MDPI

Article

Strapdown Celestial Attitude Estimation from Long Exposure Images for UAV Navigation

Samuel Teague [1,*] and Javaan Chahl [1,2]

1   School of Engineering, University of South Australia, Mawson Lakes, SA 5095, Australia
2   Joint and Operations Analysis Division, Defence Science and Technology Group, Melbourne, VIC 3207, Australia
*   Correspondence: samuel.teague@mymail.unisa.edu.au

Abstract: Strapdown celestial imaging sensors provide a compact, lightweight alternative to their gimbaled counterparts. Strapdown imaging systems typically require a wider field of view, and consequently longer exposure intervals, leading to significant motion blur. The motion blur for a constellation of stars results in a constellation of trails on the image plane. We present a method that extracts the path of these star trails, and uses a linearized weighted least squares approach to correct noisy inertial attitude measurements. We demonstrate the validity of this method through its application to synthetically generated images, and subsequently observe its relative performance by using real images. The findings of this study indicate that the motion blur present in strapdown celestial imagery yields an a posteriori mean absolute attitude error of less than 0.13 degrees in the yaw axis, and 0.06 degrees in the pitch and roll axes ($3\,\sigma$) for a calibrated wide-angle camera lens. These findings demonstrate the viability of low-cost, wide-angle, strapdown celestial attitude sensors on lightweight UAV hardware.

Keywords: celestial; stellar; navigation; strapdown; attitude

1. Introduction

The use of stabilized celestial navigation sensors for uncrewed aerial vehicle (UAV) attitude determination is well documented [1]. With recent demand for size, weight and power constrained systems, strapdown celestial sensors have become more common. A strapdown [2] celestial navigation sensor is rigidly mounted to the airframe, causing imagery to be subjected to motion artefacts from the aircraft, such as actuation, vibration and turbulence. The length of the exposure window is the primary factor governing the severity of the resultant motion blur. For wide-angled lenses, it is necessary to use longer exposure windows so as to increase the total light energy incident on the sensor. Under stable flight conditions, longer exposure windows enable the detection of higher magnitude stars, and consequently provide a more accurate attitude estimate. Under motion, however, the longer exposure window results in "smearing" of star images, leaving a trail as seen in Figure 1. This image shows a region of interest (ROI) containing a single star trail captured in-flight from a strapdown celestial imaging system. We can see from this ROI that the resultant trail tends to be noisy, and the angular velocity tends to change throughout the exposure interval.

The premise for this research comes from the hypothesis that the observed star trails contain high-resolution information pertaining to the attitude of the aircraft during the exposure window. We present a method which estimates high-resolution attitude data from long-exposure images, provided availability of a low resolution approximation from the autopilot (e.g., from an inertial measurement unit). This method makes use of long-exposure strapdown imagery simulation, presented in [3], to provide an initial approximation of the star trail location and orientation, and corrects for attitude and attitude rate errors from the

inertial unit. We demonstrate the effectiveness of this method in estimating aircraft attitude under high motion conditions.

Figure 1. A region of interest containing a single star trail, captured from a strapdown celestial imaging sensor (Pi Camera HQ, 500 ms exposure interval). The shape of the star trail indicates that the camera was subjected to significant changes in attitude throughout the exposure interval.

The method presented in this paper is unique in the domain of drone navigation. Similar research has been conducted in the field of satellite navigation for the removal of motion blur. A common technique for the removal of motion blur includes estimation of the blur kernel. Knowledge of the kernel enables methods such as inverse filtering and Wiener filtering to correct the motion-blurred image. Estimation of the blur kernel is typically straightforward when analyzing stellar imagery, due to the sparsity of the stars; thus, it is is common practice to estimate the kernel parameters as seen in [4] and apply filters that utilize this kernel, as seen in [5,6]. While offering an effective means of removing motion blur in satellite imagery, such approaches tend not to capture rotational motion about the optical axis, due to the assumption that the blur kernel is spatially invariant. The method presented in [7] does explicitly take into account rotation about the boresight; however, the linear approximation used in this method is not applicable for UAV applications.

The work presented in [8] estimates the motion parameters of a given image, and uses this information to estimate the centroid of a blurred star. This assumes that there exists a global degradation function, which tends not to be the case for UAV navigation. The method presented in [9] uses an attitude correlated frames approach, in which the attitude between frames is measured with a gyroscope to within 1 arcsecond precision. This work also assumes that the rotation about the optical axis is negligible. The approach in [10] identifies correlation between subsequent short-exposure frames, and superimposes these frames to generate a long exposure image. This approach is limited by the sensor sensitivity, however, due to the need to identify stars from the shorter exposure images.

We can see in Figure 1 that the assumption of a spatially invariant blur kernel does not hold true when factors such as turbulence and aircraft control are taken into consideration. The aircraft will typically experience rotation about the yaw axis to some extent, leading to spatial variations in the motion blur on the imaging plane. Additionally, correlated frames approaches, such as those seen in [9,10], require levels of sensitivity from the imaging equipment that are not achievable with lightweight hardware at low altitudes. Furthermore, inertially aided approaches to noise removal tend to assume that gyroscope measurements contain negligible error, as seen in [9]. We observe that, especially with low-cost UAV hardware, angular rate measurements tend to be subjected to multiple sources of noise, and thus do not offer the level of precision required to perform image stabilization.

The method presented here is unique in that no attempt is made at denoising the image. Rather, we detect points from each star trail that are correlated, and infer the attitude by using a least-squares approximation that corrects the noisy inertial navigation system (INS) attitude measurements. This approach gives rise to potential use cases such

as in-flight magnetometer calibration and thermal photogrammetry. Additionally, this may be used to correlate points from different star trails within a single frame, enabling the use of traditional point-source celestial imaging techniques, such as camera calibration [11] and star identification/tracking [12].

2. Methods

The methodology outlined in this section assumes that the camera orientation, relative to the aircraft body frame, is calibrated and fixed throughout the flight. It also assumes that the camera calibration matrix is known. These calibrations are conducted prior to takeoff.

The following steps are performed in attitude estimation:

1. Estimate the theoretical curve of the star trail on the image plane by using INS measurements.
2. Apply a smoothing filter, morphological operations, and clustering to extract the star trail for each star with brightness above a given magnitude threshold.
3. Apply a thinning algorithm on each star trail to remove the effects of Gaussian point-spread diffusion.
4. Identify the endpoints of each star trail given the INS-simulated approximation.
5. Use the endpoints of the thinned star trails, along with the endpoints of the INS approximation, to compute the weighted least squares approximation for the mean attitude offset throughout the exposure window.
6. For each point in the mean-error corrected INS approximation, compute the least squares approximation of the precise attitude offset.

This method is valid for star trails which form a simple curve with observable endpoints on the image plane. Complex curves create ambiguity in the apparent motion of the airframe.

