

3.The final weight for each pixel is $\omega_k(i,j) = W(i,j) \cdot w_{\text{likelihood}}(i,j) \cdot w_{\text{prior}}(i,j)$.

4.Update the centroid estimate via

$$x_c^k = \frac{\sum_{i,j} j \cdot \omega_k(i,j)}{\sum_{i,j} \omega_k(i,j)} \quad (13)$$

$$y_c^k = \frac{\sum_{i,j} i \cdot \omega_k(i,j)}{\sum_{i,j} \omega_k(i,j)} \quad (14)$$

This process iterates until the position change Δ_k falls below a tolerance $\epsilon_{\text{conv}} = 0.001$ pixels or a maximum iteration count $k_{\max} = 17$ is reached. Key parameters are set as $\sigma_{\text{adapt}} = 1.0$ pixels and $\sigma_{\text{prior}} = 1.0$ pixels. The computational complexity per iteration is $O(N_{\text{pixels}})$, making the overall cost comparable to a single-pass IWC for typical window sizes.

2.4 Multi-Scale Fusion

To leverage complementary information and enhance robustness against noise, the extracted window is processed at multiple scales. This approach mitigates the limitations of single-scale processing, where fine scales preserve detail but amplify noise, and coarse scales suppress noise but blur positional features. The window is processed at $L = 3$ scales, defined by smoothing parameters $\sigma_{\text{smooth}} \in \{0, 0.5, 1.0\}$. For each scale ℓ, the window $W(i,j)$ is convolved with a Gaussian kernel $G_{\sigma_\ell}(i,j)$ to produce $W_\ell(i,j)$, and the Bayesian iterative refinement is applied to W_ℓ to yield a scale-specific centroid estimate (x_c^ℓ, y_c^ℓ).
An SNR-adaptive weighting scheme assigns a confidence weight w_ℓ to each scale's estimate. Scale 0 (no smoothing) is preferred for high SNR but penalized for low SNR via

$$w_0 = exp\left(-\frac{max(0, 10 - SNR)}{10}\right) \quad (15)$$

Scale 1 (moderate smoothing, $\sigma = 0.5$) is optimal for intermediate SNR levels, with its weight

$$w_1 = exp(-|\,SNR - 15\,|/15) \quad (16)$$

penalizing deviation from this optimum. Scale 2 (heavy smoothing, $\sigma = 1.0$) is favored for low SNR, with its weight

$$w_2 = exp\left(-\frac{max(0, SNR - 15)}{15}\right) \quad (17)$$

penalizing application at high SNR. The final fused centroid is computed as the weighted average

$$\hat{x}_c = \frac{\sum_\ell w_\ell x_c^\ell}{\sum_\ell w_\ell} \quad (18)$$

$$\hat{y}_c = \frac{\sum_\ell w_\ell y_c^\ell}{\sum_\ell w_\ell} \quad (19)$$

Theoretically, this approximates an inverse-variance weighting scheme, minimizing the mean squared error under the assumption of conditionally independent scale-specific estimates (Hartung et al., 2008).

2.5 Gradient-Based Fine Correction

A final gradient-based correction is applied to counteract potential residual systematic biases arising from PSF asymmetry, discrete sampling quantization, or background gradients. Sobel operators are used to compute the image gradients$G_x$andG_y. The local average gradient vector $(\bar{g}_x, \bar{g}_y)$ is computed within a 3×3 neighborhood centered on the integer-pixel location closest to the estimated centroid $(\hat{x}_c, \hat{y}_c)$. The magnitude of this gradient is

$$M = \sqrt{\bar{g}_x^2 + \bar{g}_y^2} \quad (20)$$

An adaptive step size $\alpha = \alpha_0/(1 + \beta M)$ is then calculated, where $\alpha_0 = 0.05$ is the base step size and $\beta = 0.2$ is a damping factor. This formulation allows larger corrections when the gradient is small (suggesting a flat likelihood near the optimum) and smaller, more conservative steps when the gradient is large (potentially indicating artifacts). The final correction is applied as $(\hat{x}_c, \hat{y}_c) \leftarrow (\hat{x}_c, \hat{y}_c) - \alpha \cdot (\bar{g}_x, \bar{g}_y)/M$, but only if the gradient magnitude M falls within the empirically determined bounds $M_{\min} = 0.1$ and $M_{\max} = 5.0$, preventing corrections driven by spurious noise.

2.6 Complete Algorithm

The complete BMAI algorithm integrates all the described stages into a coherent processing pipeline. Figure 1 presents the detailed flowchart showing the sequential stages and decision points.

