

Model-based Visual Tracking of Orbiting Satellites Using Edges

Manolis Lourakis and Xenophon Zabulis

Institute of Computer Science, Foundation for Research and Technology - Hellas
N. Plastira 100, Vassilika Vouton, GR 700 13, Heraklion, Greece

Abstract—Estimating the pose of orbiting satellites is a key prerequisite for supporting autonomous proximity operations in space. This work presents a monocular 3D tracking algorithm that tracks edges with the aid of an arbitrary 3D mesh model assumed to capture a satellite’s shape. The proposed tracker propagates a pose hypothesis between successive frames, using it first to render a depth image and then refining it according to partial matches established between depth and intensity edges. Edge matching relies on fast, local 1D searches along the depth gradient direction. The tracker does not require any pre-processing of the 3D model nor does it make any assumptions regarding its characteristics, as is often the case for other approaches. It is also robust to parts of the tracked satellite being out of view, occluded, shadowed or visually undetected. Experimental results evaluating the accuracy of the tracker and comparing it with established techniques are also included.

I. INTRODUCTION

Autonomous proximity operations in space refer to the ability of a robotic satellite to conduct close maneuvers to other orbiting satellites with the purpose of rendezvous, inspection, servicing or docking [1]. A critical enabling technology for such missions concerns automatic techniques for estimating the relative pose between a service and a target spacecraft [2].

Existing pose tracking techniques are characterized by the type of sensor they employ, their operating range and by whether the target spacecraft is cooperative, i.e. carries special markers that are installed on it with the purpose of facilitating pose estimation by the service spacecraft. Owing to their compact size, simple hardware and low power consumption, cameras are particularly attractive as sensors for estimating pose in space. Visual tracking in such a setting has to overcome challenges such as rapidly changing and often inadequate illumination, hard shadows, tumbling targets, lack of strong texture, occasional background clutter, slow radiation hardened hardware, etc.

This paper is concerned with the problem of using a single camera to track the pose of an un-cooperative orbiting spacecraft at medium and close ranges, i.e. at distances below 100 m. It proposes a model-based, monocular tracking technique which relies on edges to yield accurate pose estimates with moderate computational overhead. The rest of the paper is organized as follows. Section II provides an overview of related work. The proposed tracking algorithm is presented in Section III and is compared with established approaches in Section IV. The paper concludes in Section V.

Work partially supported by the European Space Agency under the HIPNOS activity, ref. 4000117700/16/NL/LF. Image sequences used for evaluation are courtesy of GMV S.A.

II. PREVIOUS WORK

Edges are moderately robust against noise and illumination or viewpoint changes. They are also accurately and rapidly localized in images and can be extracted even from weakly textured objects, to which techniques relying on photometric local patch detectors and descriptors such as [3] cannot be applied. Owing to these properties, edges have been used as the primary type of feature by several object tracking algorithms.

For instance, Lowe [4], [5] tracks an object by extracting line segments from its image contours and fitting them to a known object model. The RAPiD tracker proposed by Harris [6] operates in a reverse direction by first projecting the model using an approximate pose and then matching it with image edges. RAPiD relies on the assumption that the difference between the actual pose and its predicted estimate is small. Therefore, data association can be efficiently established using 1D local search for an image edge along the direction that is perpendicular to a predicted edge. To keep the computational complexity low, perpendicular matching is limited to a sparse set of predetermined control points. Linearization about the current pose estimate allows each pair of orthogonally matched points to yield a linear constraint on the six parameters defining the incremental change in object pose. RAPiD is very efficient and was historically the first 3D tracker to successfully run at high frame rates on general purpose hardware. As a result, its paradigm of using a local search around a prior pose has been retained in subsequent trackers. Despite its effectiveness, however, the basic RAPiD algorithm lacks robustness to mismatches and occlusions and requires that control points and their visibility are provided externally.

Several authors have proposed improvements to the basic RAPiD algorithm. Armstrong and Zisserman [7], for example, address robustness by grouping control points using geometric primitives such as lines and conics and use RANSAC to identify and discard incorrect edge matches. Drummond and Cipolla [8] use a Lie group formalism to represent the linearized relationship between image motion and pose parameters. They define control points on object lines, the visibility of which is determined at runtime with binary space partition trees. Robustness is also attained by employing an M-estimator computed with iteratively reweighted least squares (IRLS) to estimate the pose parameters. Comport et al. [9] treat pose computation as the dual problem of 2D visual servoing and track points normal to the projections of

object lines. M-estimation is again used to obtain robustness.