2.1. Image Processing

We denote the series of n INS attitude measurements (pertaining to a long-exposure image) as

$$\mathbf{r}_i = \begin{bmatrix} \phi \\ \theta \\ \psi \end{bmatrix} \tag{1}$$

for attitude measurement i with roll ϕ, pitch θ and yaw ψ. We compute the theoretical curve of the star trails on the image plane following the methodology in [3], denoted $\mathbf{s}_j$, for star j. Initial corrections to right ascension and declination are applied given annual proper motion, precession, nutation, and aberration given the location of the aircraft and time of flight. These corrections are only applied once per flight, prior to further calculations.

Given the hour-angle of a star, ω, the local elevation El of a celestial body given declination δ and latitude Lat is computed:

$$El = \operatorname{asin}\big(\sin(\delta)\, \sin(Lat)\, +\, \cos(\delta)\cos(Lat)\cos(\omega)\big) \tag{2}$$

and the local azimuth, Az, is given by:

$$Az = \operatorname{atan2}\big(\sin(\omega)\,,\, \cos(\omega)\sin(Lat)\, -\, \tan(\delta)\cos(Lat)\big) + \pi. \tag{3}$$

We subsequently correct for refraction, given by

$$R = \frac{1.02}{\tan\left(El + \dfrac{10.3}{El + 5.11}\right)}, \tag{4}$$

where R is the refractive distance, expressed in arcminutes; thus, the apparent elevation of a star El' is then given by

$$El' = El + R. \tag{5}$$

Given the azimuth and elevation of a star, the corresponding unit vector in northeast-down (NED) coordinates is given by

$$\begin{aligned} X &= \cos(Az)\cos(El') \\ Y &= \sin(Az)\cos(El') \,. \\ Z &= -\sin(El') \end{aligned} \tag{6}$$

The rotation from the local NED frame to the aircraft frame is computed from the roll, pitch, and yaw of attitude measurement i, represented as matrix $\mathbf{C}_{a/ned}$. The transformation from aircraft to camera frame, $\mathbf{C}_{c/a}$, remains constant, and is determined by the orientation of the camera with respect to the inertial unit. The unit vector in local NED coordinates $\mathbf{v}^l$ is then transformed to the camera frame of reference, $\mathbf{v}^c$:

$$\mathbf{v}^c = \mathbf{C}_{c/a}\mathbf{C}_{a/l}\mathbf{v}^l. \tag{7}$$

For components x, y and z of unit vector $\mathbf{v}^c$, we compute the homogeneous point $\mathbf{P}$ in the camera frame of reference:

$$\mathbf{P} = \begin{bmatrix} \dfrac{X}{Z} \\ \dfrac{Y}{Z} \\ 1 \end{bmatrix}. \tag{8}$$

The camera intrinsic matrix, $\mathbf{K}$, is assumed to be known, given by

$$\mathbf{K} = \begin{bmatrix} f_x & 0 & x_0 \\ 0 & f_y & y_0 \\ 0 & 0 & 1 \end{bmatrix}, \tag{9}$$

where f_x and f_y are the focal lengths in the x and y axes, respectively, and x_0 and y_0 are the x and y locations of the principal point on the image plane. Thus, the pixel location, $\mathbf{s}_j[i]$, of the star for attitude measurement i is computed:

$$\mathbf{s}_j[i] = \mathbf{KP}. \tag{10}$$

Each $\mathbf{s}_j$ contains n two-dimensional points on the image plane, corresponding with the n attitude measurements. The prior calculations are performed n times for each star visible in the theoretical camera field of view to produce the array $\mathbf{s}_j$. Reference stars are selected based on their intensity, such that the brightest o stars in the frame are chosen. For each theoretical star trail, we apply a series of image-processing operations on the real image, as shown in Figure 2.

We first extract the ROI from the real image, given the theoretical star trail. A buffer is applied to the height and width of the ROI to allow for INS errors, typically corresponding to 2–3° of angular deviation. Gaussian blur is applied to the image with a 3×3 kernel, so as to reduce the magnitude of the image noise. A binary threshold is applied at five standard deviations above the average pixel value, as measured from the original full-scale image. This threshold is typical of stellar imaging sensors [13]. Image opening is performed with a 3×3 kernel to remove any remaining noise. The remaining contours are clustered, originating with the centre-most contour, and accumulating additional contours which are within 0.25° angular separation of the clustered set. Finally, the clustered contours are thinned by using the method presented in [14], so as to extract the centre-line from the star trail. Disjoint sections in the thinned image are connected by a straight line segment.

The thinned image is reduced to an array of two-dimensional points that are ordered from endpoint to endpoint. We denote this array as $\mathbf{p}_j$, for star j.
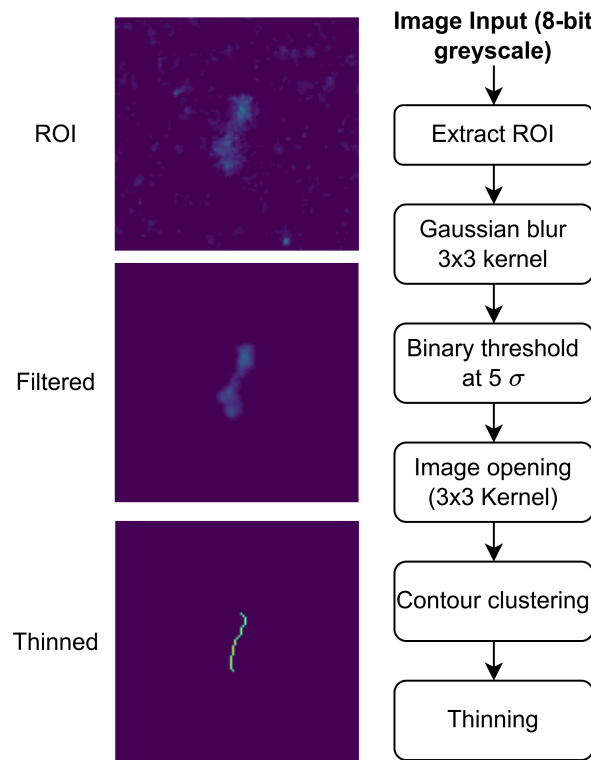
Figure 2. Flow diagram of image processing chain, with example images (black and white images converted to a perceptually uniform colour scale).

There is a possibility that the correspondence between endpoints in $\mathbf{p}_j$ and $\mathbf{s}_j$ is reversed. We use the approximate orientation from $\mathbf{s}_j$ to resolve the polarity of $\mathbf{p}_j$. We measure the angle on the image plane, θ_s, between the first and last elements of $\mathbf{s}_j$, as well as θ_p, the angle between the first and last elements of $\mathbf{p}_j$. If the magnitude of the difference between these angles exceeds $\frac{\pi}{2}$, we reverse the indices of array $\mathbf{p}_j$ to match the orientation of $\mathbf{s}_j$.

2.2. Orientation Estimation

Orientation estimation is performed in two steps. We first compute the mean attitude correction required to correlate the endpoints in $\mathbf{s}$ with the endpoints in $\mathbf{p}$. Once aligned, we compute the individual offsets for each of the n INS attitude measurements.

