Optical Engineering

[SPIDigitalLibrary.org/oe](https://www.spiedigitallibrary.org/oe)

Color star tracking II: matching

John Enright
Geoffrey R. McVittie

Color star tracking II: matching

John Enright
Geoffrey R. McVittie
 Ryerson University
 350 Victoria Street
 Toronto, Canada
 E-mail: jenright@ryerson.ca

Abstract. A novel matching algorithm is presented that can identify stars using raw images of the sky obtained from a CMOS color filter array detector. The algorithm combines geometric information with amplitude ratios calculated from the red, green, and blue color channels. Conventional algorithms that match stars based solely on inter-star geometry (and sometimes relative brightness), typically require three or more stars for a confident star match. In contrast, the presented algorithms are able to find matches with only two imaged stars in most regions of the sky. The necessary catalog preparation and a simple star-pair matching algorithm based on combined color intensity ratios and the angular spacing are discussed. Results from a large set of simulation trials and initial results from sensor field testing are presented. © The Authors. Published by SPIE under a Creative Commons Attribution 3.0 Unported License. Distribution or reproduction of this work in whole or in part requires full attribution of the original publication, including its DOI. [DOI: [10.1117/1.OE.52.1.014406](https://doi.org/10.1117/1.OE.52.1.014406)]

Subject terms: star tracker; star matching; attitude determination.

Paper 120540 received Apr. 12, 2012; revised manuscript received Nov. 30, 2012; accepted for publication Dec. 17, 2012; published online Jan. 9, 2013.

1 Introduction

A measure of robustness for autonomous star trackers is the ability to obtain an attitude estimate without *a priori* attitude information, while viewing an arbitrary part of the sky. Performance is usually limited by regions of the sky with low star densities. When solving the “lost-in-space” (LIS) problem, imaged stars must be associated with records in the star tracker’s on-board catalog. Although some star-matching routines can operate with only two stars in the field of view (FOV) once an initial solution has been found, almost all LIS matching algorithms require three visible stars for an initial fit. Unfortunately, the geometry of many pairs and triplets of stars are ambiguous and unique attitude solutions are not always assured with the minimum number of imaged stars. The most common remedy for this ambiguity is an optical design that guarantees detection of more than the minimum number of stars. Designers can meet this requirement either by enlarging the FOV, or observing dimmer stars. Our research considers a different approach to resolving ambiguous star patterns: that of using coarse measurements of stellar spectra as well as inter-star angles.

A wide variety of star identification techniques have been proposed in literature including triangle, pole-star, grid, and adaptive techniques. A survey by Spratling and Mortari¹ provides a good overview of the current state of the field. Matching routines rely on geometric arrangements between stars (e.g., arc lengths, dihedral angles, etc.). Some approaches (e.g., Servidia et al.),² use visual magnitude information to help resolve ambiguity between similar geometric patterns. Others, such as Rufino and Accardo³ argue explicitly against the use of absolute magnitude because accurate photometry is difficult with the optics and detectors found in typical star trackers.

Although spectroscopy is common practice in astronomy, prior to our current work, there seems to be no discussion of using star spectral classification to assist in star tracking. Adopting color detectors will impact many of the steps in the star tracker image processing sequence. This paper

examines how color information can be used to assist the star matching process; in our companion paper,⁴ we develop techniques to reproducibly measure position and color of the stars in a raw detector image captured by a color CMOS detector.

We begin our analysis by summarizing our recent development of a nano-satellite star tracker. This section details a reference hardware and software implementation that establishes the context for the subsequent analysis. In Sec. 3 we introduce the basic mechanisms for matching observed stars against cataloged data. We present both a basic matching strategy and an extension that allows us to match an entire scene’s worth of stars. We then discuss the challenges building catalogs suitable for color matching. Section 4 examines a number of practical and theoretical issues including: sourcing spectral information; encoding and storing the catalog based on an instrument’s optical properties; and evaluating the theoretical performance of a star tracker design. Finally, we validate the performance of our matching algorithms using simulations and night-sky tests.

2 Reference Design

This paper represents a companion study to the recent development of a very small star tracker, the S3S ST-16. A collaborative development effort between Sinclair Interplanetary (SI), the University of Toronto Space Flight Laboratory (UofT-SFL), and the Ryerson University Space Avionics and Instrumentation Laboratory (SAIL), the ST-16 (Fig. 1) is designed around a commercial CMOS detector and a simple lens assembly. This design is described at length in Enright et al.⁵ The hardware and processing architecture of the ST-16 have served as a starting point for this current study. Our performance targets and assumptions concerning accuracy, FOV size, resolution, sensitivity, update rate are derived from the ST-16 design. We recognize that this is not an optimal design for a color star tracker, but it is a very easily realized one, for the detector in the ST-16 has both monochrome and color variants.

Fig. 1 The ST-16 star tracker prototype, with penny for scale.

Fig. 2 Comparison of monochrome and color quantum efficiencies (from Refs. 6 and 7).

2.1 Optics and Performance

The ST-16 detector is the monochrome version of the Aptina MT9P031. The MT9P031 CMOS detector has a resolution of 2592×1944 pixels and a $2.2 \mu\text{m}$ pixel pitch. The monochrome version is equivalent to the color version of the detector, but lacks the color filter array (CFA) that produces distinct output color channels. The sensor's F1.2, 12 mm (dia.) lens assembly gives the ST-16 a threshold detection magnitude of $M_{\text{thresh}} = 5.75$, a 7.5 half-angle FOV, θ_{FOV} , and nominal exposure of 100 ms. The design provides at least three visible stars in view in more than 99.9% of the possible sensor orientations.

The CFA version of the detector allows distinction between sources of different colors. The CFA consists of small, absorptive filters in front of each pixel in a repeating 2×2 pattern. These filters absorb some of the incident light, so the overall sensitivity to dim sources is lower for the color detector than for the monochrome version. An illustration of the typical differences in quantum efficiency between color and monochrome detectors is shown in Fig. 2. We note that the monochrome curve is adapted from the specifications for the MT9V022 detector⁶ (the monochrome response of the MT9P031 is very similar to the MT9V022, but does not appear on the public datasheet). Although there are some differences, based on the spectrum of the source stars, the color detector generates about one third of the integrated response of the monochrome (i.e., ~ 1.25 visual magnitudes).

2.2 Baseline Processing Model

Several processing steps are necessary to convert raw images into attitude estimates (see Fig. 3). In a color star tracker, a raw sky image from a CFA detector contains the intensity readings from the red, green, and blue pixels. The initial image processing routines in the star tracker must analyze this intensity map and identify likely stars. These coarse star locations must then be refined to give precise estimates of centroid position and color. A camera model rectifies the centroid locations and predicts corresponding direction vectors for each star. Matching routines identify the stars by comparing the available data to on-board catalog of known star patterns. A final attitude estimate follows directly from the matched star information. Except for the presence of color information, this processing sequence is typical of most modern star trackers.

Color representation is an important consideration; in our companion paper,⁴ we found that integrated intensity ratios were an effective representation of star color. Our detection routines use the raw CFA output to estimate ratios between

Fig. 3 Color star tracker image processing.

the blue-green, Λ_{bg} , and red-green, Λ_{rg} , channels, together with the centroid position, (m_c, n_c) . These algorithms use a nonlinear least-squares cost function and Levenberg-Marquardt minimization. The typical root mean squared (RMS) color noise reported in the study was less than 0.1, and often about 0.05. The centroid position error was equivalent to about 0.2 pixels. These metrics will influence the effectiveness of our matching algorithms.