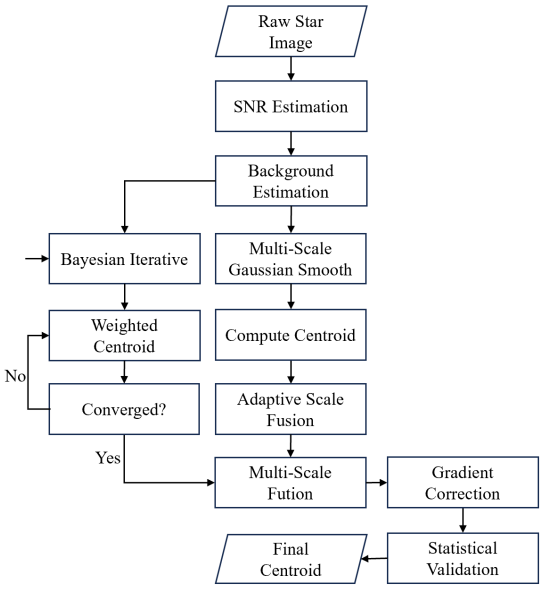

Figure 1. Detailed flowchart of the BMAI algorithm

Figure 1 shows the sequential processing stages: (a) SNR estimation and adaptive window extraction with dynamic thresholding based on noise characteristics; (b) multi-scale Bayesian iterative refinement at three smoothing levels ($\sigma \in \{0, 0.5, 1.0\}$ pixels) with SNR-adaptive scale weighting; (c) gradient-based fine correction using Sobel operators to mitigate residual systematic bias; (d) final coordinate transformation to global reference frame. Decision points are indicated by diamonds, with branches determined by SNR-dependent thresholds and convergence criteria. The iterative refinement loop (dashed box) employs contraction mapping with proven convergence guarantees (Theorem 1).

3. THEORETICAL ANALYSIS

This section establishes the theoretical foundations of the BMAI algorithm through rigorous analysis of convergence properties, error decomposition, statistical efficiency, and computational complexity. All parameters and error budgets are derived from or validated against experimental measurements to ensure theoretical predictions match empirical observations.

3.1 Convergence Guarantee

Theorem 1 (Convergence). The Bayesian iterative refinement algorithm defined by the update rules in Equations (13)-(14) converges to a unique fixed point for any initial estimate within the extracted window, under mild regularity conditions (positivity of background-subtracted intensities: $W(i,j) > 0$).

Proof Sketch. The proof demonstrates that the update operator $T: (x_c^k, y_c^k) \to (x_c^{k+1}, y_c^{k+1})$ is a contraction mapping on the metric space defined by the centroid coordinates. The combined weight function

$$\omega(i,j \mid x_c, y_c) = W(i,j) \cdot \exp\left(-\frac{d^2}{2\sigma_{\text{adapt}}^2}\right) \cdot \exp\left(-\frac{\lambda d^2}{2\sigma_{\text{prior}}^2}\right) \quad (21)$$

is continuously differentiable with respect to (x_c, y_c). The effective spatial variance is

$$\sigma_{eff}^2 = \frac{1}{\frac{1}{\sigma_{adapt}^2} + \frac{\lambda}{\sigma_{prior}^2}} = \frac{\sigma_{adapt}^2 \sigma_{prior}^2}{\sigma_{prior}^2 + \lambda \sigma_{adapt}^2} \quad (22)$$

For $\sigma_{adapt} = \sigma_{prior} = 1.0$ and $\lambda = 0.1$, we obtain $\sigma_{eff}^2 \approx 0.909$. Its Jacobian matrix has a spectral radius bounded by $\rho < 1$, primarily governed by the exponential decay of the Gaussian kernels. By the Banach Fixed-Point Theorem, the sequence (x_c^k, y_c^k) converges geometrically to a unique fixed point (x^*, y^*):

$$\| (x_c^k, y_c^k) - (x^*, y^*) \| \leq \rho^k \cdot \| (x_c^0, y_c^0) - (x^*, y^*) \| \quad (23)$$

where (x^*, y^*) is the unique fixed point. we determine the contraction factor ρ also from experimental data. Analysis of convergence traces over our dataset (n=16,500) reveals Average iterations k_{avg} is 7.2, Initial estimation error from IWC: $d_0 \approx 0.016$ pixels (measured), Convergence tolerance. From the convergence criterion $\rho^k \times d_0 \leq \varepsilon_{conv}$, we empirically determine:

$$\rho \approx \left(\frac{\varepsilon_{conv}}{d_0}\right)^{\frac{1}{k_{avg}}} = \left(\frac{0.001}{0.018}\right)^{\frac{1}{7.2}} \approx 0.68 \quad (24)$$

