Automated Astrometric Calibration of Astronomical Images with Gaia DR3 and Residual Error Decomposition

Mohammed A. H. Abozied*[1], Mohamed H. Mahgoub[1], Abdelhameed S. Eltanany[2] & Hossam Hendy[3]

[1]Guidance Department, Military Technical College, Cairo, Egypt. [2]Research Center, Cairo, Egypt. [3]Space Research Center

Abstract. Precise astrometric validation requires residuals between catalog-predicted positions and measured stellar locations, not between a prediction and its nearest integer pixel. This work presents a deterministic Flexible Image Transport System (FITS) calibration and residual-analysis pipeline that combines median-background suppression, blind plate solving with Astrometry.net, Gaia Data Release 3 (DR3) reference-star projection, Astrometry.net source-list association, bounded two-dimensional Gaussian centroid refinement, robust outlier rejection, and an affine decomposition of systematic and random residual components. Four satellite-tracking FITS frames (GAOFEN 11-3, NIGERIASAT-2, THEOS, and YAOGAN-50) were processed independently in raw and background-suppressed branches. Astrometry.net solved all eight branches. The raw images, however, produced only 3, 6, 1, and 4 validated Gaussian matches, leaving three of four raw branches below the six-match statistical threshold. Background suppression increased the corresponding match counts to 14, 43, 30, and 53; all four suppressed branches therefore supported full residual analysis. Suppressed-branch Gaia matching completeness ranged from 42.4% to 75.0% (median 43.2%). Measured total root-mean-square (RMS) residuals were 1.589–3.932 arcsec (median 3.153 arcsec), while post-affine random RMS ranged from 0.447 to 3.196 arcsec (median 2.731 arcsec). The former nearest-pixel metric yielded a much narrower 1.555–1.704 arcsec range and underestimated measured RMS in three of four fields, demonstrating why it cannot serve as an astrometric-accuracy estimator. The results show that background suppression primarily improves usable source-measurement yield in this data set, while the systematic–random decomposition provides a more defensible characterization of calibration quality. Because the local Gaia mirror lacks proper-motion and formal-error columns and the demonstration contains only four frames, the reported values are treated as a pilot validation rather than a population-wide accuracy claim.

Keywords: astrometry; Gaia DR3; FITS; Astrometry.net; sub-pixel centroiding; residual analysis.

1 INTRODUCTION

Astrometric calibration establishes the transformation between detector coordinates and celestial coordinates. For satellite tracking and other time-sensitive astronomical applications, the usefulness of a solved World Coordinate System (WCS) depends not only on obtaining a plate solution but also on quantifying how well catalog-predicted stellar positions agree with measured image centroids. Astrometry.net provides robust blind astrometric calibration by recognizing geometric star patterns and returning a standards-compatible WCS [1]. Gaia DR3 provides a dense, high-precision astrometric reference frame with positions, parallaxes, and proper motions for approximately 1.46 billion sources [2], [3]. Together, these resources support automated calibration workflows for heterogeneous optical imagery.

A central difficulty in validation is defining the residual correctly. A previous implementation projected each Gaia position through the solved WCS and then computed the difference between the floating-point pixel coordinate and its nearest integer pixel. That quantity is bounded by construction and therefore measures pixel quantization rather than the location error between a catalog prediction and an observed star. It can remain small even if the physical source is displaced, distorted, poorly centroided, or mismatched. The present work replaces that baseline with measured sub-pixel centroids and explicitly retains the old quantity only as a legacy diagnostic.

[1]* Corresponding author Mohammed A. H. Abozied, Military Technical College, Guidance Department, E-mail mohammed.abozied@mtc.edu.eg

The second objective is to separate spatially coherent calibration error from local measurement scatter. Aggregate RMS alone mixes translation, plate-scale or shear trends, residual optical distortion, and centroid noise. We therefore fit a six-parameter affine residual field and report both the measured total residual and the post-affine random component. The resulting method is designed as a deterministic and auditable analysis pipeline rather than a learned image-restoration system.

1.1 Contributions

This upgraded study makes four concrete contributions. First, it replaces nearest-pixel rounding by source-associated two-dimensional Gaussian centroid measurements. Second, it uses the source list generated by Astrometry.net during solving as the seed for Gaia association, avoiding ambiguous brightest-peak searches in noisy tracking frames. Third, it evaluates raw and background-suppressed images independently and quantifies both source-yield and residual effects. Fourth, it augments mean/RMS reporting with robust clipping, bootstrap confidence intervals, an affine systematic–random decomposition, and formal consistency diagnostics. The evaluation is deliberately limited to four frames so that claims remain tied to observed evidence.

2 RELATED WORK

Lang et al. introduced Astrometry.net as a blind calibration system that detects sources, forms geometric asterisms, searches indexed sky patterns, and returns image pointing, scale, and orientation without requiring an initial WCS [1]. Flexible Image Transport System (FITS) WCS conventions and celestial projections are standardized by Greisen and Calabretta and by Calabretta and Greisen [4], [5], providing the mathematical basis for converting Gaia celestial coordinates into detector coordinates. Astropy implements these FITS/WCS transformations and related coordinate utilities in an open scientific-Python framework [6].

Gaia DR3 builds on the astrometric solution of Gaia EDR3 and the Gaia mission reference framework, providing the reference data required for high-precision catalog validation [2], [3], [7]. The official release documentation states a reference epoch of 2016.0 for DR3 astrometry [2]. This is important for later-epoch imagery because high-proper-motion sources should be propagated to the observation epoch when proper-motion information is available. Source detection and centroid measurement are established components of astronomical image analysis; SExtractor is a classical example of automated detection and measurement software [8], while point-spread-function (PSF)-aware stellar photometry and high-precision astrometry provide established foundations for sub-pixel position estimation [9], [10]. In the present implementation, local centroid parameters are obtained by bounded nonlinear least squares using SciPy [11], while array processing is performed with NumPy [12].

The distinguishing focus of this work is not blind solving itself but validation after solving. The pipeline asks whether a Gaia source projected through the WCS has a corresponding measured image source and, if so, what residual remains after sub-pixel centroiding and after removal of a fitted spatially coherent residual field.

3 DATA AND REPRODUCIBILITY

3.1 FITS Demonstration Set

The evaluation uses four satellite-tracking FITS frames acquired on 23–24 August 2026 and identified by the tracked objects GAOFEN 11-3, NIGERIASAT-2, THEOS, and YAOGAN-50. Each frame is processed twice from the same input data: once in its raw normalized representation and once after deterministic median-background suppression. The target identity is used only as an image label; the astrometric evaluation is performed on background stars rather than on the satellite streak. The demonstration set is summarized in Table 1.