Our proposed color-based matching routines (and the monochrome matching routines on the ST-16) rely on knowing the direction vectors, s in the detector frame of reference. Each ST-16 star tracker must be calibrated to account for manufacturing tolerances and optical aberrations such as distortion. We use a 17 parameter model to relate centroid position on the detector array, i.e., (m_0, n_0) , to the star vector direction, s_D . Our calibration procedure and model is based on the scheme proposed by Heikkila and Silven.⁸ We employ a gimbaled platform, a fibre-illuminated pinhole, and an inverted telescope to produce a point source at infinity. After calibration, the typical residual centroid error is 0.2 pixels (RMS). This includes the effects of temporal noise and calibration error. Presently, we assume that the calculation of s_D does not rely on the star color; if necessary, additions to the model could correct for chromatic aberration.

3 Color-Based Matching

The challenge of solving the lost-in-space problem for star trackers rests with the matching of the observed image stars to those in the sensor's on-board catalog. Whereas other processing steps are deterministic in time and memory, the matching process involves search operations. Thus a viable algorithm must be both accurate (i.e., it identifies the correct stars), and computationally efficient (the search is bounded in time).

The star matching routines must associate the observable star characteristics (i.e., m_c , n_c , Λ_{rg} , and Λ_{bg}), with known stars in the on-board catalog. This is essentially a mapping problem: given a set of observed stars A and a catalog of stars C , we need to identify the relations that map $A_k \rightarrow C_j$ (in our notation, we use a subscript index when we need to refer to a particular star from one of these sets). Generally, the more stars that can be matched, the better the resulting attitude fix. However, not all detections made by the sensor are actually

stars—planets, other satellites, or even noise may cause false detections—so we may not be able to find a match for all the stars in A .

We first develop the basic kernel matching of a single pair of image stars against the on-board catalog. Not only does this represent the fundamental matching operation, but it also provides analysis tools that can be used to predict the performance of the sensor for a given optical design. We then extend the pairwise matching algorithm to efficiently match more than two stars. We conclude this section with a discussion of how these algorithms can be implemented efficiently.

3.1 Matching Star Pairs

The basic matching algorithm relies on measurements obtained from a pair of image stars. Each measurement set consists of five quantities: four color ratios (two for each star) and the angular separation between the pair. In monochrome star trackers, star-pairs alone are too ambiguous enough for reliable matching, but the color data adds sufficient information for a unique match in many cases.

For a particular image, we select two stars $A_k, A_l \in A$ and compute the tuple $T_{A_{kl}}$:

$$T_{A_{kl}} \equiv \{ \Lambda_{rg_k} \quad \Lambda_{bg_k} \quad \varphi_{kl} \quad \Lambda_{rg_l} \quad \Lambda_{bg_l} \}, \quad (1)$$

where φ_{kl} is computed from the two unit direction vectors:

$$\varphi_{kl} = \arccos(\hat{s}_k^T \hat{s}_l). \quad (2)$$

For notational simplicity, we have dropped the frame of reference subscript from these vectors. Our catalog contains similar tuples derived from known star locations and spectral properties. We can compare our observations to the catalog by computing the Mahalanobis distance between each pair of tuples.* In general, this has the form:

$$\Delta = \sqrt{(T_A - T_C)^T \Sigma (T_A - T_C)}, \quad (3)$$

where Σ is the measurement covariance matrix of the elements of T_A . If Σ is diagonal (i.e., errors are uncorrelated), then we can write the distance between $T_{A_{kl}}$ and $T_{C_{ij}}$ as:

$$\Delta(T_{A_{kl}}, T_{C_{ij}}) \equiv \sqrt{\left(\frac{\Lambda_{rg_k} - \Lambda_{rg_l}}{\sigma_{rg}}\right)^2 + \left(\frac{\Lambda_{rg_l} - \Lambda_{rg_j}}{\sigma_{rg}}\right)^2 + \left(\frac{\varphi_{kl} - \varphi_{ij}}{\sigma_{\varphi}}\right)^2 + \left(\frac{\Lambda_{bg_k} - \Lambda_{bg_l}}{\sigma_{bg}}\right)^2 + \left(\frac{\Lambda_{bg_l} - \Lambda_{bg_j}}{\sigma_{bg}}\right)^2}. \quad (4)$$

The k and l designations in the observed pairs are arbitrary, so we must compute both $\Delta(T_{A_{kl}}, T_{C_{ij}})$ and $\Delta(T_{A_{lk}}, T_{C_{ij}})$ and take the smaller of the two distances. The maximum likelihood match between $T_{A_{kl}}$ and the cataloged pairs is the $T_{C_{ij}}$ that minimizes Δ . Sorting C in order of increasing φ can make this search efficient.

3.2 Star Pair Ambiguity

We expect that the catalog star tuple $T_{C_{ij}}$ that minimizes Δ will be the correct match to a pair of observed stars. This assumption holds provided that centroid and color noise

does not move the image tuple $T_{A_{ij}}$ away from the correct match and closer to another catalog pair. Examining the distribution of Δ values calculated between catalog tuples allows us to assess the likelihood of measurement noise causing incorrect matches. We outline the basic framework for this analysis here and further develop the technique in Sec. 4. This analysis allows us to predict theoretical sensor performance based on catalog properties and a few optical design parameters.

*McLachlan⁹ provides a good introduction for those unfamiliar with this type of classifier.

The smaller the Mahalanobis distance between two catalog pairs, the more likely they are to be mistaken for one another. Below a certain distance threshold we consider the two tuples *ambiguous*. Using Eq. (4) to compute Δ between two catalog pairs gives:

$$\Delta_C(p, q) \equiv \Delta(T_{C_p}, T_{C_q}). \quad (5)$$

Here we have replaced the star indices i, j from each tuple with a single catalog index, p or q . Assuming normal error distributions, we can use linear discriminant analysis theory to compute the ideal error probability. Adapting the result from McLachlan et al.,⁹ this becomes:

$$P_{\text{error}} = \frac{1}{2} \left[1 + \operatorname{erf} \left(\frac{-\Delta}{2\sqrt{2}} \right) \right]. \quad (6)$$

Equivalently, we can solve for the critical separation Δ^* , for a specified success rate $P_{\text{correct}} = 1 - P_{\text{error}}$:

$$\Delta^* = -2\sqrt{2}\operatorname{erf}^{-1}(1 - 2P_{\text{correct}}). \quad (7)$$

Star pairs with a Mahalanobis distance less than Δ^* are ambiguous. The results in the remainder of this paper use $P_{\text{correct}} = 0.99$ to establish a distance threshold. The corresponding critical distance is $\Delta^* = 4.65$.

Understanding how ambiguous pairs affect system performance depends on our multistar matching strategy and the contents of the whole star-pair catalog. These concepts are developed in Secs. 3.3 and 4, respectively.

