

Preprint

D'Amico S., Benn M., and Jørgensen J.L.;

Pose Estimation of an Uncooperative Spacecraft from Actual Space Imagery;

International Journal of Space Science and Engineering, Vol.2, No.2, pp.171 - 189 (2014).

DOI: 10.1504/IJSPACESE.2014.060600

Pose estimation of an uncooperative spacecraft from actual space imagery

Simone D'Amico*

Department of Aeronautics and Astronautics,
Stanford University,
Room 262, William F. Durand Building, 496 Lomita Mall,
Stanford, CA 94305-4035, USA
E-mail: damicos@stanford.edu
*Corresponding author

Mathias Benn and John L. Jørgensen

Technical University of Denmark,
DTU Space, Elektrovej, Building 327,
DK-2800 Lyngby, Denmark
E-mail: mb@space.dtu.dk
E-mail: jlj@space.dtu.dk

Abstract: This paper addresses the preliminary design of a spaceborne monocular vision-based navigation system for on-orbit-servicing and formation-flying applications. The aim is to estimate the pose of a passive space resident object using its known three-dimensional model and single low-resolution two-dimensional images collected on-board the active spacecraft. In contrast to previous work, no supportive means are available on the target satellite (e.g., light emitting diodes) and no a-priori knowledge of the relative position and attitude is available (i.e., lost-in-space scenario). Three fundamental mechanisms – perceptual organisation, true perspective projection, and random sample consensus – are exploited to overcome the limitations of monocular passive optical navigation in space. The preliminary design is conducted and validated making use of actual images collected in the frame of the PRISMA mission at about 700 km altitude and 10 m inter-spacecraft separation.

Keywords: pose estimation; optical navigation; uncooperative spacecraft; flight results; PRISMA; space imagery.

Reference to this paper should be made as follows: D'Amico, S., Benn, M. and Jørgensen, J.L. (xxxx) 'Pose estimation of an uncooperative spacecraft from actual space imagery', *Int. J. Space Science and Engineering*, Vol. X, No. Y, pp.000–000.

Biographical notes: Simone D'Amico is an Assistant Professor of Aeronautics and Astronautics at Stanford University. He has been working as Researcher at the German Aerospace Center (DLR) from 2003 to 2013 in the fields of space flight dynamics, autonomous satellite navigation and control, spacecraft formation-flying, and on-orbit servicing. He gave key contributions to the design, development, and operations of formation-flying and rendezvous missions such as GRACE, TanDEM-X, and PRISMA for which he received several awards.

Mathias Benn is an Assistant Professor within the Measurement and Instrumentation division at the National Space Institute, Technical University of Denmark (DTU). He develops methods, software and hardware related to the DTU's microadvanced stellar compass and vision-based systems for closed-loop formation flying of spacecraft. He joined DTU in 2007. He holds a MSc EE and a PhD in Space Technology from DTU.

John L. Jørgensen is Professor and Head of Measurement and Instrumentation division at the National Space Institute, Technical University of Denmark. He is expert in space technology and systems such as rockets, satellites and advanced measurement systems. He has developed the world's first fully autonomous star tracker camera. His research focus is on the development of technology for future spacecraft formation flying as a condition for many ambitious space missions such as manned expeditions to Mars, extremely large and precise telescopes and Earth observation missions.

This paper is a revised and expanded version of a paper entitled 'Pose estimation of an uncooperative spacecraft from actual space imagery' presented at 5th International Conference on Spacecraft Formation Flying Missions and Technologies (SFFMT), Munich, Germany, 29–31 May 2013.

1 Introduction

Monocular vision navigation has been identified by various authors as enabling technology for present and future formation-flying and on-orbit-servicing missions [e.g., PRISMA – Sweden (Bodin et al., 2012), DEOS – Germany (Rupp et al., 2009), PROBA-3 – ESA (Landgraf and Mestreau-Garreau, 2013)]. Indeed simple instruments such as star trackers or infrared cameras can be modified for increased dynamic range to accurately navigate with respect to a target space vehicle at low cost from virtually zero to several tens of kilometres (Jørgensen et al., 2004).

This paper addresses the implementation of a novel vision system for estimating the pose (relative attitude and position) of a non-cooperative space resident object using single two-dimensional images given its three-dimensional geometric representation (model). Unlike most previous approaches (Ho and McClamroch, 1992; Philip et al., 1998), the estimation of the six translational and rotational degrees of freedom is accomplished without range measurements and without any a-priori information on the relative motion between the spacecraft. Furthermore, the development and verification of the pose estimator are based on actual images collected during the PRISMA mission by two cameras available onboard the servicer satellite at close separations (5–100 m) from the client.

Although this work is largely inspired by computer vision methods developed in the '80s for terrestrial applications (e.g., Lowe, 1987a; Fischler and Bolles, 1981; DeMenthon and Davis, 1995), a number of improvements have been introduced to cope with the deficiencies encountered when dealing with space navigation. Typically, the initialisation procedure of the pose estimator is not able to handle lost-in-space configurations without a-priori knowledge. Second space imagery is characterised by low carrier-to-noise ratio and high image contrasts, which cause false or partial detections of physical edges and the extraction of a restricted amount of feature points. Third most

artificial satellites have symmetric polyhedral shapes. Although these are easy to model, the perspective equations provide several ambiguous pose solutions in the presence of high symmetry which cannot be neglected in the vision system. Furthermore the unknown correspondences between image and model features result in a large search space for ambiguity resolution, and thus in a significant computational load.

A number of requirements for future monocular vision-based systems can be derived based on the aforementioned challenges. Ideally, a pose estimator shall rely on a minimum number of image features, minimise the search space for model matching, be robust to ambiguous pose solutions, compensate for image noise and model errors, not rely on the uncooperative satellite dynamics, and offer multi-stage solutions of increasing accuracy from lost-in-space all the way to fine navigation. In order to fulfil these objectives the vision system prototyped in this work makes use of three mechanisms which underlie six sequential functional modules (from image processing to final pose fitting). First a method of perceptual organisation is adopted to detect those feature groupings in the image which remain partially invariant over a wide range of viewpoints (i.e., orientation and scale) (Lowe, 1985b). These can be matched to corresponding structures in the object model in a probabilistic manner. The probability of accidental instances from randomly distributed segments is evaluated and the size of the matching search space is reduced by combining primitives into larger structures. The software complexity is reduced by the fact that the same perceptual organisation functions are used in the spacecraft modelling, image processing, and model matching modules. Furthermore the perceptual groupings of the spacecraft model can be pre-computed once before the mission launch or, if necessary, during the commissioning of the vision-system.