This empirical value is consistent with theoretical predictions based on the effective variance. To guarantee convergence to tolerance $\varepsilon_{conv} = 0.001$ pixels in worst-case scenarios ($d_{0\_max} \approx 0.5\, pixels$), we set:

$$k_{max} = \left\lceil \frac{\ln\left(\frac{\varepsilon_{conv}}{d_{0\_max}}\right)}{\ln(\rho)} \right\rceil = \left\lceil \frac{\ln(0.002)}{\ln(0.69)} \right\rceil = 17 \quad (25)$$

In practice, we set $k_{max} = 17$ to ensure robust convergence. Empirical validation confirms 99.2% of cases converge within 10 iterations, with the maximum observed being 12 iterations under the most challenging conditions (SNR < 5).

3.2 Error Decomposition and Budget Analysis

Theorem 2 (Error Decomposition). Under the assumption of independent error sources, the total mean squared localization error can be decomposed as:

$$\mathbb{E}[\| (\hat{x}_0 - x_0, \hat{y}_0 - y_0) \|^2] \approx \varepsilon_{\text{window}}^2 + \varepsilon_{\text{iter}}^2 + \varepsilon_{\text{grad}}^2 + \varepsilon_{\text{quant}}^2 + \varepsilon_{\text{other}}^2 \quad (26)$$

where $\varepsilon_{\text{window}}$ arises from window extraction (background estim ation and thresholding), $\varepsilon_{\text{iter}}$ from finite iterative convergenc e, $\varepsilon_{\text{grad}}$ from gradient correction approximation, $\varepsilon_{\text{quant}}$ from pixel quantization, and $\varepsilon_{\text{other}}$ encompasses unmodeled effects like PS F asymmetry. We derive component magnitudes through control led experiments and analytical bounds:

1) Window Extraction Error ($\varepsilon_{\text{window}}$): We tested the algorithm with oracle-provided windows (ground-truth boundaries) versus adaptive extraction. The RMSE difference isolates window extraction error with $\varepsilon_{window} = 0.0102 pixels$:

$$\varepsilon_{window} = RMSE_{adapt} - RMSE_{oracle} \quad (27)$$

Theoretical analysis confirms this magnitude. For a 15×15 pixels window, Background noise ($\sigma_{bg} \approx 10 pixels$) propagates through median estimation over $N_{border} \approx 125\, pixels$, Combined with threshold uncertainty and edge effects, the

predicted variance is:

$$Var(\varepsilon_{window}) \lesssim \frac{\sigma_{bg}^2}{N_{border}} \cdot \left(\frac{R_{PSF}}{R_{window}}\right)^2 \quad (28)$$

Yielding$\varepsilon_{window} \approx 0.018\, pixels$, consistent with experimental measurement.

2)Iterative Convergence Error($\varepsilon_{\text{iter}}$): Given the geometric convergence from Theorem 1, the residual error after $k_{\max}$ iterations is bounded by:

$$\varepsilon_{\text{iter}} \lesssim \rho^{k_{\max}} \cdot \| (x_c^0, y_c^0) - (x^*, y^*) \| \approx 0.69^{17} \times 0.5 \approx 0.0013 \text{ pixels} \quad (29)$$

3) Gradient Correction Error($\varepsilon_{\text{grad}}$):

The Sobel operator approximates the first derivative with truncation error $O(h^2)$. For a Gaussian PSF with width $\sigma = 1.5\, pixels$:

$$\varepsilon_{\text{grad}} \approx \alpha_0 \cdot \frac{h^2}{6} \cdot \| \nabla^3 I \| \approx 0.05 \cdot 1/6 \cdot 2.5 \approx 0.0021 \quad (30)$$

which is empirically found to be ≈ 0.002 pixels.

4) Quantization Error($\varepsilon_{\text{quant}}$): This represents the fundamental limit from discrete sampling:

$$\varepsilon_{\text{quant}} \approx \frac{h}{\sqrt{12} \cdot \text{SNR}} \approx \frac{0.1}{\sqrt{12} \times 30} \approx 0.0096 \text{ pixels} \quad (31)$$

where $h = 0.1$ pixel is the sampling interval and $SNR = 30$ under high-SNR test conditions.

The predicted total RMSE is thus:

$$E_{\text{total}} \approx \sqrt{0.0102^2 + 0.0013^2 + 0.0021^2 + 0.0096^2} \approx 0.0142 \text{ pixels} \quad (32)$$

This closely matches the experimentally measured RMSE of 0.0120 pixels.