3.3 Matching Whole Scenes

Our basic matching algorithm gives the best catalog match to a particular pair of image stars. Generally, we can improve the accuracy of our attitude estimate if we can find consistent matches to any additional stars in the sensor image. This additional accuracy comes at the cost of increased computational complexity. Most matching algorithms using pairs or triplets scale with $O(N_A^2)$ or $O(N_A^3)$, where N_A is the number of stars (some classes of algorithms such as Grid Algorithms can moderate this).^{10,11} For star-rich areas of the sky, we need an efficient approach to matching.

Our strategy for efficiently matching multiple star scenes is based on the random sample consensus (RANSAC) approach¹²—a popular approach in computer vision. We use a greedy heuristic to quickly assemble a self-consistent set of matches, concentrating on finding ‘inliers’ instead of rejecting outliers. This algorithm works as follows (see Fig. 4 for a pictorial representation of this process):

- From the initial set of imaged stars A [Fig. 4(a)], we select two stars, A_k and A_l [Fig. 4(b)], form the pair-tuple $T_{A_{kl}}$ and find the best pairwise match between $T_{A_{kl}}$ and the star catalog. The stars A_k and A_l form the basis for a candidate consensus set (or *conset*), S [Fig. 4(c)].
- For each star A_i , $i \neq k, l$ we select one star A_j from S (i.e., $A_j \in S$), and form the pair-tuple $T_{A_{ij}}$ [Fig. 4(d)]. We then find the best pairwise match between $T_{A_{ij}}$ and the star catalog, with the additional constraint that the mapping $j \rightarrow p$ is already known. If we can find such a

Fig. 4 Graphical illustration of consensus matching.

match where $\Delta < \Delta^*$, then we tentatively accept the mapping $i \rightarrow r$ between A_i and the catalog star C_r . If we cannot find an acceptable match for A_i we move on to the next star in A .

- If we have a tentative match for A_i , we need to check to ensure that the match is consistent with the other stars in the consensus set S [Fig. 4(e)]. Although we could evaluate the consistency using the full color information, it is faster to perform this check based on geometry alone. Thus, for each $A_s \in S$, where we have the tentative mapping $i \rightarrow r$ and confident mapping $s \rightarrow k$, we must have:

$$|\delta\varphi| \equiv |\varphi_{A_{is}} - \varphi_{C_{rk}}| < \delta\varphi^*, \quad (8)$$

where the maximum acceptable distance error can be derived from the critical Mahalanobis distance, assuming the color errors are zero:

$$\delta\varphi^* = \sigma_\varphi \Delta^* / 2. \quad (9)$$

The factor of two is needed because Δ^* was defined in terms of the separation between two catalog entries.

- If the tentative match, $i \rightarrow r$, satisfies the consistency check outlined in Step-3, then we add A_i to S [Fig. 4(f)]. A false star in the scene, may return a valid pairwise match [Fig. 4(g)], but is very unlikely to pass the geometric consistency checks [Fig. 4(h)]. If the tentative match does not meet the consistency criteria, A_i is not added to S .
- Once all of the stars in A have been evaluated [Fig. 4(i)], we consider the number of stars we have added to S , i.e., N_S . If the number of matched stars exceeds a specified fraction of the total number of imaged stars, i.e., $N_S \geq \eta N_A$, this indicates that we were able to assemble a good consensus set from the initial matched pair. In this case, the algorithm can terminate early. Conversely, a small consensus set suggests a poor initial match. In our example, this might happen if the false star was included in the initial pair of stars. In this case, we select a new initial pair and rebuild the conset. The overall matching algorithm terminates when we have found a good consensus set, or after a specified number of matching attempts has completed.

Like many RANSAC-derived heuristics, good matching performance requires some algorithm tuning. We have found that a matching threshold fraction in the range $0.6 \leq \eta \leq 0.7$ combines good tolerance to false stars (i.e., low η), with good rejection of bad matches. We must also set an upper bound on the number of RANSAC iterations we allow; for our trials, we have used $3 \cdot N_A$.

The matching algorithm keeps track of the largest conset found throughout all of its iterations. Upon termination, we can classify the results as one of three cases: a high-confidence match (conset size exceeds the η , threshold), a low-confidence match (conset is smaller than the η threshold), or no match at all (no acceptable starting pairs found). In the first and last cases our next actions are well defined. In the first, we proceed with the processing sequence and calculate the sensor attitude; in the last case, we report a matching failure and abort. The remaining case—the low-confidence match—may require extra logic to categorize.

In a flight sensor, we may want to formalize more elaborate decision rules, but for this study, we merely log the result and complete the attitude calculation.

4 Star Catalog Design

Our matching algorithms rely on comparing the star tracker observations to catalog values. Thus, we must pay particular attention to the source material, content, and organization of C . Much of the recent innovation in star tracker processing has involved different approaches to catalog generation. Calculations that can be performed off-line (areas, distances, etc.) reduce the computational (and often power), requirements on the star tracker processor. However, precomputing geometric metrics often increases the need for storage: a catalog of a few thousand stars routinely produces tens of thousands of star-pairs, and possibly millions of star-triples.

Color star catalogs share many of the same characteristics as monochrome catalogs, but present a number of unique challenges. In this section, we detail the content and organization of our baseline star catalog. We consider both conventional requirements for star geometry as well as the extra preparation required to tabulate color information.

4.1 Preliminary Catalog Processing

The source catalog for this research is the 5th Revised Edition of the Yale Bright Star (YBS) catalog.¹³ We denote this catalog by Y . We subscript this symbol, by the catalog index, i.e., Y_i to refer to a particular star. The on-board catalog is denoted by C . Working from the source catalog Y , we select a subset of stars for use in C . We select stars for C based on the following criteria:

- **Brightness.** Our primary criterion for inclusion in the on-board catalog is visual magnitude M_v (i.e., brightness). Because the larger values of M_v indicate dimmer stars we select those stars with $M_v \leq M_{\text{thresh}}$. M_{thresh} represents the brightness of the dimmest stars that a particular star tracker can reliably detect under nominal operating conditions.
- **Double stars.** Binary stars or stars whose apparent separations are too small to be distinctly resolved by the sensor optics are treated as a single star. The brighter star in each case is included in C .
- **Transient Objects.** A number of the entries in the YBS are attributed to transient novae. They have been retained in Y to preserve the numbering scheme, but are unnecessary when compiling C .

Each star in the source catalog, Y , has a variety of general astronomical data associated with it. Not all of this information is necessary in the on-line catalog. For convenience we store some basic information about each star (i.e., its Henry-Draper identifier and visual magnitude), but the star's position is most important.

Most astronomical catalogs tabulate star position in terms of right ascension α , and declination, δ . High precision applications may also include proper motion information (we neglect the effect of star motion for this study). To avoid the need for unnecessary trigonometric calculations, it is convenient to store a star's position as a unit vector, expressed in the Earth Centred Inertial (ECI) coordinate system, $\mathbf{s}_i$.

4.2 Cataloging Color Data

Assembling an on-board catalog requires compilation of color information for each star in C . We must consider what color information will be stored; where it will come from; and whether any off-line processing is necessary to tailor our data sources to the specific performance of our sensor.

In astronomy, color is an ambiguous term related to an object's emission spectrum. In addition to standard identifiers and celestial coordinates, conventional catalogs, like the YBS catalog, often include several quantities related to the spectrum of the star. These include: spectral type, color index, and effective temperature. Other collated data-sets contain detailed spectra for individual stars.