Table 1. Four-frame demonstration set and solved-WCS scale.

Frame	UTC	Gaia	Scale (″/px)
G11-3	08-23 18:22	33	3.877
NIG-2	08-24 01:23	98	3.872
THEOS	08-23 19:47	40	3.874
YAO-50	08-24 02:00	125	3.871

3.2 Reproducible Configuration

The run used Python 3.12.4 with NumPy 2.5.1, SciPy 1.18.0, Astropy 8.0.1, Matplotlib 3.8.4, and OpenCV 5.0.0. Astrometry.net was constrained to a 3.7–3.9 arcsec/pixel scale interval with a 10° search radius and downsample factor 2. Gaia sources were queried from a local SQLite mirror with G < 10.0. The code recorded SHA-256 digest 4f7fd1c129f433e4089d321532d0ca9a6e7580fd027bfbb2e8b8f6c95b3f0918. Paths are machine-specific and are excluded from the scientific configuration. The principal evaluation settings are listed in Table 2.

Table 2. Principal evaluation configuration.

Parameter	Value
Gaia magnitude limit	G < 10.0
Plate scale search	3.7–3.9 arcsec/pix
Background filter	median k=7
Background gain / clamp	2.5 / 50
Astrometry-source match radius	5 pix
Gaussian centroid window	half-width 5 pix
Minimum centroid signal-to-noise ratio (SNR)	3.0
Robust rejection	iterative 3σrobust
Minimum matches for statistics	6
Bootstrap	2000 resamples

The local Gaia table contained right ascension, declination, and G magnitude, but not source identifiers, proper motions, parallaxes, or formal position uncertainties. Consequently, proper-motion propagation and Gaia formal-error propagation were not applied in this run. This limitation is carried explicitly into the interpretation in Section 8.

3.3 Observation Epoch and Reference-Epoch Handling

Observation epochs were parsed from the FITS filenames and recorded as UTC timestamps. The four frames span approximately eight hours across two nights. Gaia DR3 astrometry is referenced to epoch 2016.0, whereas the present frames were acquired in 2026. In a complete implementation, Gaia positions should be propagated to the observation epoch using proper motion (and, where relevant, parallax and radial-velocity information) before WCS projection. Because the local database supplied only ra, dec, and G, the current run necessarily evaluates catalog coordinates without that propagation. The effect is expected to be negligible for many low-proper-motion stars but can be material for individual high-proper-motion references; this is one reason that the measured residuals are not interpreted as a pure detector/WCS noise floor.

The pipeline records whether proper-motion propagation was applied on a per-source basis. For this run the flag was false for every source. This explicit provenance avoids silently mixing propagated and unpropagated reference positions and provides a direct migration path when the Gaia mirror is upgraded.

4 METHODOLOGY

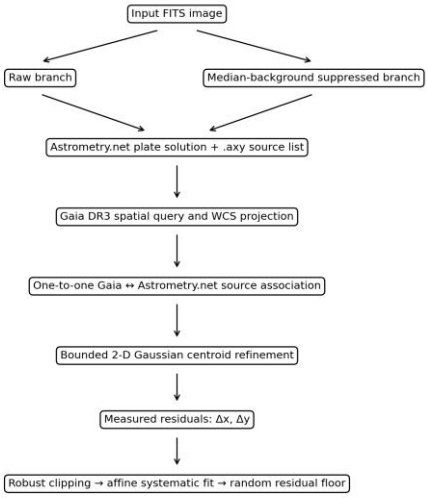


Fig. 1. Deterministic dual-branch calibration and residual-analysis workflow.

4.1 Background Suppression

Let I(x,y) be the normalized image intensity. A low-frequency background model B(x,y) is estimated with a median filter and clamped at an intensity ceiling Imax. The suppressed image is formed by subtracting a scaled background model: The complete processing workflow is shown in Fig. 1.

$$\hat{B}(x,y) = min[MedianFilter(I,k), Imax] \quad (1)$$

$$I'(x,y) = I(x,y) - G\,\hat{B}(x,y) \quad (2)$$

For the evaluated configuration, k=7, G=2.5, and Imax=50 in the normalized 8-bit working representation. The raw and suppressed images are then solved independently, rather than using suppression only as a fallback. This permits direct comparison of source-yield and residual behavior while preserving the operational distinction between plate-solving success and measurement quality.

4.2 WCS Solution and Gaia Projection

Astrometry.net solves each branch and writes a WCS along with an .axy table of sources extracted during plate solving. Gaia sources are queried over the solved field of view and projected through the Astropy WCS transformation. Only catalog projections falling within the detector boundaries are retained. Pixel scale is derived from the solved WCS rather than from a fixed hard-coded conversion, consistent with FITS celestial-coordinate conventions [4]–[6].

4.3 Astrometry-Source-Seeded Association

The final association begins from Astrometry.net source coordinates rather than from a global brightest-peak detector. For each Gaia projection, the nearest Astrometry.net source is considered within a 5-pixel radius. One-to-one uniqueness is enforced so that the same detected source cannot be assigned to multiple Gaia references. Because Astropy uses zero-based pixel coordinates while FITS-derived source tables may use a one-based convention, the implementation evaluates the candidate origin shifts 0 and −1 pixel and retains the convention producing the strongest physically valid association. In all eight branches of the present run, the selected Astrometry.net source shift was −1 pixel. Algorithm 1 summarizes the complete automated procedure.

4.4 Two-Dimensional Gaussian Centroid Refinement

Each associated Astrometry.net source seeds a local bounded fit of an elliptical two-dimensional Gaussian plus background. With local image coordinates (x,y), the model is

$$\hat{I}(x,y)=A\,exp\{-½[a(x-\hat{x})^2+2b(x-\hat{x})(y-\hat{y})+c(y-\hat{y})^2]\}+D \quad (3)$$

where $(\hat{x},\hat{y})$ is the measured sub-pixel centroid, A is source amplitude, D is local background, and the quadratic terms parameterize the PSF shape. PSF-based stellar-position measurement is well established in precision astronomical photometry and astrometry [9], [10]. Fits are constrained in centroid shift and width, require SNR ≥ 3, and reject extreme PSF axis ratios. The measured residual is then formed against the unmodified Gaia/WCS projection:

$$\Delta x = \hat{x} - xGaia\ , \quad \Delta y = \hat{y} - yGaia\ , \quad r = \sqrt{(\Delta x^2+\Delta y^2)} \tag{4}$$

Conversion to arcseconds uses the solved WCS pixel scales along each detector axis. This is the primary astrometric residual analyzed in the remainder of the paper.