The rotation which transforms a vector from real-world coordinates to camera coordinates via a yaw-pitch-roll Euler sequence, is given by

$$\mathbf{R} = \begin{bmatrix} \mathrm{c}(\theta)\mathrm{c}(\psi) & \mathrm{c}(\theta)\mathrm{s}(\psi) & -\mathrm{s}(\theta) \\ -\mathrm{c}(\phi)\mathrm{s}(\psi) + \mathrm{s}(\phi)\mathrm{s}(\theta)\mathrm{c}(\psi) & \mathrm{c}(\phi)\mathrm{c}(\psi) + \mathrm{s}(\phi)\mathrm{s}(\theta)\mathrm{s}(\psi) & \mathrm{s}(\phi)\mathrm{c}(\theta) \\ \mathrm{s}(\phi)\mathrm{s}(\psi) + \mathrm{c}(\phi)\mathrm{s}(\theta)\mathrm{c}(\psi) & -\mathrm{s}(\phi)\mathrm{c}(\psi) + \mathrm{c}(\phi)\mathrm{s}(\theta)\mathrm{s}(\psi) & \mathrm{c}(\phi)\mathrm{c}(\theta) \end{bmatrix}, \quad (11)$$

where $\mathrm{c}(x)$ and $\mathrm{s}(x)$ represent $\cos(x)$ and $\sin(x)$, respectively.

Stars are framed in the celestial coordinate system; thus, translation is negligible. Therefore, the image coordinates, $\mathbf{x}$, for infinitely distant objects, is given by

$$\mathbf{x} = \mathbf{KRX}, \quad (12)$$

where **X** is the vector containing the local NED world coordinates at a given point:

$$\mathbf{X} = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}. \tag{13}$$

We expand the vector **x** into its components to get the image coordinates x, y, and z:

$$\begin{aligned} x = f_x \big[& (\cos\theta\cos\psi)X + (\cos\theta\sin\psi)Y - (\sin\theta)Z \big] \\ & + p_x \big[(\sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi)X \\ & + (-\sin\phi\cos\psi + \cos\phi\sin\theta\sin\psi)Y + (\cos\phi\cos\theta)Z \big] \end{aligned} \tag{14}$$

$$\begin{aligned} y = f_y \big[& (-\cos\phi\sin\psi + \sin\phi\sin\theta\cos\psi)X \\ & + (\cos\phi\cos\psi + \sin\phi\sin\theta\sin\psi)Y + (\sin\phi\cos\theta)Z \big] \\ & + p_y \big[(\sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi)X \\ & + (-\sin\phi\cos\psi + \cos\phi\sin\theta\sin\psi)Y + (\cos\phi\cos\theta)Z \big] \end{aligned} \tag{15}$$

$$\begin{aligned} z = \big(& \sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi \big)X + \big(-\sin\phi\cos\psi + \\ & \cos\phi\sin\theta\sin\psi \big)Y + \big(\cos\phi\cos\theta \big)Z. \end{aligned} \tag{16}$$

The two-dimensional homogeneous image pixel coordinates, are subsequently given by

$$\mathbf{u} = \begin{bmatrix} u \\ v \end{bmatrix}; \quad u = \frac{x}{z}, \ v = \frac{y}{z}. \tag{17}$$

The Jacobian containing partial derivatives of pixel location with respect to changes in orientation is then given by (see Appendix A.1 for partial derivative equations):

$$\mathbf{J} = \begin{bmatrix} \dfrac{\partial u}{\partial \phi} & \dfrac{\partial u}{\partial \theta} & \dfrac{\partial u}{\partial \psi} \\[2mm] \dfrac{\partial v}{\partial \phi} & \dfrac{\partial v}{\partial \theta} & \dfrac{\partial v}{\partial \psi} \end{bmatrix}. \tag{18}$$

Thus, the first order Taylor series expansion gives

$$\mathbf{u}' = \mathbf{u} + \mathbf{J}\Delta\mathbf{r}, \tag{19}$$

where $\Delta\mathbf{r}$ is the vector containing the change in roll, pitch, and yaw required to translate pixel **u** to $\mathbf{u}'$:

$$\Delta\mathbf{r} = \begin{bmatrix} \Delta\phi \\ \Delta\theta \\ \Delta\psi \end{bmatrix}.$$

The linearized relationship between change in pixel location and change in orientation can be expressed as

$$\Delta\mathbf{u} = \mathbf{J}\Delta\mathbf{r}. \tag{20}$$

We extend this notation for multiple observations, $\Delta\hat{\mathbf{u}}$, where each observation $\Delta\mathbf{u_i} = [\Delta u_i, \ \Delta v_i]^T$ is vertically stacked to give a vector of length $2m$:

$$\Delta\hat{\mathbf{u}} = \begin{bmatrix} \Delta u_1, \ \Delta v_1, \ \dots \ \Delta u_m, \ \Delta v_m \end{bmatrix}^T \tag{21}$$

and similarly the Jacobian $\mathbf{J_i}$ for each observation is vertically stacked to give the matrix of size $[2m \times 3]$:

$$\hat{\mathbf{J}} = \begin{bmatrix} \mathbf{J_1} \\ \vdots \\ \mathbf{J_m} \end{bmatrix}. \tag{22}$$

Thus, provided a minimum of $m = 2$ points, we can apply the weighted least squares solution for $\Delta \mathbf{r}$,

$$\Delta \mathbf{r} = \left(\hat{\mathbf{J}}^T \mathbf{W} \hat{\mathbf{J}}\right)^{-1} \hat{\mathbf{J}}^T \mathbf{W} \Delta \hat{\mathbf{u}}, \tag{23}$$

where the diagonal weight matrix $\mathbf{W}$ with size $[2m \times 2m]$ contains the weighting for each observation

$$\mathbf{W} = \begin{bmatrix} w_1 & 0 & \dots & 0 & 0 \\ 0 & w_1 & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & w_m & 0 \\ 0 & 0 & \dots & 0 & w_m \end{bmatrix}, \tag{24}$$

where w_i is calculated from the signal strength of observation i, such that salient stars are weighted more heavily,

$$w_i = \log_{10}\left(\frac{p_i - \mu}{\sigma}\right), \tag{25}$$

and p_i is the peak pixel intensity for observation i, μ is the mean value across the image ($p_i \gg \mu$), and σ is the standard deviation in pixel intensity in the image. Note that the weightings for the u and v components of a given observation are equal.

