Real-Time Convolutional Neural Network-Based Star Detection and Centroiding Method for CubeSat Star Tracker

HONGRUI ZHAO
MICHAEL F. LEMBECK
ADRIAN ZHUANG
RIYA SHAH
JESSE WEI
University of Illinois Urbana-Champaign, IL 61820, USA

Abstract— Star trackers are one of the most accurate celestial sensors used for absolute attitude determination. The devices detect stars in captured images and accurately compute their projected centroids on an imaging focal plane with subpixel precision. Traditional algorithms for star detection and centroiding often rely on threshold adjustments for star pixel detection and pixel brightness weighting for centroid computation. However, challenges like high sensor noise and stray light can compromise algorithm performance. This article introduces a Convolutional Neural Network (CNN)-based approach for star detection and centroiding, tailored to address the issues posed by noisy star tracker images in the presence of stray light and other artifacts. Trained using simulated star images overlayed with real sensor noise and stray light, the CNN produces both a binary segmentation map distinguishing star pixels from the background and a distance map indicating each pixel's proximity to the nearest star centroid. Leveraging this distance information alongside pixel coordinates transforms centroid calculations into a set of trilateration problems solvable via the least squares method. Our method employs efficient UNet variants for the underlying CNN architectures, and the variants' performances are evaluated. Comprehensive testing has been undertaken with synthetic image evaluations, hardware-in-the-loop assessments, and night sky tests. The tests consistently demonstrated that our method outperforms several existing algorithms in centroiding accuracy and exhibits superior resilience to high sensor noise and stray light interference. An additional benefit of our algorithms is that they can be executed in real-time on low-power edge AI processors.

Index Terms— Convolution, Satellite navigation systems, Training

Hongrui Zhao, Michael F. Lembeck, Adrian Zhuang, Riya Shah, Jesse Wei are with the Aerospace Engineering Department, University of Illionis Urbana-Champaign, IL 61820, USA (e-mail: hongrui5@illinois.edu; mlembeck@illinois.edu; adrianz3@illinois.edu; riyahs3@illinois.edu; jwei28@illinois.edu). (Corresponding author: Hongrui Zhao.)

I. INTRODUCTION

With the availability of compact star tracker technology, space missions demanding precise attitude determination are increasingly leveraging these devices on cost-effective CubeSat platforms. Laser communication CubeSats in Low Earth Orbit (LEO), when equipped with a star tracker, can achieve precise pointing capabilities towards ground stations, facilitating downlink rates of up to 200 Gbps [1]. For deep space CubeSat missions, star trackers can also play a crucial role in high-gain antenna pointing [2]. Star trackers can be used in conjunction with other sensors to accurately determine the orbits of other orbiting objects, enhancing space situational awareness [3].

Many CubeSat star trackers use commercial off-the-shelf (COTS) components, trading cost for performance. Star trackers equipped with low-cost imaging sensors [4] suffer from lower sensitivity resulting in noisier star images. Conventional star centroiding algorithms [5][6] struggle with high sensor noise, leading to compromised centroid computation accuracy that is pivotal for optimal star tracker performance. Noise can also introduce false positive star detections affecting attitude determination.

Stray light from bright body or optical path dust reflections can also result in false positives. Given that the intensity of stray light often surpasses the background brightness level, conventional threshold-based star detection algorithms [5][7][8][9] are unable to discriminate false detections, undermining subsequent star identification processes. To counteract stray light, large baffles are frequently placed in front of the sensor optics [10], potentially encroaching on the limited payload space available in CubeSats. Even with a baffle, star trackers remain susceptible to blinding when directly exposed to bright bodies in their field of view [11].

In this article, we tackle these deficiencies using Convolutional Neural Networks (CNNs). CNNs have demonstrated remarkable image processing capabilities relevant to star trackers across a diverse set of applications, including precision agriculture, where CNNs excel in precisely locating and counting crops from aerial images [12]. For low-light imaging scenarios, CNNs effectively denoise and enhance short-exposure images characterized by low photon counts and diminished signal-to-noise (SNR) [13]. Recent advancements in star tracker research have leveraged CNNs for star image segmentation, showcasing robust performance against high sensor noise and stray light [14][15]. Image segmentation facilitates star pixel detection, but a subsequent denoising step is still required to achieve high SNR star images for accurate centroiding. Adopting two distinct CNNs, one for star detection and another for denoising, imposes significant computational overhead, making it impractical for edge computing hardware deployment. To overcome these limitations, this article advocates for the end-to-end training of a singular CNN capable of simultaneously handling

star detection and centroiding, resulting in a streamlined and computationally efficient architecture.

The main contributions of this article are:

1) Introducing an end-to-end training approach for a singular CNN, designed to concurrently handle star detection and centroiding. This method surpasses the centroiding accuracy and detection robustness achieved by several existing techniques.
2) Generating synthetic star images, augmented with real sensor noise and stray light, for training purposes. This approach significantly reduces the reliance on manual labeling efforts.
3) Conducting a comprehensive performance evaluation of various CNN models and identifying the most suitable CNN architectures for real-time star tracker image processing.

This article is structured as follows: In Section II, we introduce traditional star detection and centroiding algorithms, as well as existing CNN models relevant to this work. Section III presents the workflow of the proposed star detection and centroiding method, the data generation process, and the CNN implementation details. Section IV explains the experiment setups and discusses the experiment results from synthetic dataset testing, hardware-in-the-loop testing, and night sky testing. Finally, we conclude this article and point out the future works in Section V.

II. Related Works

Existing star detection algorithms can be broadly classified into two categories: global threshold methods, and local threshold methods. Global threshold methods determine a single pixel intensity threshold to distinguish star pixels (foreground) from the background. Liebe's global adaptive threshold method [5] computes the mean and the standard deviation of each star image, where the standard deviation is scaled and added to the mean to derive the global threshold. While Liebe's method offers a straightforward and effective detection approach, manual tuning of the scaling factor is often required to achieve optimal performance. Another approach, using the weighted iterative threshold method (WITM) [7], computes a global threshold for each star image through iterative adjustments of two weight coefficients until the threshold difference between successive iterations falls below a predefined value. Although WITM can autonomously derive an optimal threshold and exhibits reduced sensitivity to initialization, its iterative nature considerably impacts computation speed. On the other hand, local threshold methods assign individual thresholds to each pixel, necessitating each pixel to surpass its designated threshold to be classified as a star pixel. The detection routine of the ST-16 star tracker [8] employs a 1×29 window centered on each pixel and utilizes the window's average, augmented by a predetermined constant, as the local threshold. Meanwhile, the method proposed by Sun

et al. [9] employs a sequential application of a Gaussian filter, followed by erosion, dilation, and an average filter on each star image to compute a background estimate. For a pixel to qualify as a star pixel, its intensity must exceed the computed background estimate plus a constant. Although this method demonstrates robustness against uneven background illumination, it requires considerable computations due to the sequential application of multiple filters.

Threshold-based detection methods often fail to distinguish faint stars from background noise or discriminate between stars and other bright non-stellar objects. To enhance the robustness of star detection, several approaches have leveraged heuristic criteria [8] and local gradient information [16]. In the ST-16 detection routine [8], a cluster of pixels surpassing their respective local thresholds is identified. The algorithm then quantifies the number of connected bright pixels within the cluster using 4-connectivity. Subsequently, both the count of contiguous bright pixels and their cumulative intensity are used to eliminate isolated hot pixels. Jiang *et al.* [16] suggested that gradient information from lunar images could be examined and used to manually tune gradient thresholds to differentiate stars from lunar interference. Despite their merits, these aforementioned techniques exhibit sensitivity to variations in environmental conditions and sensor characteristics, necessitating frequent parameter recalibrations. In contrast, our method addresses this limitation by incorporating training images spanning a range of exposure durations and noise levels. This design allows the neural network to adapt dynamically to varying imaging conditions without the need for recurrent recalibrations.