Our discussion of color-based star detection relies on a familiarity with spectral type classification schemes. For simplicity we use the popular Morgan-Keenan (and Kelmann) (MK) system.^{14–16} Formally, cataloged spectral types are composite classifications based mainly on star temperature, but also on specific features (e.g., emission or absorption lines, etc) of the complete spectrum.

Although spectral type is a multiattribute categorization, the notion of color index is much closer to a directly measurable quantity. Classical color indices were introduced by Johnson and Morgan and indicate magnitude difference in the integrated response across standardized spectral bands.¹⁴ Unfortunately, the color filters on our detector do not have the same response as the official Johnson and Morgan pass-bands, making direct conversion between the two difficult. Instead, we calculate our own version of these color indices for each star in the star tracker catalog. Given a star's spectrum, $I(\lambda)$, we can calculate the red-green and blue-green color indices:

$$\Lambda_{rg} \equiv \frac{\int_{\lambda_1}^{\lambda_2} \chi_r I(\lambda) d\lambda}{\int_{\lambda_1}^{\lambda_2} \chi_g I(\lambda) d\lambda}, \quad (10)$$

$$\Lambda_{bg} \equiv \frac{\int_{\lambda_1}^{\lambda_2} \chi_b I(\lambda) d\lambda}{\int_{\lambda_1}^{\lambda_2} \chi_g I(\lambda) d\lambda}. \quad (11)$$

The $\chi_{r,g,b}$ capture the detector quantum efficiency for each of the color channels.

Choosing a source of stellar spectra for our star tracker catalog presented its own challenges. Complete spectra is unnecessary in the on-line catalog; we need only store the expected color ratios. Three approaches seemed feasible: (a) continuous black-body curves based on the stars' effective temperatures; (b) reference solar spectra based on spectral type; and (c) tabulated spectra for each individual star. The reference spectra used in this study are taken from the Pickles spectral library;¹⁷ individual spectra are taken from the recent Indo-US Library of Cloud Feed Stellar Spectra,¹⁸ and to a lesser extent from an earlier catalog by Burnashev.¹⁹ Temperature estimates are based on data in the Indo-US catalog. A comparison of these approaches is shown in plots for γ -Ori (Fig. 5) and α -Ori (Fig. 6). We used the Pickles B2II spectrum for γ -Ori (there is no B2III spectrum provided), and a black-body temperature of 22,570 K. For α -Ori we use the M2I spectrum and a temperature of 3540 K.

We decided to use the Pickles catalog as our color database, eliminating the alternatives based on the following

Fig. 5 Comparison of Spectra for γ -Ori (Bellatrix). Flux normalized to 550 nm.

Fig. 6 Comparison of Spectra for α -Ori (Betelgeuse). Flux normalized to 550 nm.

rationale. First, although the black-body curve is a good match to the hot spectrum of γ -Ori, it is a very poor match to the much cooler α -Ori. Because effective temperature is based on total integrated flux, not on a match to the spectra, black-body distributions introduced inaccuracies in the calculated color ratios [i.e., from Eqs. (10) and (11)]. Second, even though the Indo-US catalog is considered to have broad spectral coverage (346–946 nm) it is not quite sufficient to cover the full sensitive range of our CMOS detector. Improved color ratios *could* be calculated by merging partial spectral catalogs for all members of C —Pickles describes such strategies in context of building his reference spectra—but this process would be laborious. Furthermore, except for small local features, the Pickles and Indo-US curves are in reasonable agreement. The integration in the evaluation of the Λ ratios will minimize the effect of local differences. In the balance, the Pickles' data was a pragmatic compromise.

Having selected the Pickles catalog, P , as our source of color information, we must still associate each element of C with one of the reference spectra. Assigning reference

Fig. 7 Color intensity ratios as a function of spectral type. The curl in the lower-right is caused by secondary pass-bands in the array filters at near infrared wavelengths.

spectra to the stars selected from the YBS is not always straightforward. The spectral types given in the latter catalog do not always have exact matches in the former. Rules for selecting the closest match depend of the stellar classification: for hot stars, the class and subclass dominate; for cooler stars, luminosity can be quite important as well. We have developed a set of heuristic rules for this process, but further cross-checking with the available spectra in the Indo-US catalog is necessary. Figure 7 shows the predicted color ratio response across different spectral types for our baseline detector. The hook in the bottom right of the figure is caused by the second passband in the blue and green curves around 800 nm (see Fig. 2).

4.3 Star Pair Catalog

The Pickles catalog, P , provides a mechanism for computing the expected color ratios for each star in C . To complete our catalog, we compile a secondary table of feasible star tuples $T_{C_{ij}}$. Our matching algorithm is based on these pairs of stars and this table provides the basis for finding matches. The angular separation between two stars, $k, l \in C$, φ_{kl} can be computed from their direction vectors:

$$\varphi_{kl} = \arccos(\mathbf{s}_k^T \mathbf{s}_l). \quad (12)$$

To compile the table of star-pairs, we evaluate all the pair combinations and select those that satisfy

$$\varphi_{\min} < \varphi_{kl} \leq 2\theta_{\text{FOV}}. \quad (13)$$

The upper bound ensures that all catalog star-pairs fit within the minor axis of the sensor FOV; the lower bound is chosen so that the star peaks can be resolved separately. Typically $\varphi_{\min} = \theta_{\text{FOV}}/500$, an empirically chosen distance, equivalent to about four pixels.

It is unnecessary to store the entire five-element tuple online. The color information can be stored in a separate table indexed by the star identifier or reference spectrum identifier.

When we need to find the best match to a star-pair tuple T , we do not have to calculate the Mahalanobis distances for all

Fig. 8 Pair look-up tables.

tuples in the catalog. Rather we recognize that each components of Δ must lie in a bounded range. Thus, if the measured image arc-length is φ_{kl} , we need only check those catalog entries whose arc-length, φ_p , satisfies:

$$\varphi_{kl} - \delta\varphi^* \leq \varphi_p \leq \varphi_{kl} + \delta\varphi^*. \quad (14)$$

If we sort the catalog entries in order of increasing φ , then all of the feasible matches will be sequential in the pairs catalog T . We use a hash-based index scheme to find the start of the appropriate section of the catalog (see Fig. 8). We allocate storage for a pairs index $P_h, 0 \leq h < N_h$ (typically we use $N_h = 1024$, but this choice was somewhat arbitrary). For an observed φ_{kl} we calculate the separation hash h :

$$h = \lfloor H \cdot \varphi_{kl} \rfloor, \quad (15)$$

where the constant scaling factor H is chosen such that $H = N_h/\varphi_{\max}$. The value of P_h is the smallest catalog index p , where

$$\varphi_p \geq \frac{h}{H} - \delta\varphi^*. \quad (16)$$

The matching routine then starts at table index P_h and moves forward sequentially until the catalog value of φ exceeds the upper bound of Eq. (14). The hash h corresponds to a range of different φ values, so the initial φ_{P_h} may actually lie outside the lower bound of Eq. (14). In practice, this is a negligible concern because there are typically only a handful of extra pairs, and these can all be skipped quickly based on their arc-length.