The mean-offset corrected camera attitude is computed iteratively, such that

$$\Delta \mathbf{r}_k = \Delta \mathbf{r}_{k-1} + \left(\hat{\mathbf{J}}_{k-1}^T \mathbf{W} \hat{\mathbf{J}}_{k-1}\right)^{-1} \hat{\mathbf{J}}_{k-1}^T \mathbf{W} \Delta \hat{\mathbf{u}}_{k-1} \tag{26}$$

until $|\Delta \mathbf{r}_k - \Delta \mathbf{r}_{k-1}| \leq 10^{-6}$ rad. Values $\Delta \hat{\mathbf{u}}_{k-1}$ and $\hat{\mathbf{J}}_{k-1}$ are recomputed at each iteration from Equations (17) and (18), given the updated attitude:

$$\mathbf{r}_i' = \mathbf{r}_i - \Delta \mathbf{r}_k. \tag{27}$$

Equation (26) yields the mean attitude offset throughout a given exposure window. In some cases, this level of precision is satisfactory (for example, for online magnetometer calibration). The long exposure image typically contains higher-resolution attitude information pertaining to the aircraft orientation throughout the exposure window. This information can be obtained by aligning elements from the mean-corrected array $\mathbf{s}_j'$ with elements from observation $\mathbf{p}_j$. An example demonstrating the difference between mean and fine alignment can be seen in Figure 3.

A similar process to the mean attitude correction is followed for the high-resolution attitude estimation. For each star, the theoretical curve of the trail is recomputed with the mean attitude offset applied to obtain $\mathbf{s}_j'$, and the polarity once again checked to ensure that the elements of $\mathbf{p}_j$ are in the correct order. We make the following assumptions when mapping $\mathbf{s}_j'$ to $\mathbf{p}_j$:

1. The INS sampling period, T_s, is constant.
2. The photon flux density incident on the sensor from a given luminary is constant.
3. The path taken by the airframe results in a simple curve on the image plane (i.e., the star trail does not cross itself at any point).

From these assumptions, it is evident that for each successive point in the fine-attitude corrected set, $\mathbf{s}''$, the rate of increase in the cumulative intensity must be constant,

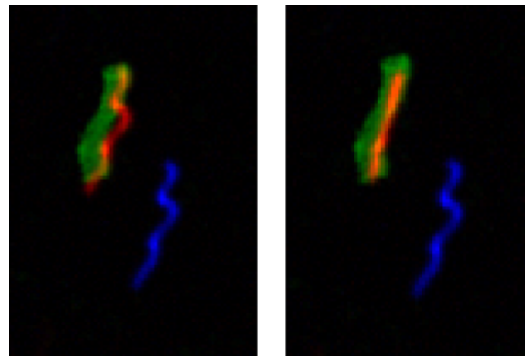$$\mathbf{I}\left(\mathbf{s}_j''[i]\right) - \mathbf{I}\left(\mathbf{s}_j''[i-1]\right) = C, \tag{28}$$

given pixel intensity $\mathbf{I}(x)$ at location x in the frame. Thus, the location of pixel $\mathbf{s}''[i]$ should be chosen to satisfy

$$\mathbf{I}(\mathbf{s}''_j[i]) = \mathbf{I}(\mathbf{s}''_j[i-1]) + C, \tag{29}$$

where the first element, $\mathbf{s}''_j[0]$ is equal to $\mathbf{p}_j[0]$, and C is chosen to ensure the last element in $\mathbf{s}''_j$ is equal to the last element in $\mathbf{p}_j$:

$$C = \frac{1}{n}\sum_{i=0}^{p}\mathbf{I}(\mathbf{p}_j[i]) \tag{30}$$

given n elements in $\mathbf{s}_j$, and p elements in $\mathbf{p}_j$.

Figure 3. An example of attitude correction, displaying a region of interest for a single star. Greyscale images are overlaid onto a three-channel image. **Left**: mean-only alignment, **Right**: fine attitude alignment. Green, real image; blue, synthetic image from INS; red, reprojection after corrections.

The candidates for $\mathbf{s}''_j[i]$ are contained in the ordered set of thinned points from the real image, $\mathbf{p}_j$. We use a cumulative sum of real image intensities along the skeleton to solve for $\mathbf{s}''_j$ empirically, as shown in Algorithm 1:

Algorithm 1 Mapping from INS points to real image points.

$\quad\mathbf{s}''_j[0] \leftarrow \mathbf{p}_j[0]$
$\quad n \leftarrow 1$
$\quad sum \leftarrow 0$
$\quad$**for** $i = 1; i < p$ **do**
$\quad\quad sum \leftarrow sum + \mathbf{p}_j[i]$
$\quad\quad$**if** $sum \geq C$ **then**
$\quad\quad\quad overshoot \leftarrow (sum - C)$
$\quad\quad\quad \mathbf{s}''_j[n] \leftarrow Interpolate(\mathbf{p}_j[i], \mathbf{p}_j[i-1], overshoot)$
$\quad\quad\quad sum \leftarrow overshoot$
$\quad\quad\quad n \leftarrow n + 1$
$\quad\quad$**end if**
$\quad$**end for**

For each observed star, there exists a set of mean attitude corrected points, $\mathbf{s}'_j$, and a set of fine attitude corrected points, $\mathbf{s}''_j$. For each of the n INS attitude measurements, we apply the least squares solution from Equation (26). The Jacobian $\mathbf{J}$ is computed for each attitude by using Equation (18), and the change in pixel location, $\Delta\mathbf{u}$, is computed as $\mathbf{s}'_j - \mathbf{s}''_j$. The updated attitudes, $\mathbf{r}''_i$, are given by

$$\mathbf{r}''_i = \mathbf{r}'_i - \Delta\mathbf{r}_{ik}. \tag{31}$$

The updated direction cosine matrices (DCM), $\mathbf{R}''_i$, are constructed from the Euler angles $\mathbf{r}''_i$ similarly to Equation (11). Given the rotational transformation $\mathbf{C}_{c/a}$ which relates the aircraft frame of reference to the camera frame of reference, the aircraft DCM for each attitude, expressed in NED coordinates is calculated:

$$\mathbf{C}_{a/ned}[i] = \mathbf{C}^T_{c/a}[i]\mathbf{R}''_i. \tag{32}$$

Thus, we use Equation (32) to compute the updated aircraft DCM for each attitude measurement from the INS.

In theory, two endpoints from a single star are sufficient to estimate attitude. In practice, however, there does not exist enough angular resolution to accurately correct for aircraft yaw (under normal flight conditions, with the camera facing upward). For this reason, we limit estimation to images containing three or more salient stars, separated by at least one third of the image width.

3. Results

Due to limitations in the accuracy of INS estimation with the available hardware, we are unable to acquire a ground-truth attitude reference for images captured in flight at the level of precision required. Consequently, we use the methodology presented in [3] to generate high-quality simulation images from real flight data for quantitative analysis. We treat the INS attitude measurements as ground truth, and apply two forms of noise to these measurements:

1. A random-valued constant offset, and
2. Perlin noise.

The random valued constant offset is applied to every attitude measurement throughout the exposure interval, and is representative of attitude/estimation bias. The gradient-based Perlin noise is generated for each individual measurement, and is representative of attitude drift from the INS estimator. We selected Perlin noise due to its gradient-based nature, and zero-crossing properties, which are typical of iterative estimators. Both sources of noise are applied together. We measure the efficacy of the methodology presented in Section 2 based on its ability to correct for these sources of noise and recover the true attitude of the aircraft.