Second the true perspective projection equations are expressed as a function of the six pose state unknowns and the available measurements (McIvor, 1988). The partial derivatives are rigorously computed from the reference state to compute the modelled-image error. Each matched image segment provides three measurements, namely the coordinates of the endpoint and the angle subtended by the corresponding line. This selection of data-types is shown to be beneficial in the presence of unexpected line terminations. The linearisation of the image perspective equations is used in two stages of the estimator. During initialisation a complete linear system can be solved for the unknown corrections to the reference state through a multi-dimensional Newton-Raphson method (Ostrowski, 1966). Full convergence with pixel level accuracy is achieved with three to four iterations, provided that the initial pose is within approximately 40 deg error. The paper describes an extensive search method which is able to provide such initial pose estimate relying only on the visibility specification attached to the spacecraft model. After the extension of the initial set of matches through projection of the model onto the image, the estimation problem becomes over-determined. The newly available measurements are processed in a weighted iterative batch least-squares estimator with a-priori information.

Third a random sample consensus method is used in combination with the least squares fit to improve robustness against false model-image matches (Torr and Zisserman, 2000). In contrast to conventional smoothing techniques, this approach uses as small an initial data set as feasible and enlarges this set with consistent data when possible. The measurement/model compatibility is based on perceptual organisation, or equivalently on the probability of non-accidentalness of the considered groups of

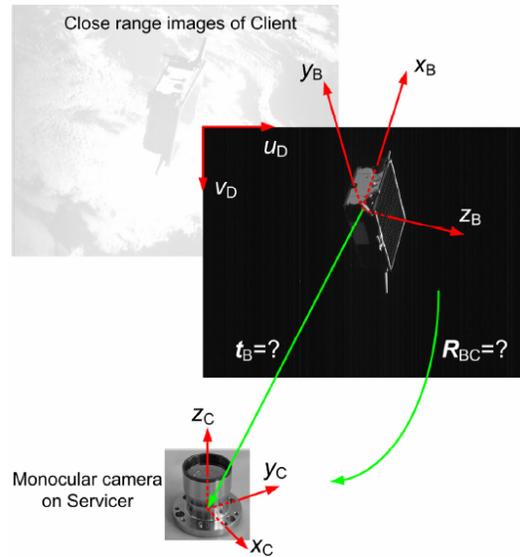
features. The quality of the estimated state is judged based on the post-fit residuals and the formal output standard deviation of the state parameters.

In addition to the detailed description of the vision-system, the paper focuses on flight results obtained from the processing of actual images taken in orbit. Reference flight dynamics products based on precision carrier-phase differential GPS (relative position) and coarse sun-magnetic attitude estimates (relative attitude) are compared with the pose estimator results (Ardaens et al., 2011; D'Amico et al., 2012). Deviations of the order of 10 cm in position and 10 deg in rotation can be inferred at 10 m separation. The paper is concluded by a discussion on the achieved robustness, the potential extensions of the vision-based approach, and the suggested improvements in terms of system architecture and navigation algorithms.

2 Problem statement

The goal of this research is to estimate the relative position, \mathbf{t}_B , and orientation, \mathbf{R}_{BC} , (pose) of a non-cooperative spacecraft (no markers or other specific supportive means) from a single two-dimensional (2D) image given its three-dimensional (3D) geometric representation (model). As depicted in Figure 1, the relative position vector is expressed in the body or structural frame (subscript B) attached to the client passive spacecraft, whereas the relative attitude defines the rotation matrix from the client body frame to the coordinate frame attached to the vision camera head unit (subscript C).

Figure 1 Coordinate frames and unknowns of pose estimation process (see online version for colours)



The close-range vision camera is embarked on the servicer spacecraft and delivers 2D low-resolution (752×580 pixels), gray-scale (8 bit per pixel), monocular (effective focal

length, $f = 20,187 \cdot 10^{-6}$ m) images of the client. Two examples are shown in Figure 1, where the image frame is indicated by the subscript D. The algorithm prototyping and testing described in this paper is based on such images collected in October 2011 in the frame of the PRISMA mission. During the considered flight experiments, the servicer (Mango) approaches the client (Tango) satellite at about 10 m separation in forced motion control based on relative GPS navigation (Bodin et al., 2012; Ardaens et al., 2011; D'Amico et al., 2012).

Each image point ρ_D can be expressed as a function of the unknown pose (\mathbf{t}_B , \mathbf{R}_{BC}) according to the following 3D-2D true perspective projection equations

$$\begin{aligned} \mathbf{r}_C &= (x_C, y_C, z_C)^T = \mathbf{R}_{BC} (\mathbf{p}_B - \mathbf{t}_B) \\ \rho_D &= (u_D, v_D) = \left(\frac{x_C}{z_C} \frac{f}{du}, \frac{y_C}{z_C} \frac{f}{dv} \right) \end{aligned} \quad (1)$$

where \mathbf{p}_B represents the corresponding point in body frame (or the matched model point), u_D and v_D are its coordinates in the image frame, du and dv are the different pixel sizes $8.6 \cdot 10^{-6}$ m and $8.3 \cdot 10^{-6}$ m respectively.

Several considerations can be done on equation (1). First of all, the relationship between image points and pose parameters is highly non-linear. Although it is straightforward to describe the projection from a 3D scene to a 2D image, the inverse problem is much more difficult and can have infinite solutions if under-constrained (Fischler and Bolles, 1981). Second, the correspondence between \mathbf{p}_B and ρ_D is unknown. Finally, the image coordinates need to be corrected for non-quadratism and lens distortion before being used in equation (1). Despite these obstacles, the pose estimation approach builds on equation (1), which constitutes a rigorous observation model and can be easily derived for the computation of partial derivatives w.r.t. the pose state.

Equation (1) suggests the main ingredients of a pose estimation system. One needs the capability of extracting client features from the available image (image processing) to obtain measurements, ρ_D . These features have to be matched to correspondent elements of a client model (model matching), \mathbf{p}_B . Finally, the unknown pose translation and rotational parameters need to be estimated based on the available measurements (pose estimation). The following sections are dedicated to each of these fundamental tasks which are highly interconnected.

