

Calibrating the Pointing and Optical Parameters of the STEREO Heliospheric Imagers

D.S. Brown · D. Bewsher · C.J. Eyles

Received: 28 July 2008 / Accepted: 12 October 2008 / Published online: 7 November 2008
© Springer Science+Business Media B.V. 2008

Abstract The Heliospheric Imager (HI) instruments on the *Solar TERrestrial RELations Observatory* (STEREO) observe solar plasma as it streams out from the Sun and into the heliosphere. The telescopes point off-limb (from about 4° to 90° elongation) and so the Sun is not in the field of view. Hence, the Sun cannot be used to confirm the instrument pointing. Until now, the pointing of the instruments have been calculated using the nominal preflight instrument offsets from the STEREO spacecraft together with the spacecraft attitude data. This paper develops a new method for deriving the instrument pointing solutions, along with other optical parameters, by comparing the locations of stars identified in each HI image with the known star positions predicted from a star catalogue. The pointing and optical parameters are varied in an autonomous manner to minimise the discrepancy between the predicted and observed positions of the stars. This method is applied to all HI observations from the beginning of the mission to the end of April 2008. For the vast majority of images a good attitude solution has been obtained with a mean-squared deviation between the observed and predicted star positions of one image pixel or less. Updated values have been obtained for the instrument offsets relative to the spacecraft, and for the optical parameters of the HI cameras. With this method the HI images can be considered as “self-calibrating,” with the actual instrument offsets calculated as a byproduct. The updated pointing results and their by-products have been implemented in SolarSoft.

D.S. Brown (✉)

Institute of Mathematics and Physics, Aberystwyth University, Penglais, Aberystwyth, SY23 3BZ, UK
e-mail: dob@aber.ac.uk

D. Bewsher · C.J. Eyles

Space Science and Technology Department, Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, UK

D. Bewsher

e-mail: d.bewsher@rl.ac.uk

C.J. Eyles

e-mail: cje@star.sr.bham.ac.uk

C.J. Eyles

Grupo de Astronomía y Ciencias del Espacio, ICMUV, Universidad de Valencia, Valencia, Spain

Keywords Instrumentation and data management · Instrumental effect

1. Introduction

The *Solar Terrestrial Relations Observatory* (STEREO) is a NASA-funded mission that is intended, in particular, to study the propagation of coronal mass ejections (CMEs) from the Sun into the heliosphere out to around 1 AU and beyond (Kaiser, 2005). It is made up of two nearly identical spacecraft at about 1 AU from the Sun in heliocentric orbits, one ahead of the Earth in its orbit (the A spacecraft) and one behind (the B spacecraft), with the separation between each spacecraft and the Earth increasing by about 22° each year. The Sun–Earth Connection Coronal and Heliospheric Investigation (SECCHI) instrument package (Howard *et al.*, 2008) comprises a set of five remote-sensing instruments on each probe, one imaging the Sun in extreme ultraviolet light (the Extreme UltraViolet Imager – EUVI), two coronagraphs (COR-1 and COR-2) with fields of view extending out to around 4° elongation from the Sun, and a pair of telescopes known as the Heliospheric Imagers (HI). HI-1 has a field of view of about 20° diameter with its optical axis pointing nominally in the ecliptic plane around 14° elongation from the Sun-centre direction, whilst HI-2 has a field of view of about 70° with its optical axis pointing at around 53.7° elongation from Sun centre (Harrison, Davis, and Eyles, 2005). These two fields overlap somewhat and cover the range of approximately 4° – 90° along the Sun–Earth line. Each STEREO spacecraft has an almost identical version of the SECCHI package (giving four HI cameras, HI-1A, HI-2A, HI-1B, and HI-2B), so Earthward-propagating CMEs will be observed from two widely separated viewpoints on either side of the Sun–Earth line.


The HI instrument is intended to primarily study CMEs as they propagate out into the heliosphere, but it also provides measurements of the F and K corona, comets, near-Earth objects, and interplanetary dust (Davis and Harrison, 2005). The background star field in the HI images can additionally be used for stellar variability studies (Bewsher *et al.*, in preparation).

One of the first CMEs observed by HI was described by Harrison *et al.* (2008), and more advanced techniques for identifying and tracking CMEs and corotating interaction regions (CIRs) have been developed (Davies *et al.*, 2008; Rouillard *et al.*, 2008a; Sheeley *et al.*, 2008a). There have also been examples of CMEs/CIRs interacting with the Earth (Sheeley *et al.*, 2008b) and other planets such as Venus (Rouillard *et al.*, 2008b).

Investigation into the behaviour of comets has also been carried out by using HI observations, in particular on the composition of the tail of Comet McNaught (Fulle *et al.*, 2007) and the interaction of a CME with Comet Encke, which caused the disconnection of the comet tail (Vourlidas *et al.*, 2007).

To determine CME positions and directions, as well as for some of the other scientific objectives of the HI instruments (*e.g.*, stellar variability studies), the pointing attitudes of the HI cameras must be accurately determined. Initial attempts to do this have made use of the provided spacecraft attitude solutions, together with the nominal preflight offsets between the HI cameras and spacecraft coordinates. One of the spacecraft axes ($+X$) is maintained accurately aligned with Sun centre by means of a guide telescope (GT), which is part of the SECCHI instrument package (Howard *et al.*, 2008). The spacecraft attitude solutions are derived from the error signals from the GT, together with a roll angle derived from the spacecraft star trackers. However, overplotting the expected positions of known stars onto a HI image using attitude solutions derived in this way (*e.g.*, Figure 1) shows an obvious discrepancy with the background star field observed by HI.

Figure 1 HI-1A image from May 2007 with catalogue stars overplotted as diamonds. Ideally, each diamond should contain a “dot” representing a star in the image. The F corona has been subtracted from the data.



Various factors can contribute to the inaccuracies in the attitude solutions obtained in this way, including errors in the spacecraft pointing solutions (particularly the roll angle) and deviations of the nominal instrument offsets from the spacecraft axes, together with changes in these offsets since launch.

The aim of this paper is to report on an alternative method for deriving attitude solutions for the HI telescopes that avoids the problems just discussed. Essentially, each camera can be used as a highly sensitive star tracker with many stars detectable in the field of view, so that an attitude solution can be derived for each image by using the stars in that image, independent of the spacecraft attitude data. The method involves matching known catalogue stars with those identified in the background star field, and perturbing the pointing parameters to optimise the fit between the predicted and observed star positions.

The methodology developed for these attitude determinations is described in Section 2. The results and their implications will be discussed in Section 3, with the specific results for HI-1A, HI-1B, HI-2A, and HI-2B contained in the [Appendix](#). Discussion on making the results of this work available to the solar community via the Solar Software distribution (SolarSoft) can be found in Section 4.

2. Methodology

To accurately deduce the pointing of a HI instrument, three steps are required. First, stars imaged in the HI field of view must be identified. Second, the locations of known stars from a catalogue must be transformed into the HI field-of-view coordinate system (which requires an estimate of the instrument pointing). Third, the identified stars and the catalogue stars must be compared, and a measure of goodness of fit must be calculated. These last two steps can be repeated with different pointing estimates to minimise the measure of fit. A flowchart showing these steps can be seen in Figure 2. An algorithm for predicting the next pointing estimate can improve the efficiency of the minimisation. Although Figure 1 shows an image

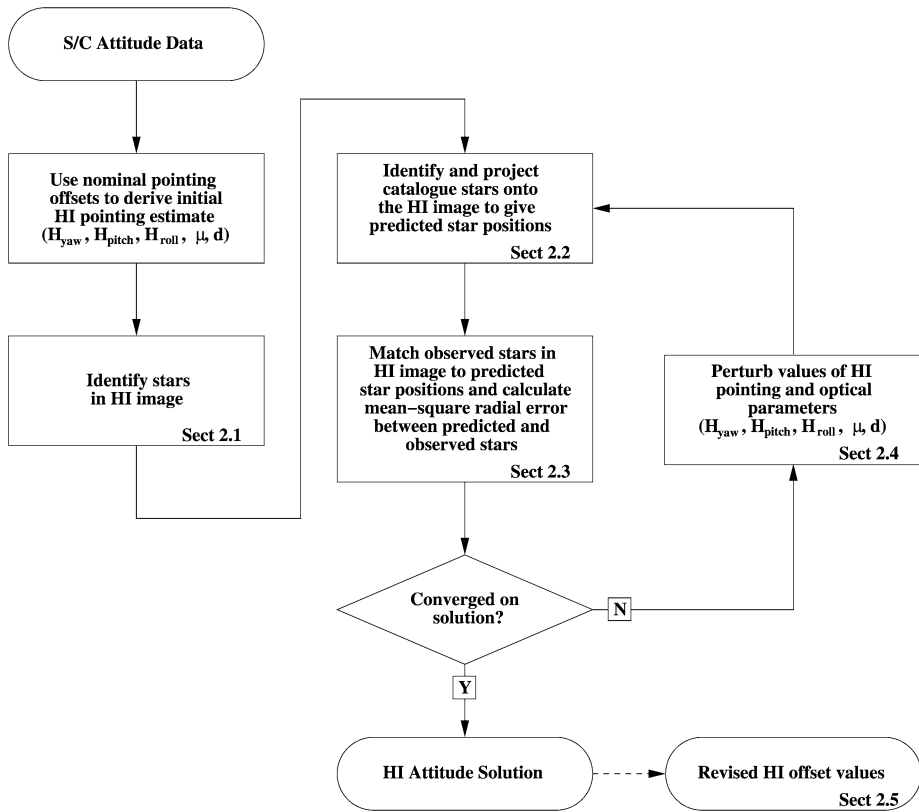


Figure 2 Flowchart outlining the minimisation algorithm described in this paper.

with the F corona removed, it is not necessary to remove the F corona before applying this methodology.

2.1. Identifying Stars in the HI Field of View

Stars appear in the HI field of view as compact spots of bright intensity. The brighter the star, the larger the bright spot. The location of bright stars can be found by locating strong, localised peaks in intensity.

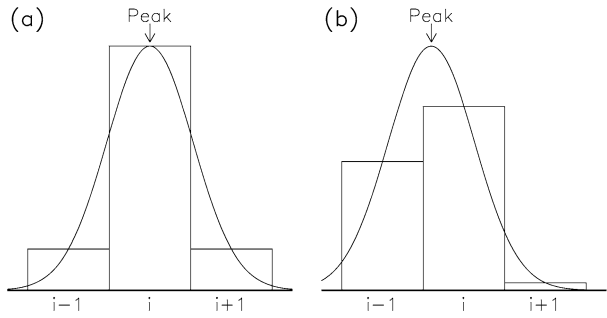
From a practical point of view, all pixels in the image are tested to see whether they are a star location. A subsquare of $(2n_p + 1) \times (2n_p + 1)$ pixels centred on the pixel to be tested is defined. The pixel is identified as a star if it has the maximum intensity in the subsquare and its intensity is greater than the local background by some threshold I_{thresh} . So

$$I_{\text{cen}} \geq \max_{-n_p \leq i, j \leq n_p} I_{i, j}, \tag{1}$$

$$I_{\text{cen}} - I_{\text{bg}} \geq I_{\text{thresh}}, \tag{2}$$

where I_{cen} is the intensity of the centre pixel of the subsquare (*i.e.*, the one that is being tested) and I_{bg} is the local background intensity, which will be defined as the average inten-

Figure 3 Intensity values of pixels next to the peak intensity can be used to deduce the subpixel location of the intensity peak. If neighbouring pixel intensities are about equal (a), then the peak will be located near to the centre of the middle pixel. If intensities differ (b), then the peak will be located away from the centre of the middle pixel towards the higher intensity pixel.



sity of the pixels around the edge of the subsquare, that is,

$$I_{bg} = \frac{1}{8n_p} \left[\sum_{i=-n_p}^{n_p} (I_{i,-n_p} + I_{i,n_p}) + \sum_{j=-(n_p-1)}^{n_p-1} (I_{-n_p,j} + I_{n_p,j}) \right]. \tag{3}$$

This gives two parameters, I_{thresh} , which controls the minimum brightness of identified stars, and n_p , which is akin to a smoothing length to reduce the effect of noise.

For data processed by the SolarSoft routine `SECCHI_PREP` (the SECCHI data preparation and calibration program), values of $I_{\text{thresh}} = 10\text{--}15 \text{ DN s}^{-1}$ and $n_p = 2$ pixels are generally sufficient to identify all stars of magnitude 6 or brighter in the red part of the spectrum, as both of the HI-1 and HI-2 instruments are most sensitive to red light (Harrison, Davis, and Eyles, 2005).

This method gives the location of a star to the nearest pixel; however, it is possible to improve on this by analysing the pixel intensities neighbouring the peak. This intensity distribution will indicate whether the peak is near the centre of the pixel or skewed towards either side of the pixel (as shown in Figure 3). A simple way to simulate this effect is with an average of pixel locations weighted by using the square of pixel intensities. So the location of the peak is approximated by

$$x = \frac{(i - 1)I_{i-1,j}^2 + iI_{i,j}^2 + (i + 1)I_{i+1,j}^2}{I_{i-1,j}^2 + I_{i,j}^2 + I_{i+1,j}^2}, \tag{4}$$

$$y = \frac{(j - 1)I_{i,j-1}^2 + jI_{i,j}^2 + (j + 1)I_{i,j+1}^2}{I_{i,j-1}^2 + I_{i,j}^2 + I_{i,j+1}^2}. \tag{5}$$

Weighting by the square of intensity, rather than intensity alone (*i.e.*, a “centre-of-mass” type calculation) is preferable as it is more robust at approximating the location of the intensity peak in asymmetric cases since a greater emphasis is put on pixels with higher intensity. For asymmetric intensity profiles, the centre of mass is unlikely to coincide with the location of peak intensity. Asymmetric intensity profiles in HI observations may arise from stars that are close to other stars, that coincide with regions of changing corona (*e.g.*, F corona, CME fronts, *etc.*), or that have asymmetric point-spread functions (*e.g.*, HI-2B; see Section 3.8). Additionally, by weighting with intensity squared, accurate results may be obtained using fewer pixels (as with the three pixels just used, for example) than would be required for comparable results using a centre-of-mass calculation. However, overall it is unlikely that either method would have an appreciable effect on the final pointing results.