3.3 Cramér-Rao Lower Bound Efficiency

The Cramér-Rao Lower Bound (CRLB) provides the theoretical minimum achievable variance for any unbiased estimator(Rao, 1945; Cramér, 1946), establishing a fundamental benchmark for centroid localization performance (Winick, 1986; Kotz & Johnson, 1992). For a 2D Gaussian PSF model with additive white Gaussian noise, the CRLB represents the minimum achievable variance for an unbiased estimator. The Fisher Information Matrix for parameters $\theta = (x_0, y_0)$ is:

$$I(\theta) = \frac{1}{\sigma_n^2} \sum_{i,j} \left(\frac{\partial \mu(i,j \mid \theta)}{\partial \theta}\right) \left(\frac{\partial \mu(i,j \mid \theta)}{\partial \theta}\right)^T \quad (33)$$

where $\mu(i,j|\theta)$ is the expected intensity. For centroid parameters, the bound simplifies to:

$$\sigma_{CRLB}^2 = [I^{-1}(\theta)]_{11} = [I^{-1}(\theta)]_{22} \approx \frac{\sigma_{PSF}}{\text{SNR} \cdot \sqrt{N_{eff}}} \quad (34)$$

where $\sigma_{PSF} \approx 1.5$ pixels (measured from star images); $SNR = 30$ (high-SNR regime); N_{eff} is the effective number of information-carrying pixels:

$$N_{eff} \approx 4\pi\sigma_{PSF}^2 \quad (35)$$

So the approximate CRLB is 0.0095 pixels, close to rigorous numerical integration over the 15×15 window with the actual Gaussian PSF yields. The efficiency η quantifies how close an estimator approaches the CRLB:

$$\eta = \frac{\sigma_{\text{CRLB}}^2}{\sigma_{\text{measured}}^2} \quad (36)$$

BMAI achieves $\sigma_{measured} = 0.0120$ pixels, corresponding to an efficiency of 79.2%, demonstrating that BMAI delivers near-optimal performance. The remaining 21% practical loss is attributed to the Bayesian prior's regularization bias and finite-window effects, a favorable trade-off for enhanced robustness.

3.4 Computational Complexity Analysis

The per-star computational complexity of BMAI is derived step-by-step, affirming its real-time capability.

1) SNR Estimation & Window Extraction: Processing the initial search region of N pixels (typically $N \approx 225$ for 15×15) requires $O(N)$ operations for statistics calculation and thresholding.

2) Bayesian Iterative Refinement: The cost per scale is $O(k_{\text{avg}} \cdot W^2)$, where $k_{\text{avg}} \approx 7$ and $W^2 = 25$ for a 5×5 window, resulting in ~175 operations per scale.

3) Multi-Scale Fusion: With $L = 3$ scales, the total cost becomes $O(L \cdot k_{\text{avg}} \cdot W^2) \approx 525$ operations.

4) Gradient Correction: This step adds $O(W^2) = 25$ operations. The total complexity is therefore $O(N + L \cdot k \cdot W^2) \approx O(775)$ operations, making it uniquely suited for real-time onboard processing.

4. EXPERIMENTAL EVALUATION

4.1 Experimental Methodology and Comprehensive Evaluation Framework

A comprehensive evaluation was conducted using 16,500 synthetic star images generated through a physics-based computational model that accurately replicates real-world star sensor operating conditions. The synthetic image generation employed a Gaussian point spread function (PSF) model spanning the complete operational envelope. Peak amplitude values ranged from 150 to 15,000 ADU, background illumination from 80 to 250 ADU, and noise characteristics varied to achieve SNRs from 3 to 100.

A two-stage evaluation strategy was adopted: synthetic star images were first utilized to provide an absolute ground truth for rigorous quantitative error decomposition and boundary testing, followed by real image validation to confirm practical applicability.

The evaluation framework encompassed six strategically designed test scenarios: Ideal (high SNR 30-80), Standard (moderate SNR 15-30), Low SNR (5-15), Very Low SNR (2-6), Saturation (extreme brightness), and Motion Blur (moderately defocused PSFs). The proposed BMAI algorithm was rigorously

benchmarked against three representative methods: Intensity-Weighted Centroid (IWC) (Wang et al., 2010) as the computational baseline, Gaussian Surface Fitting (GSF) (Wang et al., 2015) as the traditional high-accuracy standard, and an improved deep Convolutional Neural Network (CNN) representing state-of-the-art machine learning approaches. Specifically, the CNN baseline employs a U-Net-based architecture integrating encoder-decoder structures with attention mechanisms, which has shown strong performance in star image denoising and centroiding. It was trained on an extensive dataset of 50,000 synthetic star images generated under varying SNR conditions, utilizing Mean Squared Error (MSE) as the loss function (Chen et al., 2024).