3 Image and model processing

3.1 Perceptual organisation

Most model-based vision systems involve a matching process where a set of k segments extracted from the image need to be matched to a set of j segments of the model. The resulting search tree contains $(j + 1)^k$ nodes and requires an efficient strategy to reduce the computational effort. A promising technique is given by the method of perceptual grouping, which is inspired by mechanisms and constraints at the basis of the human visual system (Lowe, 1985b). Instead of considering only single image points or segments, these are combined into perceptually relevant structures known to be present in the model, and called perceptual groups. Only those perceptual groups which have the highest probability of not being accidental (i.e., not arising from randomly distributed

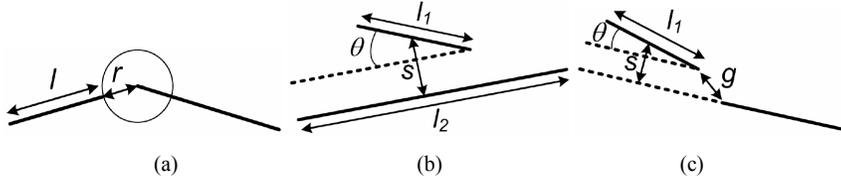
segments), are considered in the matching process. This leads to a reduction of the search space and to an efficient matching process.

The simplest perceptual groups are adopted in this research, namely segments with endpoints in close proximity, segments which are parallel and collinear segments. These primitive features are then combined into larger structures such as parallelograms and trapezoids in order to reduce the search space even more. The probability, P , of accidental instances from arbitrarily located segments can be evaluated according to the following scale independent expressions (Lowe, 1985a)

$$\begin{aligned} P_{PR} &= 2\pi r^2 / I^2 \\ P_{PA} &= 4\theta s I_2 / (\pi I_1^2) \\ P_{CO} &= 4\theta s (g + I_1) / (\pi I_1^2) \end{aligned} \quad (2)$$

where PR, PA, and CO stands for instances of proximity, parallelism, and collinearity as depicted in Figure 2 (please refer to figure for symbol explanation). The significance value for the higher-level groups is computed by multiplying the probability of non-accidentalness for each constituent.

Figure 2 (a) Geometry of proximity, (b) parallelism and (c) collinearity



Source: Adapted from Lowe (1985a)

The matching process considers only those perceptual groups which have the lowest probability of being accidental (i.e., $P \ll 1$). In addition all perceptual groups are ranked according to their significance to simplify the subsequent matching stage.

3.2 Image processing

The goal of the image processing stage is to extract line segments which correspond to true edges or at least to edges available in the spacecraft model. Unfortunately, the conditions of visual navigation in space make this task difficult and error prone. This is mainly due to the low carrier-to-noise ratio and to the high-contrast illumination which cause the detection of a small number of segments or features. Quite frequently the detected segments identify double edges or partial edges which can induce severe problems in the pose estimation scheme if not accounted for.

Despite these difficulties, here well known off-the-shelf techniques for line detection have been adopted (<http://www.opencv.org>). The approach follows three main steps:

- 1 low-pass filtering
- 2 Canny-edge detection (Canny, 1986)
- 3 Hough transform (Duda and Hart, 1972).

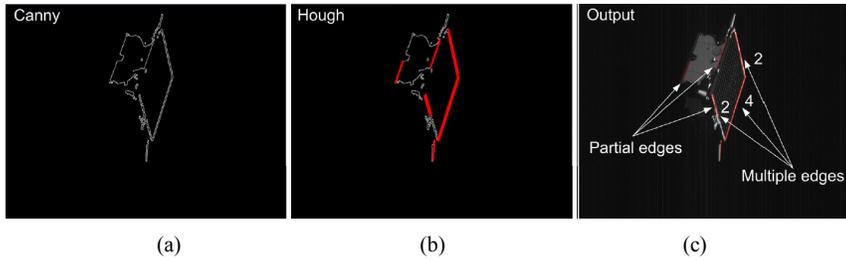
Each of these steps involve a number of process parameters to be set, such as

- 1 the kernel size
- 2 the hysteresis and Gaussian smoothing parameters
- 3 the thresholding and segment length parameters.

In general, these functions are difficult to tune for broad applicability, thus the numerical examples shown hereby can be seen as conservative test cases for the pose estimation strategy.

Figure 3 illustrates the aforementioned image processing steps applied to an image of the client spacecraft taken in orbit. Due to the illumination conditions, the Canny-edge detector can only extract features which basically reside on the two sides of the spacecraft illuminated by the Sun. The subsequent Hough transform can only identify line segments which correspond to multiple and truncated edges [see Figure 3(c)]. Unfortunately a considerable number of false edges have been detected, such as the output segments within the solar panel or on the cylindrical radio-frequency antenna (bottom of image).

Figure 3 Steps of image processing process, (a) Canny edge detector, (b) Hough transform and (c) output segments superimposed to original image (see online version for colours)



Considering the design case shown in Figure 3, the output of the image processing is a list of ten segments identified by 20 endpoints through their pixel coordinates. Perceptual groups are computed for approximate instances in 2D (i.e., $P \leq P^{\max} = 0.1$). For simplicity and to reduce ambiguities, multiple edges are removed by considering only the most perceptually significant segments (i.e., with the most meaningful connections). This leaves four image segments which are combined in three parallel and one proximity perceptual groups. For completeness, the perceptual groups output of this stage are listed in Table 1 together with their significance scorings. It is noted that a camera calibration model is applied to the coordinates of the extracted segments to correct for distortion and non-quadratism (Benn and Denver, 2012). This is essential before comparison with the spacecraft model.

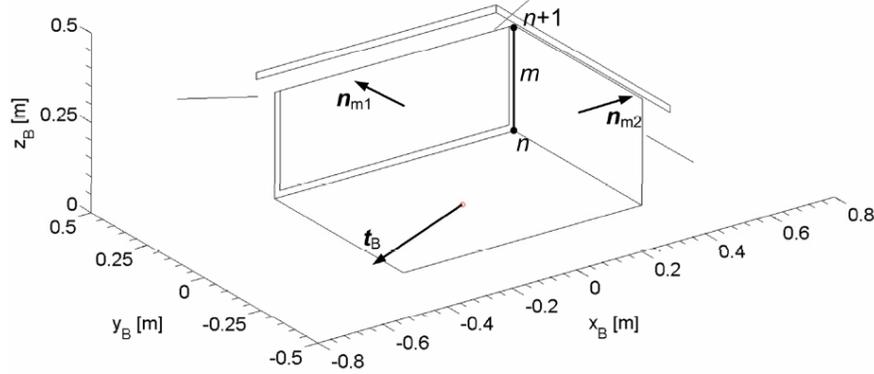
3.3 Spacecraft model

The definition of a proper spacecraft model is a fundamental step of the pose estimation strategy. On one hand, the spacecraft model has to be as minimalist as possible to reduce the system complexity and the search space for matching. On the other hand, the model must be representative to be able to find image matches among the available perceptual groups. A reasonable realism of the spacecraft model is also required to reduce the

number of ambiguities associated with the high symmetry of man-made polyhedral shapes.