Existing star-centroiding algorithms utilize the intensity distribution of star pixels to compute star centroids with subpixel accuracy. The center of gravity method stands out as a common approach due to its simplicity and computational efficiency [5]. This method derives star centroids by averaging the coordinates of a cluster of star pixels, each weighted by its brightness. However, the center of gravity method yields only an approximate centroid position and proves susceptible to background noise. In pursuit of enhanced accuracy and robustness, Delabie *et al.* introduced the Gaussian grid algorithm [6]. This approach employs a Gaussian distribution function to model the detected star pixels within a designated window, subsequently deriving the centroid as the center of this distribution. The Gaussian grid algorithm surpasses the center of gravity method in terms of accuracy and robustness while maintaining computational efficiency. However, the Gaussian grid algorithm can fail when confronted with strongly non-Gaussian distributions of star pixels, such as in the presence of stray light or sensor motion. In contrast, our method is trained for both star detection and centroiding. This integrated approach ensures robust centroiding performance across varying noise levels, thereby enhancing overall accuracy and robustness.
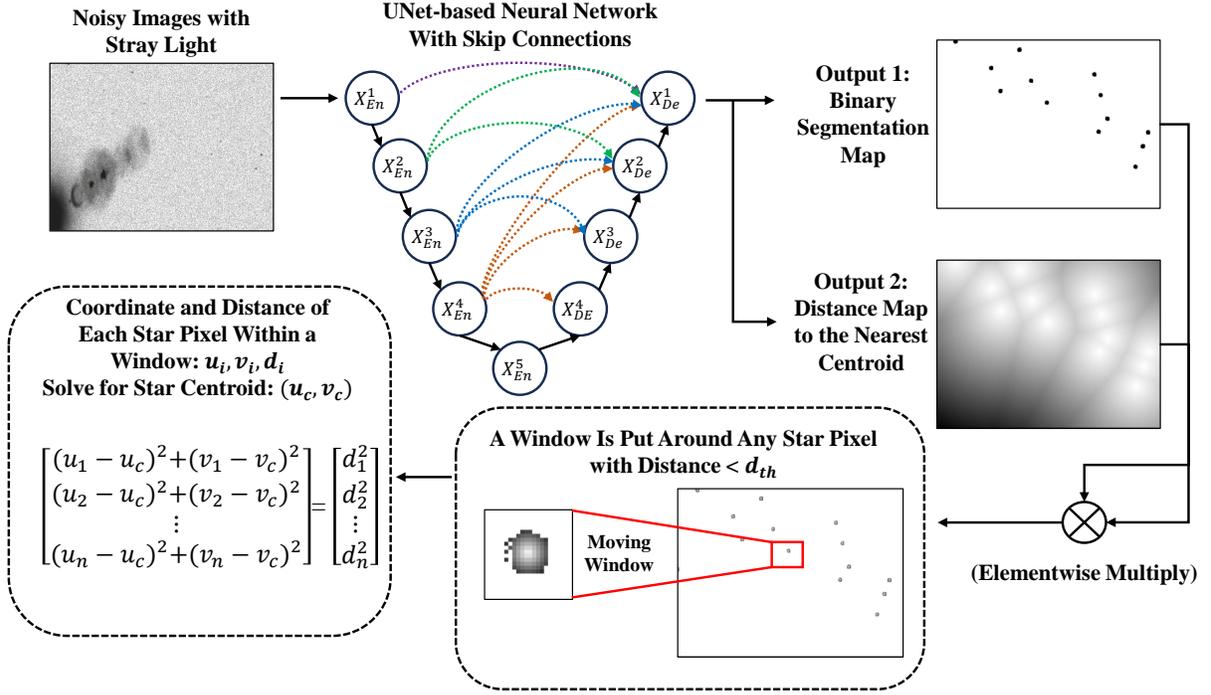
Fig. 1. Flowchart of our star detection and centroiding method. Noisy star images with stray light are fed into a UNet-based CNN. The neural network outputs both a binary segmentation map of star pixels versus background and a distance map of each pixel to the nearest star centroid. Using the distance information and the pixel coordinates, star centroiding now becomes a trilateration problem and can be directly solved using a least squares solution.

The UNet, originally designed for medical image segmentation, incorporates skip connections to retain fine details from the input image [17]. These skip connections merge feature maps containing high spatial frequency details from early convolutional layers with those containing high-level semantic information. While UNet has demonstrated impressive performance in star image segmentation with a notably low false positive rate [14][15], its original architecture, comprising approximately 31 million parameters, demands significant computational resources and executes in seconds on a standard personal computer. This computational intensity poses a challenge for real-time star tracker processing on a low-power processor. To address this limitation, various efforts have been undertaken to enhance UNet's computational efficiency and reduce its model size. MobileUNet [18] integrates inverted residual blocks from MobileNetV2 to replace the VGG-16 blocks in conventional U-Net, resulting in marked computational efficiency and a reduction in model size. ELUNet [19] introduces deep skip connections that effectively reduces channel numbers without sacrificing performance. Finally, SqueezeUNet [20] employs fire modules to decrease the number of channels, offsetting this reduction with an inception stage comprising two parallel convolutions. In this article, we undertake an exhaustive analysis of these UNet variants to determine the most suitable architectures for star tracker applications.

III. PROPOSED METHOD

Fig. 1 illustrates the flowchart detailing our CNN-based method for star detection and centroiding. The neural network receives input star images that may contain stray light and varying noise levels. Leveraging a UNet-based CNN, our method produces both a binary segmentation map and a distance map. Within the segmentation map, a "one" indicates a star pixel has been detected, while a "zero" signifies a background pixel. Concurrently, the distance map assigns each pixel a value corresponding to its distance from the nearest star centroid. These two maps are combined using element-wise multiplication. In the resultant map, pixels with values below a specified threshold, d_{th}, suggest potential star centroid locations, prompting the centering of a window at each corresponding pixel. Within each window, every pixel i is identified by its coordinates (u_i, v_i) and stores its distance d_i to the star centroid (u_c, v_c). This transforms star centroiding into a trilateration problem, determining a position based on distances from at least three known points, that can be resolved using a least squares solution. The following sections will delve into the generation of training data, the procedure for star detection and centroiding, and the specifics of our CNN implementation.

A. Training Data Generation

To generate training images for the neural network, the initial step involves acquiring synthetic noiseless star

images. The Hipparcos catalog [21] is used to provide the position and magnitude information of stars. Considering the camera parameters and the actual orientation of a star tracker, determining actual centroids and identifications of stars within the image plane is straightforward. Subsequent efforts accurately capture the distribution patterns and magnitudes of real stars.