4.5 Robust Clipping and Systematic–Random Decomposition

Spurious matches and poorly fitted centroids are removed by iterative clipping based on the median absolute deviation (MAD), a robust scale statistic widely used for contamination-resistant estimation [13], [14]. For residual magnitude r, $\sigma robust=1.4826\cdot MAD(r)$, with a $3\sigma robust$ criterion iterated up to five times. The surviving residual vectors are modeled by a six-parameter affine field centered on the mean source position:

$$[\Delta x,\Delta y]sys^T = [tx,ty]^T + M\,[x-\bar{x},y-\bar{y}]^T \tag{5}$$

$$[\Delta x,\Delta y]rand^T = [\Delta x,\Delta y]^T - [\Delta x,\Delta y]sys^T \tag{6}$$

The translation norm quantifies a field-wide offset, while the Frobenius norm of M quantifies first-order residual gradients. The RMS of the post-affine random vectors is reported as the random residual floor. A nonparametric 2000-resample bootstrap [15] provides 95% intervals for total and random RMS. A formal reduced-χ^2 diagnostic is also computed from centroid-fit covariance, but because Gaia formal errors and the full WCS/atmospheric covariance are unavailable, χ^2 is interpreted only qualitatively.

4.6 Legacy Nearest-Pixel Baseline

For continuity with the earlier implementation, the pipeline also reports

$$rlegacy = s\,\sqrt{\{[xGaia-round(xGaia)]^2 + [yGaia-round(yGaia)]^2\}} \tag{7}$$

where s is pixel scale. This quantity contains no measured source coordinate. It is therefore labeled legacy_quantization_rms_arcsec and is retained only to demonstrate how the former validation statistic differs from a true image-to-catalog residual.

4.7 Algorithmic Specification

Algorithm 1. Automated Gaia-referenced astrometric validation pipeline.

Step	Operation
1	Normalize FITS image for raw working branch.
2	Estimate median background; form suppressed branch.
3	Run Astrometry.net independently on raw and suppressed images.
4	Read solved WCS and Astrometry.net .axy source coordinates.
5	Derive pixel scale from the solved WCS.
6	Query Gaia DR3 over the solved detector footprint with $G < 10$.
7	Project Gaia coordinates to detector pixels; discard off-detector sources.
8	Test Astrometry.net source-origin convention (0 or −1 pixel).
9	Build one-to-one Gaia↔.axy associations within 5 pixels.
10	Fit bounded 2-D Gaussian centroids around associated sources.
11	Reject low-SNR, extreme-shape, duplicated, or displaced fits.
12	Compute measured Δx, Δy and radial residual in arcseconds.
13	Apply iterative MAD-based 3σ robust clipping.
14	Fit six-parameter affine systematic residual field.
15	Compute post-affine random residuals and bootstrap RMS intervals.
16	Archive per-star, per-image, and batch-level diagnostics.

4.8 Computational Structure

The computationally expensive blind-recognition task is delegated to Astrometry.net, after which association is local. Let M be the number of Gaia sources in the field and N the number of Astrometry.net source candidates. A nearest-neighbor data structure reduces source association from a naive $O(MN)$ pairwise scan to approximately $O(N \log N + M \log N)$, while each accepted association incurs only a small fixed-size nonlinear centroid fit. The affine residual model is linear in six parameters and therefore negligible compared with image preprocessing, plate solving, and repeated Gaussian fitting. This structure is suitable for batch processing because every FITS frame and every raw/suppressed branch can be analyzed independently.

5 EVALUATION METRICS

The evaluation separates plate-solving performance from residual-analysis performance. Solve success is the fraction of branches for which Astrometry.net returns a valid WCS. Matching completeness is $N_{matched}/N_{Gaia,FOV}$. Match precision after clipping is $N_{inlier}/N_{matched}$. Aggregate astrometric accuracy is summarized by mean, median, and RMS measured residuals. Systematic error is represented by the translation norm and the affine-gradient norm, while random error is represented by post-affine RMS. Bootstrap intervals quantify sampling uncertainty within each field. The raw and suppressed branches are evaluated independently with the same Gaia magnitude limit and source-association logic.

6 RESULTS

6.1 Plate Solving and Validated Source Yield

Astrometry.net solved all four raw frames and all four background-suppressed frames, giving 100% plate-solving success for both branches in this small data set. The decisive difference appeared after solving. Raw frames produced only 3, 6, 1, and 4 validated Gaussian centroid matches for GAOFEN 11-3,

NIGERIASAT-2, THEOS, and YAOGAN-50, respectively. With a six-match minimum for full statistics, only the raw NIGERIASAT-2 branch was analyzable. The suppressed branches produced 14, 43, 30, and 53 matches, so all four passed the statistical threshold. The branch-level diagnostics are reported in Table 3, and representative suppressed-frame overlays are shown in Figs. 2-5.

Table 3. Raw/suppressed plate-solve and source-association diagnostics. Frame codes: G11 = GAOFEN 11-3, NIG = NIGERIASAT-2, Y50 = YAOGAN-50. R/S = raw/suppressed; Solve = solve time (s); .axy = extracted sources; Seed = Gaia–.axy seed pairs; Match = validated Gaussian matches.

Frame	Solve R/S s	.axy R/S	Seed R/S	Match R/S
G11	23.3/23.0	654/1331	32/31	3/14
NIG	7.4/4.7	1065/1808	97/97	6/43
THEOS	7.3/5.1	561/1702	39/39	1/30
Y50	8.4/5.3	1972/2801	124/124	4/53

The .axy source lists themselves were not sparse: raw/suppressed source counts were 654/1331, 1065/1808, 561/1702, and 1972/2801. Gaia-to-.axy seed pairs were likewise abundant (31–124 per suppressed frame). The main attrition therefore occurred during Gaussian quality validation rather than during WCS solving or initial geometric association. This observation is important for future optimization because loosening the Gaia association radius would not address the principal loss mechanism.