5 Experimental Validation

We conducted a number of tests to assess the overall feasibility of the color star matching concept. First, studying the distribution of ambiguous tuples (Sec. 3.2) in a color star catalog measures the effectiveness of the Δ -based classifier and allows us to compare the performance of sensor point-designs. Second, we examine the performance of the full matching algorithm using a combination of scene simulation and night-sky testing.

5.1 Sensor Design Trade Studies

One of the key performance metrics for a star tracker is its availability, A . This quantity represents the fraction of time when an image will provide a correct attitude solution. Stray light or occlusion from the sun, Earth or moon, can reduce the effective availability for a given mission, but more

fundamental is the instrument's raw ability to image enough stars to provide an attitude fix. In this section, we examine the sensor performance as predicted from the properties of the star-pair catalogs (from Sec. 4). We use these results to evaluate the suitability of the baseline ST-16 optical design and identify evolutionary design changes that may be better suited for color-matching.

Sensor availability is strongly tied to the dimensions of the FOV, θ_{FOV} , the threshold detection magnitude, M_{thresh} , and the star matching scheme. For our color-based matching, availability will be limited by instrument orientations that meet one of the following criteria:

- Zero or one visible stars. This is strictly a function of θ_{FOV} and M_{thresh} .
- Views that only image ambiguous star-pairs. For example, if two stars are visible, then that pair must be ambiguous; if three stars are visible then all three pairs combinations must be ambiguous. Thus, ambiguous pairs are only harmful if no unambiguous pairs are available.[†]

To predict the performance of a particular star tracker design, we must consider both of these cases. Although the two situations appear to be similar, there is an important distinction between them. In views with an insufficient number of stars, the matching routines are aware they cannot provide an attitude solution. In contrast, when only ambiguous pairs are visible, the algorithms may give the wrong solution. Thus, the second situation is much more serious than the first.

Our analysis is broken down into the following steps:

1. An exhaustive off-line search identifies all ambiguous pairs in the catalog. These results are stored for reference.
2. We discretize the celestial sphere using a generalized spiral.²⁰ This spiral generates a set of N , uniformly distributed bore-sight vectors, $\mathbf{r}_k$. For these trials we use $N = 10^5$ —this puts about 430 sample points in an area size of the FOV ($\theta_{\text{FOV}} = 7.5$ deg). The separation between points is approximately 0.7.
3. For each bore-sight vector, we calculate the number of visible stars (i.e., $M_i \geq M_{\text{thresh}}$), that lie closer than θ_{FOV} to $\mathbf{r}_k$. If we have fewer than two stars, we record this view direction in a list of bad views.
4. If there are two or more stars in the FOV, then we choose pairs of stars until we find a pair that is unambiguous. If all the pair combinations are ambiguous, then make note of this view.

The availability, A , is evaluated as a spatial fraction (rather than a temporal fraction):

$$A = \frac{\text{Number of Good Views}}{\text{Total Number of Views}}. \quad (17)$$

To make our results more readily comparable from one design to another, we tabulate the fraction of *bad* views, i.e., $1 - A$.

[†]This criterion is actually more conservative than it has to be. A variety of additional inferences involving the distribution of ambiguous star-pairs could help resolve many ambiguous observations.

These analyses were repeated for a range of different sensor designs. The choice of specific cases was motivated by the characteristics of the ST-16 design (see Sec. 2.1), and the laboratory estimates of the centroid and color index noise (i.e., σ_{pix} , σ_{rg} , and σ_{bg} obtained from McVittie and Enright).⁴ The integrated difference in quantum efficiency between color and monochrome detectors is equivalent to a visual magnitude difference of approximately 1.25. Thus, the design threshold of the ST-16, $M_{\text{thresh}} = 5.75$, is equivalent to a color sensor with $M_{\text{thresh}} = 4.50$. We also consider more sensitive yet plausible designs. Manual inspection of ST-16 images reveal a number of magnitude 6.25 stars. Such a sensor would have a corresponding color threshold $M_{\text{thresh}} = 5.0$.

The centroid position error, expressed in pixels, is related to the arc-length error σ_ϕ by the following approximation:

$$\sigma_\phi \approx 2 \frac{\gamma}{f} \sigma_{\text{pix}}. \quad (18)$$

The γ/f factor converts the error from pixels to radians and the factor of two computes the diagonal error between two measured centroids. Centroid noise captures the performance of the centroid estimation routines. The two selected values for the represent nearest-pixel accuracy ($\sigma_{\text{pix}} = 0.4$), and a modest degree of subpixel centroid refinement ($\sigma_{\text{pix}} = 0.2$).

The results from these trials can be seen in Table 1. The most immediate observation is that no view has only ambiguous pairs. This suggests that system availability is limited by the sensitivity of the detector and optical design, not the pair-based matching scheme. For each test case, we repeated the analysis assuming that only monochrome information was available. Virtually all of the catalog pairs are ambiguous when compared solely on their separation. More importantly, these tests indicate a small but significant number of unresolvable, ambiguous views. Additionally, increasing the catalog size (e.g., by increasing M_{thresh} or θ_{FOV}), or increasing the noise values, will also increase the number of ambiguous star-pairs.

Several secondary observations can be made from these data. The first few rows represent color sensors very close to the reference optical design. Sensor performance for these test cases is not terrible, but substantial areas of the sky (16%) are still subject to bad views. The design solutions with around 2% bad views offer much better availability with optical designs only slightly different from the baseline sensor. In particular, because we have successfully detected magnitude 6.25 stars with the monochrome detector on the ST-16, the solutions with $M_{\text{thresh}} = 5$ and $\theta_{\text{FOV}} = 7.5$ may be plausible without any hardware changes. Although a bad view fraction of 2% is not negligible, that level of performance may be acceptable for some missions. The best availability is encountered with designs a little further from the reference values. These cases demand an increase in both θ_{FOV} and M_{thresh} . Such changes would require substantial alterations to the current optics (and possibly the detector, too).

More generally, the relatively low number of ambiguous pairs suggests that our color-based classifier provides good discrimination between catalog tuples. Ambiguity rises with both larger catalogs and increasing errors, indicating areas of the design-space that should be avoided.

Table 1 Color matching performance.

Test parameters				Color results					
M_{thresh}	θ_{FOV}	σ_{pix}	$(\sigma_{rg}, \sigma_{bg})$	Catalog pairs ($\times 10^3$)	Bad views (%)	Ambiguous pairs (%)	Unresol. pairs	Monochrome ambiguous pairs (%)	Results unresol. pairs
4.5	7.5	0.200	0.020	7.8	15.92	0.26	0	99.77	44
4.5	7.5	0.408	0.020	7.8	15.92	0.84	0	99.99	187
4.5	7.5	0.408	0.050	7.8	15.92	3.86	0	99.99	187
4.5	10.0	0.200	0.020	13.6	2.75	0.19	0	99.88	9
4.5	10.0	0.408	0.100	13.6	2.75	8.11	0	99.99	9
5.0	7.5	0.200	0.020	24.7	2.19	0.71	0	100	8
5.0	7.5	0.408	0.020	24.7	2.19	2.41	0	100	154
5.0	10.0	0.200	0.020	42.9	0.05	0.60	0	100	0
5.0	10.0	0.200	0.100	42.9	0.05	5.79	0	100	0
5.0	10.0	0.408	0.100	42.9	0.05	20.63	0	100	0

5.2 Simulation Experiments

Simulation of the entire matching process provides complementary results to the static catalog analysis. Applying the pair-wise and scene matching algorithms to synthetic observation data allows us to generate estimates of real-world performance.

