

Temporal Distribution Star Detection in Degraded Infrared Video

Per-Pixel Z-Scoring Against Cardinal Neighbors

Heath Hunnicutt Mavchin

Ruach Tov Collective

Abstract

We present a method for detecting stars in severely degraded infrared video by exploiting the temporal distribution of per-pixel brightness across aligned frames. Rather than applying spatial filters or template matching to individual frames, we compute a per-pixel z-score that compares each pixel’s temporal median against the temporal statistics of its cardinal neighbors at a fixed distance. The method is parameter-light (one distance parameter D), requires no PSF model, and naturally suppresses vignetting, sensor noise, and transient artifacts. Applied to 296 frames of publicly-released FLIR video (PR46), the method detects 15 of 16 visually-confirmed stars at z-scores of $13\text{--}29\sigma$, from data where individual frames show no discernible stars above the noise floor.

1 Introduction

Detecting point sources in degraded imagery is a fundamental problem in observational astronomy, surveillance, and remote sensing. When the degradation includes sensor noise, optical aberration, atmospheric turbulence, and lossy compression, single-frame detection methods struggle because the signal-to-noise ratio (SNR) of individual point sources may be well below unity.

Multi-frame approaches improve SNR by $\sqrt{N}$ through averaging, but require accurate frame alignment and assume stationary backgrounds. Spatial filtering (matched filters, wavelets) requires knowledge of the point spread function (PSF), which may be unknown or spatially varying.

We propose a method that exploits neither spatial structure nor PSF knowledge, but instead uses the **temporal distribution** of each pixel across N aligned frames. A star is a **temporally persistent, locally anomalous** brightness—it appears in the same pixel across all frames (persistent) and is brighter than its spatial neighbors (anomalous). Transient noise, by contrast, varies across frames and does not systematically exceed its neighbors.

2 Method

2.1 Input and Alignment

The input is a sequence of N aligned frames $\{I_1, I_2, \dots, I_N\}$, each of size $H \times W$ pixels. Frame alignment is performed via cross-correlation with sub-pixel refinement prior to this stage.

2.2 Per-Pixel Temporal Statistics

For each pixel position (x, y) , we compute the temporal median and standard deviation across all N frames:

$$\mu(x, y) = \text{median}\{I_k(x, y) : k = 1, \dots, N\} \quad (1)$$

$$\sigma(x, y) = \text{std}\{I_k(x, y) : k = 1, \dots, N\} \quad (2)$$

The median is preferred over the mean for robustness to outliers (hot pixels, cosmic rays, transient artifacts). With $N = 296$ frames, each pixel has 296 independent brightness samples—a substantial statistical basis.

2.3 Cardinal Neighbor Statistics

For each pixel (x, y) , we sample the temporal median and standard deviation of its four cardinal neighbors at distance D :

$$\mathcal{N}_D(x, y) = \{(x, y - D), (x, y + D), (x - D, y), (x + D, y)\} \quad (3)$$

The neighbor statistics are:

$$\bar{\mu}_{\mathcal{N}}(x, y) = \frac{1}{|\mathcal{N}_D|} \sum_{(x', y') \in \mathcal{N}_D} \mu(x', y') \quad (4)$$

$$\bar{\sigma}_{\mathcal{N}}(x, y) = \frac{1}{|\mathcal{N}_D|} \sum_{(x', y') \in \mathcal{N}_D} \sigma(x', y') \quad (5)$$

The distance D is chosen to be larger than the PSF extent but smaller than the spatial scale of vignetting or background gradients. At $D = 15$ pixels, the cardinal neighbors share the same vignette and background level as the central pixel but do not overlap its PSF.

2.4 Z-Score and Detection

The z-score measures how many neighbor-standard-deviations the pixel's temporal median exceeds the neighbor average:

$$z(x, y) = \frac{\mu(x, y) - \bar{\mu}_{\mathcal{N}}(x, y)}{\bar{\sigma}_{\mathcal{N}}(x, y)} \quad (6)$$

A pixel is classified as a star candidate if:

1. $z(x, y) > z_{\min}$ (significantly brighter than neighbors), and
2. $\mu(x, y) > \bar{\mu}_{\mathcal{N}}(x, y)$ (brighter, not just different).

The second condition eliminates dead pixels, dark defects, and negative outliers that would have large $|z|$ but negative sign.