The Hipparcos catalog provides the Johnson V magnitude in the Johnson UBV photometric system. The V band filter is centered at $\lambda_0 = 0.55\,\mu$m with a full width at half maximum (FWHM) bandwidth of $\Delta\lambda_0 = 0.089\,\mu$m, and has the zero magnitude spectral flux density $e_0 = 3.92 \times 10^{-8}\,\text{W}\cdot\text{m}^{-2}\cdot\mu\text{m}^{-1}$ [22]. Therefore, we can obtain the average spectral flux density of a star over the V band in photons$\cdot\text{s}^{-1}\cdot\text{m}^{-2}\cdot\mu\text{m}^{-1}$ when its magnitude m_V from the Hipparcos catalog is given

$$\bar{e}_V = \frac{e_0\lambda_0}{h\cdot c}\cdot 10^{-\left(\frac{m_V}{2.5}\right)} = 1.085356\times 10^{11}\cdot 10^{-\left(\frac{m_V}{2.5}\right)} \quad (1)$$

where h is the Planck constant and c is the speed of light. For a star with a magnitude m_V, the average photon arrival rate received by an image sensor over the V band is [23]

$$\bar{n} = \bar{e}_V \cdot AP \cdot \bar{\epsilon} \cdot \bar{E} \cdot \Delta\lambda_0 \quad (2)$$

where AP is the area of the image sensor aperture given in m², $\bar{\epsilon}$ is the average optical efficiency factor over V band, $\bar{E}$ is an average correction factor for extinction due to the airmass over the V band. For on-orbit star tracker processing, $\bar{E} = 1$. Also recall that $\Delta\lambda_0 = 0.089\,\mu$m for the V band.

The image of a star is subject to optical aberrations and spreads over several pixels, which is modeled by a point spread function (PSF). The Gaussian distribution serves as a favorable approximation to the PSF of an ideal lens. It is expressed in both the width and height directions across the image plane, as described in [6]

$$g_{width}(u) = \frac{1}{\sqrt{2\pi}\sigma_{\text{psf},u}}\exp\left[-\frac{(u-u_c)^2}{2\sigma_{\text{psf},u}^2}\right] \quad (3)$$

$$g_{height}(v) = \frac{1}{\sqrt{2\pi}\sigma_{\text{psf},v}}\exp\left[-\frac{(v-v_c)^2}{2\sigma_{\text{psf},v}^2}\right] \quad (4)$$

where the coordinates of the real star centroid are u_c, v_c. $\sigma_{\text{psf},u}$ and $\sigma_{\text{psf},v}$ are the PSF width in width and height direction in pixels. We choose $\sigma_{\text{psf},u} = \sigma_{\text{psf},v} = \sigma_{\text{psf}}$ in our simulation. In our experiments, we established that uniformly sampling σ_{psf} from $[0.5, 1]$ produces star images with realistic distribution patterns, as illustrated in Fig. 2. The PSF of a real star captured by a camera set to focus at infinity typically falls within this range. We also found that training the neural network with a variety of σ_{psf} values can improve its capacity to process stars distorted by lenses or motion, which may deviate from the Gaussian PSF.

Using the Gaussian PSF, the average rate of arrival of photons received by a pixel with coordinate (u_i, v_i) can

Real Star **Simulated, σ_{PSF}=0.5**

Simulated, σ_{PSF}=1 **Simulated, σ_{PSF}=2**
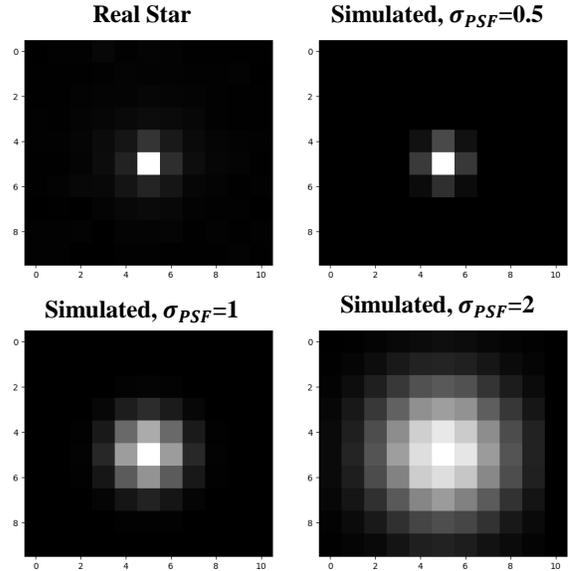
Fig. 2. Our experiments find that uniformly sampling σ_{psf} from $[0.5, 1]$ produces simulated star images with realistic distribution patterns.

be computed as

$$\bar{n}_{pixel} = \bar{n}\int_{u_i-0.5}^{u_i+0.5} g_u(u)\,du \int_{v_i-0.5}^{v_i+0.5} g_v(v)\,dv$$

where each pixel has a non-dimensional width of 1.

The arrival of photons is a stochastic process $\boldsymbol{x}(t)$ made up of Poisson impulses [22]. The probability of a pixel obtaining p photons during a exposure time T (In our data generator, we uniformly sample T from $[0.1s, 1s]$) is given by

$$P\left[\int_T \boldsymbol{x}(t)\,dt = p\right] = \frac{(\bar{n}_{pixel}T)^p\, e^{-\bar{n}_{pixel}T}}{p!} \quad (5)$$

Sampling from this Poisson distribution will give us the total number of photons N received by the pixel (u_i, v_i) within the exposure time.

Finally, the pixel bit counts $V_{(u_i,v_i)}$ can be calculated [6]

$$V_{(u_i,v_i)} = floor\left(N\cdot QE\cdot\frac{2^{N_{bits}}-1}{FWC}\right)\cdot\frac{FWC}{2^{N_{bits}}-1} \quad (6)$$

where $floor()$ outputs the greatest integer less than or equal to the input value, QE is the quantum efficiency of the image sensor, FWC is the full well capacity of the image sensor in electrons, and N_{bits} is the resolution of the ADC. A window with half width $= floor\left(3\sigma_{\text{psf}} + 0.5\right)$ centered at the centroid pixel is used to simulate the defocused star image. $V_{(u_i,v_i)}$ is calculated for every pixel within the window.

The pristine simulated star images are subsequently merged with noisy frames containing stray light captured by actual cameras, as illustrated in Fig. 3. Current star detection studies typically apply Gaussian noise to simulated star images [14][15]. Monakhova $et\ al.$ [24] demonstrated that simplistic Gaussian or Poisson-Gaussian noise models fail to adequately represent the frequently encountered
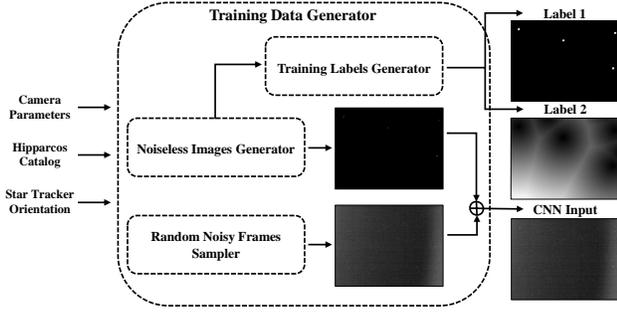
Fig. 3.   The noiseless simulated star images are merged with noisy frames containing stray light captured by actual cameras to create training inputs to CNN. Additionally, two sets of training labels are generated: training labels 1, consisting of binary segmentation maps, and training labels 2, comprising distance maps.

non-Gaussian, non-linear, and sensor-specific noise inherent in low-light images. Inspired by their observations, we acquire noisy dark frames from three distinct cameras by covering their lenses with lens caps. For each camera, we captured 100 frames at exposure times of 100 ms, 300 ms, 500 ms, 700 ms, and 900 ms, respectively. An additional set of 1000 frames is acquired by situating a camera within a darkroom environment and illuminating it with a flashlight from 12 distinct angles to simulate the presence of stray light. During the training image generation process, a noisy frame is randomly selected from this pool of 2500 captured frames to be fused with a noiseless star image. Random flips are applied to the noisy frame to facilitate data augmentation.