Figure 4 depicts the client spacecraft model which has been traded-off in the frame of this research. Basically, it consists of two stacked symmetric and convex polyhedra – one representing the solar panel ($56 \times 75 \times 2.5$ cm), and one representing the spacecraft body ($56 \times 55 \times 30$ cm). The centres of the two polyhedra are not aligned, and three additional segments (20.4 cm) represent the radio-frequency antennae. The origin of the body frame is located at the centre of the bottom face of the spacecraft body.

Figure 4 Minimalistic 3D model of client spacecraft with visibility specification for segment m (endpoints n and $n + 1$) and possible viewpoint t_B (see online version for colours)



The spacecraft model must be well organised to boost the efficiency of the search algorithms during matching and pose estimation. To this end, an indexing scheme is introduced which facilitates the finding and manipulation of all segments, including endpoints, orientation, and lengths. In addition a visibility specification is attached to each segment. This is required at various stages of the pose estimation process, including the lost-in-space initialisation (Section 4) and the extension of the first set of matches (Section 5.2). The visibility is specified by one or two unit vectors normal to the surfaces which contain the considered segment. An example is visualised in Figure 4, where segment m can be viewed from the directions indicated by the unit vectors $\mathbf{n}_{m1} = (0, 1, 0)^t$ and $\mathbf{n}_{m2} = (1, 0, 0)^t$. For simplicity, no partial visibility has been implemented, thus a generic segment m is visible under the following algebraic condition

$$t_B \cdot \mathbf{n}_{m1} \geq 0 \parallel t_B \cdot \mathbf{n}_{m2} \geq 0 \quad (3)$$

As illustrated by Fig.4, the non-rigorous implementation of the visibility constraint makes several segments not visible even if they are partially visible. This is a drawback of the adopted approach in view of the image processing output.

As for the spacecraft model, perceptual groups can be pre-computed once before mission start. In contrast to the image processing output, here exact instances in 3D of the perceptual groups are searched for (i.e., $P = P^{\max} = 0$). It is noteworthy that the same perceptual organisation functions can be used without changes for image and model processing. Table 1 lists the results of the perceptual organisation functions applied to the design case image (Figure 3), and to the spacecraft model (Figure 4). The first four

columns show the organisation tree and the size of the available structures, whereas the last column lists the perceptual group rankings or the probability of accidentalness. By definition, this is identically zero for the model groups. The only image segments which are found to be in close proximity are also the most perceptual ones among the available groups ($P = 0.01$). It is noted that the model groups ‘proximity’ are less numerous than the groups ‘parallel’, ‘collinear’, and ‘III-shaped’, and offer the same amount of measurements (two endpoints). As a consequence, the usage of these structures can improve the computational efficiency of the search process at no cost. On the other hand, the groups ‘trapezoid’, ‘parallelogram’, and ‘Y-shaped’ offer three to four measurements and are of even reduced quantity (13 to 20). These latter’s shall be prioritised in a pose estimation process, but are difficult to extract from space imagery.

Table 1 Perceptual grouping output for design case image and model

		<i>Organisation tree</i>		<i>Size</i>	<i>Probability, P</i>
Image	Segments	Endpoints		20×2	-
		Length		10×1	
	Groups	Proximity		1×2	0.01
		Parallel		3×2	(0.03, 0.03, 0.06)
Model	Segments	Endpoints		82×3	
		Length		41×1	-
		Visibility		41×6	
	Groups	Proximity		64×2	
		Parallel		189×2	0.0
		Collinear		235×2	
		Trapezoid		13×4	
		Parallelogram		13×4	
		Y-shaped		20×3	
		III-shaped		626×3	

4 Initial pose

The estimate of the initial pose is certainly the most challenging task of the pose estimation process. Many authors assume coarse a-priori knowledge of the relative position and orientation to aid the vision navigation system (Cropp, 2001), whereas others refer to this task as the weakest and least robust in the estimation chain (McIvor, 1988). Aside the client spacecraft 3D model, here no assumption on the relative translational and rotational motion is made. The lost-in-space initial pose estimation is solved in two steps and aims at achieving a three-axis attitude error below 30–40° for the subsequent refinement process.

A first guess of the pose follows the approach suggested by Lowe (1985a). In particular the rotation matrix \mathbf{R}_{BC} is computed

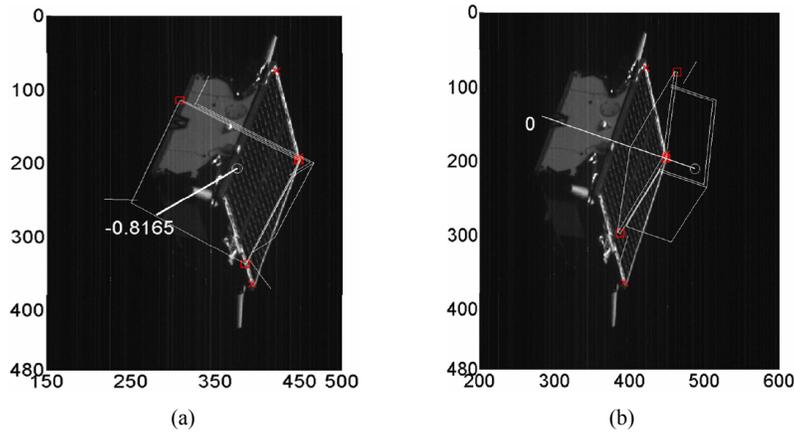
- 1 from the most likely viewing direction of the considered perceptual group
- 2 by causing one of the perceptual group model segments to be parallel to the correspondent image segment.

The most likely viewing direction is computed from the sum of the unit vectors which specify the visibility of the perceptual group. The translation t_B is obtained

- 1 as a scale factor from the segment pair of the considered perceptual group with minimum ratio model to observed line length
- 2 by aligning the endpoints of model and image segment.

Since the correspondence between observed and model segments is unknown, this initial guess has to be computed for all available pairing combinations 'model-image segment' of the considered perceptual group. This first step is completed by the computation of the initial pose error as the root-mean-square of the differences between the considered image points and the projected model points.

Figure 5 Two out of 32 initial guess pose solutions with total error below 100 pxl (rms) (see online version for colours)



Notes: Crosses indicate original measurements. Squares indicate model points.