2.5 Properties

The method has several desirable properties:

Vignette cancellation. Vignetting produces a smooth brightness gradient across the field. At distance D , the cardinal neighbors experience nearly the same vignette as the central pixel, so the vignette contribution cancels in the difference $\mu - \bar{\mu}_N$.

No PSF model required. The method does not convolve with or deconvolve from any PSF template. It detects anomalous brightness regardless of the point spread function’s shape.

Noise suppression via temporal statistics. With $N = 296$ independent samples per pixel, the temporal median is robust and the temporal standard deviation provides a per-pixel noise estimate. Transient artifacts (popcorn noise, compression artifacts, AC fan turbulence) vary across frames and are suppressed by the median.

Single parameter. The only free parameter is the neighbor distance D . The detection threshold $z_{\min}$ controls sensitivity but does not affect the z-map computation itself.

3 Application: PR46 FLIR Video

3.1 Data

PR46 is a publicly-released infrared video from the PURSUE collection [17], captured by a FLIR sensor. The available data consists of video that has been filmed off a projection screen, introducing screen grain, ambient lighting variation, and H.264 compression artifacts. An AC fan was blowing across the screen during capture. 296 frames were extracted and aligned via cross-correlation. The composite result is shown in Figure 1.

3.2 Results

With $D = 15$ and $z_{\min} = 3$, the method detects 11,344 candidate pixels (0.55% of the image). At $z_{\min} = 12$, 15 confirmed stars remain with z-scores ranging from 13σ to 29σ . One of 16 visually-marked stars was missed—it lies on a local brightness gradient where the cardinal neighbors are brighter than the star itself.

Threshold	Star pixels	Confirmed stars
$z > 3$	11,344	—
$z > 6$	~ 800	16
$z > 9$	~ 200	16
$z > 12$	~ 80	15

Table 1: Detection counts at various z-score thresholds.

4 Iterative Refinement via Star-Anchored Alignment

The method as described uses a pre-computed frame alignment. However, the detected stars themselves provide sub-pixel alignment anchors for an iterative refinement:

1. **Initial alignment:** Align frames via full-image cross-correlation (pixel-level).
2. **Initial detection:** Compute the z-map and detect stars at $z > z_{\min}$.

3. **Star-anchored re-alignment:** Use detected star positions as control points. For each frame, compute the sub-pixel offset that minimizes the residual between that frame’s star positions and the reference frame’s positions.
4. **Recompute z-map:** With improved alignment, recompute temporal statistics and the z-map. Formerly-smeared stars sharpen; new fainter stars may emerge above threshold.
5. **Iterate** until the alignment offsets converge to sub-pixel precision.

This is structurally analogous to Expectation-Maximization: the E-step discovers stars in the current z-map; the M-step refines alignment using those stars. Each iteration improves both the alignment and the detection, converging when the offsets stabilize.

With careful noise rejection—excluding pixels whose temporal variance exceeds a threshold from the alignment computation—the iterative scheme can achieve sub-pixel alignment and correspondingly stronger star detection.

5 Related Work

Astronomical source extraction. SExtractor [1] and DAOPHOT [2] are the standard tools for stellar photometry; both require PSF models and operate on single frames or co-added stacks. Our method requires no PSF model and operates on the temporal distribution across frames.

Lucky imaging. Fried [3] established the theoretical basis for short-exposure selection; Law et al. [4] demonstrated practical lucky imaging at visible wavelengths. Lucky imaging selects the best frames; our method uses all frames equally via the temporal median.

Multi-frame stacking. Drizzle [5] reconstructs undersampled images via shift-and-add. Our temporal median serves a similar alignment-and-combine purpose but uses the median (not the mean) for robustness.

CFAR detection. The method is structurally a CA-CFAR (Cell-Averaging Constant False Alarm Rate) detector [6], with temporal medians replacing single-look intensities and sparse cardinal neighbors replacing dense annular reference cells. Rohling’s OS-CFAR [7] uses ordered statistics (our median provides similar robustness). The one missed star—located on a brightness gradient where the reference cells straddle a clutter edge—is the classic CA-CFAR failure characterized by Gandhi & Kassam [8].

Temporal persistence. Track-before-detect methods [11] exploit temporal persistence of dim moving targets; our method exploits temporal persistence of stationary point sources against variable noise.