2.2. Converting Catalogue Stars to the HI Field of View

Known stars from star catalogues usually have their positions given as Right Ascension (RA) and Declination (Dec) in the Geocentric Equatorial Inertial (GEI) system. These denote the locations of stars viewed from the solar system barycentre. The STEREO spacecraft follow orbits around the Sun similar to that of the Earth at distances that are small compared to the distances to other stars. So the small subarcsecond effects of parallax are negligible for HI observations, and the catalogue RA–Dec coordinates can reasonably be used for transformation into the HI field of view.

The star coordinates are transformed to the instrument field of view in two stages. First, the positions of the stars must be transformed from RA–Dec to an instrument frame of reference, which requires an estimate of the instrument bore-sight pointing in RA and Dec, together with the instrument roll. This initial estimate is made by using the spacecraft attitude solution, together with the nominal offset between the spacecraft and instrument reference frames. The second stage involves projecting the star locations relative to the instrument frame of reference into the HI perspective (*i.e.*, predicting the positions of the stars in the camera image plane).

Traditionally, almost all solar imaging instruments, including coronagraphs, have used the gnomonic or “tan” projection. In this projection, a star at an angle α to the instrument optical axis, or bore sight, is projected onto the image plane at a radial distance R from the bore-sight position given by

$$R = F \tan(\alpha), \quad (6)$$

where F is the focal length of the optical system and angular symmetry about the optical axis is assumed. The position of a star in the image plane is given by

$$x = F \tan(\alpha) \cos(\beta), \quad (7)$$

$$y = F \tan(\alpha) \sin(\beta), \quad (8)$$

where β is the angular position of the star about the optical axis.

During preflight testing and calibration of the HI cameras, it soon became evident that this representation is not adequate for HI because of its wide-angle optics and the resultant distortion at the edge of the field of view. The distortions at the edge of the nominal circular fields of view of HI-1 and HI-2 are $\sim 1\%$ and 4% , respectively, rising to nearly 2% and 10% if the corners of the square fields of view are considered.

Although the possibility of using a polynomial relationship was considered, this complication was avoided when it was demonstrated during instrument calibration that the projection could be accurately represented by the relationship

$$R = F_p \frac{(\mu + 1) \sin(\alpha)}{\mu + \cos(\alpha)}, \quad (9)$$

where F_p is the paraxial focal length and μ is a distortion parameter. Inspection of this equation shows that for $\mu = 0$ the relationship reverts to the simpler tan projection.

It should be noted that this relationship characterises the Azimuthal Perspective (AZP) projection as discussed by Calabretta and Greisen (2002). There is no particular physical basis for applying this projection to the HI cameras; the relationship is used simply because it is a suitable two-parameter model that can accurately represent the behaviour of the optics, including the distortion at the edges of the field of view.

Although this relationship was formulated in terms of F_p and μ , two alternative parameters may be used. Firstly, F_p is related to the paraxial plate scale, f_p , by

$$f_p = \frac{t}{F_p} \quad \text{radians per pixel,} \quad (10)$$

where t is the linear dimension of a CCD pixel. By converting to practical units, and taking the size of a pixel to be $t = 13.5$ microns, this becomes

$$f_p = 2785/F_p \quad \text{arcsecs per CCD pixel,} \quad (11)$$

where F_p is measured in millimetres. It is this value of f_p that is provided in the FITS header supplied with HI data.

Also, the diameter of the (circular) field of view, d , may be related to f_p and μ by

$$f_p n_{\text{pix}} = \frac{2(\mu + 1) \sin(d/2)}{\mu + \cos(d/2)} \quad \text{radians,} \quad (12)$$

or

$$f_p n_{\text{pix}} = \frac{412530(\mu + 1) \sin(d/2)}{\mu + \cos(d/2)} \quad \text{arcsecs,} \quad (13)$$

where n_{pix} is the total number of CCD pixels along each axis (the CCD is square).

In this paper d and μ have been chosen as the two independent parameters characterising the optics, although the results in terms of F_p and f_p will also be quoted where appropriate.

In principle, it should be possible to determine the offset between the instrument optical axis (*i.e.*, the direction of symmetry implicit in this relationship) and the centre of the image (on the CCD) as part of the attitude-fitting procedure. However, initial tests showed that this offset was very strongly coupled to the instrument pointing, with a very shallow minimum in parameter space located near the centre of the CCD. Consequently, it was decided to fix the optical axis at the centre of the CCD. Consideration of the instrument design, together with the topology of the fitted parameter space, indicated that this approach was justified; however, this aspect will be revisited in Section 3.5 after the main results have been discussed.

Thus, there are five instrument parameters to be optimised in the fitting procedure for each image, the instrument pointing and roll (H_{yaw} , H_{pitch} , and H_{roll}) and the optical parameters (μ and d). Although the optical parameters do have nominal (preflight) values, they are also optimised as part of the initial fitting process. In later sections of the paper, improved values for these parameters are derived and the effects of fixing them at these refined values are considered.

As RA, Dec, and spacecraft roll are usually quoted in decimal degrees, all angles in this paper will also be in decimal degrees.

The star catalogue used for this analysis is NOMAD, the Naval Observatory Merged Astrometric Dataset (Zacharias *et al.*, 2004).

2.2.1. Transforming Stars to the Instrument Frame of Reference

First, the RA and Dec, (ϕ, θ) , of the stars are mapped to a unit sphere in a Cartesian frame, where $(0, 0, 0)$ is centred on the spacecraft, the x -axis points to the first point of Aries, the

z -axis points to the celestial pole, and the y -axis completes a right-handed system. So the location of a star is given by

$$\mathbf{u} = \begin{pmatrix} x \\ y \\ z \\ w \end{pmatrix} = \begin{pmatrix} \sin(90^\circ - \theta) \cos(\phi) \\ \sin(90^\circ - \theta) \sin(\phi) \\ \cos(90^\circ - \theta) \\ 1 \end{pmatrix}, \tag{14}$$

where w is a weighting for use with 4×4 transformation matrices as described by Watt (2000). For this particular analysis, $w = 1$ always holds and has no bearing on the analysis: however, the weighting variable is particularly useful for transformations of translation, scaling, and projection, and so the full 4×4 matrix method is described here with future extensions in mind.

To convert to an instrument frame of reference, the transformation

$$\mathbf{A} = \mathbf{RT}, \tag{15}$$

can be used, where

$$\mathbf{T} = \begin{bmatrix} \sin(H_{yaw}) & -\cos(H_{yaw}) & 0 & 0 \\ -\cos(P) \cos(H_{yaw}) & -\cos(P) \sin(H_{yaw}) & \sin(P) & 0 \\ \sin(P) \cos(H_{yaw}) & \sin(P) \sin(H_{yaw}) & \cos(P) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{16}$$

with

$$P = 90^\circ - H_{pitch},$$

adjusts for the instrument bore-sight pointing (H_{yaw} , H_{pitch}) in the RA–Dec frame of reference, and

$$\mathbf{R} = \begin{bmatrix} \cos(H_{roll}) & -\sin(H_{roll}) & 0 & 0 \\ \sin(H_{roll}) & \cos(H_{roll}) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{17}$$

corrects for the instrument roll, H_{roll} . The coordinate system now has its z -axis pointing in the direction of the instrument and its x - and y -axes being the horizontal and vertical focal-plane coordinates. The coordinate system has now become left-handed so that the x -axis goes from “screen” left to right, and the y -axis goes from bottom to top.


Note that stars in front of the observer will have positive z and stars behind will have negative z . This can be used to filter out unseen stars before they are projected into azimuthal perspective.

2.2.2. Projection into Azimuthal Perspective

The theoretical geometry of the azimuthal perspective is shown in Figure 4 and is described in further detail by Calabretta and Greisen (2002). The basic transformation equations are described in the following text. For each star, the radial distance from the centre of the focal plane, R , and the depth along the line of sight, Z , are defined as

$$R = \sin(\alpha) = \frac{1}{r} \sqrt{x^2 + y^2}, \tag{18}$$

Figure 4 Theoretical geometry of the azimuthal perspective. Stars on the unit sphere are projected onto the image plane according to a point of projection, P, which is a distance μ away from the centre of the sphere. Note that this diagram illustrates the mathematical conversion to AZP and does not represent the actual HI optical layout.



$$Z = \cos(\alpha) = \frac{z}{r}, \tag{19}$$

where

$$r = \sqrt{x^2 + y^2 + z^2}, \tag{20}$$

although in this case the stars are defined to be on a unit sphere, so $r = 1$.

The projected radial distance from the centre of the focal plane is given by

$$R_\alpha = \frac{(\mu + 1)R}{Z + \mu}, \tag{21}$$

where μ is the previously defined distortion parameter. This can be translated to projected screen coordinates by

$$x_\alpha = R_\alpha \cos(\beta) = \frac{x R_\alpha}{R}, \tag{22}$$

$$y_\alpha = R_\alpha \sin(\beta) = \frac{y R_\alpha}{R}. \tag{23}$$

These coordinates can finally be translated into pixel position by using the instrument field-of-view diameter, d , and the number of pixels across the image, n_{pix} , to scale the projected positions (assuming that the CCD image is square). The scaling factor is given by

$$S = \frac{[\cos(d/2) + \mu]n_{\text{pix}}}{2(1 + \mu) \sin(d/2)}. \tag{24}$$

The projected screen coordinates can then be translated, so that $(0, 0)$ denotes the bottom left of the image rather than the centre, and scaled to pixel number by

$$x_s = (x_\alpha + 1)S, \tag{25}$$

$$y_s = (y_\alpha + 1)S. \tag{26}$$

Note that for this analysis the pixel position $(0, 0)$ is defined to be the bottom-left corner of the image, rather than the centre of the bottom-left pixel as is often used (e.g., in IDL). Translating each of the pixel coordinates, (x_s, y_s) , by 0.5 would correct for this.

2.3. Comparing Identified and Catalogue Stars

To judge the accuracy of the pointing and optical parameters, the identified stars and catalogue stars must be compared. Each catalogue star is matched with the nearest star identified in the image and the pixel distance between the two is measured. The square of this quantity is averaged over all catalogue stars in the image to give R_{avg} , a measure of the accuracy of the fit. So

$$R_{\text{avg}}^2 = \frac{1}{n_{\text{stars}}} \sum [(x_s - x)^2 + (y_s - y)^2], \quad (27)$$

where n_{stars} is the number of catalogue stars in the image.

From a practical point of view, it may be beneficial to ignore catalogue stars that are not well matched to a star identified in the image. Stars may not always be identified in an image, for example, if they are behind a planet, occulter, saturated part of the image, or some other feature. This can be done by defining an image mask to eliminate known problems such as the occulter and saturated pixels, and setting an upper limit for the distance between a catalogue star and its matched identified star for counting the pair as being associated. For this work the maximum distance was set to 30 image pixels.

2.4. Minimising R_{avg}

The next step is to find the best set of the five pointing and optical parameters that minimise R_{avg} . Starting with an initial estimate, the parameters can be systematically perturbed to converge on the minimum.

The algorithm used is as follows. Each parameter is looped through in turn. For each parameter, a small perturbation is made in both a positive and negative direction, so

$$p_j^+ = p_j + \delta p_j, \quad (28)$$

$$p_j^- = p_j - \delta p_j, \quad (29)$$

where p_j is the parameter in question and δp_j is the step size. R_{avg} can then be calculated for both these parameter perturbations, and if either of these are lower than the previous R_{avg} , then the tested parameter takes this new value.

This process is then iterated, reducing the parameter step sizes if no better parameter step can be taken in a particular iteration (*i.e.*, taking finer steps when approaching a minima). This iteration can be continued for a fixed number of steps, or stopped when the step sizes fall below some threshold (*e.g.*, numerical precision on the computer). For the minimisations performed in this paper, a maximum of 500 iterations are performed with the loop being exited if the step sizes fall below 10^{-6} . Other applications may require different values.

This is basically a “walking downhill” algorithm to find the local minimum. This type of method depends on the local topography of $R_{\text{avg}}(p_1, p_2, p_3, p_4, p_5)$ and the initial guess to find the global minimum. Use of more advanced methodology, such as random mutation or genetic algorithms, would better find the global minimum regardless of the initial guess. However, the nominal pointing and optical parameters should provide a good initial guess, which makes the use of a more intensive method unnecessary.

To minimise the effects of unfavourable local topography, for each image, the algorithm is run multiple times from a parameter set that is randomly perturbed slightly from the nominal values. Any optimised parameter set that does not find the “best” minimum can be eliminated, and the remaining solutions can be averaged to improve the best fit parameters. For this work, 20 runs generally provide enough statistics to adequately identify the best minimum.

2.5. Calculating the Instrument Offset to the Spacecraft

Given the instrument pointing (calculated by using the aforementioned method) and the spacecraft pointing (as supplied in the FITS headers associated with the data), the offset pointing of the HI instruments with respect to the spacecraft can be deduced. The x - and y -offsets can be calculated by using only the position of the centre of the HI image, but the roll offset requires an additional (arbitrary) reference point in the HI image (for example, one of the corners).

The centre point and the reference point can be transformed from the instrument frame to the spacecraft frame by using a variant of the transformation matrix (Equation (15)) given in Section 2.2.1. The matrix \mathbf{A}_I is defined as being this transformation calculated with the optimised instrument pointing, and \mathbf{A}_{SC} is defined as being the transformation calculated using the given spacecraft pointing.

So a point $\mathbf{u} = (x, y, z, 1)$ in the instrument frame is transformed to the spacecraft frame by

$$\bar{\mathbf{u}} = \mathbf{A}_{SC} \mathbf{A}_I^{-1} \mathbf{u}. \tag{30}$$

The direct transformation from spacecraft space to instrument space is given by

$$\mathbf{B} = \bar{\mathbf{R}} \bar{\mathbf{T}}, \tag{31}$$

where

$$\bar{\mathbf{T}} = \begin{bmatrix} -\cos(X_0) & 0 & \sin(X_0) & 0 \\ -\cos(Q) \sin(X_0) & \sin(Q) & -\cos(Q) \cos(X_0) & 0 \\ \sin(Q) \sin(X_0) & \cos(Q) & \sin(Q) \cos(X_0) & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{32}$$

with

$$Q = 90^\circ - Y_0, \tag{33}$$

corrects for the instrument offsets, (X_0, Y_0) , and

$$\bar{\mathbf{R}} = \begin{bmatrix} \cos(R_0) & -\sin(R_0) & 0 & 0 \\ \sin(R_0) & \cos(R_0) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{34}$$

gives the rotation for the roll offset, R_0 . Note that this transformation differs in structure from the one given in Equations (15)–(17) as \mathbf{B} transforms between two left-handed systems (the spacecraft and instrument frames, where the z -axis is along the line of sight), whereas \mathbf{A} transforms between right- and left-handed systems (the celestial Cartesian frame, where the z -axis is “up,” and the spacecraft/instrument frame).