The application of this procedure to the aforementioned design case provides 32 pose solutions with an error below 100 pixels (rms). All possible pairing combinations provided by the most highly ranked perceptual group (i.e., proximity with $P = 0.01$) have been considered (i.e., 4×128), cf. Table 1. Figure 5 illustrates just two of the available pose solutions with fit errors equal to 39.7 [Figure 5(b)] and 62.3 [Figure 5(a)] pixels. The vector normal to the projected client solar panel is also depicted with the indication of its z_C component to improve the readability of the figure. As anticipated this initial solution is affected by a very large error which prevents any ambiguity resolution or refinement through equation (1).

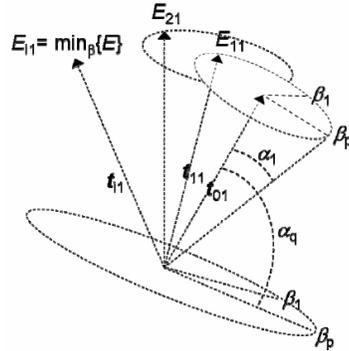
On the contrary this initial guess can be used for a more exhaustive search of the visibility domain. In particular, the most likely viewing direction t_{01} is used as a starting point for an efficient grid search whose mechanism is illustrated by Figure 6. The search

is conducted iteratively through cones defined with respect to the viewing direction. These cones are parameterised through their aperture angles, α and β . The chosen search grid and iteration parameters are listed in Table 2. At the first iteration $\alpha = \alpha_1$ is set with respect to t_{01} and the viewing direction t_{11} is searched which minimises the error E_{11} among all $\beta_k = (\beta_1, \dots, \beta_p)$. This is repeated using t_{11} as new viewing direction until the error $E_{i1} \leq E_{i-1,1}$ or until the maximum number of iterations is reached (i.e., $i = n$). The obtained error represents the minimum achievable pose error starting from α_1 as initial search cone. The same procedure can be repeated for different initial cone apertures $\alpha_j = (\alpha_1, \dots, \alpha_q)$ in order to span the visibility domain as exhaustively as possible. The obtained pose solution is affected by the minimum possible error for all considered α and β .

Table 2 Initial pose search grid and iteration parameters

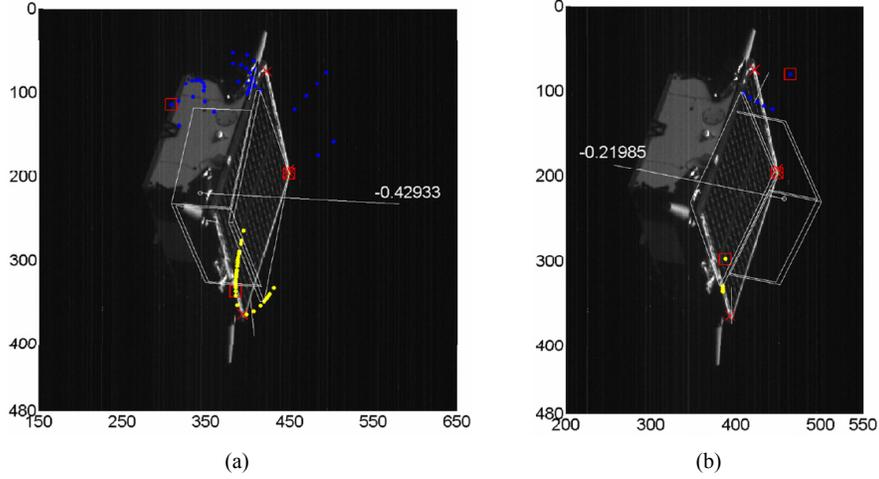
Item	Value
p	6
$\Delta\beta = \beta_k - \beta_{k-1}$	10°
q	5
$\Delta\alpha = \alpha_j - \alpha_{j-1}$	20°
n	4

Figure 6 Initial pose search algorithm



Representative results of this second step are illustrated in Figure 7. In particular, four pose solutions out of the 32 selected at the previous step show an error below 25 pixels (rms). The solutions addressed in Figure 5 are re-proposed here for better comparison. The projection errors have been reduced to 19.4 [Figure 7(a)] and 23.5 [Figure 7(b)] pixels. This has required a rotation as large as $\alpha = 100^\circ$ (i.e., $q = 5$) in the first case, whereas no rotation was necessary in the second case. Similar results are obtained when processing other images of the client spacecraft. The resulting initial pose is considered good enough to proceed with a refinement based on equation (1). In addition the search space for ambiguity resolution has been reduced to four possible pose solutions for the considered perceptual group.

Figure 7 Two out of four initial pose solutions with total error below 25 pxl (rms) (see online version for colours)



Notes: Crosses indicate original measurements. Squares indicate model points. Blue and yellow dots represent intermediate iterations of search algorithm.

5 Pose estimation

5.1 Newton-Raphson refinement

The fit error E between modelled and observed points ρ_D can be linearised about the current pose estimate as follows

$$E = \frac{\partial \rho_D}{\partial r_C} \frac{\partial r_C}{\partial t_B} \Delta t_B + \frac{\partial \rho_D}{\partial r_C} \frac{\partial r_C}{\partial \phi_{BC}} \Delta \phi_{BC} \quad (4)$$

where the partial derivatives with respect to the relative position t_B and to the Euler angles ϕ_{BC} (rotation matrix R_{BC}) can be simply computed from equation (1) as

$$\frac{\partial \rho_D}{\partial r_C} = \begin{pmatrix} f_u/z & 0 & -f_u x/z^2 \\ 0 & f_v/z & -f_v y/z^2 \end{pmatrix}; \quad \frac{\partial r_C}{\partial t_B} = -R_{BC}; \quad \frac{\partial r_C}{\partial \phi_{BC}} = \frac{\partial R_{BC}}{\partial \phi_{BC}}. \quad (5)$$

Each image point provides two error equations [i.e., equation (4)] in six unknowns which express small variations $\Delta(\cdot)$ with respect to the reference state. As a consequence a complete linear system of six equations and six unknowns can be derived from just three point correspondences. This is the amount of observables given by a proximity perceptual group composed of two segments. The unknown corrections to the reference state can be computed through a multi-dimensional Newton-Raphson method (Ostrowski, 1966), provided that the initial pose is affected by an error below about 40° . As shown in Figure 3, using only point correspondences is weak in the presence of partial edges. In order to cope with uncertain line terminations, the angle θ subtended by the considered

segment can be used as observable as well. Fortunately the partial derivative of θ can be expressed as a linear combination of the partial derivatives of u and v as follows