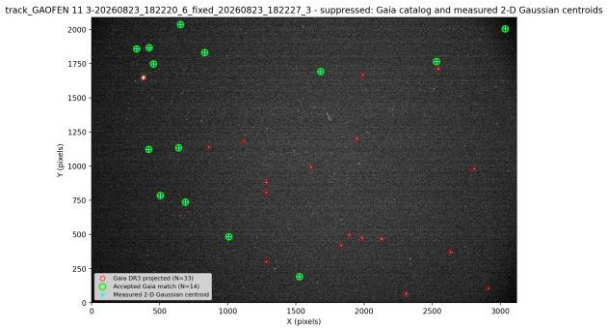


Fig. 2. GAOFEN 11-3 suppressed frame: Gaia DR3 projections (red), validated matches (green), and measured 2-D Gaussian centroids (cyan); Nmatch=14.

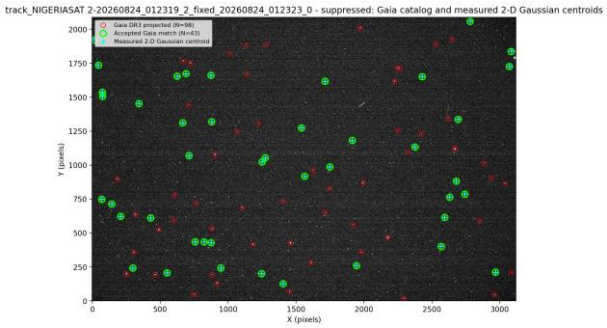


Fig. 3. NIGERIASAT-2 suppressed frame: Gaia DR3 projections (red), validated matches (green), and measured 2-D Gaussian centroids (cyan); Nmatch=43.

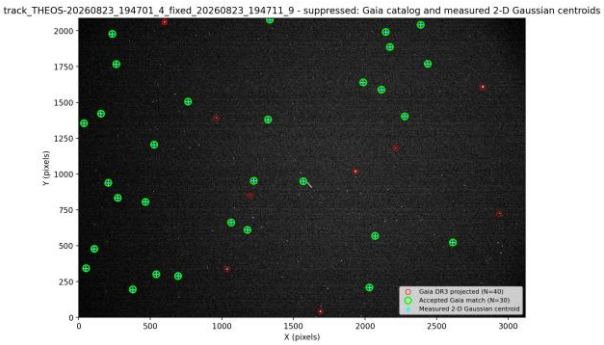


Fig. 4. THEOS suppressed frame: Gaia DR3 projections (red), validated matches (green), and measured 2-D Gaussian centroids (cyan); Nmatch=30.

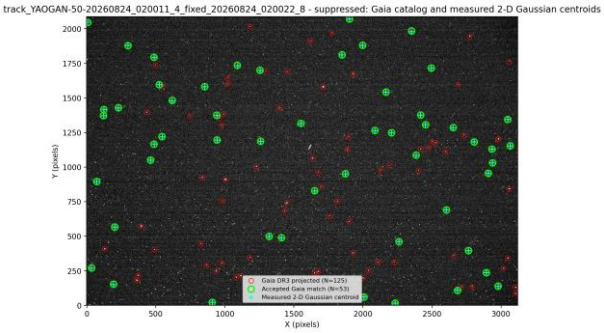


Fig. 5. YAOGAN-50 suppressed frame: Gaia DR3 projections (red), validated matches (green), and measured 2-D Gaussian centroids (cyan); Nmatch=53.

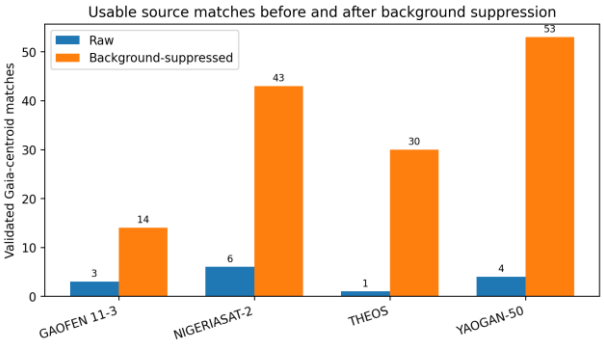


Fig. 6. Validated Gaia-centroid matches in raw and background-suppressed branches.

The median match count increased from 3.5 in raw images to 36.5 after suppression, a factor of 10.4. Suppressed-branch completeness was 42.4%, 43.9%, 75.0%, and 42.4%, with a median of 43.2%. The corresponding inlier counts after robust clipping were 6, 35, 29, and 50. These results indicate that, for the evaluated tracking frames, suppression did not improve the already-successful blind WCS solve rate but did substantially increase the number of sources that could support stable Gaussian centroid measurement. Figures 6 and 7 summarize the source-yield improvement and matching completeness.

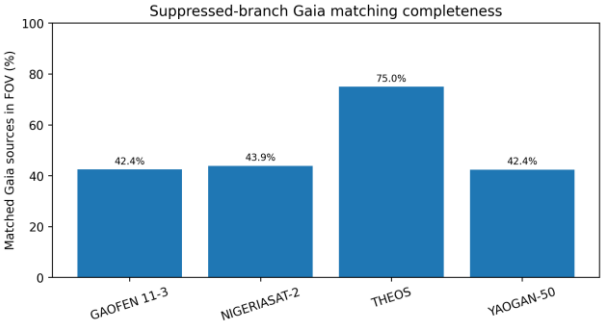


Fig. 7. Fraction of Gaia G < 10 sources in the solved field that produced validated suppressed-branch centroid matches.

6.2 Measured Astrometric Residuals

Table 4. Suppressed-branch measured residual statistics. M = validated matches; In = inliers after clipping; C% = Gaia completeness; RMS = total RMS; Rnd = post-affine random RMS; Sh = fitted affine translation magnitude.

Frame	M	In	C%	RMS″	Rnd″	Sh″
G11	14	6	42.4	1.589	0.447	1.452
NIG	43	35	43.9	2.935	2.565	0.848
THEOS	30	29	75.0	3.370	2.896	1.701
Y50	53	50	42.4	3.932	3.196	1.600

Measured total RMS ranged from 1.589 arcsec for GAOFEN 11-3 to 3.932 arcsec for YAOGAN-50; the four-frame median was 3.153 arcsec. NIGERIASAT-2 and THEOS produced total RMS values of 2.935 and 3.370 arcsec. At solved pixel scales near 3.87–3.88 arcsec/pixel, these values correspond to approximately 0.41–1.02 pixel RMS. The 95% bootstrap intervals were 1.217–1.914 arcsec for GAOFEN 11-3, 2.513–3.326 arcsec for NIGERIASAT-2, 2.806–3.949 arcsec for THEOS, and 3.341–4.543 arcsec for YAOGAN-50. Table 4 reports the suppressed-branch statistics, and Fig. 8 compares total and post-affine random RMS.