Infrared change detection. The STLC family [12] applies temporal pixel analysis to infrared surveillance video for change detection, but targets moving objects rather than stationary astronomical point sources.

Iterative registration. The detect→align→redetect loop of Section 4 follows the iterative back-projection pattern of Irani & Peleg [13], applied here to astronomical frame registration.

We note that the z-score with $\bar{\sigma}$ estimated from four reference cells follows a Student- t -like distribution rather than a Gaussian; the CFAR literature handles this via threshold multipliers [6]. At our observed margins of $13\text{--}29\sigma$, this distinction does not affect detection conclusions.

6 Future Directions

Several extensions merit investigation.

Hyperparameter optimization. The neighbor distance $D = 15$ was chosen heuristically. A systematic sweep of D across the PSF extent, background correlation length, and vignette scale would characterize the sensitivity surface. At small D , reference cells overlap the star’s PSF and suppress detection; at large D , the vignette-cancellation assumption weakens. The optimal D likely depends on the imaging system and can be estimated from the data.

HUD and object mask interaction. The current pipeline masks HUD overlay and the central object prior to star detection, excluding those pixels from the z-map. However, stars that transit behind the masked region are lost entirely. An iterative approach could detect stars outside the mask, establish a stellar coordinate frame, and then predict star positions **within** the masked region—testing whether the expected stellar signal is present beneath the overlay. This would recover occluded stars and provide additional alignment anchors.

Vignette as a screen geometry diagnostic. The vignetting pattern in the PR46 data may encode the geometry of the air-gap copy chain: a FLIR sensor projected onto a screen, filmed by a second camera. The vignette combines the FLIR’s native vignette with the projection geometry and the recording camera’s optics. If the detected star positions are compared against a catalog (via plate solving), any systematic displacement of stars as a function of field position would reveal the screen distortion function. In particular, stars near the central object may exhibit apparent displacement due to either (a) screen warping from projector heat or physical deformation, or (b) refractive effects in the original FLIR optical path. Distinguishing these two sources—screen distortion versus atmospheric or optical refraction in the original capture—would constrain the provenance chain of the video.

Iterative star-anchored alignment. As described in Section 4, detected stars can anchor sub-pixel re-alignment, with noise rejection excluding high-variance pixels from the alignment computation. Each iteration sharpens the temporal statistics, potentially revealing fainter stars and enabling plate solving of fields that are currently too degraded for astrometry.

7 Conclusion

The per-pixel temporal z-score against cardinal neighbors is a simple, effective, and parameter-light method for detecting stars in degraded multi-frame imagery. It requires no PSF model, naturally cancels vignetting, and suppresses transient noise through temporal statistics. Applied to 296 frames of degraded FLIR video, it detects 15 of 16 confirmed stars at $13\text{--}29\sigma$ significance. The iterative refinement via star-anchored alignment promises further improvement by exploiting the mutual reinforcement between detection and alignment.




Figure 1: Top: temporal distribution heatmap from 296 aligned FLIR frames. Color encodes z-score: red ($z \geq 5.1$), orange ($z \geq 6$), yellow ($z \geq 9$), white ($z \geq 12$). Bottom: a single raw FLIR frame (frame 200 of 296, contrast enhanced) with green circles marking the 15 confirmed star positions. Individual frames show no discernible stars above the sensor noise; the temporal z-map recovers all 15 as persistent, locally anomalous point sources.

II. PSF-Matched Star Detection in Handheld Phone Photography

1 Introduction

The temporal distribution method of Part I requires many aligned frames. A complementary approach extracts the point spread function (PSF) from identified stars in a **single** long-exposure frame, then uses that PSF as a matched filter to detect every star in the image. This method was developed for handheld phone photographs of the night sky—4-second exposures at ISO 6400 with a consumer OnePlus Nord N20 5G, no tripod, no tracking mount.

The key insight is that the hand motion during a long exposure produces a characteristic **motion trace** that is identical for every point source in the frame. Stars, being point sources at infinity, all experience the same hand shake. The motion trace IS the PSF, and it can be extracted from the brightest stars, then used as a correlation kernel to find every other star.