A common thread in the works of [7], [8], [9] is that they assume simplified object models, in which all modeled edges must give rise to visible image edges that are sampled for defining the control points. Thus, models suitable for tracking must retain only the most prominent edges and should be primarily comprised of line segments. Such requirements are typically not met by CAD models represented by densely tessellated polygonal 3D meshes (cf. Fig. 2), clearly limiting the flexibility of the aforementioned tracking algorithms. To alleviate this, more recent approaches leverage the processing capabilities of modern GPUs in order to dynamically identify the edges of a model that are visible from a particular viewpoint. For example, Reinke et al. [10] propose a technique for hidden line removal which relies on rendering a CAD model to automatically extract visible edges that are used for tracking. Additionally, they recommend a random distance sampling strategy for defining the control points on visible edges. Petit et al. [11] also rely on rendering an object model to determine edge visibility but avoid any model line processing by extracting edges from discontinuities of the rendered depth image.

A shortcoming of purely edge-based methods is that unrelated edges might locally look very similar and thus can give rise to erroneous edge correspondences that will cause tracking to fail. To deal with edges' lack of distinctiveness, certain works suggest maintaining multiple pose hypotheses, e.g. [12], [13], [14], [15]. In particular, as [15] will be latter used for comparative evaluation, it is noted that it performs tracking in a particle filtering framework, using chamfer matching to form pose hypotheses and initialize particles that are subsequently evolved with standard edge-based tracking. Aiming to further increase robustness and reduce drift, techniques such as [16], [17], [18] which combine edge and point features have also been proposed. These define keyframes that are used for anchoring with the aid of point features.

This work puts forward a RAPiD-like tracking algorithm that can accommodate any triangle mesh model without pre-processing or manual intervention for determining the control points. This is achieved by using the object model in combination with rasterization rendering to produce a depth image. Rendering automatically handles self-occlusions, thus permitting control points to be defined on edges extracted from the depth image. The 6D pose is estimated by maintaining a single hypothesis which is evolved from frame to frame using robust regression techniques. Our approach is thoroughly described in the next section.

III. PROPOSED METHOD

A triangle mesh model for the tracked object and its pose in the initial frame are assumed to be externally provided. At each subsequent frame, an approximate object pose, which might simply be the one estimated for the previous frame, is also assumed to be available. For each incoming image, Canny's algorithm [19] detects its strongest edges. Independently, the object model is projected on the

image at the approximate pose, using rasterization rendering. Rendering relies on the approximate pose to yield a depth image (Sec. III-A). Edges are then detected in the depth image and used to determine the control points by local, 1D searches that start at depth edge pixels and extend in directions parallel to the depth gradients. Specifically, a control point is defined for every match established between depth and intensity edges (Sec. III-B). Each control point provides a linear constraint on the incremental change in pose (Sec. III-C). The set of all such linear constraints is used in a robust regression framework which estimates pose after removing the influence of outliers (Sec. III-D). This pose is used to repeat the pose estimation process for a new intensity frame.

The contributions of the method are threefold. First, control points are defined automatically by rendering unrestricted mesh models, without any manual pre-processing, and partially matching edges of different modalities. Second, outliers are handled with a combination of robust regression techniques that do not call for the definition of arbitrary outlier thresholds. Third, perpendicular matching is made more robust by maintaining multiple candidate matches for each control point. More details are provided in the following subsections.

A. Rendering

Given a triangle mesh model of an object and a camera pose, depth rendering simulates an image whose pixel values are distances rather than intensities. More specifically, every pixel in the rendered depth image contains the distance to the nearest point on the model's surface that projects on the pixel in question. Depth images are rendered with the aid of rasterization rendering [20]. Rasterization involves only geometry (i.e., no texturing/shading) and has been entirely implemented in software. Nevertheless, rendering can easily be accelerated with a GPU implementation.