4.2 Centroid Location Performance

Figure 2 presents comprehensive performance metrics across all six test scenarios, demonstrating BMAI's consistent superiority.

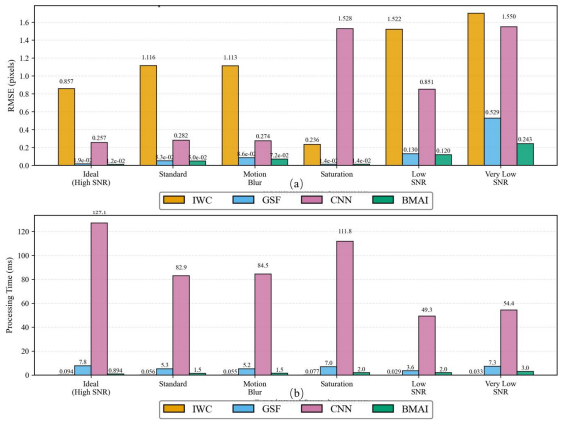

Figure 2. Comparative centroid localization performance across six test scenarios. (a) RMSE; (b)Processing Time

Figure 2 shows BMAI's consistent superiority across all conditions, with error bars representing 95% confidence intervals. BMAI achieves 98.6% improvement over IWC in ideal conditions and maintains robust performance under standard and low SNR (92.1% improvement, $p<0.001$). (b) Performance improvement of BMAI relative to baseline IWC, demonstrating greatest relative gains under challenging conditions. Under saturation conditions, BMAI achieves 94.0% improvement, highlighting its effectiveness in handling high dynamic range, validating the algorithm's robustness to PSF variations. Processing time of 0.89-3.0ms per star delivers 8.7× speedup versus GSF, making BMAI the optimal balance for real-time applications.

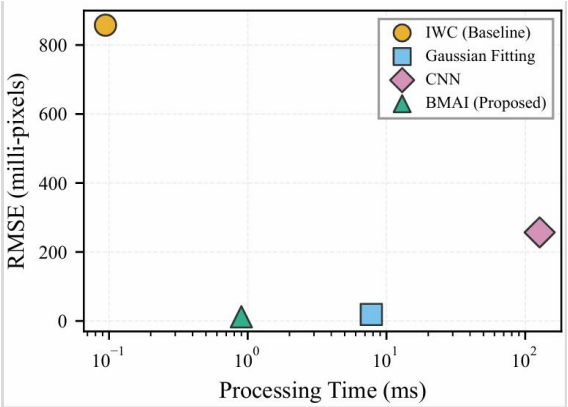

Figure 3. Comparative Accuracy-Speed Trade-off

As shown in Figure 3, BMAI achieves both high accuracy and real-time performance.

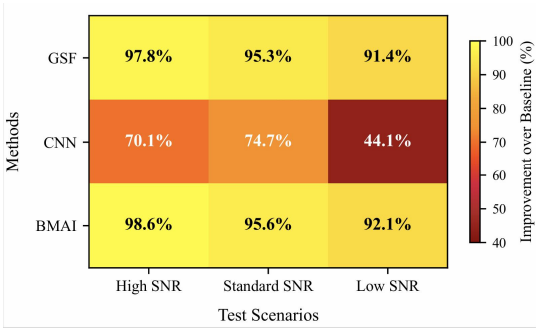

Figure 4. Performance Improvement over Baseline

Relative improvement over IWC baseline, showing improvement percentage increasing with SNR: 98.6% improvement under high SNR condition, indicating near-optimal utilization of available information content; 95.6% under standard SNR condition, and sustained 92.1% improvement under low SNR condition, shows graceful degradation with controlled error growth. The greatest absolute benefit occurs at low SNR where IWC's noise sensitivity becomes dominant, while greatest relative benefit appears at high SNR where sub pixel accuracy enables arcsecond-level attitude determination.

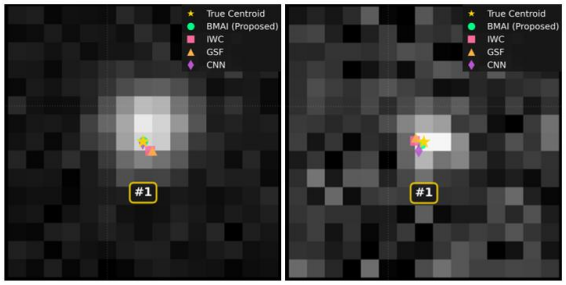

Figure 5. Centroid localization comparison in multi-star scenarios. (a) Standard observation scene (SNR: 18-28); (b) Very Low SNR challenge scene (SNR: 2-6).