So the transformation

$$\mathbf{u} = \mathbf{B} \bar{\mathbf{u}} \tag{35}$$

will return the previously calculated point $\bar{\mathbf{u}}$ back to \mathbf{u} .

As the centre of the HI image is at $(0, 0, r)$ in instrument space (where r is the radius of the projection sphere whose exact value is irrelevant for this calculation) the roll does

not matter (as this is the reference point that the instrument is rolled about). So the direct transformation

$$\mathbf{u} = \bar{\mathbf{T}}\bar{\mathbf{u}} \tag{36}$$

can be rearranged to give

$$\tan(X_0) = \frac{\bar{x}_0}{\bar{z}_0}, \tag{37}$$

$$\tan(90^\circ - Y_0) = \frac{\bar{x}_0 \sin(X_0) + \bar{z}_0 \cos(X_0)}{\bar{y}_0}, \tag{38}$$

where

$$\bar{\mathbf{u}}_0 = \begin{pmatrix} \bar{x}_0 \\ \bar{y}_0 \\ \bar{z}_0 \\ 1 \end{pmatrix} \tag{39}$$

is the centre of the HI image in the spacecraft frame.

The roll offset can be calculated by transforming the reference point, $\bar{\mathbf{u}}_r$, in spacecraft coordinates by $\bar{\mathbf{T}}$ to get $\hat{\mathbf{u}}_r = \bar{\mathbf{T}}\bar{\mathbf{u}}_r$ and manipulating the matrix equation

$$\mathbf{u}_r = \bar{\mathbf{R}}\hat{\mathbf{u}}_r \tag{40}$$

to get

$$\tan(R_0) = \frac{x_r\hat{x}_r - y_r\hat{y}_r}{x_r\hat{x}_r + y_r\hat{y}_r}, \tag{41}$$

which completes the offset parameters.

The offsets are calculated with respect to a celestial frame of reference. The primary coordinate system in the associated FITS headers is the HelioProjective Cartesian (HPC) frame, which has the opposite sign convention for longitude (*i.e.*, it is a left-handed system rather than a right-handed one). To obtain the offsets with respect to the HPC frame, the x -offset and the roll offset should be given the opposite sign. Note that general conversion between celestial and HPC frames is much more complicated than this.

In this method it is assumed that the spacecraft pointing is accurate. The guide telescope guarantees spacecraft bore-sight pointing to within at least 70 arcseconds (Howard *et al.*, 2008), but the roll is more uncertain and of unknown accuracy. If any aspect of the spacecraft pointing were to be inaccurate, then this methodology will give the instrument offset relative to the assumed spacecraft pointing rather than the actual spacecraft pointing. This means that variations in the accuracy of the spacecraft pointing will contribute to variations in the instrument offset.

2.6. Calculating Instrument Pointing in HPC Coordinates

As previously mentioned, the default coordinate system for HI data in the associated FITS headers is that of HPC coordinates, as described in Thompson (2006), which are provided along with celestial RA and Dec coordinates. The HPC system has the Sun at 0° in longitude and latitude on a sphere centred on the observer (*i.e.*, the spacecraft). The HPC roll is from the solar equator.

Once the celestial instrument pointing is known, the HPC instrument pointing can be easily calculated by using variations of the transformation matrices (Equations (15)–(17)) that have already been encountered. $\mathbf{A}_{I,C}$ is defined to be the transformation matrix in Section 2.2.1 with the celestial instrument pointing (calculated from the minimisation). $\mathbf{A}_{SC,C}$ is this matrix with the celestial spacecraft pointing, $\mathbf{A}_{SC,HPC}$ is the matrix with the spacecraft HPC pointing, and $\mathbf{A}_{I,HPC}$ is the matrix with the unknown HPC instrument pointing. For a point \mathbf{u} in instrument space, the following circular relation should hold:

$$\mathbf{u} = \mathbf{A}_{I,HPC} \mathbf{A}_{SC,HPC}^{-1} \mathbf{A}_{SC,C} \mathbf{A}_{I,C}^{-1} \mathbf{u}. \tag{42}$$

So by setting

$$\bar{\mathbf{u}} = \mathbf{A}_{SC,HPC}^{-1} \mathbf{A}_{SC,C} \mathbf{A}_{I,C}^{-1} \mathbf{u}, \tag{43}$$

where $\bar{\mathbf{u}}$ is the location of \mathbf{u} in the HPC frame, this becomes

$$\mathbf{u} = \mathbf{A}_{I,HPC} \bar{\mathbf{u}}. \tag{44}$$

By considering the point $(0, 0, r)$, the centre of the HI image in instrument space (where instrument rotation does not affect the transformation) and an additional reference point, the matrix equation can be analytically manipulated to show that

$$\tan(H_{yaw}) = \frac{\bar{y}_0}{\bar{x}_0}, \tag{45}$$

$$\tan(90^\circ - H_{pitch}) = \frac{\bar{x}_0 \cos(H_{yaw}) + \bar{y}_0 \sin(H_{yaw})}{\bar{z}_0}, \tag{46}$$

where H_{yaw} and H_{pitch} are the instrument pointing in HPC, \mathbf{u}_0 and $\bar{\mathbf{u}}_0$ refer to the centre of the HI image, and the vector elements of $\bar{\mathbf{u}}$ are given by

$$\bar{\mathbf{u}}_0 = \begin{pmatrix} \bar{x}_0 \\ \bar{y}_0 \\ \bar{z}_0 \\ 1 \end{pmatrix}. \tag{47}$$

The instrument roll angle, H_{roll} , in HPC can be calculated by applying the transformation

$$\hat{\mathbf{u}}_r = \mathbf{T}_{I,hpc} \bar{\mathbf{u}}_r, \tag{48}$$

where $\bar{\mathbf{u}}_r$ refers to the additional reference point in HPC coordinates (calculated with the transformation just outlined), $\mathbf{T}_{I,hpc}$ is the translation part of $\mathbf{A}_{I,hpc}$ (given by Equation (16)), and $\hat{\mathbf{u}}_r$ is a calculable point that differs from the original point \mathbf{u}_r in the instrument frame by the HPC instrument roll, that is,

$$\mathbf{u}_r = \mathbf{R}_{I,hpc} \hat{\mathbf{u}}_r. \tag{49}$$

This can be manipulated to give

$$\tan(H_{roll}) = \frac{x_r \hat{x}_r - y_r \hat{y}_r}{x_r \hat{x}_r + y_r \hat{y}_r}, \tag{50}$$

which completes the pointing parameters.

This calculation uses the spacecraft pointing in both the celestial and HPC frames. So long as the spacecraft pointing is accurate (or at the very least any inaccuracies in the celestial and HPC values are consistent), then the HPC instrument pointing will also be accurate. This pointing information can be calculated alongside the celestial pointing.

2.7. Some Notes on Implementation

The main implementation issue is identifying stars in the image. A large majority of images are good and provide no problems; however, some images have bright features (*e.g.*, stars and planets) that saturate the detector and prevent stars in these saturated pixels from being detected. This can be solved by generating a mask of saturated pixels and eliminating catalogue stars that fall on these pixels.

A related problem is that of stars that lie very near the edge of the image, or a mask. Perturbation of the parameter set can push the catalogue star in and out of the acceptable image, whereas its theoretical accompanying identified star may be outside of the acceptable image and so not found. The catalogue star will then be incorrectly matched, causing an incorrect increase in R_{avg} . This can be corrected by adding a border mask near the edge of the image, around occulters and around saturated pixels (of about 10 bins seems to be sufficient). Stars may be identified in the image within this mask, which can be matched with catalogue stars as they come into and out of the acceptable image.

Some images are heavily saturated (particularly near the beginning of the mission when the Earth and Moon are in the field of view), or otherwise corrupted, such that no sensible minimisation can be done. When no stars are identified in the image an error is flagged, but as these images are unlikely to be of scientific use, accurate pointing should not be necessary. Caution should also be taken with images that are optimised by using a small number of stars.

When identifying stars in an image, a subsquare size, n_p , and intensity threshold, I_{thresh} , were chosen. For most of the HI cameras the point-spread function (PSF) is narrow, so a small subsquare (*e.g.*, $n_p = 2$) is satisfactory for star identification. However, for images with a wider PSF (especially HI-2B; see Section 3.8), a larger value of n_p may be preferable. The ideal value of I_{thresh} is one that is as large as possible while still identifying all of the catalogue stars that are to be matched. Too high a value will fail to identify stars that are in the catalogue, but a value too low may lead to many dimmer noncatalogue stars being identified, which can confuse the minimisation. For example, tests with good HI-1A data suggest that $I_{\text{thresh}} = 15 \text{ DN s}^{-1}$ is suitable to identify all 6th magnitude stars.

All of the HI instruments are most sensitive to red light, in particular HI-1, which has a 630–730 nm bandpass filter (Harrison, Davis, and Eyles, 2005). As such, it is preferable to use a catalogue of red stars, and look at the magnitude of red light, rather than the total or blue magnitudes.

A slight modification to the walk downhill method is employed in practice. Instead of a single positive and negative perturbation for each parameter in a single iteration, the algorithm tries a large step followed by a small step. So in a single iteration a parameter may take no steps, a large step, a small step, or a large step followed by a small step (which may be in the opposite direction to the large step). This is to tension for different rates of topographic variation in R_{avg} ; for example, it may be beneficial in valley-like features to take small steps in one direction, but large steps in another.

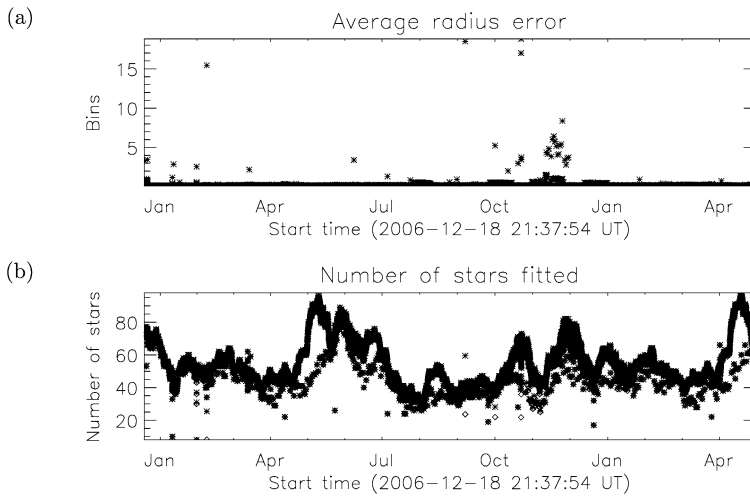


Figure 5 Plots showing (a) the error in the fits and (b) the number of stars used for the HI-1A minimisation where μ and d are allowed to vary over time.

3. Results

The operational phase of the mission began with the opening of the protective doors covering the HI cameras on 13 December 2006 and 11 January 2007 for HI-A and HI-B, respectively, around the time of the final lunar swingby and injection into heliocentric orbit for each spacecraft. There followed a period of a few months during which various commissioning and calibration activities were carried out and the final synoptic observing programmes for the HI instruments were developed and tested.

The main science operations phase began around April 2007. The HI-1 camera provides 36 images per day at 40-minute cadence, whilst HI-2 provides 12 images at 2-hour cadence. Each image consist of a large number of individual exposures of the 2048×2048 pixel CCD, which are cleaned of cosmic-ray hits, 2×2 binned, and then summed onboard to produce the 1024×1024 pixel science images, which are downlinked through the telemetry. In addition, there is normally one single exposure with the full 2048×2048 pixel resolution of the CCD downlinked every day for each camera.

The algorithm outlined in the previous section has been run on all HI data from the beginning of the mission to the end of April 2008, excluding images in the known bad data lists (see <http://www.ukssdc.rl.ac.uk/solar/stereo/info>).

The goodness of fit for each minimisation from the HI-1A camera can be seen in Figure 5. Figure 5(a) shows the average bin (image pixel) error per star fitted, R_{avg} . For the overwhelming majority of data (more than 99%), this falls below 1 bin. Figure 5(b) shows the number of catalogue stars in the instrument field of view (and not behind any bad pixel mask) that are used for the minimisation. For every fit, there is an asterisk that shows the number of catalogue stars in the image and a diamond that shows the number of stars that are accurately fitted (*i.e.*, the catalogue star was matched to an identified star within a 30-bin radius). For the majority of images these points coincide (and so are not discernible in the plot). In general, between 20 and 100 catalogue stars are successfully fitted for each HI-1A image.

For some periods of time there appears to be two populations on the lower panel (*e.g.*, May 2007). The lower population corresponds to the daily 2048×2048 pixel high-resolution

calibration image, whilst the upper population corresponds to the summed 1024×1024 pixel science images. The single-exposure images seem to be more sensitive to saturation, although this is likely due to identification of saturated pixels being optimised for the 1024×1024 image rather than an instrumental effect. Hence, in these periods there are saturation masks for the 2048×2048 images that hide stars from the fitting algorithm. Note that the location of saturated pixels is provided from the prepared image rather than being determined as part of the minimisation, and the construction of saturation masks performed in this work is done using this information.

These plots are characteristic of the goodness-of-fit plots for the other cameras, although HI-2B is not quite as accurate owing to the larger PSF of that camera (see Section 3.8).

3.1. The Causes of Poor Minimisation

A reasonable criterion is that an R_{avg} less than 1 bin is a good fit, if R_{avg} is between 1 and 5 bins it is acceptable, and if R_{avg} is greater than 5 bins then it is not a good fit. In the majority of cases when there is not a good fit, the increase in error is due to some issue with the observations themselves. Figure 6 shows four examples of HI images with increased errors. If the spacecraft is rolling during the exposure [Figure 6(a)], then the stars do not maintain a fixed position in the image and star trails are evident. This causes problems in identifying stars in the image, and there is no unique pointing solution to be obtained. When particularly bright features are in the field of view [*e.g.*, the Earth in the early part of the mission, Figure 6(b)], then ghosting can occur in the camera optics, which can interfere with the star identification algorithm. Telemetry data dropout can cause empty blocks in an image [Figure 6(c)], which can “hide” stars that would otherwise be identified in the image. This can leave catalogue stars without an identified star to be matched with, so the catalogue star may be matched incorrectly, leading to a higher R_{avg} . Other features such as the Milky Way [Figure 6(d)] can also cause problems, since stars in the Milky Way can be closely packed together and not easily identified in the image, which can lead to matching errors.