B. Selection of Control Points

Most of the existing edge-based approaches, e.g. [7], [8], [9], require that the object to be tracked is modeled with a simple wireframe that consists of a small number of straight edges. In effect, the object's image contours should be described by piecewise linear segments. Control points are defined by sampling either image [7], [8] or model [9] edges. While such a choice ensures computational efficiency, it imposes strong constraints on acceptable models, thus limiting applicability. This is because detailed CAD object models that might be available need to be manually redesigned so that they are made suitable for tracking by retaining only their most salient edges. Furthermore, object models often include curved parts such as spheres or cylinders, which are not faithfully represented with straight edges. To deal with arbitrary 3D meshes, [18], [10] use wireframe rendering to determine which model edges are visible from a particular viewpoint, whereas [11] uses full depth rendering of the model and finds edges with the Laplacian operator as discontinuities in the depth image. In both cases, rendering

requires a predicted object pose, which might simply be the object pose in the previous frame. In this work, we follow an approach similar to [11], but employ a more reliable alternative to detecting depth edges.

The Laplacian edge detector looks for zero crossings of the depth's second derivative, i.e. pixels where the Laplacian changes sign. As this produces many spurious edges, in practice edges are detected at pixels which in addition to being adjacent to zero crossings, have absolute Laplacian magnitudes differing by more than a threshold. However, absolute thresholding makes the detection of depth edges sensitive to the choice of threshold.

In this work, we have employed [21], which primarily detects occluding (i.e., jump) edges. The algorithm operates in two passes, examining depth image pixels along rows and then columns. At each pass, it detects depth edges by comparing the difference in the depth values of neighboring pixels against a proportional threshold. The edge detector of [21] was further adapted to detect crease edges based on the magnitude of the contained angle of the normal vectors corresponding to neighboring depth pixels. More specifically, a crease edge is detected whenever the absolute value of the dot product between two neighboring normal vectors is less than a threshold. It is noted that this calculation is independent of the coordinate system used to express the normal vectors, which are estimated once when an object model is loaded. The jump and crease edges resulting from the two applications of [21] are then merged into a single depth edge map. Control points are defined on the detected depth edge pixels, after matching them with intensity edge pixels. This process makes no assumptions about the shape of the object's contour and is explained next.

Given two edge maps and their corresponding gradient orientations quantized in eight 45° wide bins, edge matching concerns the establishment of perpendicular correspondences between edge pixels. This is achieved by examining each edge pixel in the source (i.e., depth) edge map and moving along the gradient direction in the target (i.e., intensity) edge map, until either an edge pixel is found or a maximum distance from the starting pixel has been traced (see also Fig.1). To declare a match, the edge pixel found in the target map has to have a gradient orientation compatible with that of the source pixel (cf. edge polarity [6]). In accordance with the aperture problem, this procedure determines a partial rather than a full match, since local analysis can only provide the displacement component perpendicular to the edge. The search for a perpendicular match along the edge normal has linear rather than quadratic complexity, this being a crucial enabling factor for efficient performance. The search for corresponding edge pixels has to be performed in both orientations, as it is not possible to know in advance which side of the source edge the target edge has moved. In the case that matching candidates are found for both orientations, the one closest to the source pixel is retained. In all cases, the visited target pixels are determined with Bresenham's line drawing algorithm which involves integer coordinates only.

As described up to this point, edge matching seeks for

the closest edge pixel in the target map along the gradient direction, similarly to previous edge-based trackers like [6], [7], [8], [9]. This can be improved by considering multiple candidate matches for each control point, as follows. During matching, all edge pixels from the target map along the gradient and up to a maximum distance are retained. A pose estimate is first computed by relying on the closest candidate match for each control point. The control points are then backprojected to the model using the depth image and then reprojected on the image with the estimated pose. The candidate match closest to each reprojected control point is then found and pose estimation is repeated, using the revised matches. After a new pose has been estimated, the process can be repeated. In practice, we have observed that one or two iterations are sufficient for the assignment of matches to stabilize.

C. Pose Constraints

This section relies on findings from [6] and the related description in [22]. Let $\mathcal{SE}(3)$ denote the special Euclidean group comprised of the six-parameter family of proper rotations and translations in the 3D Cartesian space. Assume an object coordinate system with its axes aligned with those of the camera coordinate system and its origin at $\mathbf{T} = (T_x, T_y, T_z)$ in camera coordinates. A control point's camera coordinates \mathbf{M} and its object coordinates $\mathbf{P} = (P_x, P_y, P_z)$ are related with $\mathbf{M} = \mathbf{T} + \mathbf{P}$. Let \mathbf{m} be this point's normalized image projection.¹ Note that the camera coordinates \mathbf{M} can be determined by backprojecting \mathbf{m} with the aid of the rendered depth image. Assume now that the object moves with a rigid motion $\xi \in \mathcal{SE}(3)$, consisting of an incremental rotation \mathbf{R} and translation \mathbf{t} which express the object's pose in the next camera frame with respect to the current one. The control point's new camera coordinates are thus given by