Real imagery and INS attitude data was captured from a single test flight. A Pixhawk version 2 autopilot was used for vehicle control and attitude estimation. Attitude data was logged from the autopilot's extended Kalman filter (EKF) at a rate of 30 Hz. The camera was mounted to the autopilot via a rigid plastic 3D-printed structure, such that all autopilot attitudes were coupled with the aerial imagery. We used a Raspberry Pi 4 companion computer for image storage, and a Raspberry Pi high-quality camera sensor fitted with the official 6-mm wide-angle lens for image capture. The sensor resolution was set to 3280×2464, with an ISO of 800. The flight was conducted at a height of 150 m above ground level with an exposure interval of 500 ms. No clouds were present, and the trial was conducted under moonless conditions. The airframe used for this study was a Zeta Science FX61 with 1.5 m wingspan, with an approximate mass of 1.5 kg, as seen in Figure 4.
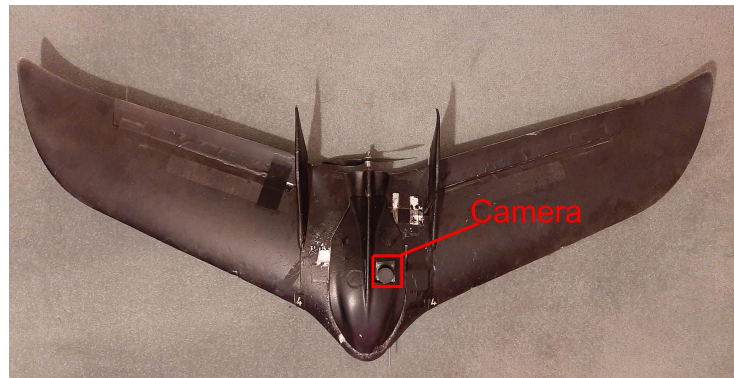
Figure 4. Zeta Science FX61 airframe used for capturing in-flight imagery.

3.1. Simulation Results

For each image captured, a series of attitude measurements were stored to disk corresponding with the exposure interval. We used a static ground image to calibrate the simulation, and subsequently generate each synthetic image from the log data. A uniformly distributed random offset between $-1°$ and $1°$ was applied to the roll and pitch channels, and a uniformly distributed random offset between $-3°$ and $3°$ was applied to the yaw channel. We applied a greater offset to the yaw channel to replicate the magnetometer bias typically seen in low-cost INS systems. A sequence of Perlin noise of length n was generated for the roll, pitch, and yaw channels, with number of octaves uniformly randomly selected between $0.1°$ and $2°$, and magnitude less than $0.5°$. This noise is representative of angular rate errors, as can be seen in Figure 5.
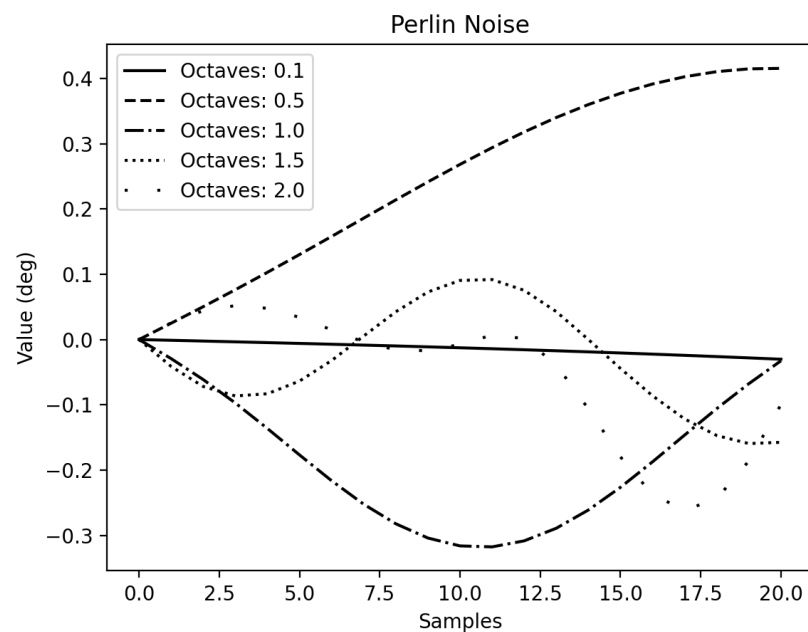

**Figure 5.** An example of Perlin gradient-based noise generation across various octaves (frequencies).

The methodology from Section 2 was applied, given the synthetic images (ground truth) and noisy attitude data (INS). Points $\mathbf{p}_j$ were extracted by using the presented image processing methods, and points $\mathbf{s}_j$ were computed by using the noisy attitude data.

A total of 110 images were simulated and processed. From these, 56 images contained a sufficient number of salient stars to perform attitude estimation (in accordance with the constraints set in Section 2). Each image had $n = 14$ noisy attitude estimates. Figure 6 shows histograms of mean absolute errors. Table 1 shows detailed results from this test. For each image captured, the mean, mean absolute, and max errors were recorded from the

attitude correction output. We can see from these results that the mean errors were close to zero, indicating that bias in the estimation process is low. The three standard deviation limit indicates that 99.7% of yaw errors are less than 0.2828°, pitch errors less than 0.1453°, and roll errors less than 0.1366°. Similarly, 99.7% of mean absolute yaw errors fell within $\pm 0.1294°$, pitch errors within $\pm 0.0591°$, and roll errors within $\pm 0.0604°$. The residual mean absolute error in the least squares approximation was 1.246 pixels, and an average of 6.036 stars were used for each estimation.
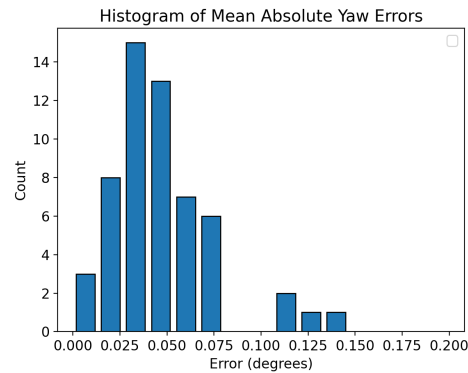

(**a**)
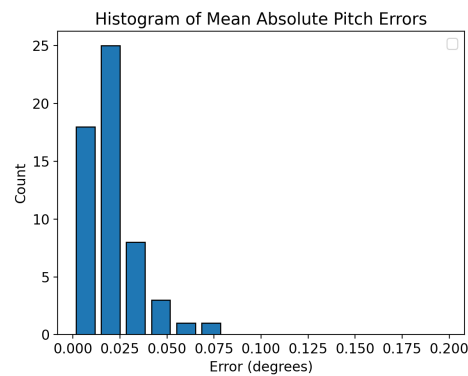
(**b**)
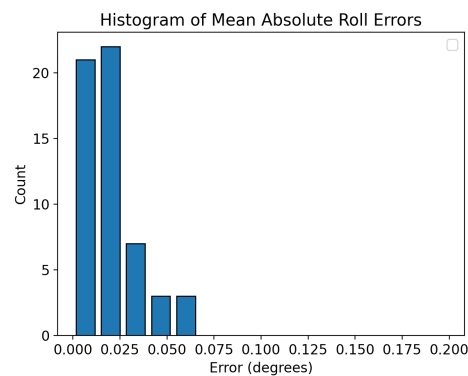
(**c**)

Figure 6. Histogram of mean absolute errors from each simulated image containing $n = 14$ attitude references. (**a**) Yaw. (**b**) Pitch. (**c**) Roll.