As shown in Fig. 3, two distinct label types are produced for each training image: binary segmentation maps and distance maps. A binary segmentation map, mirroring the resolution of the training image, assigns each pixel a value of either "one" or "zero." Pixels with a value of "one" denote star signals, while those with a value of "zero" represent background areas. Additionally, distance maps, sharing the same resolution as the training image, are segmented into Voronoi cells. Each Voronoi cell $\mathcal{V}_k$ represents a cluster of pixels within the image I that are nearest to a star k with a centroid coordinate C_k. Mathematically, a Voronoi cell is defined as:

$$\mathcal{V}_k = \{q \in I \mid \|q - C_k\| \leq \|q - C_j\|, \forall j\} \qquad (7)$$

Here, each pixel q within a Voronoi cell $\mathcal{V}_k$ stores the distance $\|q - C_k\|$ from its center to the centroid of the corresponding star. During training, a CNN is tasked with generating both segmentation and distance maps that accurately correspond to the labels, given a star image as input. This comprehensive training regimen enables the CNN to effectively learn the tasks of star detection and centroiding simultaneously.

B. Star Detection and Centroiding Procedure

Our star detection and centroiding procedure from image input to centroid output is described in algorithm 1. Given a captured star image I_{in} as the input, neural

Algorithm 1

1: **procedure** DETECT AND CENTOID(I_{in})
2:      $S, D = f(I_{in})$
3:      $I_{out} = S \otimes D$
4:      **while** $q \in I_{out}$ with $d < d_{th}$ exists **do**
5:          Find a coarse centroid pixel with $d < d_{th}$
6:          Put a 5×5 window around the pixel
7:          **for** q_i within the window **do**
8:              Obtain (u_i, v_i) and d_i
9:              remove q_i from I_{out}
10:         **end for**
11:         Obtain A and B using Eq. 9
12:         (u_c, v_c) = DGELSY(A, B)
13:         store (u_c, v_c) into a list
14:     **end while**
15:     **return** the list of centroids
16: **end procedure**

network $f(I_{in})$ outputs both a segmentation map S and a distance map D. To filter out the background, we perform an element-wise multiplication for both maps $I_{out} = S \otimes D$. To first find coarse centroid positions, we find every pixel $q \in I_{out}$ that has a distance d value smaller than a threshold d_{th}. A pixel q with smaller d indicates it is closer to the corresponding star centroid. In our experiments, we found that our method is not sensitive to the selection of d_{th} and $d_{th} = 0.5\sqrt{2}$ gives good results. A 5×5 window is put around a coarse centroid pixel to obtain the coordinate (u_i, v_i) and the distance value d_i of every pixel q_i within.

To find the coordinate (u_c, v_c) of the corresponding star centroid, we first obtain the following equation

$$\begin{bmatrix} (u_1 - u_c)^2 + (v_1 - v_c)^2 \\ (u_2 - u_c)^2 + (v_2 - v_c)^2 \\ \vdots \\ (u_n - u_c)^2 + (v_n - v_c)^2 \end{bmatrix} = \begin{bmatrix} d_1^2 \\ d_2^2 \\ \vdots \\ d_n^2 \end{bmatrix} \qquad (8)$$

We want to solve for the centroid coordinate using distances from a list of known coordinates, which is essentially a trilateration problem. Eq. 8 can be rewritten in the form of $Ax = B$

$$\begin{bmatrix} 2(u_n - u_1) & 2(v_n - v_1) \\ 2(u_n - u_2) & 2(v_n - v_2) \\ \vdots \\ 2(u_n - u_{n-1}) & 2(v_n - v_{n-1}) \end{bmatrix} \begin{bmatrix} u_c \\ v_c \end{bmatrix} =$$
$$\begin{bmatrix} d_1^2 - d_n^2 - u_1^2 - v_1^2 + u_n^2 + v_n^2 \\ d_2^2 - d_n^2 - u_2^2 - v_2^2 + u_n^2 + v_n^2 \\ \vdots \\ d_{n-1}^2 - d_n^2 - u_{n-1}^2 - v_{n-1}^2 + u_n^2 + v_n^2 \end{bmatrix} \qquad (9)$$

We can obtain the centroid coordinates (u_c, v_c) by passing matrices A and B into an off-the-shelf least-squares solver. We use the DGELSY solver based on complete orthogonal factorization from the LAPACK library for its fast computation speed and numerical stability [25]. After a region of pixels has been processed and its star centroid

TABLE I

CNN Models

Model	Input Dim	Encoder Dim	Out Dim
UNet	1	64-128-256-512-1024	2
MobileUNet	1	16-24-32-96-1280	2
ELUNet	1	8-16-32-64-128	2
SqueezeUNet	1	64-128-256-512-1024	2

has been obtained, we remove all of these pixels from I_{out} to avoid repetitive computations. The process described above repeats until no pixel $q \in I_{out}$ with $d < d_{th}$ exists.

C. CNN Implementation Details

We implemented and compared four UNet-based CNN models, as detailed in Table I: the vanilla UNet [17], MobileUNet [18], ELUNet [19], and SqueezeUNet [20]. Each CNN model is designed with an input dimension of one for monochrome star images and an output dimension of two for segmentation and distance maps. The term "Encoder Dim" refers to the output channel dimension of the encoder blocks, with the dimensions of the decoder blocks adjusted accordingly. For instance, in the vanilla U-Net model with encoder dimensions of 64-128-256-512-1024, the decoder blocks have output dimensions of 512-256-128-2. The original ELUNet employs bilinear interpolation for upsampling. However, our chosen edge AI processor does not support high-resolution bilinear upsampling operations. Consequently, we substituted all bilinear upsampling layers with transposed convolution layers in our implementations. All CNN models were developed using PyTorch, and we based our implementations on existing public repositories of these models.

We employed binary cross-entropy loss, denoted as L_S, for training the segmentation map and mean squared error loss, denoted as L_D, for the distance map. The total loss function is formulated as a weighted sum of these two individual loss terms

$$L_S =$$
$$- mean \left(\sum_i^N S_i \otimes \log \hat{S}_i + (1 - S_i) \otimes \log \left(1 - \hat{S}_i \right) \right) \tag{10}$$

$$L_D = mean \left(\sum_i^N \left(\hat{S}_i \otimes \hat{D}_i - S_i \otimes D_i \right)^2 \right) \tag{11}$$

$$L_{total} = \alpha L_S + \beta L_D \tag{12}$$

Here, the $mean()$ function computes the average value of an array, N represents the batch size, S_i is the actual binary segmentation map, $\hat{S}_i$ is the estimated segmentation map with values ranging between zero and one, D_i is the actual distance map, and $\hat{D}_i$ is the estimated distance map. Notation $\otimes$ denotes element-wise multiplication. For our training procedure, we set $\alpha = 1$ and $\beta = 2.5$ to balance the magnitude of the two loss terms.

Our dataset consists of 2,500 images for training and an additional 500 images for evaluation. CNN models

undergo training for 100 epochs with a batch size set to 10 to leverage batch normalization layers effectively. We employ the Adam optimizer with a weight decay of 5×10^{-4} and an initial learning rate of 1×10^{-3}. The learning rate is decreased by a factor of two every 20 epochs to facilitate better convergence. All training processes are conducted on Google Colab using a single NVIDIA V100 GPU. Our codes are available at the project website https://hongruizhao.github.io/CNNStarDetectCentroid/.