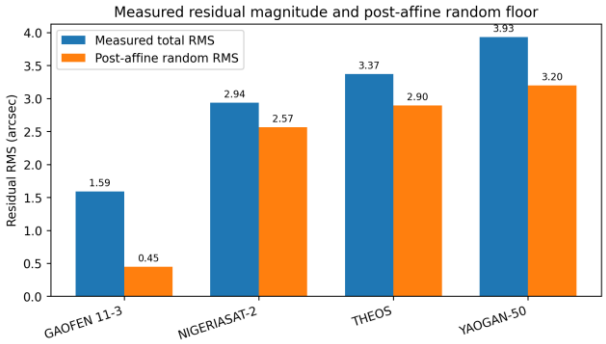


Fig. 8. Measured total RMS and residual RMS after removal of the fitted affine systematic component.

6.3 Systematic and Random Components

The fitted translation magnitude ranged from 0.848 to 1.701 arcsec. GAOFEN 11-3 showed the strongest benefit from affine decomposition: total RMS 1.589 arcsec was reduced to a random RMS of 0.447 arcsec, a 71.9% reduction in RMS magnitude. The other frames retained larger random components: 2.565 arcsec for NIGERIASAT-2, 2.896 arcsec for THEOS, and 3.196 arcsec for YAOGAN-50. Their total-to-random reductions were 12.6%, 14.0%, and 18.7%, respectively. Thus a first-order coherent field component explains much of the GAOFEN residual structure, whereas the remaining three frames are dominated by residual scatter not removed by a six-parameter affine model. A representative post-affine vector field is shown in Fig. 9.

The affine-gradient Frobenius magnitudes were small in absolute terms (4.62×10^{-4} to 2.08×10^{-3} arcsec/pixel) but act across detector scales of several thousand pixels. These terms are diagnostic of low-order spatial structure; they should not be interpreted as a complete optical-distortion solution because no higher-order SIP refit was performed.

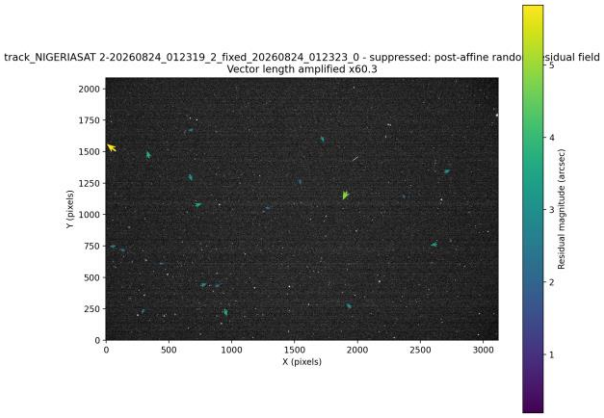


Fig. 9. NIGERIASAT-2 post-affine random residual field; vectors are amplified for visibility and colored by residual magnitude.

6.4 Per-Frame Interpretation

GAOFEN 11-3 produced the smallest suppressed total RMS and the smallest post-affine random RMS. Fourteen catalog associations survived Gaussian fitting, but only six remained after robust clipping. The reduction from 1.589 to 0.447 arcsec after affine removal indicates a strong coherent field component relative to the remaining random scatter. Because the inlier count is small, the result is best viewed as evidence of a removable low-order offset rather than a definitive estimate of the instrument noise floor.

NIGERIASAT-2 provided the strongest raw-versus-suppressed comparison because both branches met the six-match threshold. The raw branch contained only six validated matches and exhibited 7.441 arcsec total RMS. Suppression increased the match set to 43 and produced 35 inliers, while total RMS fell to 2.935 arcsec. The post-affine random term remained near 2.56 arcsec, showing that the large raw error was largely associated with sparse measurement and coherent field displacement rather than a dramatic change in the residual random floor.

THEOS yielded the highest catalog completeness of the four suppressed frames: 30 of 40 Gaia sources in the detector field were measured, with 29 surviving clipping. Its 3.370 arcsec total RMS and 2.896 arcsec random RMS imply that most of the residual was not removed by the first-order affine model. This frame therefore provides the clearest demonstration that high match completeness does not automatically imply a low residual floor.

YAOGAN-50 contained the densest bright-star field in the demonstration, with 125 Gaia G < 10 projections. Fifty-three produced validated Gaussian centroids and 50 remained inliers. The total RMS of 3.932 arcsec was the largest suppressed-branch value, while the post-affine RMS was 3.196 arcsec. The large source count makes this field particularly useful for future tests of higher-order distortion and spatially varying PSF models.

6.5 Legacy Quantization Metric Versus Measured Error

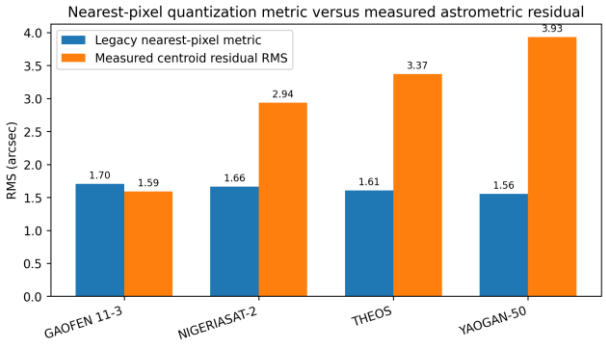


Fig. 10. Legacy nearest-pixel quantization RMS compared with measured Gaia-to-centroid residual RMS.

The legacy nearest-pixel values clustered tightly between 1.555 and 1.704 arcsec, with median 1.636 arcsec. In contrast, measured suppressed-branch RMS spanned 1.589–3.932 arcsec. The measured-to-legacy

RMS ratios were 0.93, 1.76, 2.10, and 2.53 for GAOFEN 11-3, NIGERIASAT-2, THEOS, and YAOGAN-50, respectively; the median ratio was 1.93. The legacy value happened to be slightly larger than the measured total RMS for GAOFEN, but it substantially underestimated the measured residual in the other three fields. This behavior is expected because rounding error is constrained by pixel phase and does not respond to physical image displacement. Figure 10 compares the legacy quantization metric with the measured centroid residuals.