2 PSF Extraction from Starlight

2.1 Star Identification

Stars are detected as local maxima in a Gaussian-smoothed ($\sigma = 3$) grayscale image, thresholded above background, with a roofline mask excluding the lower half of the frame. For the Orion field, 8 named stars were identified via a belt collinearity algorithm: three collinear, equally-spaced stars with spacing ratio > 0.96 and collinearity deviation < 0.003 uniquely identify Orion's belt.

2.2 Brightness-Binned PSF

Each identified star's neighborhood (radius $R = 75$ px) is extracted, background-subtracted (annular median), and centroid-aligned to sub-pixel precision. Stars are sorted by total disc brightness and grouped into brightness bins. A sliding window of 3 sequential stars produces PSFs at each brightness level.

The critical finding: **bright stars have wider PSFs than faint stars**. At high brightness, the sensor saturates and the signal blooms into adjacent pixels, inflating the apparent PSF. The medium-brightness stars (FWHM ≈ 13 px) give the truest motion trace; the faint stars (FWHM ≈ 7 px) are cleanest but noisiest. This brightness-dependent PSF variation is consistent across all sliding windows (all sequential triangle closures < 2 px), confirming that the shape is real and the alignment is correct.

2.3 Motion Trace Characterization

The extracted PSF (Figure 2) records the photographer's hand motion during the 4-second exposure with sub-pixel precision. Cross-section analysis reveals an asymmetric two-lobed structure: the hand paused briefly at two positions with motion between them, producing FWHM of 10 px (horizontal) $\times$ 14 px (vertical). The vertical axis shows more motion than the horizontal—consistent with the biomechanics of a standing handheld camera.




Figure 2: Sliding window PSF extraction from 8 identified Orion stars, sorted by brightness (left to right = bright to faint). The PSF shrinks as stars get fainter: bright stars bloom due to sensor saturation (FWHM = 31 px), medium stars show the true motion trace (FWHM = 13 px), faint stars are tightest but noisy (FWHM = 6 px). The shape is consistent across all windows.

3 Matched Filter Detection

The extracted PSF is used as a correlation kernel across the entire image. For each pixel position (x, y) , the correlation with the PSF kernel h is computed via FFT:

$$C(x, y) = \mathcal{F}^{-1} \{ \mathcal{F}\{I\} \cdot \mathcal{F}^*\{h\} \} \quad (7)$$

The correlation map is z-scored:

$$z(x, y) = \frac{C(x, y) - \bar{C}}{\sigma_C} \quad (8)$$

Stars are detected as local maxima above a threshold $z \geq z_{\min}$. The matched filter is the optimal detector for a known signal in Gaussian noise (the Neyman-Pearson lemma [16]); here the “known signal” is the PSF extracted from the data itself.

3.1 Results

Applied to a single 4-second handheld exposure of Orion (ISO 6400, $f/1.7$, 27 mm equivalent), the matched filter detects 23 sky stars at $z \geq 4$ and 29 at $z \geq 3$. The brightest detection—Sirius—reaches $z = 84.3\sigma$, confirming that the matched filter achieves extraordinary sensitivity from a single consumer phone photograph.

3.2 True-Hue False-Saturation Rendering

Detected stars are rendered as colored discs on a black background. The color is the **true hue** from the original image (the phone sensor captures real stellar color at ISO 6400), with saturation and value maximized ($S = V = 1$ in HSV space). Disc radius encodes detection significance via a log scale:

$$r = 3.90 \cdot \ln(z) - 1.29 \quad (9)$$

mapping $z = 3 \rightarrow r = 3$ px, $z = 4 \rightarrow r = 4$ px, and $z = 84 \rightarrow r = 16$ px. This visualization reveals stellar color differences invisible at natural brightness levels: Betelgeuse appears warm, Rigel and Sirius appear blue-white.

4 Plate Solving and Validation

The original 4-second photograph was submitted to astrometry.net [1] and successfully plate-solved: RA = 81.35, Dec = −17.94, confirming Orion. The pixel scale (59.5 arcsec/px) and field orientation (352.9) were recovered. This independently validates the star positions detected by our pipeline.

5 Wiener Deconvolution

With the PSF known, the entire image can be deconvolved via the Wiener filter:

$$\hat{F}(\omega) = \frac{F_{\text{img}}(\omega) \cdot H^*(\omega)}{|H(\omega)|^2 + \lambda} \quad (10)$$

where λ is a regularization parameter balancing sharpness against noise amplification.