Occasionally, there are some images for which no result is obtained. This usually happens when the image is heavily saturated and the saturation mask covers all catalogue stars in the field of view. In this situation, no minimisation can be done.

3.2. Fixing μ and the Field-of-View Diameter

The minimisation process provides values for the optical parameters μ and the field-of-view diameter, d . As the algorithm is numerical, these optimised values will have slight random variations about the “actual solution” over different minimisations.

Figure 7 shows scatter plots of the optimised μ value against the field-of-view diameter for all cameras. The diagonal orientation of the data suggests a correlation between these parameters, so they can trade off to some degree and still maintain a good accuracy of fit in the minimisation.

From an instrumental perspective, the values of μ and the field-of-view diameter should be fixed, since they are parameters of the optical system. Significant variation of these parameters would imply changes in the dimensions or alignment of the optics. For example, such effects could arise from thermal distortion or motion of mechanical components within clearances. There is no evidence of any long-term changes in either parameter through the mission, and short-term changes are probably ruled out by the fact that they would produce other effects (*e.g.*, blurring of the images). Furthermore, it is difficult to envisage any mechanism that could cause short-term changes in these parameters.




Figure 6 A selection of example images where R_{avg} is above 5 bins. They are (a) a HI-1A image where the spacecraft is rolling so the moving stars produce trails in the image, (b) a HI-1B image where the Earth is so bright that, in addition to the broad vertical band of saturated pixels, “ghosting” within the camera optics produces artifacts in other parts of the image, (c) a HI-2A image that has missing telemetry data blocks that stars can “hide” behind, and (d) a HI-2B image that contains the Milky Way where stars packed closely together in the image are difficult to separate and identify. The F corona has not been removed from any of these images, and the intensity scaling may vary between images.

Thus, the variation in the solution for these parameters is likely to be due to the coupling between them, together with random errors introduced within the minimisation. To improve the estimates for μ and the field-of-view diameter, d , the assumption that they should not change over time will be made, and a weighted average (in favour of low R_{avg}) from all the fitted attitude solutions will be calculated. The averages for all HI instruments can be seen in Table 1. The errors on the weighted means are very low, supporting the assumption that these values are accurate and are not significantly changing.

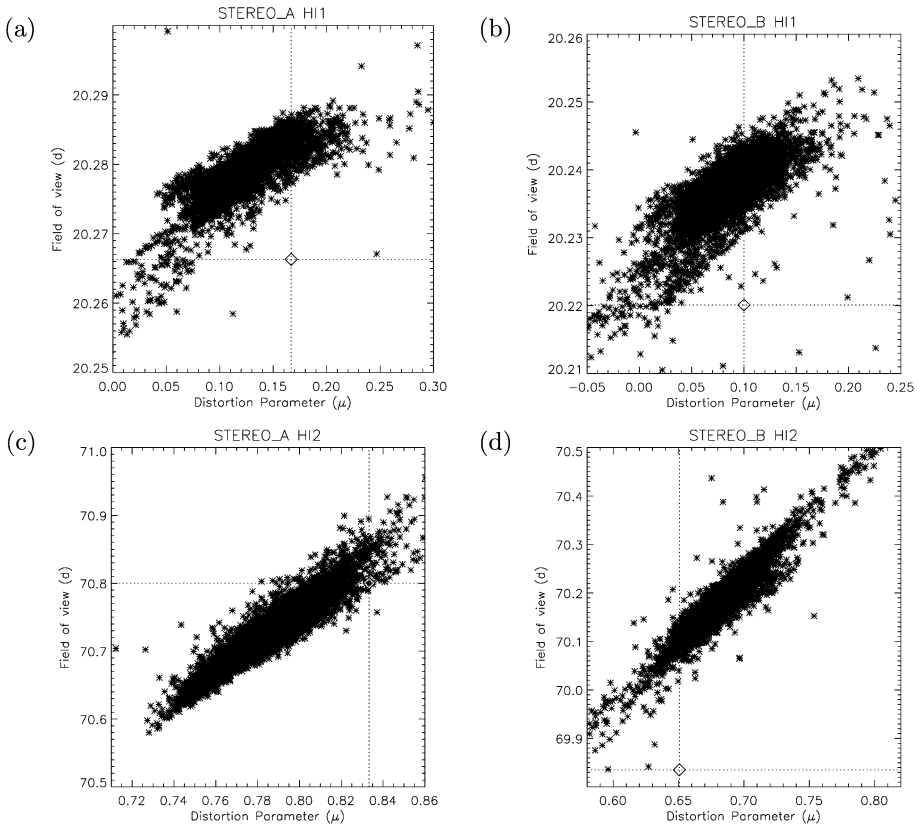


Figure 7 Scatter plots of μ against d in (a) HI-1A, (b) HI-1B, (c) HI-2A, and (d) HI-2B. The dotted lines and diamonds indicate the values of the nominal preflight optical parameters. There is clearly correlation between μ and d for all cameras. Note that because of the axis scaling of the plot, a few points lie outside the plotted region.

Table 1 The averages for μ and the field-of-view diameter, d , of the fitted data. Errors on each are displayed (though for the HI-1 field-of-view diameters, the calculated weighted error falls below computational numerical precision), along with the number of frames used for the calculation. The preflight nominal values for each camera are also given.

Instrument	μ	σ_μ	d (°)	σ_d (°)	N_{frames}
HI-1A	0.102422	2×10^{-6}	20.27528	$<1 \times 10^{-5}$	16085
Preflight	0.16677		20.2663		
HI-2A	0.785486	9×10^{-6}	70.73507	2×10^{-5}	6299
Preflight	0.83329		70.8002		
HI-1B	0.09509	1×10^{-5}	20.23791	$<1 \times 10^{-5}$	15281
Preflight	0.10001		20.2201		
HI-2B	0.68886	1×10^{-5}	70.20152	2×10^{-5}	5796
Preflight	0.65062		69.8352		

It is clear from both Table 1 and Figure 7 that the optimised values for μ and the field-of-view diameter are significantly different from their preflight nominal values. This correction will help improve the matching of stars across the whole image.

Table 2 The paraxial plate scale, f_p , in arcseconds per CCD pixel and the focal length, F_p , in millimetres calculated using the average values for μ and d . Errors propagated through the calculation fall below computational numerical precision. Preflight nominal values are also given.

Instrument	μ	d (°)	f_p (arcsec pixel ⁻¹)	F_p (mm)
HI-1A	0.102422	20.27528	35.96382	77.43894
Preflight	0.16677	20.2663	35.91917	77.53520
HI-2A	0.785486	70.73507	130.03175	21.41785
Preflight	0.83329	70.8002	129.77177	21.46075
HI-1B	0.09509	20.23791	35.89977	77.57710
Preflight	0.10001	20.2201	35.86529	77.65167
HI-2B	0.68886	70.20152	129.80319	21.45556
Preflight	0.65062	69.8352	129.41279	21.52028

Table 3 The overall means of the average radial errors between identified stars and catalogue stars for the cases where μ and d are allowed to vary, and where they are fixed.

Instrument	$\overline{R_{avg}}$ varying	$\overline{R_{avg}}$ fixed
HI-1A	0.233	0.248
HI-2A	0.760	0.783
HI-1B	0.288	0.298
HI-2B	1.457	1.481

As discussed in Section 2.2, the HI optical systems can be parameterised in terms of the paraxial focal length, F_p , or the corresponding paraxial plate scale f_p , rather than the field-of-view diameter, d . These parameters describe the optics *per se*, whereas d describes the camera system, since it includes the dimensions of the detector. Table 2 shows the values for F_p and f_p , as calculated from μ and d , together with their preflight values. It is noted that all values for focal length have decreased from the preflight calibrations, by $\sim 0.1\%$ in the case of HI-1 and $\sim 0.2\% - 0.3\%$ in the case of HI-2. This is interpreted as a result of the preflight calibrations of the optics being done in air, whilst the values derived in these optimisations are in vacuum. The difference in magnitude between the HI-1 and HI-2 changes is interpreted as resulting from the larger number of lens elements, and consequently refracting surfaces, in HI-2.

The minimisation algorithm is now repeated using the averaged values of μ and d given in Table 1. These will not be allowed to vary in the new minimisation, which will produce only optimised pointing information.

3.3. Pointing Results for the Fixed Optical Parameters

Figure 8 shows the revised goodness-of-fit plots for HI-1A when the fixed, average values for μ and d are used in the minimisation. The accuracy seems comparable to the previous case when μ and d were allowed to vary.

Table 3 shows the mean values of R_{avg} for each camera in the cases when μ and d are allowed to vary and when they are taken as fixed. This confirms the visual observation that there is little difference between Figures 5 and 8, although for all cameras there is a slight loss of overall quality of the star fits of the order $(1 - 2) \times 10^{-2}$ bins. This level of accuracy is certainly higher than the level of accuracy in identifying star positions in the HI field of view, so the assumption that μ and the d do not vary is deemed to be acceptable.

Table 4 shows the number of good, acceptable, and bad minimisations given the definition made in Section 3.1 that R_{avg} less than 1 bin is a good fit, R_{avg} between 1 and 5 bins is acceptable, and R_{avg} greater than 5 bins is not a good fit.

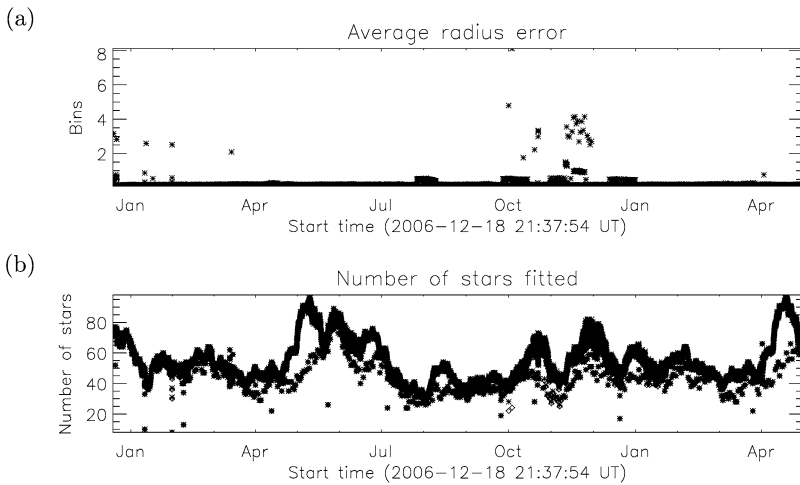


Figure 8 Plots showing (a) the error in the fits and (b) the number of stars used for the HI-1A minimisation where μ and d are fixed at their averaged values.

Table 4 The number of images with a good ($R_{\text{avg}} \leq 1$), acceptable ($1 < R_{\text{avg}} \leq 5$), and bad ($R_{\text{avg}} > 5$) radial error. Also shown is the number of images that produce no result.

Instrument	$R_{\text{avg}} \leq 1$	$1 < R_{\text{avg}} \leq 5$	$R_{\text{avg}} > 5$	No result
HI-1A	15973 (99.29%)	52 (0.32%)	60 (0.38%)	2 (0.01%)
HI-2A	5725 (90.86%)	497 (7.89%)	77 (1.22%)	2 (0.03%)
HI-1B	15032 (98.36%)	173 (1.13%)	76 (0.50%)	2 (0.01%)
HI-2B	2531 (43.21%)	3164 (54.02%)	101 (1.73%)	61 (1.04%)

For HI-1A, HI-2A, and HI-1B, the vast majority of images produce good pointing solutions with a healthy proportion of the remainder being acceptable. The majority of HI-2B images are either good or acceptable with the lower accuracy being due to the PSF (see Section 3.8). The 61 HI-2B cases where no result was obtained all occur in the first month after opening the HI-B door, when the spacecraft was still close to the Earth, with the Earth and Moon in the HI-2B field of view, causing large-scale saturation and ghosting in the images.

3.4. Visual Check of the Pointing Solutions

The accuracy of the optimised pointing solutions can be visually checked by plotting catalogue stars over the image and checking how they align with the background star field. Calculation of the positions of catalogue stars in the HI field of view has been discussed in Section 2.2.

The plots for the spacecraft-derived pointing using the nominal instrument offsets (previously shown in Figure 1) and the optimised pointing for a HI-1A image from May 2007 can be seen in Figure 9. The diamonds indicate where the catalogue stars should appear. For a good fit, a star (dot) should appear at the centre of each diamond. It is clear that the nominal pointing [Figure 9(a)] does not match catalogue stars with the background star field and there is an offset (which can be quite pronounced near the edges of the image) between the

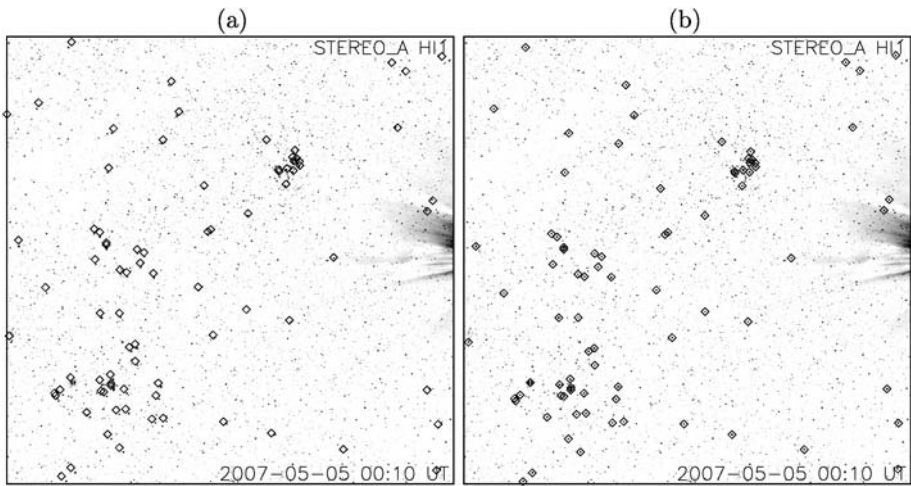


Figure 9 HI-1A images from May 2007 with catalogue stars overlotted as diamonds. Ideally, each diamond should contain a “dot” representing a star in the image. Image (a) uses the nominal pointing to find star locations, and image (b) uses the optimised pointing. For both images, the F corona has been subtracted.

two, whereas the optimised pointing [Figure 9(b)] shows very good correspondence between the background stars and the catalogue stars.