6.6 Raw–Suppressed Paired Case

Only NIGERIASAT-2 supplied at least six matches in both branches, so it is the sole frame supporting a direct raw-versus-suppressed residual comparison under the fixed threshold. Raw total RMS was 7.441 arcsec and dropped to 2.935 arcsec after suppression, while the post-affine random RMS changed only from 2.538 to 2.565 arcsec. The suppressed branch also increased matched sources from 6 to 43 and reduced the fitted systematic translation from 3.821 to 0.848 arcsec. This case suggests that the major gain was improved source measurement and reduced coherent mismatch rather than a lower stochastic centroid floor.

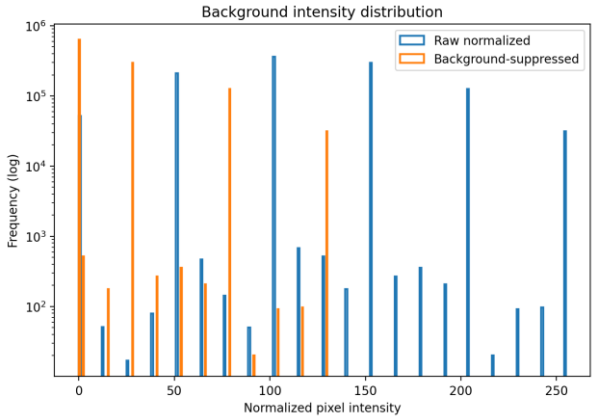


Fig. 11. NIGERIASAT-2 normalized intensity distributions before and after median-background suppression (log-frequency axis).

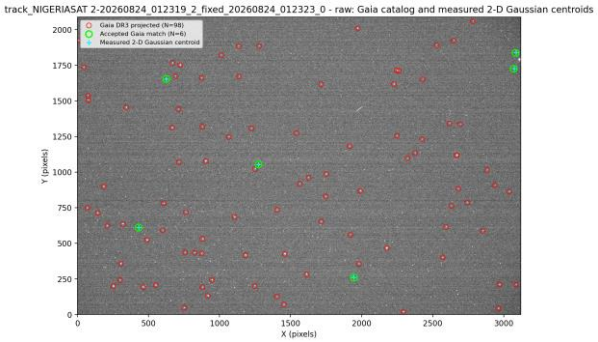


Fig. 12. NIGERIASAT-2 raw frame: 98 Gaia projections but only six validated Gaussian centroid matches; compare with the 43 suppressed-frame matches in Fig. 3.

The intensity histogram illustrates that the chosen preprocessing is not a simple global brightness scaling: the median-background operation redistributes low-frequency sky structure and changes the contrast landscape used by the local centroid fits. The raw overlay in Fig. 12 visually confirms that many Gaia locations are correctly projected inside the image even when the local Gaussian-quality criteria accept only a small subset. This is why catalog-overlay density and validated-centroid count must be reported separately. Figure 11 shows the corresponding raw and suppressed intensity distributions.

6.7 Formal Consistency Diagnostic

Reduced-χ^2 values for the four suppressed branches were 0.419, 7.080, 11.282, and 9.029. Three fields therefore show residual scatter far larger than the centroid-fit covariance model alone predicts, while GAOFEN is formally under-dispersed. These values should not be used as calibrated probabilities: the local Gaia mirror supplied neither formal astrometric uncertainties nor proper motions, and the covariance model

omits WCS-solution uncertainty, atmospheric refraction, timing, detector systematics, and PSF model mismatch. The diagnostic instead indicates that a centroid-only uncertainty model is incomplete for most evaluated frames. The confidence intervals and formal consistency diagnostics are reported in Table 5.

Table 5. Bootstrap intervals and formal reduced-χ^2 diagnostics.

Frame	Total RMS 95% CI″	Random RMS 95% CI″	$\chi^2\nu$
G11	1.22–1.91	0.30–0.57	0.42
NIG	2.51–3.33	2.12–3.02	7.08
THEOS	2.81–3.95	2.46–3.31	11.28
Y50	3.34–4.54	2.68–3.69	9.03

7 DISCUSSION

7.1 Solving Success Is Not the Same as Measurement Success

A notable result is that all raw images were successfully plate-solved even though three raw branches could not support six validated Gaussian matches. Astrometry.net can establish a global WCS from a subset of suitable source patterns, whereas the residual-analysis stage demands that individual Gaia projections be associated with stable local centroids under stricter shape and SNR constraints. The distinction is operationally important: a FITS image can possess a valid WCS yet still provide too few robust source measurements for a statistically meaningful residual field.

7.2 Why the Astrometry.net Source List Matters

The final pipeline uses Astrometry.net's extracted .axy coordinates as association seeds. This choice exploits information already generated by the successful solver and avoids searching a broad neighborhood for the brightest local peak, which can be unreliable in noisy, crowded, or streak-contaminated tracking imagery. In the current run, the seed-pair counts were 31–124 in suppressed branches before Gaussian-quality filtering, demonstrating that the plate solver had already identified many geometrically plausible sources. Gaussian refinement then acts as a measurement-quality stage rather than as a second blind detector.

7.3 Interpretation of Background Suppression

The data do not support the claim that suppression is required for plate solving, because both raw and suppressed solve rates were 4/4. They do support a narrower and useful claim: the evaluated median-filter suppression substantially increased the number of source associations that survived Gaussian centroid quality checks. Median validated matches increased by roughly an order of magnitude, and all four suppressed frames reached the statistical-analysis threshold. This distinction should be preserved in any operational implementation: preprocessing may be valuable even when WCS solving already succeeds, because downstream centroid precision and source-yield can remain limiting.

7.4 Centroid Yield Versus Catalog Density

The number of Gaia stars inside the detector is not the same as the number of stars that can be used for precision residual statistics. The suppressed branches contained 33, 98, 40, and 125 Gaia $G < 10$ projections, yet only 14, 43, 30, and 53 produced validated Gaussian centroids. Several mechanisms can reduce the usable fraction: faint or blended sources near the adopted magnitude threshold, local streak contamination, PSF asymmetry, low local SNR, saturation, edge effects, and disagreement between the Astrometry.net source list and the Gaia-projected location. The explicit completeness metric makes these losses visible instead of silently treating every catalog projection as a measurement.