Our simulations are designed to measure matching performance over the entire sky. Instead of random orientations, we start with the uniform grid of bore-sight directions, and consider multiple roll angles for each case. A typical simulation involves 10^6 inertial orientations, generated using 10^4 unique bore-sight vectors. This corresponds to a density of approximately 70 different bore-sight angles within the FOV of the ST-16.

For each simulated orientation, we determine the stars that would be visible to the sensor, and determine the observed detector-frame star vectors s_{D_i} . This process considers the actual rectangular FOV of the sensor. In this study, we are primarily interested in evaluating the matching performance of our algorithms rather than detection behavior. Thus instead of pixel level image simulation, we abstract the detection process and generate star measurement vectors directly (with appropriate direction and color index noise). We use the following error model:

$$\tilde{s}_{D_i} = \frac{s_{D_i} + \mathbf{n}_\lambda}{\|s_{D_i} + \mathbf{n}_\lambda\|} \quad (19)$$

$$\tilde{\Lambda}_{rg} = \Lambda_{rg} + n_{rg} \quad (20)$$

$$\tilde{\Lambda}_{bg} = \Lambda_{bg} + n_{bg}. \quad (21)$$

The star vector noise, $\mathbf{n}_\lambda$, is distributed in a Rayleigh distribution around the true star vectors (uniform in direction,

effectively Gaussian in both axes perpendicular to the true star direction). color index noise (i.e., n_{rg} , n_{bg}) is normally distributed and zero-mean.⁴ A summary of the simulation parameters is shown in Table 2.

If the simulated scene produces an acceptable match, the sensor orientation is estimated using the Davenport's q-method,²¹ giving the estimate of the quaternion attitude $\hat{\mathbf{q}}_{\text{DI}}$. The net error in the estimate, ϕ_{err} is the angle derived from the error quaternion:

$$\mathbf{q}_{\text{err}} = \mathbf{q}_{\text{DI}} \otimes \hat{\mathbf{q}}_{\text{DI}}^{-1}, \quad (22)$$

$$\phi_{\text{err}} = 2 \arccos(q_{s_{\text{err}}}), \quad (23)$$

where $q_{s_{\text{err}}}$ is the scalar part of the error quaternion.

Figure 9 shows a plot of cumulative distribution in ϕ_{err} for the entire trial. About 90% of the tests have an error of 0.01 deg or less; 99.3% of the trials had 1 deg or less.

Table 2 Simulation parameters.

Parameter	Value
Total poses	10^6
Unique bore-sight vectors	10^4
σ_λ	$55 \times 10^{-6} \mu\text{rad}$ (≈ 0.2 pixels)
σ_{rg}, σ_{bg}	0.05
M_{thresh}	5.0
Half-angle FOV (minor axis)	7.5 deg

Fig. 9 Simulation results showing cumulative distribution of error angle.

While not particularly noteworthy performance for a star tracker, it is consistent with the simulated centroid accuracy. Understanding where the algorithm performs poorly offers valuable insight. Table 3 provides a summary of the problematic orientations (i.e., no match, or a match with high error). These cases represent about 0.7% of the total trials. Most of high-error solutions result from incorrect matching assignments—only one test had both high error, and no matching mistakes. A few cases exist in which a single correct star was matched, but these merely represent two similar star-pairs with a common star. No cases appeared with two or more correct matches, combined with one or more bad matches, suggesting that the geometric consistency check of Eq. (9) is functioning correctly.

Table 3 Simulation results.

Total problem orientations	7078
No solution	897
Poor match (<40% of available stars)	1763
Incorrectly matched stars	
2 (incorrect) of 2 (total stars)	1013
2 of 3	1244
2 of 4	1247
2 of 5	905
3 of 4	2
3 of 6	4
3 of 7	3

We can identify several distinct types of problematic orientations. Generally, matching failures that are diagnosable as such are preferable to potentially incorrect matches:

- No solutions. The sensor matching routines were unable to calculate a match that satisfied the consistency criteria. The most common cause of this response were scenes with fewer than two stars in the FOV.[‡] This class of problem is of least concern to us, since the algorithms are aware of their inability to find a match. Although there were 897 distinct orientations that yielded no-match results there were only 49 unique bore-sight vectors represented in these results. Thus, no-match problems are strongly correlated with star-poor areas of the sky.
- Poor matching. This arises from trials that returned a candidate match solution using fewer than 40% of the available stars. In almost all cases this was an incorrect, two-star match result. We suspect that noise is the primary cause of these problems, but this type of result requires further study. This type of result is more serious than the no-match case, but the low match percentage offers a convenient confidence metric by which we can discard the estimates.
- Wrong matches (two and three stars). This class of response results from an incorrect star-match solution returned by the matching algorithm. Most of these solutions report one matched pair, but a small number of views report three matched stars. This type of response would be very serious on-orbit, and the frequency of occurrence (0.44%) is too high to be neglected. Minimizing these errors is the focus of ongoing research, but we note that raising the poor-match threshold to 50% would halve the number of results in this class.

Figure 10 shows the spatial distribution of problematic orientations. The coordinates in the figure correspond to the direction of the bore-sight axis in celestial (inertial) coordinates. Each point represents one of the bore-sights evaluated and the data aggregated over the whole set of bore-sight angles. “Good” results have low error characteristics for all roll angles. Any other marker represents the worst result for that direction. Most matching difficulties occur in regions of low star density (the band of the Milky Way follows a rough sinusoid through the green region, starting from high declination at the left side of the figure).

5.3 Field Testing

As a complement to our whole-sky simulations, we also conducted a set of night-sky tests, evaluating the performance of the color star matching algorithms with real detector images. These tests rely on the complete processing chain from image acquisition through attitude estimation.

Proper calibration is vital to the collection and interpretation of precise star tracker measurements. Not only must we remove the effects of optical aberrations using our conventional camera calibration (see Sec. 2.2), but we must also compensate for the influence of the Earth’s atmosphere on stellar spectra and positions.

[‡]Or two stars separated by more than the minor-axis arc-length.

Fig. 10 Mollweide projection of spatial distribution of matching problems.

5.3.1 Color calibration

Ground tests of any star tracker will be affected by atmospheric attenuation and refraction. These phenomena affect detection thresholds and distort the inter-star geometry in the images. Color matching routines must also deal with wavelength-dependent atmospheric attenuation of the apparent color indices of the observed stars.

As part of the field trials described by McVittie et al.,⁴ we measured color ratios from 21 bright stars, three from each primary spectral class. Figure 11 compares the measured ratios to the expected values based on the theoretical sensitivity of our detector and the reference spectra from the Pickles catalog. The endpoints of each 'bar' on the figure represent the observed and catalog values for the stars in our survey. As a general trend, atmospheric attenuation and dispersion reduce the observed response, particularly at shorter wavelengths.

Fig. 11 Observed color ratios.