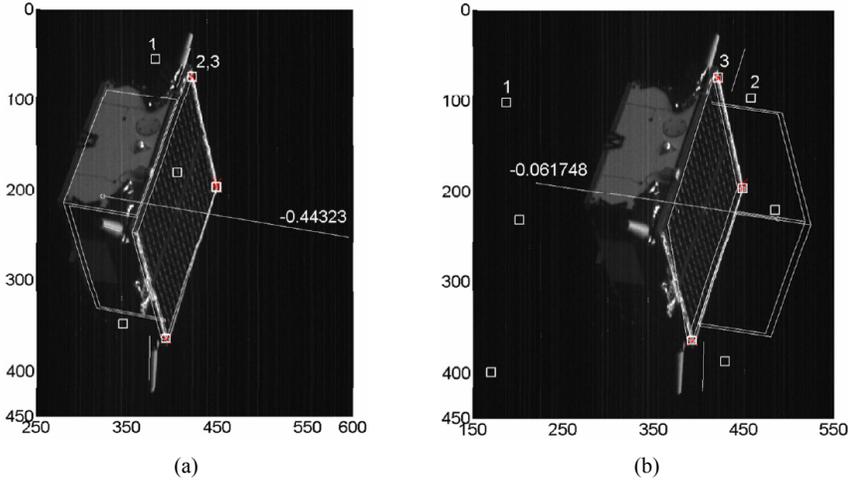
$$\frac{\partial \theta}{\partial r_c} = \left(\frac{\partial v_2}{\partial r_c} - \frac{\partial v_1}{\partial r_c} \right) \frac{1}{I \cos \theta} = \left(\frac{\partial u_2}{\partial r_c} - \frac{\partial u_1}{\partial r_c} \right) \frac{-1}{I \sin \theta} \quad (6)$$

where 1 and 2 represent the segment start- and end-points. Equation (6) is derived from the normal form of the equation of a straight line parameterised through its inclination angle θ and its distance from the origin. Independent from the choice of the observables, at each iteration the Newton-Raphson method solves for a vector of six corrections using equation (4)

$$\Delta \mathbf{x} = (\Delta t_{Bx} \quad \Delta t_{By} \quad \Delta t_{Bz} \quad \Delta \phi_{BCx} \quad \Delta \phi_{BCy} \quad \Delta \phi_{BCz})^t \quad (7)$$

and given six computed model-image error components E_i . After each iteration, $\Delta \mathbf{x}$ shrinks by about one order of magnitude, and no more than four iterations are needed for pixel level accuracy. The powerfulness of this approach is clearly shown by Figure 8. All four initial pose solutions available from the previous step can be made accurate to below 2.5 pixels (rms). The price to pay is the inversion of a 6×6 Jacobian [composed from equation (5)] at each iteration. Here only point correspondences have been used.

Figure 8 Solutions of Figures 5 and 7 refined to below 2.5 pxl error (rms) (see online version for colours)



Notes: Crosses indicate original measurements. Squares indicate model points upon successive iterations (numbered from 1 to 3).

5.2 Model matching

Although only based on an initial set of hypothesised matches, the initial pose estimate and its further refinement through the projection equations can be used to predict the locations of other model features in the image and extend the match. The selection of matches in the image which are consistent with predictions from the available pose

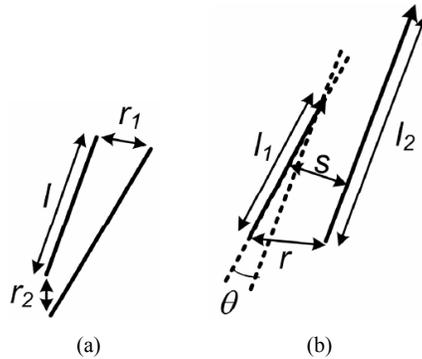
follows a probabilistic approach based on the formalism of perceptual organisation (see Section 3.1). The same perceptual grouping functions first applied to the 3D model (Section 3.3) and to the 2D image (Section 3.2), can be used to process a new 2D object resulting from the superposition of the detected image segments and the projected 3D model.

Structures equivalent to perceptual groups are formed by comparing the projected segments of the model which are visible with the image segments. Each potential match is assigned a probability of non-accidentalness based on simultaneous instances of proximity and parallelism. In particular, two segments are considered as correct matches if one of the following applies:

- 1 both endpoints are in proximity [see Figure 9(a)]
- 2 one endpoint is in proximity and the segments are parallel with overlapping legs [see Figure 9(b)].

The goodness of match is given by the product of the probabilities of proximity and parallelism as computed from equation (2).

Figure 9 Criteria used for proper model-image matching

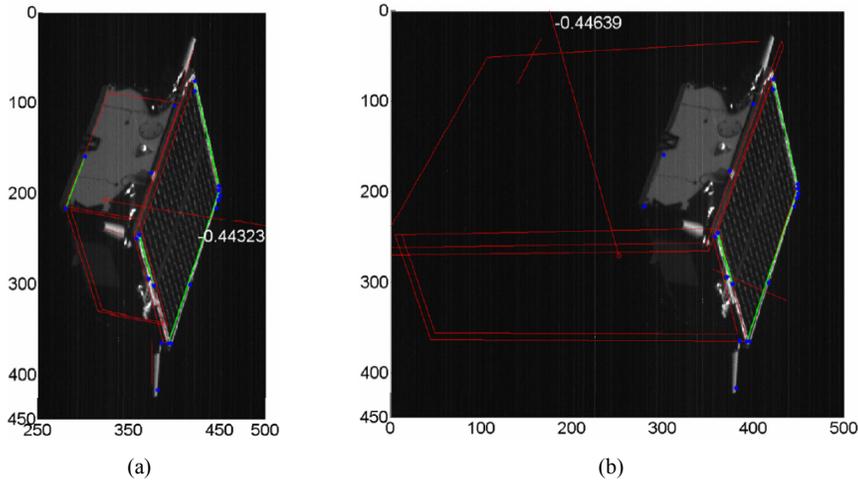


The application of the method described above to the four pose solutions available from the previous step gives the possibility to finally resolve the remaining ambiguity. Only two preliminary solutions give new matches as shown in Figure 10. The solution with the most highly ranked matches is considered for the final stage [Figure 10(a)], being the one which best fits the image in a probabilistic sense.