This distinction also explains why overlay figures should show both populations. Red circles document all Gaia sources predicted within the detector, while green circles and cyan crosses identify the subset that is actually associated and measured. A visually dense catalog overlay is therefore not evidence that every marked star contributes a residual. Only the validated measured subset enters the statistical tables and affine fit.

7.5 Implications for Accuracy Claims

The measured results do not support a universal sub-arcsecond total astrometric-accuracy claim for these four frames. Total RMS is between 1.6 and 3.9 arcsec, and the post-affine random floor exceeds 2.5 arcsec in three fields. GAOFEN 11-3 is the exception, with a 0.447 arcsec post-affine random RMS, but that value is based on six inlier sources and should therefore be interpreted cautiously. The scientifically stronger contribution is methodological: the reported residuals now refer to actual measured source centroids, and the separation of total, systematic, and random components prevents a quantization artifact from being presented as calibration accuracy.

7.6 Relevance to Satellite-Tracking Astrometry

For optical satellite tracking, the background stars establish the celestial reference frame used to convert image detections into angular measurements; the quality of those angular measurements directly affects optical orbit determination and tracking performance [16]. A residual field of several arcseconds can propagate directly into right-ascension/declination uncertainty and may become a significant error source for orbit determination or reacquisition, particularly when the detector scale is only a few arcseconds per pixel. The current workflow provides a per-image quality-control layer: frames with poor match completeness, coherent translation, large random RMS, or incomplete uncertainty calibration can be flagged before their astrometry is passed to orbit-processing stages.

8 LIMITATIONS AND REQUIRED NEXT WORK

The first limitation is sample size. Four frames are sufficient to validate the corrected residual definition and demonstrate the source-yield effect, but they are not sufficient to claim instrument-wide or site-wide accuracy. A journal-level generalization study should repeat the same protocol on at least tens of frames spanning multiple nights, sky regions, background levels, seeing conditions, target elevations, and stellar densities. Raw-versus-suppressed paired residual statistics require enough raw branches to satisfy the same minimum-match rule.

Second, the local Gaia mirror lacks pmra, pmdec, ref_epoch, and formal uncertainty columns. Gaia DR3 astrometry is referenced to epoch 2016.0 [2], with the broader Gaia mission providing the astrometric framework on which these releases are built [7]. The present observations are in 2026, so high-proper-motion stars can accumulate non-negligible catalog-to-epoch displacement. The database should therefore be regenerated with source_id, ra, dec, pmra, pmdec, parallax when available, ref_epoch, and the associated uncertainties. Positions can then be propagated to the FITS observation epoch before WCS projection, and the χ^2 statistic can include a defensible Gaia uncertainty contribution.

Third, the affine model captures only first-order spatial structure. Residual fields that remain position-dependent after the affine fit should be tested against higher-order distortion models such as the Simple Imaging Polynomial (SIP) convention [17], provided the source count and field coverage are sufficient. Fourth, the current centroid covariance is model-based and does not include WCS-solution uncertainty, atmospheric differential refraction [18], timing, detector effects, or correlations introduced by preprocessing. A complete uncertainty budget should be estimated empirically using repeated exposures or independent reference solutions.

A further implementation limitation is that this evaluation performs detection and Gaussian centroiding on a normalized working image used for deterministic processing. For precision metrology, the next code revision should preserve detector-native linear-intensity values (after physically justified bias/dark/flat corrections where available) for centroid estimation while allowing separately scaled copies for visualization and solver assistance. This would avoid any possibility that normalization or clipping modifies the point-spread-function weights used by the centroid fit.

Finally, the present $G < 10$ cut intentionally limits the catalog to bright reference stars but also restricts the number of possible matches. Future evaluation should sweep the magnitude threshold and quantify the trade-off between source density, centroid SNR, completeness, and residual quality. Any threshold optimization must be evaluated on held-out frames to avoid selecting parameters solely because they improve the same images used to tune them.

9 QUALITY-CONTROL INTERPRETATION FOR OPERATIONAL USE

The pipeline is intended not only to produce a single RMS number but also to expose whether a solved image is suitable for downstream angular measurement. A practical quality-control decision should therefore be hierarchical. First, the plate solution must succeed. Second, the Gaia projection must contain a sufficient number of reference sources inside the detector. Third, enough of those references must produce one-to-one Gaussian centroids after shape and SNR checks. Fourth, robust clipping must leave enough inliers to support the affine model. Only after these conditions are met should total RMS, systematic translation, random RMS, and uncertainty diagnostics be interpreted. Table 6 summarizes the proposed image-level quality-control interpretation.

Table 6. Proposed image-level quality-control interpretation for the current pipeline.

Diagnostic	Purpose	Observed status
Plate solve	Required global WCS	8/8 branches solved
Min. measured matches	Enable residual statistics	Raw 1/4; suppressed 4/4
Completeness	Expose catalog→measurement yield	42.4–75.0% suppressed
Affine translation	Flag coherent field offset	0.848–1.701 arcsec
Random RMS	Estimate post-affine scatter	0.447–3.196 arcsec
Proper-motion provenance	Reference-epoch integrity	Unavailable in local DB
Formal χ^2	Uncertainty consistency	Diagnostic only; not calibrated

9.1 Suggested Status Classes

For batch operation, four status classes are useful. SOLVE_FAILURE denotes an image for which no valid WCS is returned. INSUFFICIENT_MATCHES denotes a valid WCS with fewer than the minimum number of measured Gaia centroids; three of four raw branches fall into this category. MEASURED_UNCALIBRATED denotes a branch with enough residual measurements but with an incomplete external uncertainty budget, which describes all four suppressed branches in the present experiment because proper-motion and Gaia formal-error columns are missing. A future MEASURED_CALIBRATED class should be reserved for runs in which catalog positions are propagated to the observation epoch and the combined Gaia, centroid, WCS, and environmental uncertainty model has been validated empirically.

9.2 Separation of Visualization and Measurement

The output figures deliberately use different symbols for catalog projection and accepted measurement. A red Gaia circle answers the geometric question "where should this catalog star fall under the solved WCS?" A green circle and cyan centroid answer the measurement question "was a corresponding source actually validated and centroided?" This separation prevents a dense catalog overlay from being mistaken for a dense residual sample. The distinction is particularly important for satellite-tracking frames, where streaks, nonuniform background, local detector artifacts, and saturation can make some reference locations unusable even when the global plate solution is correct.