This demonstrates that there is a nontrivial positioning error with the preflight nominal pointing offsets, particularly for fine position analyses, which is dramatically reduced when the attitude solutions are derived from the star fields in the individual images as described in this paper.

3.5. Offset of the Optical Axis from the Centre of the HI Image

In Section 2.2, the assumption was made that the optical axis of the instrument lines up with the centre of the HI image (or CCD). The attitude solutions calculated in this paper are based on this assumption. To accept these pointing solutions as valid, it is important to review this assumption and quantify the accuracy of this assumption.

Take the offset of the optical axis from the centre of the HI image to be given by $(x_{\text{off}}, y_{\text{off}})$. This gives two new parameters that are used when converting catalogue star positions to their positions in the HI image by modifying Equations (25) and (26) so they become

$$x_s = (x_\alpha + 1)S + x_{\text{off}}, \tag{51}$$

$$y_s = (y_\alpha + 1)S + y_{\text{off}}. \tag{52}$$

The minimisation process can now be performed for seven parameters (the two offsets plus the original five parameters) and the results analysed.

Recall from Section 2.4 that for each HI image the minimisation procedure was run 20 times from random perturbations from the nominal pointing. To investigate how the offsets vary in more detail, the minimisation will be run 100 times for each image, randomly perturbing from the best-fit pointing already calculated.

Figure 10 shows scatter plots of (a) RA against x_{off} and (b) Dec against y_{off} for a HI-1A image from 1 June 2007 at 09:00 UT. It is clear that, in both cases, the quantities are




Figure 10 Scatter plots showing (a) how the instrument RA varies with the optical axis x -offset and (b) how the instrument Dec varies with the optical axis y -offset. Only points where $R_{avg} \leq 1$ are plotted.

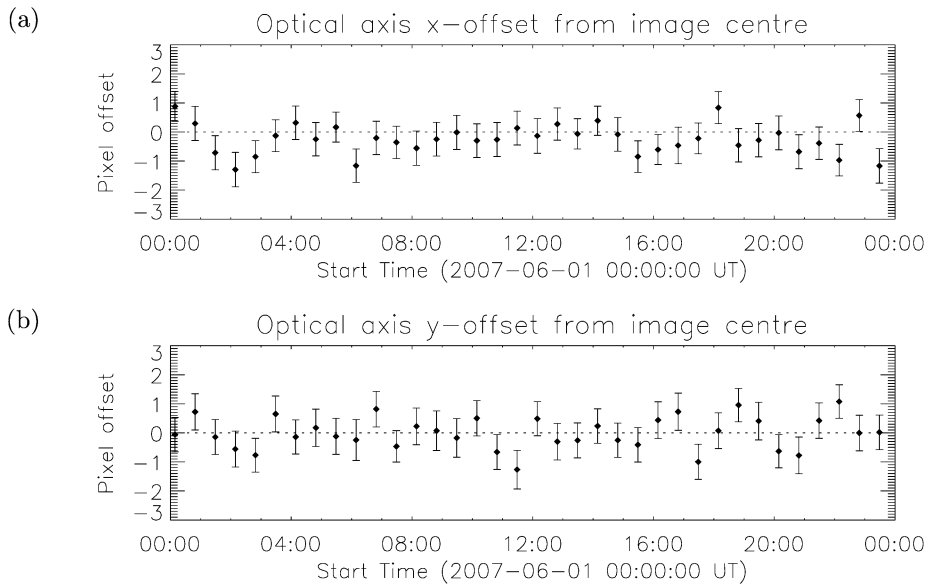


Figure 11 Plots showing the variation of (a) the average optical axis x -offset and (b) the average optical axis y -offset for each 1024×1024 HI-1A image taken on 1 June 2007.

strongly correlated. For a nonzero instrument roll, some correlation between RA and y_{off} , and between Dec and x_{off} , would be expected – to the extreme of a $\pm 90^\circ$ roll where correlation would be completely this way around. However, the instrument roll typically remains closer to 0° and so the original correlation pairings are dominant.

The mean of these x - and y -offsets can be found to give averages for the image, and this process can be repeated for multiple images to see how the average offsets vary over time. Figure 11 shows the average offsets for each 1024×1024 HI-1A image taken on 1 June

Table 5 The daily weighted average x - and y -offsets, along with the errors on the averages, for all four HI cameras on 1 June 2007.

Instrument	\bar{x}_{off}	σ_x	\bar{y}_{off}	σ_y
HI-1A	-0.3	0.1	0.0	0.1
HI-2A	0.23	0.05	-0.29	0.05
HI-1B	0.12	0.09	0.3	0.1
HI-2B	1.05	0.06	0.32	0.06

2007. The average deviates from zero offset by no more than a couple of pixels, with many of the error bars overlapping the zero-offset line.

From these a daily weighted average for the offsets can be found, which for HI-1A on 1 June 2007 gives $\bar{x}_{\text{off}} = -0.3 \pm 0.1$ and $\bar{y}_{\text{off}} = 0.0 \pm 0.1$. However, the subpixel error term is comparable with other sources of error not included in the calculation (such as the accuracy in identifying stars in the HI image), so some care must be taken when interpreting the high degree of significance of this figure as the true error may be larger than stated (although still subpixel).

The evidence from 1 June 2007 suggests that the assumption of fixing the optical axis to the centre of the HI image is a reasonable one for HI-1A. Table 5 shows the daily weighted average offsets for all four HI cameras; the corresponding plots can be seen in the Appendix. In almost every case, the averages are a subpixel offset from the zero-offset point, so by using the same argument as for HI-1A, the offsets for HI-1B and HI-2A may be reasonably set at zero. The exception is HI-2B, which has a 1-pixel x -offset; however, this is small given the issues with the HI-2B PSF (see Section 3.8), so it may also be reasonable to take zero offset for HI-2B, although this will need to be revisited when the HI-2B PSF has been compensated for.

Note that only 1024×1024 images have been used in evaluating the optical axis offset. If setting the offsets to zero is reasonable for the 1024×1024 images then it follows that it is also reasonable for the 2048×2048 images, as the difference between the two is the binning of pixels.

3.6. Instrument Offsets from the STEREO Spacecraft

The nominal pointing solutions are calculated by using the preflight instrument pointing offset from the spacecraft, together with spacecraft attitude solutions. The attitude solutions derived from background stars in the HI images along with the spacecraft pointing contained in the associated FITS headers can be used to calculate the post-launch instrument offsets (see Section 2.5).

The calculated offsets for HI-1A can be seen in Figure 12. The offsets clearly deviate from the nominal values ($X_0 = 13.98^\circ$ for the yaw offset and $Y_0 = R_0 = 0^\circ$ for the pitch and roll offsets, in terms of spacecraft coordinates). The yaw offset is about 0.02° from the nominal and, generally, varies by 0.005° . The pitch offset is about 0.12° from the nominal and, again, generally varies by 0.005° . The roll offset is the most deviated from the nominal by about 1.02° with a variation of 0.01° .

The deviations from the nominal instrument offsets may arise from a number of causes, including prelaunch alignment tolerances, changes to alignments resulting from the launch environment, or changes since launch, possibly from thermal distortion or outgassing of CFRP (carbon fibre reinforced plastic) structures. Changes of alignment could occur within the HI instrument itself, in the mounting of HI onto the spacecraft, or within the spacecraft structures.




Figure 12 Plots showing the variation in offset of the HI-1A instrument from the spacecraft (in terms of spacecraft coordinates). Dashed line shows the nominal offset when it falls within the range of the plot.

The spacecraft attitude data are derived in part from the SECCHI guide telescope, which maintains the Sun-pointing SECCHI instruments aligned accurately with Sun centre (to better than 70 arcseconds), but there is also a roll angle derived from the spacecraft star trackers that is subject to much larger uncertainties, which could also contribute to deviations in the offset.

The origin of the fluctuations in the offsets of HI-1A (and similar fluctuations in HI-2A) during the early part of the mission is at present not fully understood but may be associated with commissioning activities prior to the start of science mission operations in April 2007.

However, it must be emphasised that whatever the origin of the deviations and fluctuations in the offsets may be, the process of using the background stars in each image to derive an optimised attitude solution for that image completely avoids problems in the science data analysis associated with uncertainties in the instrument offsets. The HI images are in effect self-calibrating.

Table 6 shows average values of the pointing offsets between the instruments and spacecraft, along with errors on the averages and preflight nominal offsets. HI-1A and HI-2A are averaged from April 2007 onwards to capture the general trend, after what appears to be an anomalous period early in the mission. HI-1B and HI-2B are averaged over the entire mission, and the larger variation in their offsets is evidenced by the larger errors on the averages.

Table 6 The average instrument offsets from the STEREO spacecraft. Errors on the averages are given, and the nominal preflight values are given on a separate line.

Instrument	x_0	σ_{x_0}	y_0	σ_{y_0}	r_0	σ_{r_0}
HI-1A	14.0037	4×10^{-4}	0.1159	2×10^{-4}	-1.0215	5×10^{-4}
Preflight	13.98		0.0		0.0	
HI-2A	53.4075	2×10^{-4}	0.0662	2×10^{-4}	-0.1175	2×10^{-4}
Preflight	53.68		0.0		0.0	
HI-1B	-14.10	1×10^{-2}	0.022	4×10^{-3}	-0.37	1×10^{-2}
Preflight	-13.98		0.0		0.0	
HI-2B	-53.690	2×10^{-3}	0.213	3×10^{-3}	0.052	7×10^{-3}
Preflight	-53.68		0.0		0.0	

The average offsets can differ from the nominal values by up to a degree in the case of roll and almost all of the differences are significant.

3.7. Pointing Discontinuities in HI-1B

Although average values for the instrument offsets are shown in Table 6, closer analysis of the HI-1B data indicates an additional effect. Figure 13 shows the pointing of HI-1B over an example period of three days. The pointing is clearly discontinuous at several times. This is consistent with pointing discontinuities detected visually in HI-1B data movies. These discontinuities can also be seen in the HI-1B offset plots in the Appendix.

The reason for the discontinuities is not known, but discussion within the SECCHI instrument team suggests that as HI-B is facing the direction of travel of STEREO-B, bombardment of the instrument by dust particles may be perturbing the position of HI-B relative to the spacecraft.

The procedure described in this paper for deriving an attitude solution for each image based on the background stars within that image completely resolves the problem of the pointing discontinuities. A valid attitude solution is obtained for every valid image. In a small proportion of HI-1B images the pointing discontinuities occur during the exposure, resulting in smeared, trailed, or even split star images. However, these images are flagged as “bad” and will not normally be used for science data analysis. The high proportion of HI-1B images in Table 4 that give good solutions confirms the validity of this approach.

These pointing discontinuities may also affect HI-2B observations, depending on whether the pointing of the whole instrument is being perturbed, or just that of the HI-1B camera. However, the degraded PSF of HI-2B (see Section 3.8) makes it difficult to determine which is the case, although there are indications of a larger spread in the offsets for this camera.

3.8. Issues with the HI-2B PSF

As has been previously mentioned, the typical pointing errors in HI-2B data are higher than has been seen with the other three cameras. This results from the PSF of HI-2B, which is larger and more asymmetric than for the other cameras. A comparison of how the PSF affects the HI-2B images as opposed to the HI-2A images can be seen in Figure 14.

The diffuse intensity pattern produced by the PSF at star locations will cause inaccuracies when identifying stars in the HI-2B field of view. These inaccuracies lead to the larger inaccuracies in pointing solutions.




Figure 13 Plots showing the optimised attitude solutions for HI-1B over three days in May 2007. Discontinuities in the pointing can be clearly seen. These correspond to the dashed vertical lines, which denote times where a visible discontinuity is seen in HI-1B data movies.

The derived attitude solutions are not unreasonable and generally give acceptable results for the camera. However, once the PSF of HI-2B has been adequately modelled, the minimisation will be rerun taking it into account, which should give improved attitude solutions.

3.9. Stellar Light Curves

With accurate pointing and optical parameters, it becomes much easier to produce light curves of stars from the background star field. The star-fitted attitude solutions can accurately be used to follow a fixed point in RA and Dec over multiple images and provide a centre point for performing image photometry. Figure 15 shows the light curves for a catalogue star tracked by using both the nominal spacecraft-derived attitude solutions and the fitted attitude solutions. The nominal pointing is unable to follow a specific star, as the catalogue stars do not necessarily coincide with the stars in the HI image (see Section 3.4) and the distance between the predicted position and the actual position in the image can change as it crosses the instrument field of view. This produces the unsatisfactory light curve seen




Figure 14 Comparison of how the PSF affects (a) the HI-2A images and (b) the HI-2B images for an almost identical field of view (containing the Milky Way). The stars in HI-2B are more spread out, in an asymmetric way, than in HI-2A.




Figure 15 Light curves from the binary system HD22766 (RA: 3.677° , Dec: 28.773°) (a) calculated using the nominal pointing and (b) calculated using the optimised pointing. Note that the light curve has only been intensity corrected by using the preflight nominal large-scale flat field as the postlaunch correction is not yet available.

in Figure 15(a). Using the optimised pointing means that the star can be followed with much more confidence, and the resulting light curve [Figure 15(b)] is much more satisfactory and submagnitude variations can be observed. This means that good-quality light curves can easily be produced for stars over periods of up to 20 days for HI-1 data and 70 days for HI-2 data, these being the approximate times for stars to drift across the field of view. A more detailed description of obtaining light curves will be presented by Bewsher *et al.* (in preparation).

4. SolarSoftWare Implementation

To make the best pointing and optical parameters available to the solar community, the attitude solutions calculated in this paper, along with the information for future data as they

become available, have been contributed to the SolarSoftWare (SSW) tree (<http://www.lmsal.com/solarsoft/>; also see Freeland and Handy, 1998) as a set of pointing files.

An IDL subroutine has been provided to update header structures of HI data that have been loaded into IDL. The usage is

```
IDL> hi_fix_pointing, header
```

where `header` is a structure variable (or array) of the HI fits header. This code finds the appropriate pointing file, extracts the pointing and optical information for the provided header, and substitutes that information into the structure. This routine is called as part of `secchi_prep`, the SSW SECCHI preparation and calibration routine.