It is noted that the alternative solution [Figure 10(b)] gets high fit scores although obviously wrong to a human eye. This is due to the very limited amount of segments identified by the image processing. In order to increase the robustness of this ambiguity resolution, an extra image processing stage could be considered which aims at finding specific features of the model in reduced areas of the image according to the available pose solution. In the presence of more image segments, instead of using all gained matches at once in the least-squares fit (described in the next section), one could consider to add these new matches to the best fit solution gradually. By the time a sufficient number of reliable matches have been found, the pose estimation is accurate enough to choose between the few ambiguous alternatives. The set of matches is repeatedly extended in this way until no more can be found. This iterative matching procedure has

the appealing property of using the easy cases to provide better pose estimates to disambiguate the more difficult situations.

Figure 10 Two solutions which provide new matches (green segments) when projected onto the image plane (see online version for colours)



Note: Blue dots indicate the original measurements.

5.3 Least-squares fit

After the matching step, more correspondences between model and image are available than strictly necessary (i.e., ≥ 2 segments). The pose estimation problem is now over-determined with more equations (measurements) than unknowns (pose state). This can be used to compensate for measurements and model errors or to estimate extra parameters (e.g., camera parameters). All available measurements are processed in a weighted iterative batch least-squares estimator with a-priori information (Montenbruck and Gill, 2001). For each matched segment, three measurements are derived, namely the coordinates, (u, v) , of the endpoint in closest proximity and the angle, θ , subtended by the line segment with the u -axis. The sensitivity matrix is computed using the same equations as for the Newton-Raphson method [i.e., equations (4)–(6)]. The a-priori state is inherited from the initial pose estimate (refined through Newton-Raphson). The a-priori state covariance matrix is diagonal with 1 m and 10° standard deviation for the components of \mathbf{t}_B and ϕ_{BC} respectively. A process noise of 10 cm and 1° standard deviation is injected to keep the filter recipient to measurements at each further iteration. The measurement noise is assumed Gaussian and uncorrelated with standard deviations (du, dv) which amount to the different pixel size for (u, v) and du/l for θ . The quality of the estimated state is judged based on the total rms of the post-fit residuals and the formal output standard deviation of the state parameters. After three to four iterations no further improvement of the obtained solution is observable.

For the solution under consideration, a total of four segments are used providing 12 measurements. The post-fit residuals amount to 0.57 and 1.12 pixels for (u, v) , and 0.33° for θ , whereas the formal state standard deviation is (3.9, 4.7, 3.3) cm for \mathbf{t}_B , and

$(0.40 \ 0.20 \ 0.32)^\circ$ for ϕ_{BC} after four iterations. A closer look at the post-fit residuals shows that the error decreases by 50% in u and θ , whereas it remains constant in v (from first to last iteration). Although the error trend can be largely improved by correcting for the offsets which affect most detected edges (see Figure 3), it was not possible to remove the bias of 1 pixel in the available v -axis measurements. It is believed that such an error could be absorbed through the identification, matching, and processing of image segments mostly aligned with the u -axis. The availability of more matched segments could greatly improve the standard deviation of the pose error.

In order to get an independent assessment of the pose estimation accuracy, flight dynamics products from the PRISMA mission have been used for comparison, namely the post-facto on-ground precise relative orbit determination based on GPS (accurate to about 2 cm rms) (Ardaens et al., 2011), and the on-board coarse client attitude estimate from sun-sensors and magnetometers (accurate to about 3° rms). As compared with these operational products, the pose estimate is affected by errors of $(14.0 \ -12.7 \ -6.6)$ cm and $(3.8 \ -10.6 \ 2.9)^\circ$ along and around the camera axis respectively. As expected from the post-fit residuals the largest error components reflect an uncertainty due to rotations around the u -axis. This is due to the particular viewing geometry where the client spacecraft is mostly aligned with the v -axis. The obtained results are very promising since the pose estimation error lies within the typical error budget of the applied reference products. It is also noted that the error computation is affected by uncertainties in the location and orientation of the camera. Furthermore other authors have obtained similar errors at about 10% of the inter-spacecraft separation in more optimistic conditions based on pure software simulations (Cropp, 2001).

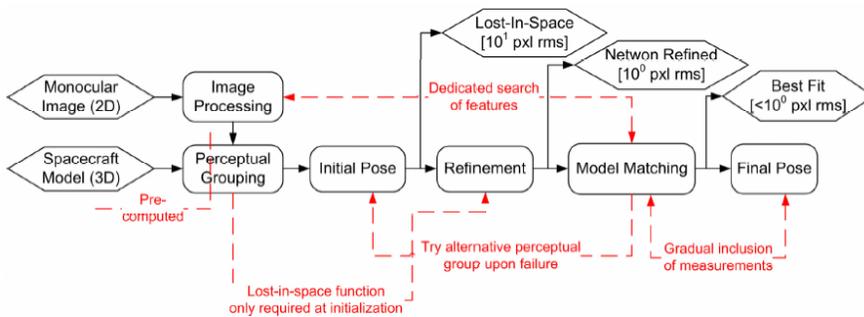
6 Discussion and conclusions

This paper has described the foundations of a spaceborne model-based pose estimation methodology whose overall layout is depicted in Figure 11. The detailed description of the applied algorithms has been accompanied by a representative design case to show the functional applicability and the accuracy potential of the developed approach. The architecture of the vision-based navigation system comprises six main functions (cf., Figure 11), namely image processing (i.e., Canny edge detector and Hough transform, Section 3.2), perceptual grouping (applied to model, image and during matching process, Section 3.1), initial pose (lost-in-space, Section 4), Newton-Raphson refinement (based on the true perspective projection equations, Section 5.1), model matching (based on the projection of the visible model onto the image plane, Section 5.2), and the final pose (least-squares fit, Section 5.3). The resulting multi-stage estimator is capable of delivering intermediate pose solutions of increasing accuracy (post-fit rms residuals from 10^1 to below 10^0 pixels) and decreasing ambiguity level (possible number of pose solutions from 10^2 to 10^0) throughout the process. Comparisons with independent flight dynamics operational products have shown pose accuracies at the level of 10 cm and 10° .