$$\mathbf{M}' = \mathbf{R}\mathbf{P} + \mathbf{T} + \mathbf{t}. \quad (1)$$

Assuming that motion ξ is infinitesimal, \mathbf{R} can be approximated as $\mathbf{I} + [\omega]_\times$ with \mathbf{I} being the 3×3 identity matrix and $[\omega]_\times$ the skew-symmetric matrix associated with the cross product. Then, \mathbf{M}' from Eq. (1) can be approximated as

$$\mathbf{M}' \approx \boldsymbol{\Omega}\mathbf{P} + \mathbf{M} + \mathbf{t}. \quad (2)$$

After some algebraic manipulation, the right side of Eq. (2) is expanded as

$$\mathbf{M}' \approx \begin{pmatrix} \omega_y P_z - \omega_z P_y + M_x + t_x \\ \omega_z P_x - \omega_x P_z + M_y + t_y \\ \omega_x P_y - \omega_y P_x + M_z + t_z \end{pmatrix}. \quad (3)$$

The normalized projection \mathbf{m}' of \mathbf{M}' on the image can be obtained by dividing the first two elements in the right side of Eq. (3) with the third. Multiplying the numerators and denominators by $M_z - (\omega_x P_y - \omega_y P_x + t_z)$, expanding and ignoring second order terms, yields

$$\begin{aligned} m'_x &= m_x + (t_x + \omega_y P_z - \omega_z P_y - m_x(\omega_x P_y - \omega_y P_x + t_z)) \cdot (T_z + P_z)^{-1} \\ m'_y &= m_y + (t_y + \omega_z P_x - \omega_x P_z - m_y(\omega_x P_y - \omega_y P_x + t_z)) \cdot (T_z + P_z)^{-1}. \end{aligned} \quad (4)$$

¹A normalized image projection refers to a projection on an ideal pinhole camera, i.e. if $\mathbf{M} = (X, Y, Z)$ then $\mathbf{m} = (\frac{X}{Z}, \frac{Y}{Z})$. In other words, the effects of the camera intrinsics \mathbf{K} on the projection have been removed.

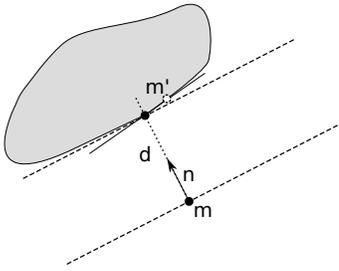


Fig. 1: Illustration of perpendicular edge matching. \mathbf{m} is the predicted projection of a control point and \mathbf{m}' the actual projection of this control point in the next frame. The edge through \mathbf{m}' is approximately parallel to that through \mathbf{m} (solid and dashed lines). Due to the aperture problem, the location of \mathbf{m}' cannot be fully determined. Instead, a linear search performed along the direction of \mathbf{n} provides distance d .

The above two equations are linear in the elements of $\xi = (\omega_x, \omega_y, \omega_z, t_x, t_y, t_z)^t$, therefore they can be written more compactly as

$$\mathbf{m}' = \mathbf{m} + \mathbf{A}\xi, \quad (5)$$

with \mathbf{A} being a 2×6 matrix defined by the coefficients of the elements of ξ from Eq. (4).

As shown in Fig. 1, let \mathbf{n} be the unit vector perpendicular to the edge at \mathbf{m} . The distance between points \mathbf{m} and \mathbf{m}' in the direction of \mathbf{n} is given by the projection of $\mathbf{m}' - \mathbf{m}$ on \mathbf{n} , i.e. $\mathbf{n}^t(\mathbf{m}' - \mathbf{m})$, and is approximately equal to d . Combined with Eq. (5), this yields $\mathbf{n}^t\mathbf{A}\xi = d$. In simple terms, every control point provides one constraint on the motion ξ , therefore six such points determined by searching along different directions suffice to estimate ξ . If more than six constraints are available, ξ can be estimated in a least squares manner by solving

$$\hat{\xi} = \arg \min_{\xi} \sum_i (\mathbf{n}_i^t \mathbf{A}_i \xi - d_i)^2. \quad (6)$$

Finally, the estimated $\hat{\xi}$ is used to incrementally update the pose. For future reference, the differences $\mathbf{n}_i^t \mathbf{A}_i \xi - d_i$ are called residuals.