We can see from testing that the corrected attitude contained significantly reduced errors. Yaw errors with average magnitude 1.5° were typically corrected to within approximately 0.05°, and roll and pitch errors with average magnitude 0.5° were corrected to within approximately 0.02°. The maximum errors were also significantly lower than the average Perlin noise of 0.2°, highlighting the efficacy of this algorithm in removing both sources of noise. A graphic showing the simulated attitude correction can be seen in Figure 7.

Table 1. Simulated flight results. All units in degrees.

|	Mean	Median	Std	Mean + 3σ
Max yaw error	0.1053	0.0903	0.0856	0.2828
Max pitch error	0.0503	0.0414	0.0317	0.1453
Max roll error	0.0506	0.0418	0.0286	0.1366
Mean absolute yaw error	0.0428	0.0442	0.0274	0.1294
Mean absolute pitch error	0.0205	0.0167	0.0128	0.0591
Mean absolute roll error	0.0217	0.0207	0.0129	0.0604
Mean yaw error	0.0118	0.01076	0.0322	-
Mean pitch error	−0.0080	−0.0061	0.0138	-
Mean roll error	−0.0078	−0.0061	0.02124	-
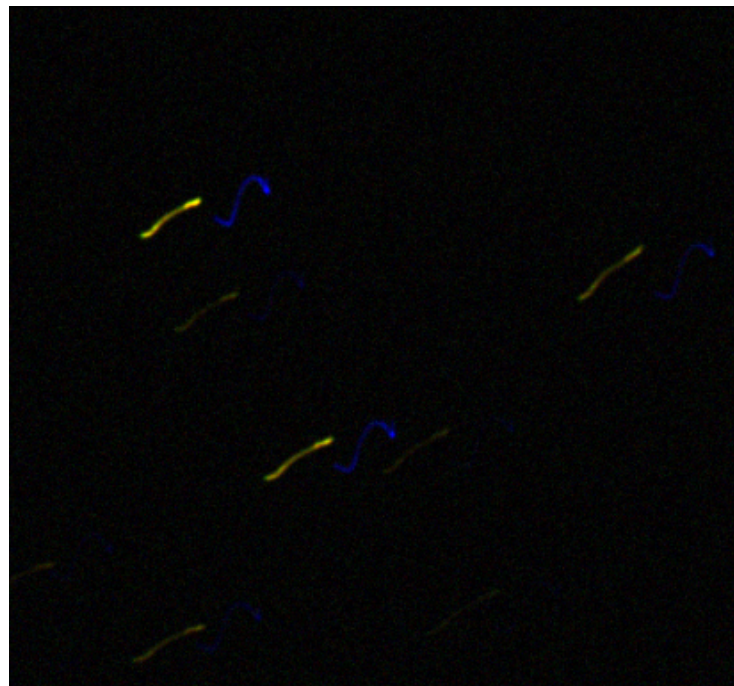
Figure 7. An example of simulation attitude correction, displaying superimposed regions of interest. Green channel, baseline simulation image; blue channel, synthetic image from noisy INS; red channel, synthetic image after corrections. Max yaw error: 0.0727°, max pitch error: 0.0286°, max roll error: 0.0226°.
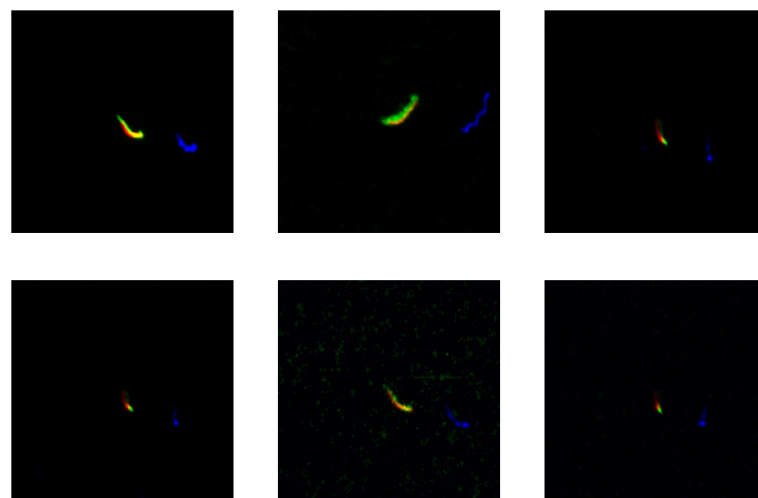
3.2. Real Imagery

To supplement the simulation results, we observe the output of our methodology when applied directly to a real image. Although no ground truth was available, we assess the resulting reprojection in a qualitative manner, so as to validate the efficacy of the image-processing techniques. The real images are subjected to small errors in the camera matrix **K**, which are most prevalent at maximal radial distance from the principal point, as well as other minor unmodeled sources of noise.