IV. Experiments

A. Synthetic Data Test

We begin our experimental evaluation by benchmarking our method against established algorithms using 500 synthetic test images. To further assess the robustness of our approach, an additional set of 500 synthetic test images, contaminated by stray light, was employed. Our evaluation encompasses various CNN models, including the vanilla UNet [17], SqueezeUNet [20], ELUNet [19], and MobileUNet [18]. For comparative analysis, we considered several existing star detection methods: Liebe's global threshold method [5], the Weighted Iterative Threshold Method (WITM) [7], the ST-16 star tracker detection routine [8], and Sun *et al.*'s local threshold method [9]. We also compared our method's centroiding performance to that of the center of gravity [5] and Gaussian grid algorithms [6]. The details of the CNN models and traditional methods are discussed in Section II.

We employed precision, recall, and the F1 score as metrics to assess the detection performance. Precision measures the algorithm's capability to exclude false detections, while recall quantifies its ability to identify all stars. The F1 score offers a balanced evaluation by combining both precision and recall. The mathematical expressions for these metrics are defined as

$$Precision = \frac{TP}{TP + FP} \times 100 \tag{13}$$

$$Recall = \frac{TP}{TP + FN} \times 100 \tag{14}$$

$$F1 = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{15}$$

In these formulas, "TP" denotes true positives, representing stars correctly identified by the star tracker. "FP" indicates false positives, and "FN" signifies false negatives, which are stars not detected by the star tracker. To evaluate centroiding accuracy, we employ the root-mean-square error (RMSE) between the actual centroid coordinates and the estimated coordinates. Additionally, we document the computational time, the total number of multiply-accumulate operations (MACs) in giga, and the number of model parameters in millions for each CNN model. All CNN models are executed on an NVIDIA RTX2060M GPU. It is worth noting that the proposed least squares centroiding process runs efficiently, requir-

ing only 4.49 milliseconds on an AMD Ryzen 7 4800H CPU.

Evaluation results are presented in Tables II, III, IV, and V. Across both detection and centroiding assessments, all CNN models consistently outperformed traditional methods. We also found that the introduction of stray light had minimal impact on the CNN models' performance. We exclusively showcase the stray light performance of our proposed method, as existing detection methods struggle with these contaminated images. Remarkably, despite having a considerably smaller model size and fewer computations, ELUNet achieves performance comparable to UNet. While MobileUNet exhibits slightly lower performance than other CNN models in both detection and centroiding, it distinguishes itself as the fastest method among the tested models. On the other hand, despite its computational intensity, SqueezeUNet does not surpass ELUNet in performance. Consequently, due to their computational efficiency coupled with high performance, we exclusively focused on MobileUNet and ELUNet for the subsequent tests.

TABLE II
Detection Evaluation with Synthetic Images

Method	Precision	Recall	F1
UNet	99.6	97.5	98.6
SqueezeUNet	99.6	97.5	98.5
ELUNet	99.7	97.5	98.6
MobileUnet	99.3	96.0	97.6
Liebe's	91.8	89.6	90.7
WITM	91.1	62.2	73.9
ST16	93.4	83.6	88.2
Sun *et al.*'s	81.1	93.9	87.0

TABLE III
Centroiding Evaluation with Synthetic Images

Method	Centroid RMSE (pixels)
UNet	0.1298
SqueezeUNet	0.1382
ELUNet	0.1363
MobileUnet	0.1695
Center of Gravity	0.6966
Gaussian Grid	0.1922

TABLE IV
Evaluation with Straylight Synthetic Images

Method	F1	Centroid RMSE (pixels)
UNet	97.3	0.1312
SqueezeUNet	96.7	0.1394
ELUNet	97.3	0.1348
MobileUnet	96.6	0.1567

B. Hardware Inference Time Test

We continued evaluating our method using an Edge AI processor. Generally, space applications require hardware with low power consumption to accommodate a

TABLE V
CNN Model Size and Computation Time

Model	Time (s)	MACs (G)	Params (M)
UNet	0.113	256.25	31.04
SqueezeUNet	0.067	70.06	2.64
ELUNet	0.028	22.50	0.8
MobileUnet	0.020	7.71	4.4

constrained power budget. The attitude update rate of a star tracker, in particular, necessitates hardware to possess at high processing speed to facilitate rapid inferences by the AI model. A few viable options on the market are identified by G. Furano *et al.* [28]: Intel Myriad X, Google Coral Edge TPU, and Nvidia Jetson Nano. Ultimately, we selected the Google Coral Edge TPU for its low power consumption of 2 Watts, relatively high processing speed of 4 TOPS, and its readily available documentation and well-supported community.

To execute MobileUNet and ELUNet on the Edge TPU and measure the inference time, the model must first undergo quantization. This process transforms the high-bit representations of weights and activations in the original model into low-bit representations, thereby reducing the model's inference time and enabling efficient execution on edge AI hardware. TinyNN [29] is employed to perform quantization on the model. Then, the model is compiled using the EdgeTPU compiler [30], resulting in an Edge TPU model. However, because the Edge TPU has limitations in terms of the model operations it supports [30], we substitute the resize bilinear operation in our model with the supported transpose convolution operation prior to compilation. The model is then ready to be loaded onto the Edge TPU for inference time measurement.

The Edge TPU is interfaced with a Raspberry Pi 4B over a USB 3.0 connection, and a sample test image file is loaded onto the device. Given that the model must initially be loaded into the Edge TPU's memory, the first inference operation requires a longer duration to complete. To ensure the consistency of data collection, a set of twenty inferences is executed prior to any inference time measurements. Subsequently, a series of 100 inferences is performed on the sample test image and the time for each inference is measured. The mean and standard deviation of all the inference times are then computed. This procedure is replicated for both ELUNet and MobileUNet, and the results are consolidated in Table VI.

TABLE VI
Edge TPU Inference Results

Model	Mean (ms)	σ (ms)	Int8 File Size (MB)
ELUNet	808.3532	83.3374	15.5
MobileUNet	265.5293	21.2383	6.7

MobileUNet performs exceptionally with a frequency of 3.77 Hz. ELUNet's mean inference time is under 1000 milliseconds with a frequency of 1.237 Hz. However, the performance of ELUNet is noticeably worse. This lower

performance in terms of inference time can be ascribed to ELUNet having to perform an increased number of computations on higher-resolution feature maps, a consequence of its deep skip connections, despite it having fewer model parameters.

C. Static Hardware-in-the-loop Test

To gather data on the performance of our models, we constructed a starfield simulator. Creating our own starfield simulator enables us to easily control the conditions of the test and conduct them at any time. Specifically, our simulator system consists of a screen, a collimating lens, and a testing camera. The screen is a 5" diagonal 1920x1080 resolution IPS monitor and the achromatic collimation lens has a focal length of 350mm, a distance that lets the monitor fit nicely in the field of view of our camera. These components are mounted on optical posts, which slide along an optical rail to align and adjust their positions. The camera specifications are detailed in Table VII. An image of the simulator setup is shown in Fig. 4.
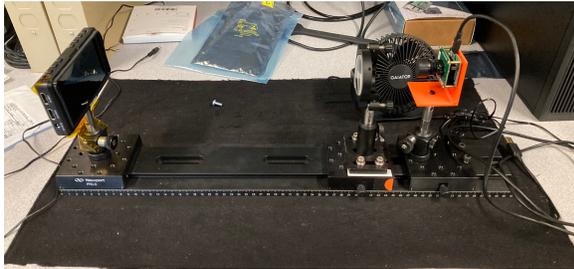
Fig. 4.   Starfield simulator allowed us to evaluate our method with real camera images.