5. Conclusions

This paper has discussed an algorithm for accurate calculation of the instrument pointing and key optical parameters for the Heliospheric Imager instruments onboard the STEREO spacecraft, by using the background star field observed in HI images.

Known stars from the NOMAD catalogue are transformed to their predicted locations in the HI image (given estimated pointing and optical parameters) and compared to stars identified in the background star field. Once catalogue and identified stars are matched, the average radial distance between all matched pairs can be calculated. By varying the pointing and optical parameters, the average radius can be minimised to find the best-fit pointing and optical parameters.

As the optical parameters should not vary greatly between different images from the same HI camera, the optimised optical parameters from the beginning of the mission to the end of April 2008 are averaged to provide a constant estimate. The minimisation process is repeated for the fixed optical parameters to deduce the final attitude solutions.

The values for the optical parameters obtained in the optimisations differ by small but significant amounts from the values obtained in prelaunch calibrations. In the case of the paraxial focal length the differences are consistent with the changes to be expected from making the measurements in air (for the prelaunch calibrations) and in vacuum.

Selected data have been discussed within the paper. A complete overview for all of the HI cameras can be found in the Appendix. The vast majority of data from all cameras are optimised with good accuracy, although HI-2B is not as good as the other cameras, owing to the broader and asymmetric point-response function of this camera. Nevertheless, most solution sets for HI-2B are acceptable, and once a correction for the PSF has been deduced, the minimisation algorithm can be run on the corrected data.

The optimised attitude solutions provide a much better correspondence between catalogue background stars, which can be demonstrated visually by overplotting catalogue stars onto the HI data. This is important for many applications, including accurately deducing the position and direction of CMEs and obtaining reliable stellar light curves from the HI data. Stars can stay within the HI field of view for around 20 days in HI-1 and 70 days in HI-2, so HI can provide long-time-span light curves that are impractical to obtain from ground-based observatories.

The attitude solutions in this paper are calculated for all HI images that are not in the bad-image list from the beginning of the mission to the end of April 2008 (although this work has identified additional observations that will be added to the bad-image list). It is intended that this will continue to be done for all future data, and the accurate pointing and optical parameters will be made available through the Solar Software tree.

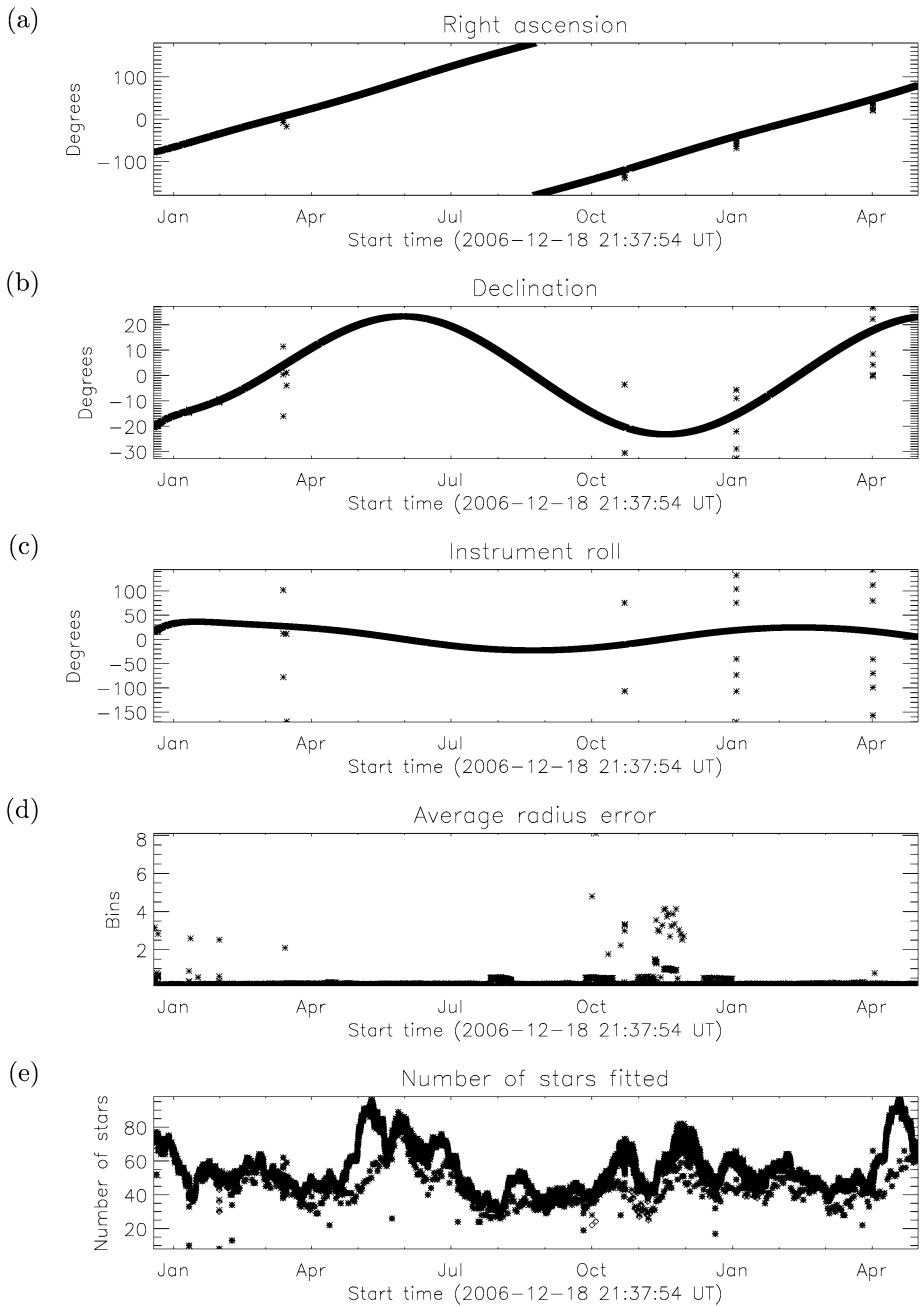


Figure 16 Plots showing the variation in pointing of the HI-1A instrument in (a) RA, (b) Dec, and (c) roll, (d) the error in the fits, and (e) the number of stars used for the fits, where μ and d take their averaged values.

Acknowledgements D.S.B. would like to acknowledge the Science and Technology Funding Council (STFC) for funding through the Aberystwyth rolling grant. J. Davies (RAL) is also acknowledged for her comments on the text. The Heliospheric Imager instrument was developed by a collaboration that included the Rutherford Appleton Laboratory and the University of Birmingham, both in the United Kingdom, the Centre Spatial de Liège (CSL), Belgium, and the U.S. Naval Research Laboratory (NRL), Washington DC, USA. The STEREO/SECCHI project is an international consortium of the Naval Research Laboratory (USA), Lockheed Martin Solar and Astrophysics Lab (USA), NASA Goddard Space Flight Center (USA), Rutherford Appleton Laboratory (UK), University of Birmingham (UK), Max-Planck-Institut für Sonnensystemforschung (Germany), Centre Spatial de Liège (Belgium), Institut d'Optique Théorique et Appliquée (France), and Institut d'Astrophysique Spatiale (France).

Appendix

A.1. HI-1A Results

Assuming that the spacecraft + X -axis (*i.e.*, the Sun pointing axis) always points to the Sun, Figure 16(a) shows that the instrument RA has gone through more than 360° , which demonstrates that the spacecraft has indeed gone through a complete orbit of the Sun in the first year. The declination [Figure 16(b)] varies smoothly between -20° and 20° as the spacecraft orbits the Sun. The instrument roll angle relative to GEI coordinates [Figure 16(c)] also




Figure 17 Plots showing the variation in offset of the HI-1A instrument from the spacecraft. Dashed line shows the nominal offset when it falls within the range of the plot.

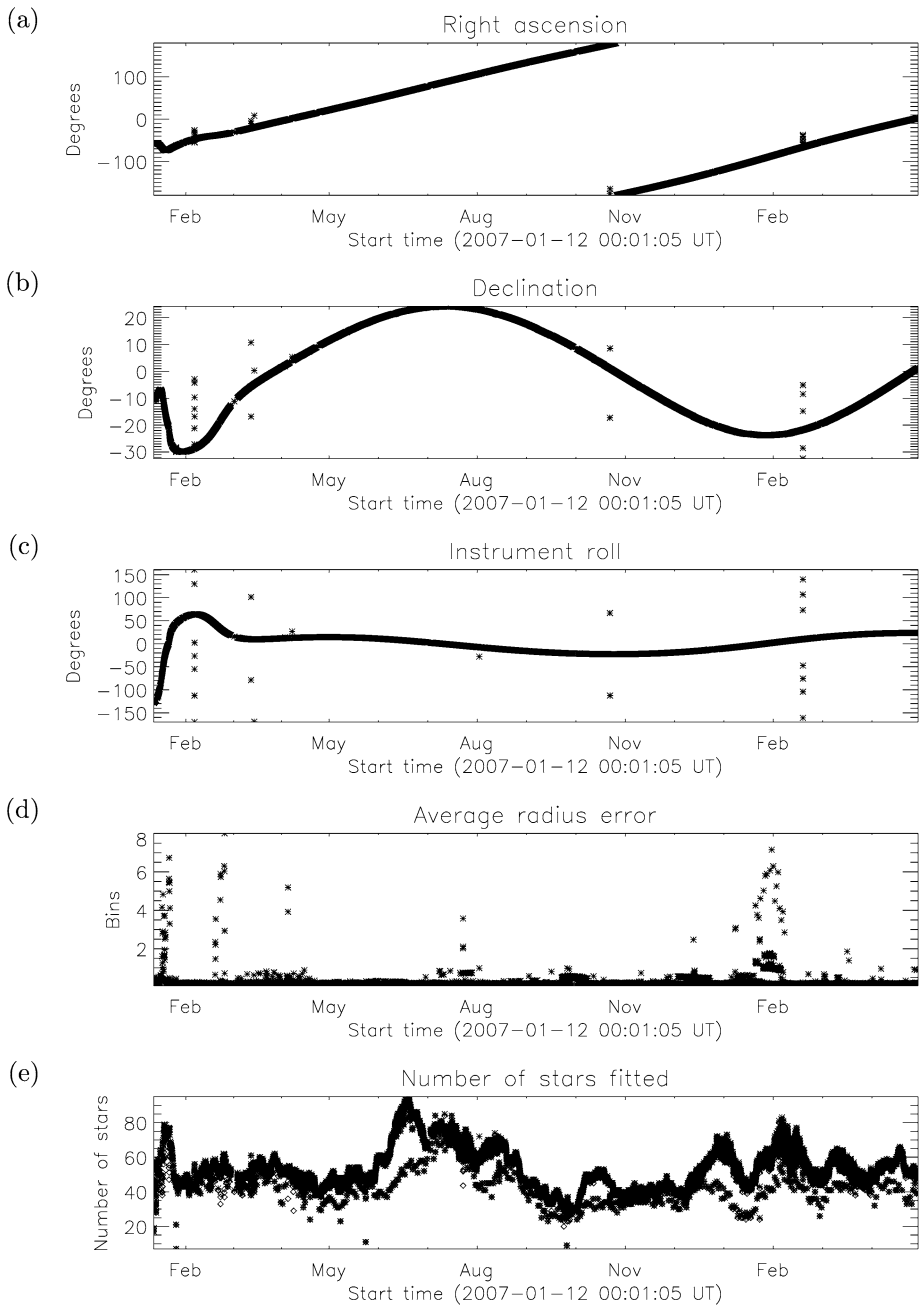


Figure 18 Plots showing the variation in pointing of the HI-1B instrument in (a) RA, (b) Dec, and (c) roll, (d) the error in the fits, and (e) the number of stars used for the fits, where μ and d take their averaged values.




Figure 19 Plots showing the variation in offset of the HI-1B instrument from the spacecraft. Dashed lines show the nominal offset when it falls within the range of the plot.

varies over the course of the year. The spacecraft axes are designed to align closely with the ecliptic plane following a settling down period at the beginning of the mission. Variation of the declination and roll relative to GEI coordinates reflects the fact that the ecliptic plane and the celestial equatorial plane do not coincide. The data spikes in Dec and roll in March and October 2007 and January and April 2008 [Figures 16(b) and 16(c)] correspond to calibration rolls by the spacecraft.

The goodness of fit for each minimisation can also be seen in Figure 16 (which are repeated from Figure 8). Figure 16(d) shows the average bin (*i.e.*, image pixel in either the 1024×1024 or 2048×2048 images) error per star fitted. For the majority of data, this falls below 1 bin. Figure 16(e) shows the number of catalogue stars in the instrument field of view that are used for the minimisation; for every fit, there is an asterisk that shows the number of catalogue stars in the image and a diamond that shows the number of stars that are accurately fit.

Figure 17 shows the calculated instrument offset from the spacecraft and is a repeat of Figure 12. Excluding fluctuations in the first three months, which are probably associated with spacecraft commissioning activities, the pointing offsets maintain a fairly constant value, which is to be expected.