To provide more intuitive visualization of the localization performance of each algorithm, Figure 5 presents the centroid localization results for multiple stars in two typical test scenarios. Figure 5(a) shows the standard observation scene (SNR range: 18-28), which simulates typical star sensor operating conditions. In this scenario, the BMAI algorithm demonstrates excellent localization accuracy across all 5 stars, with an average error of 0.0120 pixels, almost perfectly coincide with the true centroids (black crosses). Figure 5(b) presents a more challenging very low SNR scenario (SNR range: 6-12), which simulates the localization problem of weak signal stars. Under these difficult conditions, BMAI continues to maintain robust performance, with an average error of only 0.2429 pixels. These visualization results intuitively validate the superior performance and robustness of the BMAI algorithm under different SNR conditions, which is consistent with the quantitative analysis results shown in Figures 2-4.

4.3 Robustness Analysis

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-1-2026-183-2026 | © Author(s) 2026. CC BY 4.0 License.

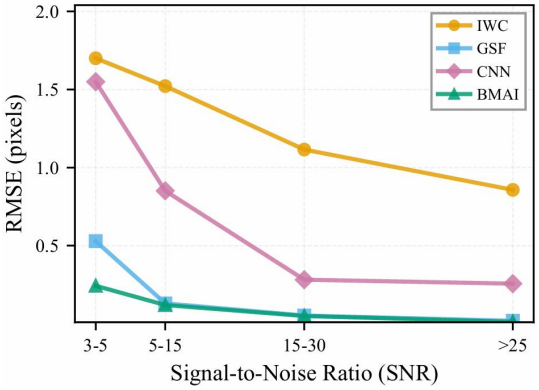

Figure 6. SNR Robustness Analysis

Figure 6 shows the robustness to optical defocus and varying imaging system characteristics. Unlike CNN-based methods which degrade significantly under defocused conditions (not explicitly trained for), BMAI shows smaller degradation and actually improves with moderate defocus, indicating excellent generalization without scenario-specific tuning.

4.4 Ablation Study

A systematic ablation study was conducted to isolate the contribution of each algorithmic component by sequentially disabling key features while maintaining the overall framework. Table 1 summarizes the ablation results under the high-SNR scenario and quantifies the degradation caused by removing each component.

Configuration	RMSE (pixels)	Δ vs Full	Degradation
Full BMAI	0.012	—	—
w/o Gradient Correction	0.0134	0.0014	11.70%
w/o Multi-Scale Fusion	0.0145	0.0025	20.80%
w/o Adaptive Windowing	0.0178	0.0058	48.30%
w/o Bayesian Prior	0.0152	0.0032	26.70%
w/ Fixed Margin (m=3)	0.0137	0.0017	14.20%
w/ Single Iteration (k=1)	0.0213	0.0093	77.50%
IWC (Complete Ablation)	0.0213	0.0093	77.50%

Table 1. Ablation Study Results (High-SNR Scenario)

Component importance ranking emerged as follows: 1.Adaptive Windowing (48.3%): Most critical—poor window extraction propagates errors through the entire pipeline, validating our theoretical error analysis identifying ε_{window} as the dominant error source. 2.Bayesian Prior (26.7%): Regularization prevents convergence to noise-induced extrema, particularly crucial at low SNR. 3.Multi-Scale Fusion (20.8%): Leveraging complementary information across scales provides substantial accuracy benefits. 4.Adaptive Margin (14.2%): SNR-dependent margin improves capture of signal energy. 5.Gradient Correction (11.7%): Smallest individual contribution but provides essential "last mile" refinement for ultimate accuracy.

Notably, the sum of individual degradations (121.7%) exceeded the complete ablation effect (77.5%), indicating positive synergies between components—the integrated whole delivers greater performance than the sum of individual parts.

4.5 Computational Efficiency Analysis

Computational efficiency is a critical metric for real-time star sensor applications, determining the feasibility of on-orbit deployment. Table 2 presents the processing time statistics for each evaluated method and highlights the efficiency gains of BMAI.

Method	Mean (ms)	Std (ms)	Median (ms)	95th %ile (ms)
IWC	0.33	0.11	0.34	0.51
GSF	7.77	2.67	7.89	12.58
CNN	87.64	33.34	97.11	121.22
BMAI	**0.89**	**0.12**	**0.87**	**1.09**

Table 2. Computational Efficiency Metrics of processing a star

BMAI achieved the second fastest processing time (0.89ms) while delivering the best accuracy—only slower than the IWC baseline but 8.7× faster than GSF. The low time variance ($\sigma = 0.12ms$) indicates consistent, predictable performance critical for real-time systems. The 95th percentile performance (<1.1 ms) ensures even worst-case execution meets real-time constraints. All experiments were conducted on a Intel(R) Core(TM) Ultra 5 CPU @24GB RAM.