TABLE VII

Camera Specifications

Sensor	CMOS Monochrome Global Shutter
Resolution	640x480
Focal Length	16mm
Full Field of View	13.8x10.3 Degrees
F/ratio	1.2

The simulator functions by displaying a set of stars on the screen from the Hipparcos Catalog, only showing stars brighter than visual magnitude 6. The camera views this visual field through the collimating lens, which is placed to focus the starfield to infinity. The images captured by the camera are then fed into the star tracking algorithms for our tests. A cover lined with light-absorbing fabric is placed over the setup to ensure minimal light leakage into the simulator. With this setup, there are some limitations. Since our screen has a limited spatial and luminance resolution, the positions and magnitude of stars are displayed with limited accuracy. The star vectors, on average, will have 12.9 arcseconds of angular distance error. The small dot pitch of our chosen monitor helps minimize this error. Based on lux measurements of the

TABLE VIII

Static Detection Results

Model	Cold			Hot		
	Precision	Recall	F1	Precision	Recall	F1
ELUnet	100	99.4	99.7	100	99.4	99.7
MobileUnet	100	96.8	98.4	99.3	97.4	98.4
Liebe's	94.9	47.5	63.3	100	43.2	60.4
WITM	100	61.4	76.1	89.2	84.1	86.6
ST16	100	37.3	54.4	100	24.1	38.8
Sun *et al.*'s	100	43.0	60.2	98.9	59.2	74.1

screen, we also estimate that the displayed stars have an apparent magnitude of 2.35.

With the simulator setup complete, we were able to run tests to compare our neural network and conventional baseline models. The first series involved a static starfield to measure performance, with our detection metrics being the precision, recall, and F1 scores defined in (13). To maintain consistency between tests, the simulator is calibrated by aligning the camera image center with the center of the monitor, establishing a fixed optical path. In addition, a fan is placed near the camera to provide cooling when desired, controlling the amount of thermal noise present in the camera images. Ten celestial fields of view were selected to provide test images for the algorithms. The camera integration time was fixed at 150 ms for all tests, and distortion correction was implemented using calculated OpenCV distortion coefficients. The parameters for each traditional method were kept the same for all tests; these were set to values with the best detection results for a real night sky image with 150 ms integration time. The results for the static detection are shown in Table VIII.

For both hot (noisy) and cold camera operating conditions, the neural networks performed comparably with near-perfect detection scores. Between the baseline methods, there is a notable difference between hot and cold operating conditions, with WITM and Sun *et al.*'s methods improving in F1 score while the Liebe and ST16 methods had lower scores. We note that the traditional method performances are significantly worse here than when we tested them using synthetic data. We hypothesize that the parameters used for each method may not be fully optimized for these circumstances. However, this illustrates the robustness of our neural network models and how they do not need continuous recalibration to achieve good detection performance, surpassing all tested traditional methods in precision, recall, and F1 scores.

Regarding the centroiding tests, we compared our neural networks to both the Center of Gravity and Gaussian Grid methods each paired with the WITM detection method, since it displayed the highest detection scores among the baseline methods. Since our static tests showed that our neural networks are not impacted by thermal noise and that the WITM model exhibited the highest detection performance in a hot camera operating condition, we tested centroiding accuracy with no cooling on the camera. To measure centroiding accuracy, we compared

TABLE IX
Static Centroiding Results

Model	Avg. Error (arcsec)	σ (arcsec)
ELUnet	56	50.8
MobileUnet	57.9	53.6
WITM + Gaussian Grid	58.2	49.2
WITM + Center of Gravity	71.2	61.1

the difference in the angular separations of stars computed by each algorithm versus their true values displayed on the simulator screen. An average of 300 star pairs were compared for all algorithms. The static centroiding results are shown in Table IX.

Due to the limited-resolution error and any potential distortion calibration errors, the centroiding results may not represent the true capability of the algorithms. In addition, some stars near the far edge of the screen appear stretched due to optical distortion, which makes centroiding inaccurate for those stars. To counter this, we removed outlier stars that had average angular errors of greater than 200 arcseconds. Overall, since they were all tested in the same system and with the same calibration, the consistency enables comparison between the algorithms. Specifically, the neural networks are shown to have significantly improved centroiding accuracy to the Center of Mass method, while remaining competitive with the Gaussian Grid method.

Finally, we also tested performance when stray light is purposefully added to deteriorate the image quality. This was done by placing a flashlight inside the simulator, resulting in light flares in the captured images. The results for detection and centroiding accuracy are shown in Table X. Only our neural network models are tabulated for stray light results, as all other conventional models failed to identify any stars and were plagued with false positives when exposed to the same lighting conditions.

TABLE X
Stray Light Performance

Model	Precision	Recall	F1	Avg Error	σ
ELUnet	94.4	96.8	95.6	60.5	51.9
MobileUnet	87.9	97.4	92.4	61.2	50.3

Even with stray light contamination, our proposed neural network methods display only slightly worse detection metrics while all other traditional models displayed a multitude of false positives. While the precision scores of the models are slightly less due to the appearance of some false positives, the neural network models are still able to consistently and correctly identify true stars in heavily contaminated images, as evidenced by their precision, recall, and F1 scores. Our models also retained their high accuracy under these conditions, averaging only about 4 arcseconds worse than the non-contaminated static tests.

D. Dynamic Hardware-in-the-loop Test

The second series of testing involved a rotating starfield to evaluate detection performance. The simulator is calibrated in the same manner as the static tests. Since the results of the static tests demonstrated the quality of performance with varying levels of thermal noise, the dynamic tests were only conducted at one temperature. Ten fields of view (the same as the static tests) were chosen for the starfields. For each starfield, three separate rates of rotation were used to dictate the manner in which the displayed starfield moved across the monitor. Each method (ELUnet, MobileUnet, WITM, and Sun *et al.*'s) was tested with a rotating starfield a total of thirty times. Fig. 5 demonstrates the star image quality depending on rotation rate. A random time was chosen to capture the stills used for analysis. By comparing the number of stars detected in the captured images to the simulated images and applying the metrics in Eq. 13, we were able to come up with an F1 score that measures the accuracy of each method in detecting stars.

(a)   (b)   (c)
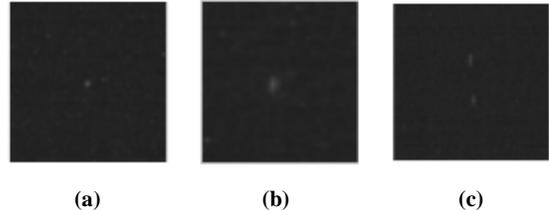
Fig. 5. Images depicting star pixels with zero rotation (a), ω=[0.5, 0, 0] (b), and ω=[1, 0, 0] (c) show that higher rotation rates produce more elongated star pixels in the starfield images.

The results for dynamic detection for ω=[0.5, 0, 0], ω=[0, 0, 6], and ω=[1, 0, 0] are shown in Table XI. For all three rotation rates, our proposed two neural network methods perform the best with F1 scores above 98. The WITM method achieved slightly lower scores due to having more false positives. Nonetheless, it still performed significantly better than Sun *et al.*'s method, which struggled notably in distinguishing between true stars and false positives.