From the 110 images captured, 49 satisfied the requirements for attitude estimation. The mean residual error of the least squares estimation was 3.945 pixels, as compared with 1.264 pixels in simulation. The average number of stars used for estimation was 5.510, as compared to 6.036 in simulation. Figures 8 and 9 demonstrate the efficacy of this method in practice. We show the ROI for each star used in the attitude estimation process. The green channel contains the real image, the blue channel contains the uncorrected INS data, and the red channel is a reprojection of the corrected INS data. We can see that our method corrected for both types of noise. Some bias remains present in some stars, which is more prominent at greater distances from the optical centre.
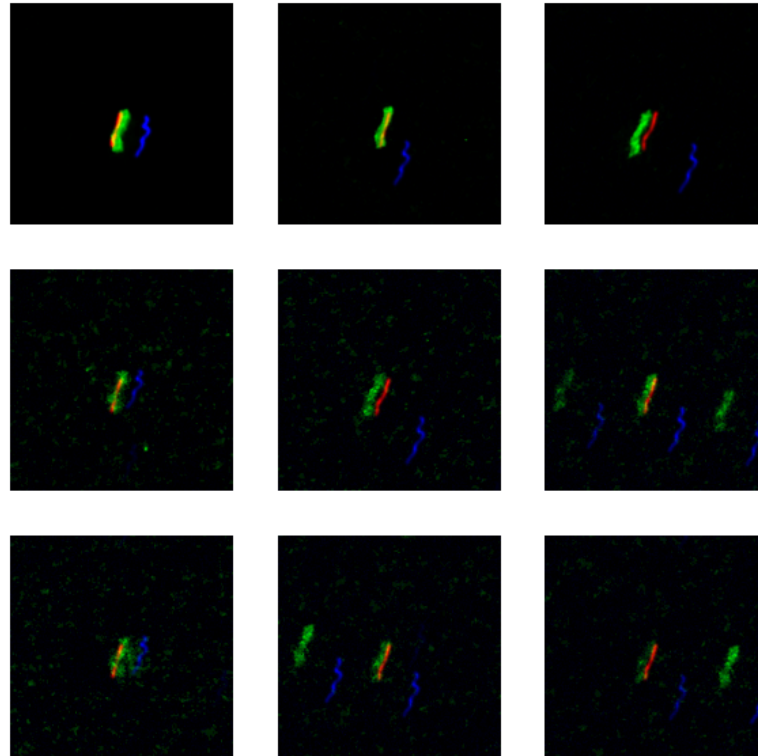
Figure 8. ROIs of stars used for attitude correction on a real image. Green channel, real image; blue channel, synthetic image from raw INS data; red channel, synthetic image from corrected INS data. The intensity of each ROI is amplified such that the peak pixel intensity is 255.

Figure 9. ROIs of stars used for attitude correction on a real image. Green channel, real image; blue channel, synthetic image from raw INS data; red channel, synthetic image from corrected INS data. The intensity of each ROI is amplified such that the peak pixel intensity is 255.

4. Discussion

The assumptions made in Section 2 may limit potential use cases for this approach. Assumption 2, constant photon flux density from a given luminary, disqualifies the use of rolling shutter cameras, and indicates that the performance will be hindered if the image is subjected to obscurities such as partial cloud cover. Assumption 3, that the star forms a simple curve on the image plane, may become limiting for imaging systems with longer exposure intervals, due to the increased likelihood of complex curves forming on the image plane. Interestingly, it may be possible to plan aircraft trajectories, which reduces the likelihood of this occurring.

It is evident in the results that some re-projection bias remains when performing attitude correction on the real images (3.945 pixels, as compared with 1.264 in simulation). There are two potential causes for this value being significantly higher than in simulation.

1. The camera calibration matrix, **K** does not perfectly characterize the camera.
2. Unmodeled sources of noise caused the image processing techniques not to transfer from simulation to reality.

Our observation is that the former is the most likely cause of this error. It is evident that certain areas on the imaging plane are defocused and do not conform with a Gaussian point-spread distribution, which may be contributing to the delocalization of pixels during the image processing stage. The image-thinning algorithm preserves the endpoints of the simulated stars (which are drawn from a sequence of Gaussian point-spread functions); however it tends to truncate the endpoints of the real stars. We postulate, however, that this effect is less significant than the effect caused by camera calibration. We see in Figures 8 and 9 that the fine-attitude correction appears to remove both the high- and low-frequency components of noise, but the bias offset remains for some stars. If the image-processing techniques were failing, we would not observe the removal of noise, particularly with the lower signal-to-noise ratio stars seen in Figure 8. We also observe that the bias tends to be greater for stars located further from the optical centre. These effects are indicative of residual nonlinear errors, such as radial and tangential distortion. Despite our efforts to remove the sources of distortion, this is a practical limitation of low-cost hardware. Interestingly, if this bias is caused by nonlinear distortion, then the residual error is not a good indicator of attitude error. This is evident when considering that a perfect attitude estimate will still yield reprojection errors.

The geometry of the camera tends toward a greater error in yaw than for pitch and roll. Under stable flying conditions (minimal roll and pitch), the yaw angle of the aircraft is measured from the angle of arc between stars about the principal point. When only a small number of observations are made, the dispersion of stars about this principal point is unlikely to be uniform, and consequently the resolution is decreased. In theory, this could be offset by applying a higher weight to stars that are further displaced from the principal point. In practice, however, the effects of lens distortion are most prominent at points farthest from the principal point, and these observations are likely to be more erroneous.

In this study we seldom encountered multiple stars within the same ROI. With more sensitive optical equipment, or less accurate attitude sensors, the occurrence of multiple stars may increase. We have not explored the possibility of outlier removal; however, we expect that the use of a random sample consensus (RANSAC) [15] algorithm may be beneficial for the selection of stars under such circumstances.

5. Conclusions

We have demonstrated a novel use for long-exposure imagery captured from a low-cost strapdown celestial sensor mounted to a lightweight, low-altitude, fixed-wing airframe. The captured imagery contains high-resolution data pertaining to the attitude of the aircraft. Standard image processing techniques were used on the long exposure images, in conjunction with a linearized weighted least squares approximation, so as to produce a corrected attitude estimate for each attitude reported by the INS. Through simulation, we demonstrated that attitude is estimated with means absolute error less than 0.13 degrees in

the yaw axis, and less than 0.06 degrees in the pitch and roll axes (3 σ). We subsequently demonstrated that this algorithm translates to real imagery, with some additional noise due to calibration. Future work will explore the use of this technique for the online calibration of magnetometer offsets.

Appendix A

Appendix A.1. Jacobian Matrix Entries

The partial derivatives of the image coordinates with respect to the aircraft Euler angles are given by:

$$
\begin{aligned}
\frac{\partial x}{\partial \phi} =& p_x \big[\cos\phi \big(\sin\psi X - \cos\psi Y \big) \\
&- \sin\phi \big(\sin\theta \cos\psi X + \sin\theta \sin\psi Y + \cos\theta Z \big) \big]
\end{aligned}
\tag{A1}
$$

$$
\begin{aligned}
\frac{\partial x}{\partial \theta} =& -f_x \big(\sin\theta \cos\psi X + \sin\theta \sin\psi Y + \cos\theta Z \big) \\
&+ p_x \cos\phi \big(\cos\theta \cos\psi X + \cos\theta \sin\psi Y - \sin\theta Z \big)
\end{aligned}
\tag{A2}
$$

$$
\begin{aligned}
\frac{\partial x}{\partial \psi} =& f_x \cos\theta \big(-\sin\psi X + \cos\psi Y \big) \\
&+ p_x \sin\phi \big(\cos\psi X + \sin\psi Y \big) \\
&+ p_x \cos\phi \big(-\sin\theta \sin\psi X + \sin\theta \cos\psi Y \big)
\end{aligned}
\tag{A3}
$$

$$
\begin{aligned}
\frac{\partial y}{\partial \phi} =& f_y \big[\big(\sin\phi \sin\psi + \cos\phi \sin\theta \cos\psi \big) X \\
&+ \big(-\sin\phi \cos\psi + \cos\phi \sin\theta \sin\psi \big) Y \\
&+ \big(\cos\phi \cos\theta \big) Z \big] + p_y \big[\big(\cos\phi \sin\psi - \sin\phi \sin\theta \cos\psi \big) X \\
&+ \big(-\cos\phi \cos\psi - \sin\phi \sin\theta \sin\psi \big) Y - \big(\sin\phi \cos\theta \big) Z \big]
\end{aligned}
\tag{A4}
$$