D. Robustification of Pose Estimation

The linearization about the current pose estimate detailed in the previous section allows each pair of orthogonally matched points to yield one linear constraint in the six parameters defining the incremental change in object pose. In practice, more than six matches will be established and some of them will be erroneous due to various sources of error. Erroneous matches will give rise to outlying constraints with large residuals, which will cause least squares to fail due to its lack of robustness. Therefore, Eq. (6) cannot be used directly, despite its simplicity. Instead, the pose refinement computation is carried out in a robust regression framework, which allows problematic measurements to be identified and discarded before corrupting the estimate. Robustness is achieved by employing the Least Median of Squares (LMedS) estimator [23], which substitutes the summation

operation in the minimization of the squared residuals in Eq. (6) with the median:

$$\hat{\xi} = \arg \min_{\xi} \text{med}_i (\mathbf{n}_i^t \mathbf{A}_i \xi - d_i)^2. \quad (7)$$

LMedS can tolerate up to 50% erroneous constraints, but unlike least squares, has no closed form solution. Despite its robustness, LMedS suffers from low precision. Therefore, the LMedS estimate of ξ is not used as such. Instead, an extra computation step uses it to identify as outliers those constraints that yield large residuals, then discards them and estimates ξ using least squares on the remaining inliers. RANSAC [24] could also be used in place of LMedS to achieve robustness. However, to discriminate between inliers and outliers, RANSAC requires that the standard deviation of valid residuals is externally supplied. LMedS, on the other hand, estimates this deviation as a byproduct, thus automatically adapts to the level of noise present in the constraints. It is noted that the median is computed with the Quickselect algorithm [25], which avoids sorting and has an average complexity of $O(n)$.

An extra level of robustness is achieved by substituting the squared L_2 norm in the least squares minimization applied to the inliers with the p -th power of the L_p norm, $p < 2$:

$$\hat{\xi} = \arg \min_{\xi} \sum_i (\mathbf{n}_i^t \mathbf{A}_i \xi - d_i)^p. \quad (8)$$

The effect of the L_p norm is to down-weight large residuals by replacing their squares with their p -th power, which increase less steeply than quadratically. A value of 1.5 is chosen for p and the minimization in Eq. (8) is carried out with the IRLS algorithm after converting it to weighted least squares [26]. It is stressed that IRLS is applied to the inliers of the LMedS algorithm and not to the set of all available constraints.

To achieve more stable tracking and tolerate larger object motions, pose estimation is integrated with a linear Kalman filter with a constant velocity model. The filter is defined by six measurements and twelve state parameters. The measurements are the 6D pose estimated by tracking, whereas the state consists of the pose plus the linear and angular velocities. To attain a linear motion model, rotations are represented using Euler angles. At each iteration, the Kalman filter is used to predict the pose of the subsequent frame and drive rendering.

IV. EXPERIMENTAL RESULTS

Results from a C implementation of the proposed tracker are presented in this section. The experiments rely on synthetic image sequences depicting ENVISAT, a defunct Earth observation satellite measuring $26 \text{ m} \times 10 \text{ m} \times 5 \text{ m}$ (see Figs. 2 & 3). The sequences were rendered using the high detail mesh model shown on the left in Fig. 2, which consists of approximately 30K faces. They realistically simulate the motion of the satellite and the solar illumination, consisting of 1024×1024 images with a field of view equal to 40° . Each image is accompanied by accurate ground truth pose,

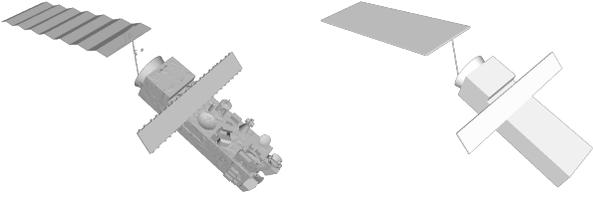


Fig. 2: 3D models of ENVISAT used for tracking: detailed model made up of 30K faces (left) and simplified model with 1K faces (right).

facilitating the quantitative measurement of a tracker’s performance. Sample images from the two ENVISAT sequences are shown in Fig. 3. A common characteristic of both test sequences is that they exhibit poor contrast due to insufficient illumination (cf. Fig. 3). To mitigate this, a low-complexity algorithm was used to enhance their contrast [27]. It is noted that this pre-processing is applied prior to applying all tracking techniques compared below.