Slant path length (related to altitude and star elevation relative to the Earth observation site), atmospheric aerosols, and humidity can cause significant variations in astronomical observations. As a full atmospheric correction model goes beyond the current scope of work, we instead present a simple system that has proven useful in validating the color matching as a proof-of-concept. Our proposed color correction is a six-parameter bilinear transformation of the form:

$$\begin{bmatrix} \Lambda'_{rg} \\ \Lambda'_{bg} \end{bmatrix} = \begin{bmatrix} \beta_1 & \beta_2 \\ \beta_3 & \beta_4 \end{bmatrix} \begin{bmatrix} \Lambda_{rg} \\ \Lambda_{bg} \end{bmatrix} + \begin{bmatrix} \beta_5 \\ \beta_6 \end{bmatrix}, \quad (24)$$

where $(\Lambda_{rg}, \Lambda_{bg})$ are the observed color ratios and β_i are the model coefficients. The model is a simple low-order curve-fit, and the model parameters are not intended to be physically meaningful. The fitting parameters are determined using a least squares minimization, with any obvious outliers removed from the dataset (see red data point in Fig. 11). Figure 12 shows the calibrated dataset compared to the catalog values. Although the fit does not provide a perfect correction, it reduces the residual RMS color-error from 0.43 (uncorrected) to 0.10 (corrected). A comprehensive atmospheric correction combined with larger sample set could see this error reduced even further.

5.3.2 Atmospheric refraction

To minimize geometric distortions, the star vectors used in these tests were corrected for atmospheric refraction. Using a manual star match and the raw star vectors, we calculate the inertial attitude C_{DI} of the sensor frame D . Refraction may bias these estimates but not by more than a degree. From the latitude, longitude, and time of the observations, we can calculate the rotation, C_{IG} , between the inertial frame I and local topocentric frame G . The precise azimuthal definition of T is immaterial, as long as the 'up' direction is clearly understood. These two matrices give us the sensor-to-topocentric rotation:

Fig. 12 Calibrated color ratios.

$$\mathbf{C}_{DG} = \mathbf{C}_{DI}\mathbf{C}_{IG}. \quad (25)$$

Once we have calculated, $\mathbf{C}_{DT}$, the star vectors, $\mathbf{s}_D$ can be rotated to G , corrected for refraction using the formulae in Bennett,²² and the resulting vector rotated back into D prior to matching.

5.3.3 Matching tests

Once the color-calibration data were processed, we collected a number of images in order to test the color -matching

Fig. 13 Stylized view of one of a field test image. Axes indicate approximate image scale and the dashed rectangle indicates stars used in the field trial.

routines. Due to limitations of the prototype sensor optics, we had to make some adjustments to the instrument focus to get consistent detection of star color. These modifications increased the defocus of the prototype in order to provide better sampling of the star point spread function (PSF). Good detections were limited, to a relatively small region near the centre of the detector. Consequently, we concentrate on a small cluster of three known stars, specifically Orion's belt. Figure 13 shows the three stars of interest.

All three stars are bright and easily identifiable to the naked-eye. As spectral-type O and B stars, they all have strong blue-green response. In fact, both δ -Ori and ζ -Ori map to the same reference spectra, and ϵ -Ori, is only slightly different. These properties make this a challenging triplet of stars to match. The observation data for the stars are shown in Table 4. The Λ_{rg} ratios have low variability, but the Λ_{bg} ratios remain higher than ideal (this catalog was designed for $\sigma_{rg} = \sigma_{bg} = 0.05$). The mean of the color-corrected ratios is fairly close to the predicted catalog values.

Table 5 shows the accuracy of the matching routines when applied to our Orion observations (a total of 70 images). All of the correct results, successfully matched all three stars; all of the incorrect matches reported an incorrect pair of stars (with the third star unmatched). The initial success rate of 50% is much lower than what would be suitable for a flight instrument. Even attaining this modest performance relied on relaxing the geometric constraint threshold [from Eq. (9)], and increasing the values of σ_{rg} and σ_{bg} used in the catalog look-ups to 0.1, from the design value of 0.05. To understand this poor performance, we took a closer look at the observed quantities we were using in the matching routines.

Our algorithms rely on accurate color and geometric measurements. Biases in either quantity lead directly to increases in Δ , and will impair our ability to match to the star catalog. The critical Mahalanobis distance, Δ^* , explicitly limits the bounds of the search space. To investigate how these errors contributed to our poor match success, we devised two idealized corrections. The first was an offset to each measured color ratio, to bring the mean color ratio error to zero. The second was correction to the measured arc-length φ between δ -Ori and ζ -Ori (the longest arc), again bringing the mean error to zero (the magnitude of this correction was approximately 18×10^{-3} deg). A final trial combined both idealized corrections. Although these corrections are artificial, they isolate the separate effects of measurement noise and systematic bias.

When the color and geometric corrections were combined, the matching performance greatly improved. The geometric correction is particularly significant because we have

Table 4 Field testing color ratios.

Star	σ_{rg}	σ_{bg}	Λ_{bg}			Λ_{rg}		
			Raw	Corrected	Catalog	Observed	Corrected	Catalog
δ -Ori	0.045	0.116	1.160	1.206	1.232	0.582	0.639	0.652
ϵ -Ori	0.038	0.096	1.116	1.156	1.161	0.598	0.663	0.674
ζ -Ori	0.030	0.081	1.123	1.166	1.232	0.580	0.652	0.652

Table 5 Field test matching results.

Corrections	Correct matches	Incorrect matches
Basic color	35	35
Ideal color	40	30
Chromatic aberration	58	12
Ideal color + chromatic	65	5

only corrected a single arc length (and not the position of all three stars). Some simple laboratory tests suggest that the geometric problems stem from chromatic aberration. Our geometric camera calibration is conducted using a fibre-coupled, 3000 K incandescent source. Replacing the broadband calibration source with a blue light emitting diode (ζ -Ori is much more blue than our broadband source), introduced differential centroid motion very similar to the observed geometric error. Some contributions to the measurement bias are attributable to testing in the atmosphere (e.g., color bias), whereas others represent limitations of the prototype hardware (e.g., chromatic aberrations to geometry). Careful attention must be paid to these factors in when developing a revised color star tracker prototype.

Atmospheric scintillation is a likely cause for some of the observed color and centroid noise. Observations made with short exposures and small apertures (as in our sensor), can see large amplitude variations between images, and significant band-to-band variability as well.²³ As an example, separate experiments with the monochrome ST-16 show that around 75% of the temporal centroid noise during field tests is attributable to scintillation. This effect makes it very difficult to design a field test that assesses both noise performance and matching effectiveness. Matching performance during terrestrial tests will always need to deal with larger color error than would be encountered on-orbit.

One final insight was gained from a careful inspection of algorithm execution. The three stars observed in the field test are very close in spectrum; two are identical and the third only slightly different. As a result, it is common for color noise to lead to a reversed star assignment, particularly with the initial matched pair. With only a single pair of stars in view, little more can be done to resolve the ambiguity. Our algorithm does not use an explicit search and hence, cannot backtrack through a star assignment. This approach is very fast if we can make one correct initial match, but does not deal well with situations where color-noise has reversed a star assignment. An alternate implementation with a limited amount of backtracking is currently under study, and is expected to resolve this shortcoming.