9.3 Recommended Acceptance Policy for Future Campaigns

The thresholds used here—$G < 10$, a 5-pixel association radius, centroid SNR ≥ 3, and six matches for statistical evaluation—should be treated as reproducible test settings rather than universal operational gates. Future campaigns should choose acceptance thresholds only after multi-night validation. A defensible policy would require a minimum inlier count and field coverage, an upper bound on post-affine random RMS, a limit on systematic translation, and successful epoch propagation for the Gaia references. The accepted limits should be derived from held-out frames and, where possible, from an independent astrometric truth source. This approach prevents the quality gate from being tuned merely to reproduce favorable residuals on the same images used to develop the algorithm.

10 CONCLUSION

This work converts an earlier nearest-pixel residual diagnostic into a measured astrometric-validation pipeline. The final method combines independent raw/suppressed plate solving, Gaia DR3 projection, Astrometry.net source-list seeding, two-dimensional Gaussian centroid refinement, robust clipping, and affine systematic–random decomposition. Across four satellite-tracking frames, Astrometry.net solved all raw and suppressed images. Background suppression did not change the 100% solve rate, but it increased median validated Gaia-centroid matches from 3.5 to 36.5 and enabled full residual statistics for all four frames. Suppressed-branch total RMS was 1.589–3.932 arcsec, with post-affine random RMS of 0.447–3.196 arcsec. The former nearest-pixel metric remained near 1.6 arcsec and underestimated measured error in three of four fields.

The principal conclusion is therefore not that the system universally achieves sub-arcsecond accuracy. Rather, the revised pipeline produces an auditable measurement of actual Gaia-to-image residuals and makes explicit whether error is dominated by coherent plate structure or by residual local scatter. That distinction is essential before calibrated star positions are used as the angular reference for satellite tracking and orbit determination. Extension of the Gaia database with proper motion and formal uncertainties, followed by multi-night validation on a larger frame set, is the immediate next step.

FUNDING

This paper is funded (For publication costs) by Egyptian Knowledge Bank (EKB).

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest regarding the publication of this paper.

DATA AND CODE AVAILABILITY

The datasets used and/or analyses during the current study are available from the corresponding author on reasonable request.

ACKNOWLEDGEMENTS

The authors acknowledge the Military Technical College Space Research Center and the developers of Gaia, Astrometry.net, Astropy, SciPy, NumPy, and the broader open-source astronomical software ecosystem.

AUTHOR CONTRIBUTIONS

M. A. H. A., M. H. M., A. S. E. and H. H. have analyzed the data, revised the manuscript, and approved the final version.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to M.A.H.A. (email: mohammed.abozied@mtc.edu.eg).

REFERENCES

[1] D. Lang, D.W. Hogg, K. Mierle, M. Blanton, S. Roweis, Astrometry.net: blind astrometric calibration of arbitrary astronomical images, Astron. J. 139 (2010) 1782-1800. https://doi.org/10.1088/0004-6256/139/5/1782.

[2] Gaia Collaboration, Gaia Data Release 3: summary of the content and survey properties, Astron. Astrophys. 674 (2023) A1. https://doi.org/10.1051/0004-6361/202243940.

[3] L. Lindegren, et al., Gaia Early Data Release 3: the astrometric solution, Astron. Astrophys. 649 (2021) A2. https://doi.org/10.1051/0004-6361/202039709.

[4] E.W. Greisen, M.R. Calabretta, Representations of world coordinates in FITS, Astron. Astrophys. 395 (2002) 1061-1075. https://doi.org/10.1051/0004-6361:20021326.

[5] M.R. Calabretta, E.W. Greisen, Representations of celestial coordinates in FITS, Astron. Astrophys. 395 (2002) 1077-1122. https://doi.org/10.1051/0004-6361:20021327.

[6] Astropy Collaboration, The Astropy Project: sustaining and growing a community-oriented open-source project and the latest major release of the core package, Astrophys. J. 935 (2022) 167. https://doi.org/10.3847/1538-4357/ac7c74.

[7] T. Prusti, et al., The Gaia mission, Astron. Astrophys. 595 (2016) A1. https://doi.org/10.1051/0004-6361/201629272.

[8] E. Bertin, S. Arnouts, SExtractor: software for source extraction, Astron. Astrophys. Suppl. Ser. 117 (1996) 393-404. https://doi.org/10.1051/aas:1996164.

[9] P.B. Stetson, DAOPHOT: a computer program for crowded-field stellar photometry, Publ. Astron. Soc. Pac. 99 (1987) 191-222. https://doi.org/10.1086/131977.

[10] J. Anderson, I.R. King, Toward high-precision astrometry with WFPC2. I. Deriving an accurate point-spread function, Publ. Astron. Soc. Pac. 112 (2000) 1360-1382. https://doi.org/10.1086/316632.

[11] P. Virtanen, et al., SciPy 1.0: fundamental algorithms for scientific computing in Python, Nat. Methods 17 (2020) 261-272. https://doi.org/10.1038/s41592-019-0686-2.

[12] C.R. Harris, et al., Array programming with NumPy, Nature 585 (2020) 357-362. https://doi.org/10.1038/s41586-020-2649-2.

[13] P.J. Rousseeuw, C. Croux, Alternatives to the median absolute deviation, J. Am. Stat. Assoc. 88 (1993) 1273-1283. https://doi.org/10.1080/01621459.1993.10476408.

[14] P.J. Huber, E.M. Ronchetti, Robust Statistics, second ed., Wiley, Hoboken, NJ, 2009. https://doi.org/10.1002/9780470434697.

[15] B. Efron, Bootstrap methods: another look at the jackknife, Ann. Stat. 7 (1979) 1-26. https://doi.org/10.1214/aos/1176344552.

[16] E. Lee, S.-Y. Park, B. Shin, S. Cho, E.-J. Choi, J. Jo, J.-H. Park, Orbit determination of KOMPSAT-1 and Cryosat-2 satellites using optical wide-field patrol network data with batch least squares filter, J. Astron. Space Sci. 34 (2017) 19-30. https://doi.org/10.5140/JASS.2017.34.1.19.

[17] D.L. Shupe, M. Moshir, J. Li, D. Makovoz, R. Narron, R.N. Hook, The SIP convention for representing distortion in FITS image headers, in: Astronomical Data Analysis Software and Systems XIV, ASP Conf. Ser. 347 (2005) 491-495.

[18] R.C. Stone, An accurate method for computing atmospheric refraction, Publ. Astron. Soc. Pac. 108 (1996) 1051-1058.