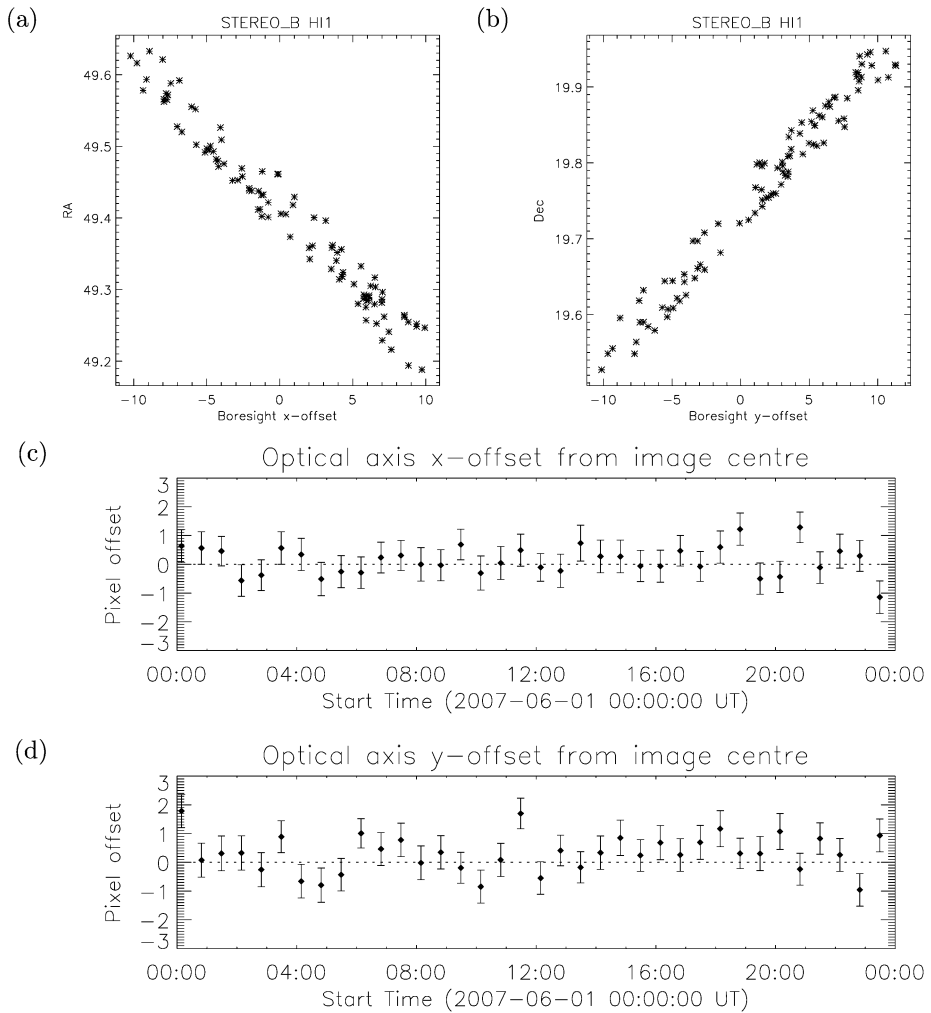


Figure 20 Plots showing the behaviour of the offset of the HI1B instrument’s optical axis. Scatter plots show (a) how the instrument RA varies with the optical axis x -offset and (b) how the instrument Dec varies with the optical axis y -offset. Only points where $R_{\text{avg}} \leq 1$ are plotted. Also shown are (c) the average optical axis x -offset and (d) the average optical axis y -offset for each 1024×1024 HI-1B image taken on 1 June 2007.

A.2. HI-1B Results

The pointing solutions calculated for HI-1B using the average μ and d values can be seen in Figure 18 along with the average pixel error of the fits and the number of fitted stars. The data are generally consistent with very few outlying values. The positions of calibration rolls can be seen, along with the more complicated trend in January–February 2007 corresponding to orbital mechanics following the lunar swingby. This spacecraft has also clearly orbited the Sun in just over the course of the first year of observations. For the majority of data R_{avg} is below 1 bin, providing a good accuracy, with typically 20–100 stars fitted for the minimisation.

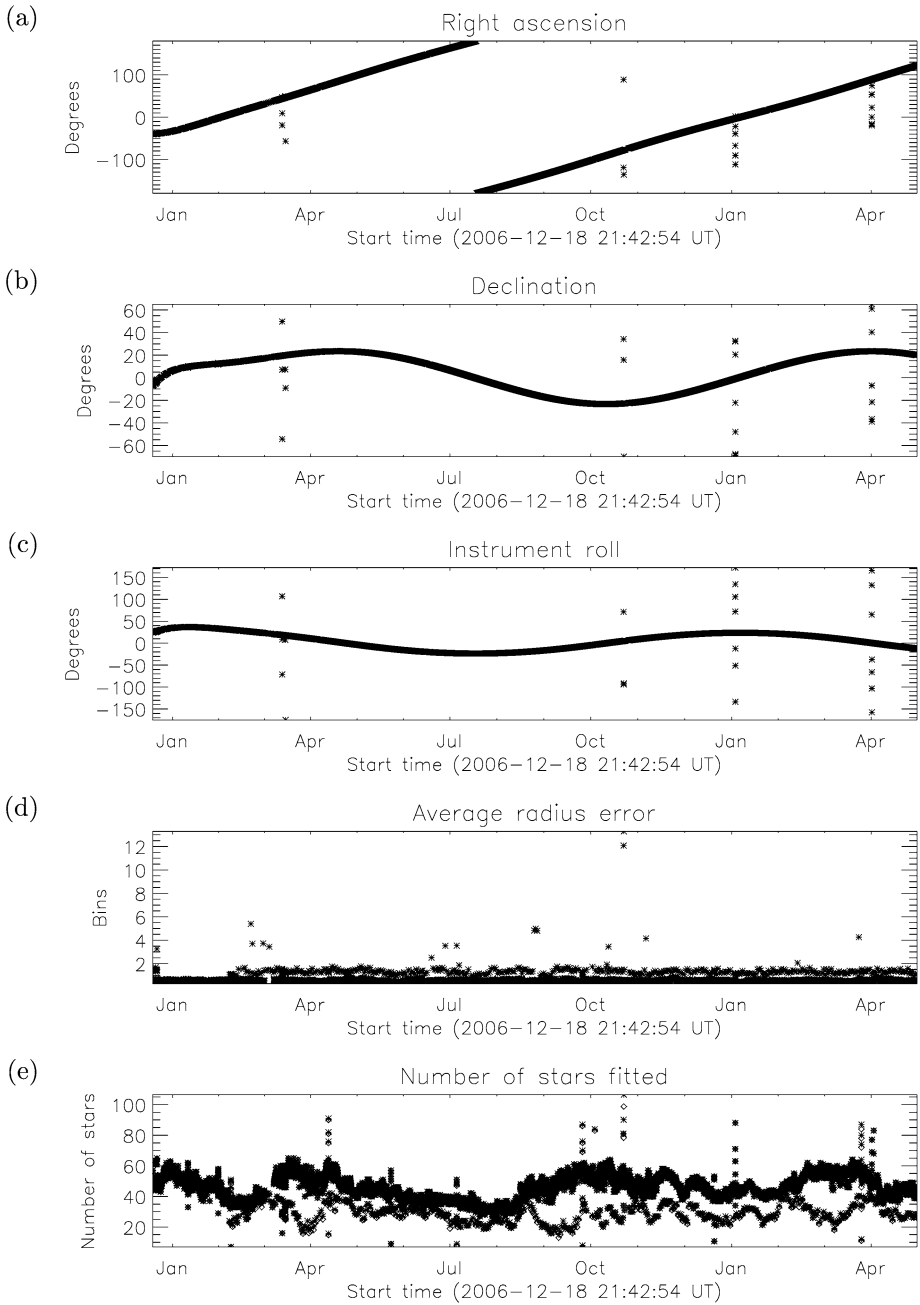


Figure 21 Plots showing the variation in pointing of the HI-2A instrument in (a) RA, (b) Dec, and (c) roll, (d) the error in the fits, and (e) the number of stars used for the fits, where μ and d take their averaged values.




Figure 22 Plots showing the variation in offset of the HI-2A instrument from the spacecraft.

The HI-1B instrument offset can be seen in Figure 19. Discontinuities in the pointing of around 0.1° can be seen (as discussed in Section 3.7). The minimisation algorithm using the background stars in each image enables an accurate attitude solution to be obtained independent of the discontinuities.

Figure 20 shows the behaviour of the optical axis offset from the centre of the HI image. The scatter plots show correlation between RA and x_{off} and between Dec and y_{off} . The image averages for 1 June 2007 are clustered around the zero-offset line with good overlap of error bars. These plots suggest that it is reasonable to set the x - and y -offsets to zero for the minimisation calculation.

A.3. HI-2A Results

The pointing solutions for HI-2A calculated using the average μ and d values can be seen in Figure 21. The data are generally consistent, with very few outlying values. Additionally, it is consistent in shape with HI-1A, which should be expected as both cameras are mounted on the same spacecraft. Similarly, the spacecraft calibration rolls are clearly visible at the same times as they are seen in the HI-1A pointing solutions.

The average pixel error of the fits along with the number of fitted stars can be seen in Figures 21(d) and 21(e). As with the HI-1A case, the vast majority of pixel errors are below 1 bin; however, there is a second population (the 2048×2048 data) that have an error of just above 1 bin. This corresponds with the fact that the width of 1 bin in the 1024×1024

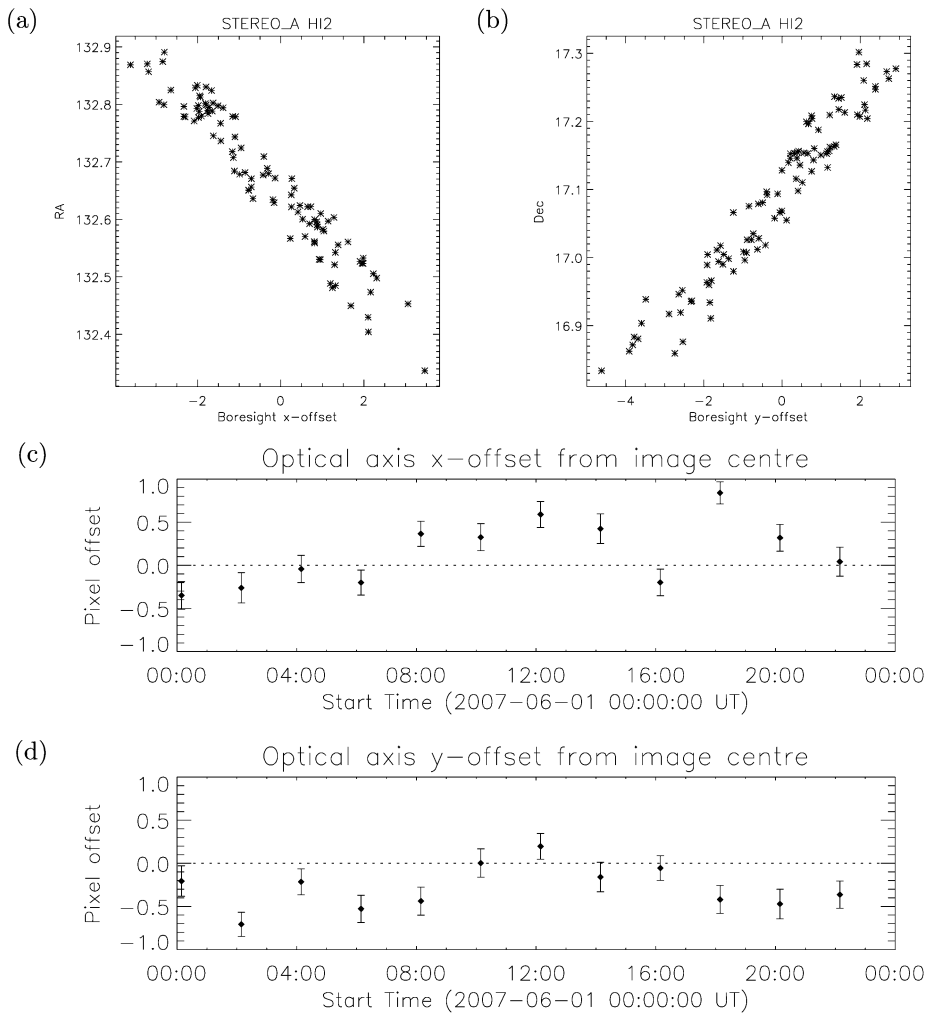


Figure 23 Plots showing the behaviour of the offset of the HI-2A instrument's optical axis. Scatter plots show (a) how the instrument RA varies with the optical axis x -offset and (b) how the instrument Dec varies with the optical axis y -offset. Only points where $R_{\text{avg}} \leq 1$ are plotted. Also shown are (c) the average optical axis x -offset and (d) the average optical axis y -offset for each 1024×1024 HI-1B image taken on 1 June 2007.

data is the same as 2 bins in the 2048×2048 data. Typically 20–80 stars are fit as part of the minimisation algorithm; however, since the HI-2 cameras have a larger field of view, it is only necessary to use catalogue stars down to 4th magnitude to achieve this, rather than 6th magnitude as used for HI-1.

The HI-2A instrument offset can be seen in Figure 22. There is evidence of fluctuations in the first three months consistent with that seen in HI-1A. The rest of the plots are noisier than their HI-1A counterparts; however, this is in part because the pixels are approximately four times larger for HI-2 than for HI-1 (in absolute terms).

Figure 23 shows the behaviour of the optical axis offset from the centre of the HI image. The scatter plots show correlation (a) between RA and x_{off} and (b) between Dec and y_{off} .