The metric used for measuring pose error concerns the average misalignment between the model’s vertices at the true and estimated pose, as proposed in [28]. More specifically, for the true pose $\{\mathbf{R}_g, \mathbf{t}_g\}$ and an estimate $\{\mathbf{R}_e, \mathbf{t}_e\}$, the alignment error is defined as

$$E = \frac{1}{N} \sum_{i=1}^N \|(\mathbf{R}_g x_i + \mathbf{t}_g) - (\mathbf{R}_e x_i + \mathbf{t}_e)\|_2, \quad (9)$$

where x_i denote the N mesh model vertices.

The proposed tracker was tested on the two ENVISAT sequences of Fig. 3. Furthermore, two recent model-based trackers were also applied to these sequences for comparison. Specifically, the first is the edge-based tracker of the ViSP library² [29] that corresponds to the implementation of [9] by its authors. The other is the EBT library³, i.e. Choi’s own implementation of the multi-hypotheses tracker from [15]. Whereas the proposed method can employ the 3D model on the left in Fig. 2 directly, this is not the case for the method of [9]. As already pointed out, this is because [9] makes strong assumptions about the type of mesh edges that are permissible in the model of the object to be tracked, and assumes that the latter consists only of edges visible in the images. Thus, to enable a comparison, the simplified model on the right in Fig. 2 that conforms to the requirements of [9] was designed. On the other hand, the EBT tracker from [15] has no such restrictions and can work with the unmodified 30K faces detailed model.

The first sequence, a frame of which is shown on the left in Fig. 3, consists of 1152 frames. ENVISAT rotates out of plane and remains approximately at a constant distance from the camera, equal to about 55 meters. The top graph in Fig. 4 plots for each frame the alignment errors computed with Eq. (9) for the poses estimated by [9] with the simplified,

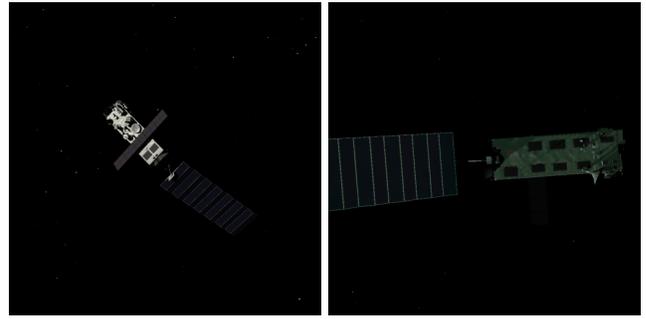


Fig. 3: Initial frames from the far (left) and close (right) tracked sequences; notice their poor contrast.

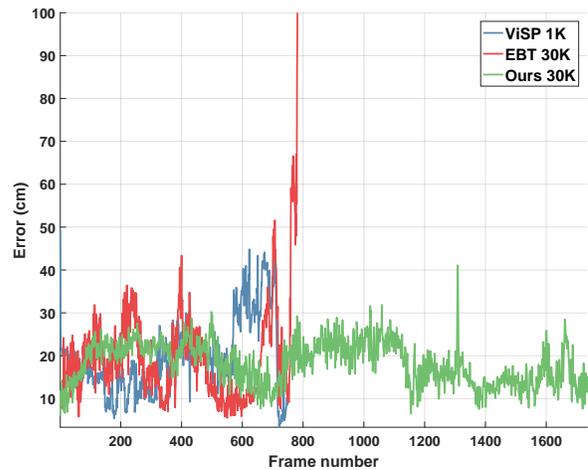
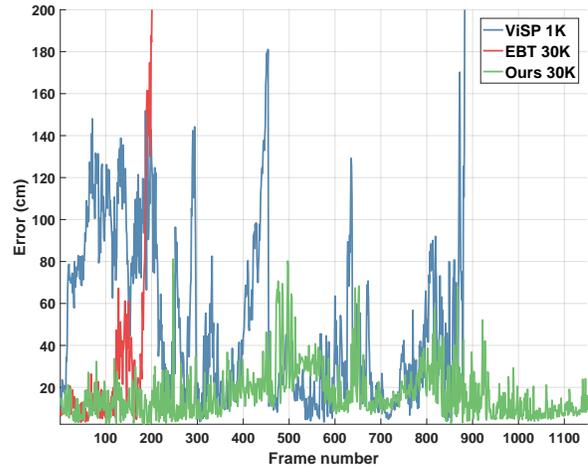