A systematic lambda search across multiple images reveals that the deconvolution landscape is **structural**—the normalized curves for sky noise, star FWHM, ringing energy, and sharpness overlay almost perfectly across different exposures, pointings, and atmospheric conditions. The structural knee at $\log_{10}(\lambda) \approx -2.0$ is a universal Pareto optimum: a property of the Wiener method, not of any particular image. At $\lambda = 0.01$, the deconvolution resolves barn siding panels, power lines, and window frames from 1.6-second handheld exposures taken in darkness.

6 Connection to Part I

Parts I and II solve the same problem—detecting stars in degraded imagery—via complementary approaches:

	Part I (Temporal)	Part II (PSF Matched Filter)
Input	296 video frames	1 phone photograph
PSF	Not needed	Extracted from stars
Detection	Temporal CA-CFAR	Matched filter correlation
Best detection	29σ	84σ
Plate solved	No	Yes
Deconvolution	No	Yes (Wiener)

Table 2: Comparison of temporal and PSF-based star detection.

The temporal method requires many frames but no PSF model; the PSF method requires only one frame but needs identifiable stars to extract the PSF. Together they cover the space from multi-frame degraded video to single-frame long-exposure photography.

References

- [1] Bertin, E. & Arnouts, S. (1996). SExtractor: Software for source extraction. **Astronomy & Astrophysics Supplement Series**, 117, 393–404.
- [2] Stetson, P.B. (1987). DAOPHOT: A computer program for crowded-field stellar photometry. **PASP**, 99, 191–222.
- [3] Fried, D.L. (1978). Probability of getting a lucky short-exposure image through turbulence. **JOSA**, 68(12), 1651–1658.
- [4] Law, N.M., Mackay, C.D. & Baldwin, J.E. (2006). Lucky imaging: High angular resolution imaging in the visible from the ground. **Astronomy & Astrophysics**, 446, 739–745.
- [5] Fruchter, A.S. & Hook, R.N. (2002). Drizzle: A method for the linear reconstruction of undersampled images. **PASP**, 114, 144–152.
- [6] Finn, H.M. & Johnson, R.S. (1968). Adaptive detection mode with threshold control as a function of spatially sampled clutter-level estimates. **RCA Review**, 29, 414–464.
- [7] Rohling, H. (1983). Radar CFAR thresholding in clutter and multiple target situations. **IEEE Trans. Aerospace and Electronic Systems**, AES-19(4), 608–621.
- [8] Gandhi, P.P. & Kassam, S.A. (1988). Analysis of CFAR processors in nonhomogeneous background. **IEEE Trans. Aerospace and Electronic Systems**, 24(4), 427–445.
- [9] Kay, S.M. (1998). **Fundamentals of Statistical Signal Processing, Vol. II: Detection Theory**. Prentice Hall.
- [10] North, D.O. (1963). An analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems. **Proc. IEEE**, 51(7), 1016–1027.
- [11] Barniv, Y. (1985). Dynamic programming solution for detecting dim moving targets. **IEEE Trans. Aerospace and Electronic Systems**, AES-21(1), 144–156.
- [12] Deng, H. et al. (2016). Infrared small target detection based on modified STLC. **Infrared Physics & Technology**, 76, 504–513.
- [13] Irani, M. & Peleg, S. (1991). Improving resolution by image registration. **CVGIP: Graphical Models and Image Processing**, 53(3), 231–239.
- [14] Irwin, M.J. (1985). Automatic analysis of crowded fields. **MNRAS**, 214, 575–604.
- [15] Starck, J.L. & Murtagh, F. (2002). **Astronomical Image and Data Analysis**. Springer.
- [16] Neyman, J. & Pearson, E.S. (1933). On the problem of the most efficient tests of statistical hypotheses. **Phil. Trans. Royal Soc. A**, 231, 289–337.
- [17] Department of Defense (2021). PURSUE: Public Unclassified Release of Selected UAP Evidence. Defense Visual Information Distribution Service (DIVIDS), <https://www.dvidshub.net>. Note: the source video was obtained via the non-authenticated public interface. The authors declined to register with DIVIDS due to concerns about the End User License Agreement, which asserts jurisdiction over “attached devices” of registered users in terms that could be read as authorizing warrantless inspection of user equipment. The public data used in this work requires no authentication and carries no such terms.