$$
\begin{aligned}
\frac{\partial y}{\partial \theta} =& f_y \big[\big(\sin\phi \cos\theta \cos\psi \big) X \\
&+ \big(\sin\phi \cos\theta \sin\psi \big) Y + \big(-\sin\phi \sin\theta \big) Z \big] \\
&+ p_y \big[\big(\cos\phi \cos\theta \cos\psi \big) X \\
&+ \big(\cos\phi \cos\theta \sin\psi \big) Y + \big(-\cos\phi \sin\theta \big) Z \big]
\end{aligned}
\tag{A5}
$$

$$\frac{\partial y}{\partial \psi} = f_y \big[\big(-\cos\phi\cos\psi - \sin\phi\sin\theta\sin\psi \big) X$$
$$+ \big(-\cos\phi\sin\psi + \sin\phi\sin\theta\cos\psi \big) Y \big]$$
$$+ p_y \big[\big(\sin\phi\cos\psi - \cos\phi\sin\theta\sin\psi \big) X$$
$$+ \big(\sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi \big) Y \big] \tag{A6}$$

$$\frac{\partial z}{\partial \phi} = \big[\cos\phi\sin\psi - \sin\phi\sin\theta\cos\psi \big] X$$
$$+ \big[-\cos\phi\cos\psi - \sin\phi\sin\theta\sin\psi \big] Y + \big[-\sin\phi\cos\theta \big] Z \tag{A7}$$

$$\frac{\partial z}{\partial \theta} = \big[\cos\phi\cos\theta\cos\psi \big] X + \big[\cos\phi\cos\theta\sin\psi \big] Y$$
$$+ \big[-\cos\phi\sin\theta \big] Z \tag{A8}$$

$$\frac{\partial z}{\partial \psi} = \big[\sin\phi\cos\psi - \cos\phi\sin\theta\sin\psi \big] X$$
$$+ \big[\sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi \big] Y \tag{A9}$$

and the 2-dimensional homogeneous image coordinates are given by:

$$u = \frac{x}{z}, \quad v = \frac{y}{z} \tag{A10}$$

then by applying the quotient rule, we compute the partial derivatives:

$$\frac{\partial u}{\partial \phi} = \frac{z\frac{\partial x}{\partial \phi} - x\frac{\partial z}{\partial \phi}}{(z)^2} \tag{A11}$$

$$\frac{\partial u}{\partial \theta} = \frac{z\frac{\partial x}{\partial \theta} - x\frac{\partial z}{\partial \theta}}{(z)^2} \tag{A12}$$

$$\frac{\partial u}{\partial \psi} = \frac{z\frac{\partial x}{\partial \psi} - x\frac{\partial z}{\partial \psi}}{(z)^2} \tag{A13}$$

$$\frac{\partial v}{\partial \phi} = \frac{z\frac{\partial y}{\partial \phi} - y\frac{\partial z}{\partial \phi}}{(z)^2} \tag{A14}$$

$$\frac{\partial v}{\partial \theta} = \frac{z\frac{\partial y}{\partial \theta} - y\frac{\partial z}{\partial \theta}}{(z)^2} \tag{A15}$$

$$\frac{\partial v}{\partial \psi} = \frac{z\frac{\partial y}{\partial \psi} - y\frac{\partial z}{\partial \psi}}{(z)^2} \tag{A16}$$

References

1. Kayton, M.; Fried, W.R. *Avionics Navigation Systems*; John Wiley & Sons: Hoboken, NJ, USA, 1997.
2. Titterton, D.; Weston, J.L.; Weston, J. *Strapdown Inertial Navigation Technology*; IET: Stevenage, UK, 2004; Volume 17.

3.   Teague, S.; Chahl, J. Imagery Synthesis for Drone Celestial Navigation Simulation. *Drones* **2022**, *6*, 207. [CrossRef]
4.   Chen, X.; Liu, D.; Zhang, Y.; Liu, X.; Xu, Y.; Shi, C. Robust motion blur kernel parameter estimation for star image deblurring. *Optik* **2021**, *230*, 166288. [CrossRef]
5.   Wei, Q.; Weina, Z. Restoration of motion-blurred star image based on Wiener filter. In Proceedings of the 2011 Fourth International Conference on Intelligent Computation Technology and Automation, Shenzhen, China, 28–29 March 2011; Volume 2, pp. 691–694.
6.   Wang, S.; Zhang, S.; Ning, M.; Zhou, B. Motion blurred star image restoration based on MEMS gyroscope aid and blur kernel correction. *Sensors* **2018**, *18*, 2662. [CrossRef] [PubMed]
7.   Zhang, W.; Quan, W.; Guo, L. Blurred star image processing for star sensors under dynamic conditions. *Sensors* **2012**, *12*, 6712–6726. [CrossRef] [PubMed]
8.   Sun, T.; Xing, F.; You, Z.; Wei, M. Motion-blurred star acquisition method of the star tracker under high dynamic conditions. *Opt. Express* **2013**, *21*, 20096–20110. [CrossRef] [PubMed]
9.   Ma, L.; Zhan, D.; Jiang, G.; Fu, S.; Jia, H.; Wang, X.; Huang, Z.; Zheng, J.; Hu, F.; Wu, W.; et al. Attitude-correlated frames approach for a star sensor to improve attitude accuracy under highly dynamic conditions. *Appl. Opt.* **2015**, *54*, 7559–7566. [CrossRef] [PubMed]
10. He, Y.; Wang, H.; Feng, L.; You, S. Motion-blurred star image restoration based on multi-frame superposition under high dynamic and long exposure conditions. *J. Real-Time Image Process.* **2021**, *18*, 1477–1491. [CrossRef]
11. Klaus, A.; Bauer, J.; Karner, K.; Elbischger, P.; Perko, R.; Bischof, H. Camera calibration from a single night sky image. In Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004, Washington, DC, USA, 27 June–2 July 2004; Volume 1, p. I.
12. Rijlaarsdam, D.; Yous, H.; Byrne, J.; Oddenino, D.; Furano, G.; Moloney, D. A survey of lost-in-space star identification algorithms since 2009. *Sensors* **2020**, *20*, 2579. [CrossRef] [PubMed]
13. Liebe, C.C. Accuracy performance of star trackers-a tutorial. *IEEE Trans. Aerosp. Electron. Syst.* **2002**, *38*, 587–599. [CrossRef]
14. Zhang, T.Y.; Suen, C.Y. A fast parallel algorithm for thinning digital patterns. *Commun. ACM* **1984**, *27*, 236–239. [CrossRef]
15. Fischler, M.A.; Bolles, R.C. Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* **1981**, *24*, 381–395. [CrossRef]