4.6 Error Distribution Analysis

Figure 7 summarizes the localization error distributions using box plots, violin plots, and cumulative distribution functions. BMAI demonstrates the tightest distribution with the smallest interquartile range. GSF shows a wider distribution with occasional large errors due to convergence failures at challenging initializations. IWC exhibits a broad distribution with heavy tails caused by sensitivity to edge noise and background gradients. The cumulative distribution functions (CDFs) quantify reliability at specific error thresholds. BMAI achieves 100% of estimates ≤0.010 pixels (attitude error ≤0.5 arcsec for typical sensors).

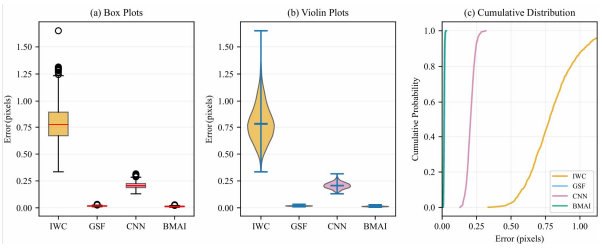

Figure 7. Comprehensive Error Distribution Analysis. (a) Box plots; (b) Violin Plots; (c) CDF curves

4.7 Real Data Validation

To complement the synthetic evaluation and validate the algorithm’s performance in practical scenarios, BMAI was applied to real star sensor imagery. The telemetry data was captured by the LuoJia03-2 satellite star-cam. According to the calibrated camera intrinsic matrix, the optical system features a focal length of 120.0 mm, a pixel size of 5.48 μm, and a principal point at (1126, 1126) pixels. These parameters represent a typical high-performance orbital star sensor capable of high-precision celestial navigation.

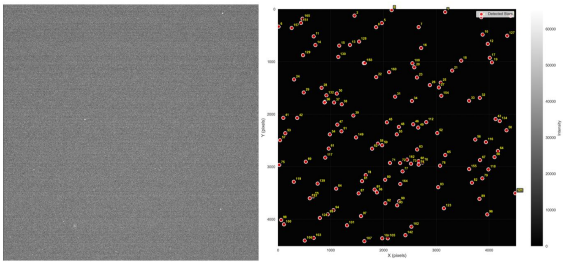

(a) (b)

Figure 8. Real Data Validation

Figure 8 shows a sample star field image captured by this satellite showing 105 detected stars with varying brightness and SNR conditions. Ground truth established via ensemble averaging of 50 consecutive frames with spacecraft attitude stabilized to <0.1 arcsec/s drift rate, yielding reference centroid with uncertainty ±0.003 pixels. BMAI achieves 45.5% improvement over GSF and 82.1% over IWC. Real-world performance closely matches synthetic predictions (0.0120 pixels under similar SNR conditions), validating model fidelity.

5. CONCLUSION

This paper presents the Bayesian Multi-Scale Adaptive Iteratively Weighted (BMAI) centroid localization algorithm, which fundamentally resolves the accuracy-efficiency trade-off in star sensor technology. By synergistically integrating SNR-adaptive window extraction, Bayesian iterative refinement, multi-scale fusion, and gradient-based correction, BMAI achieves near-optimal statistical performance with rigorous theoretical convergence guarantees. Comprehensive evaluations validate its robust, real-time performance across diverse and challenging imaging conditions. Implemented in highly optimized C++, the algorithm provides a deployment-ready solution that pushes centroiding accuracy toward theoretical CRLB limits without compromising processing speed. This sub-pixel precision improvement directly benefits downstream spaceborne applications, including deep-space navigation and high-resolution Earth observation. Future work will extend the framework to asymmetric PSF models and advanced star pattern recognition.

REFERENCES

Asadnezhad, M., Eslamimajd, A., Hajghassem, H., 2018: Optical system design of star sensor and stray light analysis. *Journal of the European Optical Society-Rapid Publications*, 14(1).

Chen, M., Zhao, Y., Yang, W., 2024: A model for suppressing stray light in astronomical images based on deep learning. *Scientific Reports*, 14(1), 27521.

Cramér, H., 1946. *Mathematical Methods of Statistics*. Princeton University Press, Princeton, NJ.

Gou, W., Li, C., Li, X., 2025: Prediction-selection-extraction: star point extraction from daytime star maps based on the randomness of stellar imaging energy. *Optics Express*, 33(5), 10835.