Despite the promising results, several aspects of the navigation system shall be further analysed and improved. The major limiting factor for accuracy and reliability has been identified in the image processing module. Future work will explore interest point detectors and descriptors such as the speeded up robust features (SURF) algorithm (Bay et al., 2008) or the scale-invariant feature transform (SIFT) of Lowe (1999). Aside its

overall improvement, efforts shall be put into a dedicated image processing step after the matching process to identify spacecraft features in limited portions of the image (where they are supposed to be). In order to gain new relevant features, the observation model based on line-to-line correspondences (Section 5.1) could be extended to match arbitrary curves rather than only straight line segments. In particular, silhouette-based approaches as adopted in the orbital express mission might be considered (Leinz et al., 2008). In addition the processing and comparison of subsequent images could be exploited. In this case, the initial pose estimate does not need to be repeated, provided that the client spacecraft relative attitude does not change by more than 30–40° between two images. The interplay between model matching and final pose estimation is also considered a major area of research. The gradual inclusion of new measurements and the data editing process need to be studied further. Biased edges have been found to largely affect the performance of the least-squares fit. It has to be understood how such uncertainties in the measurements or model errors could be properly incorporated in the filter scheme. Finally, in addition to real space imagery, simulated test image sequences will be added to the evaluation of the proposed method for future evaluations.

Figure 11 Architecture of pose estimation method, including data flow (arrows), inputs/outputs (hexagonal blocks), and key functions (rounded rectangles) (see online version for colours)



Acknowledgements

This work was done in the frame of the sabbatical program (Forschungssemester) of the German Aerospace Center (DLR) while the first author was affiliated with DLR. The authors would like to thank Jan Kleberg (DTU) for the provision of the Tango spacecraft 3D model. Space images used in this work have been collected during a GPS-based flight experiment defined by GMV/ESA and operated by OHB Sweden in the frame of the PRISMA extended mission.

References

- Ardaens, J.-S., D'Amico, S. and Montenbruck, O. (2011) 'Final commissioning of the PRISMA GPS navigation system', *22nd International Symposium on Spaceflight Dynamics (ISSFD)*, 28 February–4 March, Sao Jose dos Campos, Brazil.
- Bay, H., Ess, A., Tuytelaars, T. and Van Gool, L. (2008) 'SURF: speeded up robust features', *Computer Vision and Image Understanding (CVIU)*, Vol. 110, No. 3, pp.346–359.
- Benn, M. and Denver, T. (2012) *Advanced Stellar Compass – Lens Distortion Correction Method for the μ ASC*, 19 March, ASC-DTU-MA-3013, No. 1.0.
- Bodin, P., Noteborn, R., Larsson, R., Karlsson, T., D'Amico, S., Ardaens, J.-S., Delpéch, M. and Berges, J.-C. (2012) 'Prisma formation flying demonstrator: overview and conclusions from the nominal mission', *AAS 12-072: 35th Annual AAS Guidance and Control Conference*, 3–8 February 2012, Breckenridge, Colorado.
- Canny, J. (1986) 'A computational approach to edge detection', *IEEE Trans Pattern Recognition and Machine Intelligence*, November, Vol. PAMI-8, No. 6, pp.679–698.
- Cropp, A. (2001) *Pose Estimation and Relative Orbit Determination of a Nearby Target Microsatellite using Passive Imagery*, PhD Thesis, University of Surrey, UK.
- D'Amico, S., Ardaens, J.-S. and Larsson, R. (2012) 'Spaceborne autonomous formation-flying experiment on the PRISMA mission', *Journal of Guidance, Control, and Dynamics*, May–June, 0731–5090, Vol. 35, No. 3, pp.834–850, doi: 10.2514/1.55638.
- DeMenthon, D.F. and Davis, L.S. (1995) 'Model-based object pose in 25 lines of code', *International Journal of Computer Vision*, June, Vol. 15, Nos. 1–2, pp.123–141.
- Duda, R.O. and Hart, P.E. (1972) 'Use of the Hough transformation to detect lines and curves in pictures', *Communication of the ACM*, January, Vol. 15, No. 1.
- Fischler, M.A. and Bolles, R.C. (1981) 'Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography', *Communications of the ACM*, June, Vol. 24, No. 6, pp.381–395.
- Ho, C.J. and McClamroch, N.H. (1992) 'Autonomous spacecraft docking using a computer vision system', *Proc. 31st Conf. on Decision and Control*, December, Tucson, Arizona, pp.645–650. <http://www.opencv.org> (accessed December 2012).
- Jørgensen, J.L., Denver, T. and Jørgensen, P.S. (2004) 'Using an autonomous star tracker as formation flying sensor', *4th Symposium on Small Satellites Systems and Services*, 20–24 September, La Rochelle, France.
- Landgraf, M. and Mestreau-Garreau, A. (2013) 'Formation flying and mission design for Proba-3', *Acta Astronautica*, Vol. 82, No. 1, pp.137–145, January.
- Leinz, M.R. et al. (2008) 'Orbital express autonomous rendezvous and capture sensor system (ARCSS) flight test results', *Proc. SPIE 6958, Sensors and Systems for Space Applications II*, 69580A, 15 April, doi:10.1117/12.779595.
- Lowe, D.G. (1985a) 'Three dimensional object recognition from single two dimensional images', *Artificial Intelligence*, March, Vol. 31, No. 3, pp.355–395.
- Lowe, D.G. (1985b) *Perceptual Organization and Visual Recognition*, Kluwer Academic Publishers, MA, USA.
- Lowe, D.G. (1999) 'Object recognition from local scale-invariant features', *Proceedings of the International Conference on Computer Vision 2*, pp.1150–1157, doi:10.1109/ICCV.1999.790410.
- McIvor, A.M. (1988) 'An analysis of Lowe's model-based vision system', *Proc. Fourth Alvey Vision Conference*, August, Univ. of Manchester, pp.73–78.
- Montenbruck, O. and Gill, E. (2001) *Satellites Orbits – Models, Methods, and Applications*, Springer Verlag, Heidelberg, Germany.
- Ostrowski, A.M. (1966) *Solutions of Equations and Systems of Equations*, 2nd ed., Academic Press, New York.

- Philip, N.K., Kumar, E.K. and Ananthasayanam, M.R. (1998) 'Sliding observer for a robust relative position and attitude estimation during the final phase of an autonomous docking mission', *Proc. 49th Int'l. Astronautical Congress*, October, Melbourne, Australia.
- Rupp, T., Boge, T., Kiehling, R. and Sellmaier, F. (2009) 'Flight dynamics challenges of the german on-orbit servicing mission DEOS', *21st International Symposium on Space Flight Dynamics (ISSFD)*, 28 September–2 October 2009, Toulouse, France.
- Torr, P. and Zisserman, A. (2000) 'MLE-SAC: a new robust estimator with application to estimating image geometry', *Computer Vision and Image Understanding*, Vol. 78, No. 1, pp.138–156.