The second phase of dynamic testing involved computing centroid accuracy by comparing the actual star angular distances from the star catalog with the angular distances computed by the algorithms. This evaluation was conducted across three different starfields, each featuring three distinct rates of rotation. This resulted in nine test images for each method (MobileUnet, ELUnet, and WITM Gaussian), each containing multiple stars. Similar to the static centroiding tests, stars with an average angular error of greater than 200 arcseconds were considered outliers and removed. The results for all three rates of rotation are presented in Table XII. In all cases, the neural networks had a lower error than the WITM Gaussian method. Despite not being trained on dynamic images, the proposed method demonstrates robustness for star

TABLE XI

Dynamic Detection Results

Model	ω=[0.5, 0, 0]			ω=[0, 0, 6]			ω=[1, 0, 0]		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
ELUnet	100	100	100	100	100	100	99.3	100	99.9
MobileUnet	100	99.3	99.65	100	100	100	100	100	100
WITM	99.2	80.1	88.6	99.2	84.6	91.3	98.6	91.3	94.8
Sun *et. al's*	98.4	41.2	58.1	100	38	55.1	100	50.3	67.0

TABLE XII

Dynamic Centroiding Results

Model	ω=[0.5, 0, 0]		ω=[0, 0, 6]		ω=[1, 0, 0]	
	Error(arcsec)	σ (arcsec)	Error (arcsec)	σ (arcsec)	Error (arcsec)	σ(arcsec)
ELUnet	69.59	61.02	72.17	61.73	59.3	46.6
MobileUnet	62.56	52.52	69.03	62.53	69.1	52.5
WITM + Gaussian Grid	69.82	53.69	74.75	61.62	82.72	56.37

images with motion blur, consistently producing accurate centroids.

E. Night Sky Testing

While hardware-in-the-loop tests offer control over temperature and rotation rates, our current hardware setup has limitations. It can only provide a coarse estimation of accuracy performance and lacks the capability to simulate faint stars. To achieve a more comprehensive evaluation, we performed night sky testing at Middle Fork River Forest Preserve, IL, U.S. The camera, detailed in Table VII, was mounted on a tripod and aimed near the zenith with its exposure time set to 500 ms. To align with standard practices [5], the inertial-to-body attitude outputs were converted to 3-2-1 Euler angles representing right ascension, declination, and roll. The right ascension would be expected to change linearly at Earth's sidereal rate, while the declination and roll remain constant. The cross-boresight and around-boresight accuracy of a star tracker can be modeled as the standard deviations of the declination and roll. For star identification, we employed the geometric voting algorithm [26], and the singular value decomposition method [27] was utilized for attitude determination. Recorded videos of the test results can be found at our project website https://hongruizhao.github.io/CNNStarDetectCentroid/. We recorded two 10-minute videos to evaluate the detection and accuracy of our method. Results are presented in Table XIII. Among conventional methods, Sun *et al.*'s local threshold approach demonstrated the best detection performance and served as our comparison benchmark. Overall, both CNN models outperformed traditional approaches, with ELUNet showing a slight edge over MobileUNet. As illustrated in Fig. 6, while our method and the conventional baseline exhibit similar detection performance, our method achieved a significantly higher star identification rate due to its superior centroiding accuracy.

To assess the impact of stray light, we first recorded a 5-minute video under standard conditions, immediately

TABLE XIII

Night Sky Test

Model	Cross-Boresight (arcsecond 1-σ)		Around Boresight (arcsecond 1-σ)	
	Test 1	Test 2	Test 1	Test 2
ELUnet	12.92	13.79	76.95	151.90
MobileUnet	12.93	15.44	81.49	146.88
Sun *et al.*'s + Center of Gravity	21.26	22.99	251.72	290.84
Sun *et al.*'s + Gaussian Grid	19.84	19.34	221.09	253.34
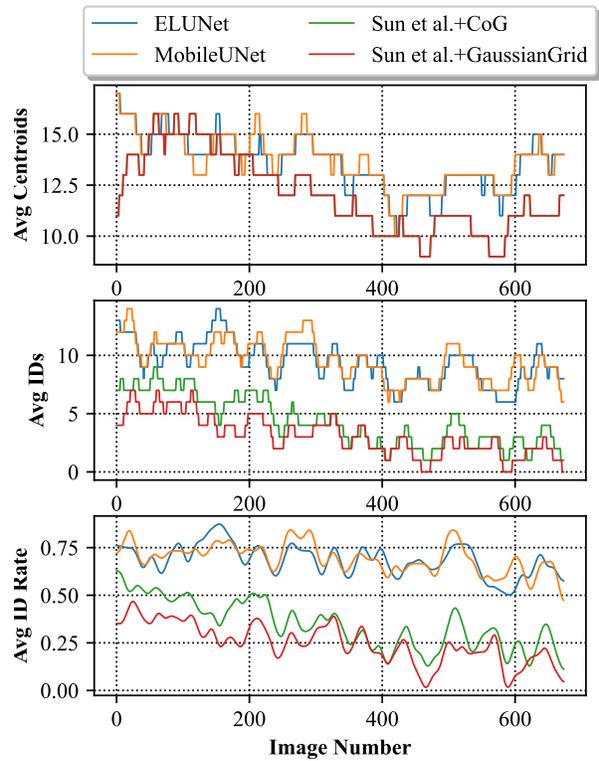
Fig. 6. Detection and identification performance in the first night sky test. With a similar number of star centroids detected, the proposed method has higher star identification rates compared to the conventional methods.

TABLE XIV
Night Sky Test, Stray light

Model	Cross-Boresight (arcsecond 1-σ)		Around Boresight (arcsecond 1-σ)	
	w/o Straylight	Straylight	w/o Straylight	Straylight
ELUnet	8.07	20.25	85.85	191.73
MobileUnet	8.07	23.09	85.85	216.96

(a)   (b)
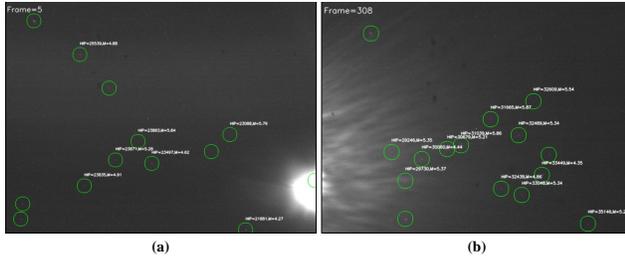
Fig. 7. Despite (a) the presence of the moon within the field-of-view or (b) severe lens flare, the proposed method leads to accurate star detection and identification.

followed by a subsequent 5-minute video with stray light introduced via a flashlight. The results are presented in Table XIV. While both models are affected by stray light, their accuracy remains comparable to typical commercial CubeSat star trackers. Under normal conditions, ELUNet and MobileUNet perform identically. However, when exposed to stray light, ELUNet demonstrated greater robustness. Fig. 7 showcases visual examples of ELUNet prcoessing results of images affected by stray light and those with the moon in the field-of-view. Detected stars are highlighted with green circles, while identified stars display their information at the top. Even when confronted with significant lens flare or the moon's interference, our method reliably rejected false detections and precisely computed star centroids, leading to successful star identification.

V. Conclusion

In this article, we introduced a CNN-based method for real-time star detection and centroiding in star tracker processing. We generated training images by blending synthetic "clean" star images with authentic sensor noise and stray light. The CNN model yields a binary segmentation map for star detection and a distance map for centroiding. By incorporating distance data with pixel coordinates, centroid calculation was reformulated as a trilateration problem, which is solved using the least squares method. The chosen CNN models operates with under 300 ms latency on a low-power Google Coral TPU. We conducted synthetic image tests, hardware-in-the-loop tests, and night sky tests to thoroughly assess our method against existing approaches. Our method not only surpassed existing methods in detection and centroiding performance but also exhibited remarkable robustness across varying thermal, rotation, and stray light conditions. Moving forward, our research will progress to

constructing a prototype star tracker to further validate our method in an actual flight-qualified setup. Additionally, a potential enhancement could involve training the CNN to process star trails resulting from higher rotation rates.