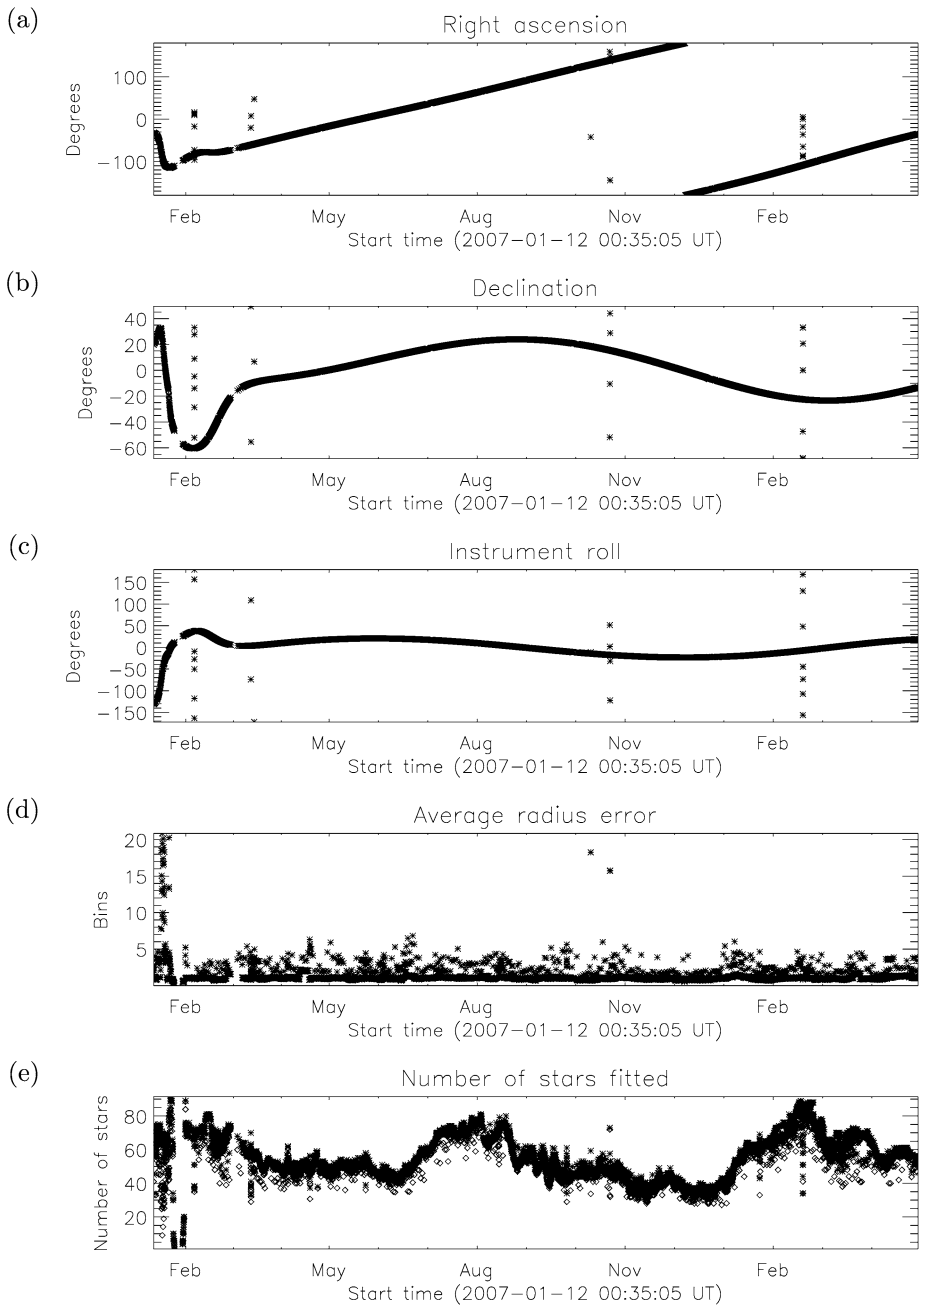


Figure 24 Plots showing the variation in pointing of the HI-2B instrument in (a) RA, (b) Dec, and (c) roll, (d) the error in the fits, and (e) the number of stars used for the fits, where μ and d take their averaged values.




Figure 25 Plots showing the variation in offset of the HI-2B instrument from the spacecraft. Dashed line shows the nominal offset when it falls within the range of the plot.

The image averages for 1 June 2007 are clustered around the zero-offset line but small error bars do not give good overlap with the zero line; however, other subpixel errors in the system that are not included in this plot may compensate for these small error bars. These plots suggest that it is reasonable to set the x - and y -offsets to zero for the minimisation calculation.

A.4. HI-2B Results

The pointing solutions for HI-2B calculated using the average μ and d values can be seen in Figure 24. The data are generally consistent, with very few outlying values. Additionally, the plots are consistent in shape with those for HI-1B, also showing the initial spacecraft maneuvering, which would be expected as both cameras are mounted on the same spacecraft. The positions of calibration rolls can be seen, and they agree with the calibration rolls for HI-1B, which is also to be expected.

The average pixel error of the fits along with the number of fitted stars can also be seen in Figures 24(d) and 24(e). The R_{avg} for HI-2B observations has not minimised as well as for other cameras owing to the PSF for this camera (see Section 3.8); however, the vast majority of optimised attitude solutions are still acceptable.

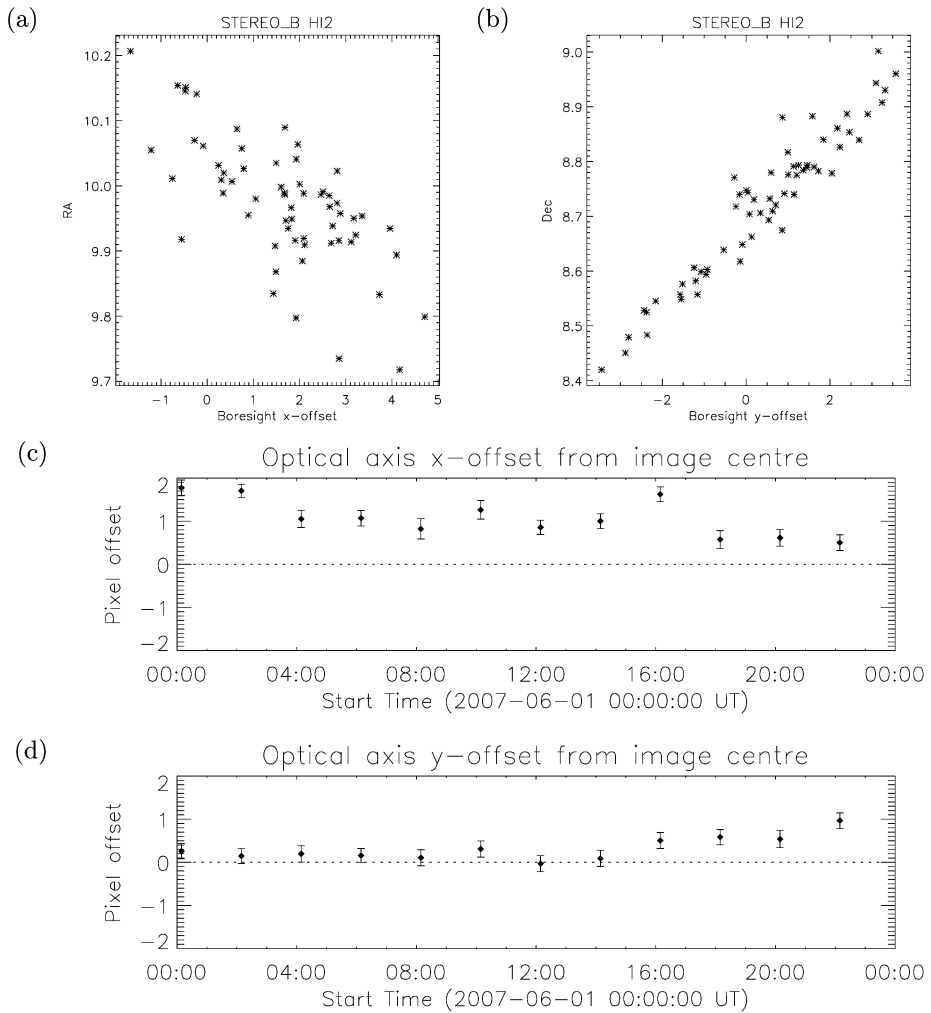


Figure 26 Plots showing the behaviour of the offset of the HI-2B instrument’s optical axis. Scatter plots show (a) how the instrument RA varies with the optical axis x -offset and (b) how the instrument Dec varies with the optical axis y -offset. Only points where $R_{avg} \leq 3$ are plotted. Also shown are (c) the average optical axis x -offset and (d) the average optical axis y -offset for each 1024×1024 HI-1B image taken on 1 June 2007.

Typically 20–80 stars are fit as part of the minimisation algorithm. Again, as this is a HI-2 camera, 4th magnitude catalogue stars have been used for the minimisation algorithm instead of 6th magnitude stars.

The HI-2B instrument offset can be seen in Figure 25. If the pointing discontinuities observed in the case of HI-1B are due to changes in alignment of the entire HI-B instrument (rather than just the HI-1B camera) then evidence for them may also be expected in the HI-2B solutions. This is not evident in Figure 25; however, a combination of the broader, asymmetric PSF and the larger pixel size may be masking the effect. Repeating the minimisation for HI-2B once the PSF has been adequately modelled may help shed light on the cause of the discontinuities in HI-1B.

Figure 26 shows the behaviour of the optical axis offset from the centre of the HI image. The scatter plots show correlation (a) between RA and x_{off} and (b) between Dec and y_{off} . Note that solutions with a larger R_{avg} are used in the calculation as the HI-2B PSF has an impact on their accuracy. The image averages for 1 June 2007 are fairly near the zero-offset line but small error bars do not give good overlap with the zero line. Again, the effect of the PSF will contribute to the accuracy here, so it is taken as being reasonable to set the x - and y -offsets to zero for the minimisation calculation. However, this will need to be revisited once the HI-2B PSF has been accounted for.

References

- Calabretta, M.R., Greisen, E.W.: 2002, Representations of celestial coordinates in FITS. *Astron. Astrophys.* **395**, 1077 – 1122.
- Davis, C.J., Harrison, R.A.: 2005, STEREO/HI from near-Earth objects to 3D comets. *Adv. Space Res.* **36**, 1524 – 1529.
- Davies, J.A., Harrison, R.A., Rouillard, A.P., Sheeley, N.R., Perry, C.H., Bewsher, D., Davis, C.J., Eyles, C.J., Crothers, S.R., Brown, D.S.: 2008, A synoptic view of solar transient evolution in the inner heliosphere using the heliospheric imagers on STEREO. *Geophys. Res. Lett.*, submitted.
- Freeland, S.L., Handy, B.N.: 1998, Data analysis with the SolarSoft system. *Solar Phys.* **182**, 497 – 500.
- Fulle, M., Leblanc, F., Harrison, R.A., Davis, C.J., Eyles, C.J., Halain, J.-P., Howard, R.A., Bockeleé-Morvan, D., Cremonese, G., Scarmato, T.: 2007, Discovery of the atomic iron tail of comet MCNaught using the heliospheric imager on STEREO. *Astrophys. J.* **661**, L93 – L96.
- Harrison, R.A., Davis, C.J., Eyles, C.J.: 2005, The STEREO heliospheric imager: how to detect CMEs in the heliosphere. *Adv. Space Res.* **36**, 1512 – 1523.
- Harrison, R.A., Davis, C.J., Eyles, C.J., Bewsher, D., Crothers, S.R., Davies, J.A., Howard, R.A., Moses, D.J., Socker, D.G., Newmark, J.S., Halain, J.-P., Defise, J.-M., Mazy, E., Rochus, P., Webb, D.F., Simnett, G.M.: 2008, First imaging of coronal mass ejections in the heliosphere viewed from outside the Sun Earth line. *Solar Phys.* **247**, 171 – 193.
- Howard, R.A., Moses, J.D., Vourlidas, A., Newmark, J.S., Socker, D.G., Plunkett, S.P., Korendyke, C.M., Cook, J.W., Hurlay, A., Davila, J.M., Thompson, W.T., St Cyr, O.C., Mentzell, E., Mehalick, K., Lemen, J.R., Wuelsel, J.P., Duncan, D.W., Tarbell, T.D., Wolfson, C.J., Moore, A., Harrison, R.A., Waltham, N.R., Lang, J., Davis, C.J., Eyles, C.J., Mapson-Menard, H., Simnett, G.M., Halain, J.P., Defise, J.M., Mazy, E., Rochus, P., Mercier, R., Ravet, M.F., Delmotte, F., Auchere, F., Delaboudiniere, J.P., Bothmer, V., Deutsch, W., Wang, D., Rich, N., Cooper, S., Stephens, V., Maahs, G., Baugh, R., McMullin, D., Carter, T.: 2008, Sun Earth connection coronal and heliospheric investigation (SECCHI). *Space Sci. Rev.* **136**, 67 – 115.
- Kaiser, M.L.: 2005, The STEREO mission: an overview. *Adv. Space Res.* **36**, 1483 – 1488.
- Rouillard, A.P., Davies, J.A., Forsyth, R.J., Davis, C.J., Harrison, R.A., Lockwood, M., Bewsher, D., Crothers, S.R., Eyles, C.J., Hapgood, M., Perry, C.H.: 2008a, First imaging of corotating interaction regions using the STEREO spacecraft. *Geophys. Res. Lett.* **35**, L10110.
- Rouillard, A.P., Davies, J.A., Forsyth, R.J., Savani, N., Sheeley, N.R., Thernisien, A., Burlaga, L.F., Zhang, T., Vourlidas, A., Howard, R.A., Wang, Y.-M., Rees, A., Anderson, B., Krimigis, T., Slavin, J., Carr, C.M., Tsang, S., Lockwood, M., Davis, C.J., Harrison, R.A., Bewsher, D., Brown, D.S., Crothers, S.R., Eyles, C.J., Hapgood, M., Perry, C.H., Whittaker, I., Jones, G.H., Coates, A.J., Grande, M., Frahm, R.A., Winningham, J.D.: 2008b, A solar storm observed from the Sun to Venus by using the STEREO, Venus Express and MESSENGER spacecraft, *J. Geophys. Res.*, submitted.
- Sheeley, N.R. Jr., Herbst, A.D., Palatchi, C.A., Wang, Y.-M., Howard, R.A., Moses, J.D., Vourlidas, A., Newmark, J.S., Socker, D.G., Plunkett, S.P., Korendyke, C.M., Burlaga, L.F., Davila, J.M., Thompson, W.T., St Cyr, O.C., Harrison, R.A., Davis, C.J., Eyles, C.J., Halain, J.P., Wang, D., Rich, N.B., Battams, K., Esfandiari, E., Stenborg, G.: 2008a, SECCHI observations of the Sun's garden-hose density spiral. *Astrophys. J.* **674**, L109 – L112.
- Sheeley, N.R. Jr., Herbst, A.D., Palatchi, C.A., Wang, Y.-M., Howard, R.A., Moses, J.D., Vourlidas, A., Newmark, J.S., Socker, D.G., Plunkett, S.P., Korendyke, C.M., Burlaga, L.F., Davila, J.M., Thompson, W.T., St Cyr, O.C., Harrison, R.A., Davis, C.J., Eyles, C.J., Halain, J.P., Wang, D., Rich, N.B., Battams, K., Esfandiari, E., Stenborg, G.: 2008b, Heliospheric images of the solar Wind at Earth. *Astrophys. J.* **675**, 853 – 862.
- Thompson, W.T.: 2006, Coordinate systems for solar image data. *Astron. Astrophys.* **449**, 791 – 803.

- Vourlidas, A., Davis, C.J., Eyles, C.J., Crothers, S.R., Harrison, R.A., Howard, R.A., Moses, D.J., Socker, D.G.: 2007, First direct observation of the interaction between a comet and a coronal mass ejection leading to a complete plasma tail disconnection. *Astrophys. J.* **668**, L79–L82.
- Watt, A.: 2000, *3D Computer Graphics*, 3rd edn., Addison-Wesley, Reading.
- Zacharias, N., Monet, D.G., Levine, S.E., Urban, S.E., Gaume, R., Wycoff, G.L.: 2004, The Naval Observatory merged astrometric dataset (NOMAD). *Bull. Am. Astron. Soc.* **36**, 1418.