Fig. 4: Errors for the poses estimated with ViSP [9], EBT [15] and the proposed tracker for the far (top) and close (bottom) ENVISAT sequences. The red and green graphs correspond to poses estimated by [15] and the proposed method using the model on the left in Fig. 2, while the blue graphs correspond to poses estimated with [9] using the simplified model on the right in Fig. 2.

²<https://visp.inria.fr/>

³https://github.com/CognitiveRobotics/object_tracking_2D

1K faces model and by [15] and the proposed method with the 30K faces model. The proposed method yields lower pose errors compared to both [9] and [15], and succeeds in tracking the target along the entire sequence. On the contrary, [9] loses track around frame 900 and [15] after 220, and yield much larger errors before they fail.

In the second sequence (cf. Fig. 3 right), which is 1735 frames long, the camera is initially at a distance of around 30 m from ENVISAT. While the satellite rotates, the camera gradually approaches it up to approximately 11 m. Due to its large size, only some parts of ENVISAT are within the camera field of view in this range of distances. This sequence presents more challenging illumination which results in low contrast images, even after the enhancement by [27]. Tracked pose errors are shown in the bottom graph of Fig. 4, from which can be seen that the tracker of [9] is more accurate than the proposed one for frames between 100 and 400 but gradually becomes worse and fails completely after around frame 750. The EBT tracker of [15] produces sharper spikes and fails after frame 780.

Concerning execution speed, this depends on the apparent size of the tracked object in the images. Thus, the proposed tracker runs at about 10Hz on an Intel Core i7 CPU @ 3.60GHz for the far sequence and about 5Hz for the close one. These frame rates include the time spent loading compressed PNG images from disk and do not benefit from any hardware acceleration. Efforts are underway to accelerate the proposed tracker with space-grade, rad-hard FPGAs [30].

V. CONCLUSION

Relative localization of orbiting satellites is an important component of future autonomous proximity operations in space. Towards this direction, this work has suggested a generic approach for model-based tracking that relies on image edges. The developed tracker does not require any pre-processing of the 3D object model and makes no assumptions about its nature. It is also robust to common tracking nuisances, such as parts of the tracked object being out of view, occluded or completely undetected. Tracking accuracy has been evaluated with the aid of simulated image sequences with known ground truth poses and has been shown to compare favorably to established tracking techniques.