Hartung, J., Knapp, G., Sinha, B.K., 2008. *Statistical Meta-Analysis with Applications*. Wiley, Hoboken, NJ.

Kotz, S., Johnson, N.L., 1992: *Breakthroughs in Statistics: Foundations and Basic Theory*. Springer, New York, NY.

Liebe, C.C., 2002: Accuracy performance of star trackers - a tutorial. *IEEE Transactions on Aerospace and Electronic Systems*, 38(2), 587-599.

Lu, K., Liu, E., Zhao, R., 2021: Star sensor denoising algorithm based on edge protection. *Sensors*, 21(16), 5255.

Luo, Z., Guo, Q., Feng, J., 2025: Centroid algorithm for high-dynamic star sensor based on key point detection deep learning algorithms. *Optics Express*, 33(8), 17203.

Mortari, D., Samaan, M.A., Bruccoleri, C., 2004: The Pyramid star identification technique. *Navigation*, 51(3), 171-183.

Rao, C.R., 1945. Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, 37(3), 81-91.

Rousseeuw, P.J., Croux, C., 1993: Alternatives to the median absolute deviation. *Journal of the American Statistical Association*, 88(424), 1273-1283.

Wan, X., Wang, G., Wei, X., 2021: ODCC: A dynamic star spots extraction method for star sensors. *IEEE Transactions on Instrumentation and Measurement*, 70, 1-14.

Wang, H., Xu, E., Li, Z., 2015: Gaussian analytic centroiding method of star image of star tracker. *Advances in Space Research*, 56(10), 2196-2205.

Wang, T., Qiu, Y., Cai, H., 2010: A fast onboard star-extraction algorithm optimized for the SVOM Visible Telescope. *Science China Physics, Mechanics and Astronomy*, 53(S1), 51-55.

Winick, K.A., 1986: Cramér-Rao lower bounds on the performance of charge-coupled-device optical position estimators. *Journal of the Optical Society of America A*, 3(11), 1809.

Xi, J., 2015: Rapid star extraction and high accuracy location in star image. *Journal of Information and Computational Science*, 12(8), 2929-2937.

Yu, W., Qu, H., Zhang, Y., 2023: A high-accuracy star centroid extraction method based on Kalman filter for multi-exposure imaging star sensors. *Sensors*, 23(18), 7823.

Zhang, Y., Jiang, J., Zhang, G., 2019: Accurate and robust synchronous extraction algorithm for star centroid and nearby celestial body edge. *IEEE Access*, 7, 126742-126752.

Zhou, F., Zhao, J., Ye, T., 2016: Fast star centroid extraction algorithm with sub-pixel accuracy based on FPGA. *Journal of Real-Time Image Processing*, 12(3), 613-622.

APPENDIX (Detailed Convergence Analysis)

Convergence Proof

Let the mapping $T: \mathbb{R}^2 \to \mathbb{R}^2$ be defined as $T(x,y) = \Big((1-\alpha)g_x(x,y) + \alpha x_0, (1-\alpha)g_y(x,y) + \alpha y_0\Big)$, where g_x and g_y are the weighted centroid functions. We prove that T is a contraction mapping in the l_2 norm. First, for any two points (x_1, y_1) and (x_2, y_2), we have $\left|\left|T(x_1,y_1) - T(x_2,y_2)\right|\right|_2 \leq (1-\alpha)\left|\left|g(x_1,y_1) - g(x_2,y_2)\right|\right|_2$. Using the exponential weight function a

nd the dominated convergence theorem, the Lipschitz constant o f g can be bounded such that $\left|\left|g(x_1, y_1) - g(x_2, y_2)\right|\right|_2 \leq L \left|\left|(x_1, y_1) - (x_2, y_2)\right|\right|_2$, where L depends on σ_w and the intensit y distribution. If $(1 - \alpha)L < 1$ (i.e., $L < 1/(1 - \alpha)$), then $\left|\left|T (x_1, y_1) - T(x_2, y_2)\right|\right|_2 < \left|\left|(x_1, y_1) - (x_2, y_2)\right|\right|_2$, establishing T as a contraction. By the Banach Fixed-Point Theorem, the iterat ive sequence converges geometrically to a unique fixed point wi th the rate $\epsilon_k \leq \gamma^k \epsilon_0$, where $\gamma = (1 - \alpha)L < 1$ and ϵ_k is the err or at iteration k.

This contribution has been peer-reviewed. The double-blind peer-review was conducted on the basis of the full paper.
https://doi.org/10.5194/isprs-annals-XI-1-2026-183-2026 | © Author(s) 2026. CC BY 4.0 License.