REFERENCES

[1] K. M. Riesing *et al.*
On-orbit results of pointing, acquisition, and tracking for the TBIRD CubeSat mission
In *Proc. SPIE LASE*, San Francisco, CA, USA, 2023.
[2] D. Sternberg, J. Essmiller, C. Colley, A. Klesh, and J. Krajewski
Attitude control system for the Mars cube one spacecraft
2019 IEEE Aerospace Conference, Big Sky, MT, USA, 2019.
[3] J. Kruger, A. W. Koenig, S. D. Amico
Starling formation-flying optical experiment (StarFOX): system design and preflight verification
J Spacecr Rockets, vol. 60, no. 6, Nov. 2023.
[4] CubeSpace
(2023, Jan).
Gen 1 Sensors Actuators. [Online]. Available: link
[5] C.C Liebe
Accuracy performance of star trackers - a tutorial
IEEE Transactions on Aerospace and Electronic Systems, vol. 38, no. 2, pp. 587-599, Apr 2002.
[6] T. Delabie, J.D. Schutter, B. Vandenbussche
An accurate and efficient Gaussian fit centroiding algorithm for star trackers
J of Astronaut Sci, vol. 61, pp. 60-84, 2014.
[7] W. Xu, Q. Li, H. Feng, Z. Xu, Y. Chen
A novel star image thresholding method for effective segmentation and centroid statistics
Optik, vol. 124, no. 20, pp. 4673-4677, Oct 2013.
[8] T. Dzamba *et al.*
Success by 1000 improvements: flight qualification of the ST-16 star tracker
Small Satellite Conference, Logan, UT, USA, 2014.
[9] T. Sun, F. Xing, Z. You, M. Wei
Motion-blurred star acquisition method of the star tracker under high dynamic conditions
Opt. Express, vol. 21, no. 17, pp. 20096-20110, 2013.
[10] Rocket Lab
Second Generation Star Tracker (ST-16RT2). [Online]. Available: link
[11] eoPortal
MarCO (Mars Cube One). [Online]. Available: link
[12] K. Dijkstra, J. van de Loosdrecht, L.R.B. Schomaker, M.A Wiering
CentroidNet: a deep neural network for joint object localization and counting
ECML PKDD, Turin, Italy, 2018.
[13] C. Chen, Q. Chen, J. Xu, V. Koltun
Learning to see in the dark
CVPR, Salt Lake City, UT, USA, Jun 2018.
[14] M. Mastrofini, I. Agostinelli, F. Curti
Design and validation of a U-Net-based algorithm for star sensor image segmentation
Applied Sciences, vol. 13, no. 3, Feb 2023.
[15] H. Liu, X. Wei, G. Wang, J. Li
A star spot extraction method based on SSA-UNet for star sensors under dynamic conditions
IEEE Sensors Journal, vol. 23, no. 7, pp. 7410-7419, Apr 2023.
[16] J. Jiang, K. Chen
FPGA-based accurate star segmentation with moon interference
Jorunal of Real-Time Image Processing, vol. 16, pp. 1289-1299, 2019.

[17] O. Ronneberger, P. Fischer, T. Brox
U-Net: convolutional networks for biomedical image segmentation
MICCAI, Munich, Germany, Oct 2015.

[18] J. Jing, Z. Wang, M. Rätsch, H. Zhang
Mobile-Unet: an efficient convolutional neural network for fabric defect detection
Textile Research Journal, vol.92, no. 1-2, 2022.

[19] Y. Deng, Y. Hou, J. Yan and D. Zeng
ELU-Net: an efficient and lightweight U-Net for medical image segmentation
IEEE Access, vol.10, pp. 35932-35941, 2022.

[20] N. Beheshti, L. Johnsson
Squeeze U-Net: a memory and energy efficient image segmentation network
CVPRW, Seattle, WA, USA, 2020.

[21] European Space Agency
The Hipparcos and Tycho catalogues
ESA SP-1200, 1997.

[22] P. Léna, D. Rouan, F. Lebrun, F. Mignard, D. Pelat
Observational Astrophysics, 3rd ed., Heidelberg, Germany: Springer Berlin, 2012.

[23] W.K. Merline, S.B. Howell
A realistic model for point-sources imaged on array detectors: The model and initial results
Experimental Astronomy, vol. 6, pp. 163-210, 1995.

[24] K. Monakhova, S.R. Richter, L. Waller, V. Koltun
Dancing under the stars: video denoising in starlight
CVPR, New Orleans, LA, USA, Jun 2022.

[25] E. Anderson *et al.*
LAPACK Users' Guide, 3rd ed.
Soc. Ind. Appl. Math., 1999.

[26] M. Kolomenkin, S. Pollak, I. Shimshoni, M. Lindenbaum,
Geometric voting algorithm for star trackers
IEEE Transactions on Aerospace and Electronic Systems, vol. 44, no. 2, pp. 441-456, Apr 2008.

[27] F.L. Markley,
Attitude determination using vector observations and the singular value decomposition
J of Astronaut Sci, vol. 38, no. 3, pp. 245-258, 1987.

[28] G. Furano *et al.*
Towards the use of artificial intelligence on the edge in space systems: challenges and opportunities
IEEE Aerospace and Electronic Systems Magazine, vol. 35, no. 12, pp. 44-56, Dec 2020.

[29] H. Ding, J. Pu, C. Hu
TinyNeuralNetwork: an efficient deep learning model compression framework
2021. [Online]. Available: link

[30] Coral
(2020).
Edge TPU Compiler. [Online]. Available: link

Michael F. Lembeck was born in Peoria, IL. He has B.S. (1980), M.S.(1981), and Ph.D. (1991) degrees in Aeronautical and Astronautical Engineering from the University of Illinois at Urbana-Champaign (UIUC).

He is a Clinical Associate Professor and Director of the Laboratory for Advanced Space Systems at Illinois (LASSI) providing students with opportunities to design, build, and fly experimental CubeSat satellites. He has led or worked on multiple government and commercial spaceflight programs, including JPL's Galileo Jupiter Orbiter and Space Industries, Inc.'s Wake Shield Facility-03 flown on STS-80, Orbital Sciences' OrbView-3 and OrbView-4/Warfighter satellites, and he was the Requirements Division Director for the Exploration Systems Mission Directorate at NASA Headquarters.

Dr. Lembeck is an Associate Fellow of the American Institute of Aeronautics and Astronautics (AIAA).

Adrian Zhuang is a third-year undergraduate student pursuing a B.S. degree in aerospace engineering at the University of Illinois at Urbana-Champaign. He is performing research at the Laboratory for Advanced Space Systems at Illinois on star tracker systems and test systems.

Riya Shah is a third-year undergraduate student studying aerospace engineering at the University of Illinois at Urbana-Champaign. She works at the Laboratory for Advanced Space Systems at Illinois (LASSI) on the star-tracker research project.

Jesse Wei is a third-year undergraduate student studying aerospace engineering at the University of Illinois Urbana-Champaign. He is working for the Laboratory for Advanced Space Systems at Illinois on the star-tracker research project.

Hongrui Zhao was born in China in 1995. He holds an M.S in aerospace engineering from the University of Illinois Urbana Champaign (UIUC), 2020, and a B.S. in aerospace engineering from Nanjing University of Aeronautics and Astronautics, 2018.

He is a PhD candidate in aerospace engineering at UIUC. He is the GNC and star tracker research lead at the Laboratory for Advanced Space Systems at Illinois. He has worked on multiple UIUC CubeSat missions as a GNC engineer, including CubeSail, SASSI2, and CAPSat. He has worked as a GNC engineer intern at CU Aerospace from 2022 to 2023. Currently, his research focuses on applying artificial intelligence techniques to address challenges in space applications.