REFERENCES

- [1] A. Flores-Abad, O. Ma, K. Pham, and S. Ulrich, "A review of space robotics technologies for on-orbit servicing," *Progress in Aerospace Sciences*, vol. 68, pp. 1–26, Jul. 2014.
- [2] C. English, G. Okouneva, P. Saint-Cyr, A. Choudhuri, and T. Luu, "Real-time dynamic pose estimation systems in space: Lessons learned for system design and performance evaluation," *Intl. Journal of Intelligent Control and Systems*, vol. 16, no. 2, pp. 79–96, 2011.
- [3] M. Lourakis and X. Zabulis, "Model-based pose estimation for rigid objects," in *Intl. Conf. on Computer Vision Systems (ICVS)*, ser. LNCS. Springer-Verlag, 2013, vol. 7963, pp. 83–92.
- [4] D. Lowe, "Robust model-based motion tracking through the integration of search and estimation," *Intl. Journal of Computer Vision*, vol. 8, no. 2, pp. 113–122, 1992.
- [5] —, "Fitting Parameterized Three-Dimensional Models to Images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 5, pp. 441–450, May 1991.
- [6] C. Harris, "Tracking with rigid objects," in *Active Vision*, A. Blake and A. Yuille, Eds. Cambridge, MA: MIT Press, 1992, pp. 59–73.
- [7] M. Armstrong and A. Zisserman, "Robust object tracking," in *Asian Conf. on Computer Vision*, vol. I, 1995, pp. 58–61.
- [8] T. Drummond and R. Cipolla, "Real-time visual tracking of complex structures," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 7, pp. 932–946, 2002.
- [9] A. I. Comport, É. Marchand, M. Pressigout, and F. Chaumette, "Real-time markerless tracking for augmented reality: The virtual visual servoing framework," *IEEE Trans. Vis. Comput. Graph.*, vol. 12, no. 4, pp. 615–628, 2006.
- [10] S. Reinke, E. Gutzzeit, B. Mesing, and M. Vahl, "Tracking technical objects in outdoor environment based on CAD models," in *Proc. Intl. Symp. on Advances in Visual Computing (ISVC)*. Springer, 2012, pp. 437–446.
- [11] A. Petit, E. Marchand, and K. Kanani, "Tracking complex targets for space rendezvous and debris removal applications," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 4483–4488.
- [12] C. Kemp and T. Drummond, "Dynamic measurement clustering to aid real time tracking," in *IEEE Intl. Conf. on Computer Vision (ICCV)*, vol. 2, 2005, pp. 1500–1507.
- [13] G. Klein and D. W. Murray, "Full-3D edge tracking with a particle filter," in *Proc. British Machine Vision Conf. (BMVC)*. British Machine Vision Association, 2006, pp. 1119–1128.
- [14] C. Teuliere, E. Marchand, and L. Eck, "Using multiple hypothesis in model-based tracking," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 4559–4565.
- [15] C. Choi and H. I. Christensen, "3D textureless object detection and tracking: An edge-based approach," in *Proc. IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 3877–3884.
- [16] L. Vacchetti, V. Lepetit, and P. Fua, "Combining edge and texture information for real-time accurate 3D camera tracking," in *Proc. Intl. Symp. on Mixed and Augmented Reality (ISMAR)*, 2004, pp. 48–57.
- [17] E. Rosten and T. Drummond, "Fusing points and lines for high performance tracking," in *Proc. IEEE Intl. Conf. on Computer Vision (ICCV)*, vol. 2. IEEE Computer Society, 2005, pp. 1508–1515.
- [18] C. Choi and H. I. Christensen, "Real-time 3D model-based tracking using edge and keypoint features for robotic manipulation," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 4048–4055.
- [19] J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 8, no. 6, pp. 679–698, Nov. 1986.
- [20] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes, *Computer Graphics: Principles and Practice*, 2nd ed. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1990.
- [21] L. Bose and A. Richards, "Fast depth edge detection and edge based RGB-D SLAM," in *Proc. IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2016, pp. 1323–1330.
- [22] V. Lepetit and P. Fua, "Monocular model-based 3D tracking of rigid objects: A survey," *Foundations and Trends in Computer Graphics and Vision*, vol. 1, no. 1, pp. 1–89, 2005.
- [23] P. J. Rousseeuw, "Least median of squares regression," *Journal of the American Statistical Association*, vol. 79, no. 388, pp. 871–880, 1984.
- [24] M. Fischler and R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography," *Commun. ACM*, vol. 24, no. 6, pp. 381–395, June 1981.
- [25] C. A. R. Hoare, "Algorithm 65: Find," *Commun. ACM*, vol. 4, no. 7, pp. 321–322, July 1961.
- [26] A. Björck, *Numerical Methods in Matrix Computations*, ser. Texts in Applied Mathematics. Springer, 2015.
- [27] T. Arici, S. Dikbas, and Y. Altunbasak, "A histogram modification framework and its application for image contrast enhancement," *IEEE Trans. Image Processing*, vol. 18, no. 9, pp. 1921–1935, Sept 2009.
- [28] S. Hinterstoisser, V. Lepetit, S. Ilic, S. Holzer, G. Bradski, K. Konolige, and N. Navab, "Model based training, detection and pose estimation of texture-less 3D objects in heavily cluttered scenes," in *Asian Conf. on Computer Vision*. Springer, 2012, pp. 548–562.
- [29] E. Marchand, F. Spindler, and F. Chaumette, "ViSP for visual servoing: a generic software platform with a wide class of robot control skills," *IEEE Robot. Autom. Mag.*, vol. 12, no. 4, pp. 40–52, Dec. 2005.
- [30] G. Lentaritis et al., "Project HIPNOS: Case study of high performance avionics for active debris removal in space," in *IEEE Comp. Soc. Ann. Symp. on VLSI (ISVLSI)*, 2017, to appear.