6 Conclusions

In this paper, we have established a framework for the star matching routines necessary for color -based star tracking. Color star tracking is a novel concept that can be used to leverage advances in commercial imaging systems to develop new types of satellite attitude sensors. This technology can be employed in dedicated star trackers, or as a contingency mode of operation for a payload camera. Our

approach combines coarse spectral information available from a CMOS CFA detector and conventional inter-star geometry. In many cases, the color and geometry of a single pair of stars is sufficient for an unambiguous match.

The color star matching algorithms use a star-pair comparison based on a Mahalanobis-distance classifier. When available, we match additional stars by employing a RANSAC-derived consensus algorithm. The RANSAC approach is robust to false stars and other outliers. We have also examined some of the practical challenges of star catalog organization, including color representation, choice of data sources, and catalog formatting. A series of simulation and field trials have validated the effectiveness of these contributions.

Our experimental results are encouraging. Although the reference point design is not well-suited for color applications, catalog analysis shows that evolutionary design changes can lead to better performance. Comprehensive simulations demonstrate good accuracy and robustness in the majority of sensor orientations; these tests also help to identify cases where the algorithm performance can be improved to prevent incorrect matches. Finally, our field-tests demonstrate that the matching routines will work with the noise properties encountered in a real sensor.

Notwithstanding the positive outcomes of our tests, our experiments do suggest the need for some crucial hardware, software, and testing refinements. First, improved optical performance is necessary for future sensors. Better achromatic optics and explicit color corrections in the sensor calibrations represent promising initial paths to achieving this goal. Second, although the bilinear color corrections improve the correspondence between catalog and measured color values, better approaches are necessary to reduce residual errors. Finally, we recognize that the need for complicated corrections for local terrestrial viewing conditions would make validation of any flight-instrument troublesome. An adaptive, on-line algorithm in which the sensor gradually learned the specific color ratios for each cataloged star or spectral type is currently under investigation.

We recognize that the algorithms and optical designs presented here are not optimal. Rather they demonstrate the essential feasibility of this novel star tracking concept. Measurements from simple color cameras can be used to match stars. Although color sensors are not as sensitive as their monochrome counterparts, the information gain from color information makes two-star lost-in-space solutions practical. We feel that performance in these tests is quite promising and that the next generation of color star trackers will provide performance competitive with traditional monochrome sensors.

Acknowledgments

The authors would like to thank Mr. Doug Sinclair from Sinclair Interplanetary for the donation of prototype hardware.

References

1. B. B. Spratling and D. Mortari, "A survey on star identification algorithms," *Algorithms* **2**(1), 93–107 (2009).
2. P. Servidia and R. S. Pena, "New robust star identification algorithm," *IEEE Trans. Aerosp. Electron. Syst.* **42**(3), 1127–1131 (2006).
3. D. Accardo and G. Rufino, "Brightness-independent start-up routine for star trackers," *IEEE Trans. Aerosp. Electron. Syst.* **38**(3), 813–823 (2002).

4. G. McVittie and J. Enright, "Colour star tracking I: star measurement," *Opt. Eng.* **51**(8), 084402 (2012).
5. J. Enright et al., "Towards star tracker only attitude estimation," in *Proc. 24th AIAA/USU Conf. on Small Satellites*, AIAA and Utah State University (USU), Logan, Utah (2010).
6. Aptina Imaging, "Image Sensors—MT9V022IA7ATM," (January 2012), http://www.aplina.com/products/image_sensors/mt9v022ia7atm/.
7. Aptina Imaging, "Image Sensors—MT9P031I12STM," (January 2012), http://www.aplina.com/products/image_sensors/mt9p031i12stm/.
8. J. Heikkila and O. Silven, "A four-step camera calibration procedure with implicit image correction," in *Proc. IEEE Comput. Soc. Conf. on Comput. Vis. and Pattern Recognit.*, pp. 1106–1112, IEEE Computer Society, San Juan, Puerto Rico (1997).
9. G. J. McLachlan, *Discriminant Analysis and Statistical Pattern Recognition*, pp. 17–19, John Wiley & Sons, Inc., Toronto (2004).
10. H. Lee, C. S. Oh, and H. Bang, "Modified grid algorithm for star pattern identification by using star trackers," in *Proc. IEEE Int. Conf. on Recent Adv. in Space Technol.*, pp. 385–391, IEEE, Istanbul, Turkey (2003).
11. H. Lee and H. Bang, "Star pattern identification technique by modified grid algorithm," *IEEE Trans. Aerosp. Electron. Syst.* **43**(3), 1112–1116 (2007).
12. M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM* **24**(6), 381–395 (1981).
13. D. Hoffleit and W. H. Warren Jr., "Bright Star Catalogue," 5th revised Ed., Astronomical Data Center, NSSDC/ADC (1991).
14. H. L. Johnson and W. W. Morgan, "Fundamental stellar photometry for standards of spectral type on the revised system of the yerkes spectral atlas," *Astrophys. J.* **117**(1), 313 (1953).
15. W. W. Morgan and P. C. Keenan, "Spectral classification," *Annu. Rev. Astron. Astrophys.* **11**(1), 29–50 (1973).
16. R. O. Gray and C. J. Corbally, *Stellar Spectral Classification*, Princeton University Press, New Jersey (2009).
17. A. J. Pickles, "A stellar spectral flux library: 115025000," *Publ. Astron. Soc. Pac.* **110**(749), 863–878 (1998).
18. F. Valdes et al., "The Indo-US library of cool feed stellar spectra," *Astrophys. J. Suppl. Ser.* **152**(2), 251–259 (2004).
19. V. I. Burnashev, "Spectrophotometry of 1588 stars," *Abastumanskaya Astrofiz. Obs. Bull.* **59**(83) (1985).
20. E. B. Saff and A. B. J. Kuijlaars, "Distributing many points on a sphere," *Math. Intellig.* **19**(1), 5–11 (1997).
21. M. D. Shuster, "Algorithms for determining optimal attitude solutions," NASA Technical Report NASA-CR-183473 (1978).
22. G. G. Bennett, "The calculation of astronomical refraction in marine navigation," *J. Navigation* **35**(2), 255–259 (1982).
23. D. Dravins et al., "Atmospheric intensity scintillation of stars II: dependence on optical wavelength," *Publ. Astron. Soc. Pac.* **109**(736), 725–737 (1997).

John Enright holds a BSc (1997) from the University of Toronto (Engineering Science: Aerospace) and a MS (1999) and a PhD (2002) from MIT in Aerospace Systems. He is currently an associate professor in Aerospace Engineering at Ryerson University in Toronto. Having joined the faculty at Ryerson University in 2003, he is now the Principal Investigator of the Space Avionics and Instrumentation Laboratory (SAIL). While at MIT (1999 to 2003), he led the software development for the SPHERES flight project, and the GFLOPS real-time spacecraft simulation testbed. His research interests include spacecraft avionics and sensor processing, systems engineering and flight software. He is a member of the AIAA, CASI, and the IEEE.

Geoffrey R. McVittie received his ME from Ryerson University, in 2008. He is currently a doctoral student studying Aerospace Engineering in the Space Avionics and Instrumentation Laboratory (SAIL) at Ryerson University. His research focuses on the design and development of new generation star tracker sensors.