

Sensitivity of automated attitude determination from ISAR radar mappings

S. Lemmens and H. Krag

European Space Agency, ESOC/HSO-GR, 64293 Darmstadt, Germany.

ABSTRACT

Inverse synthetic aperture radars (ISAR) are valuable instruments for assessing the state of a large object in low Earth orbit. The imaging capabilities of these radars can reach a sufficient quality for their products to be used during launch support or contingency operations, e.g. determining the structural integrity, or analyzing the dynamic behavior of an object. However, the direct interpretation of ISAR images can be a demanding task due to the nature of the range-Doppler space in which these images are produced. Recently, a tool has been developed by the European Space Agency's Space Debris Office to generate radar mappings of a target in orbit. These mappings are a 3D-model based simulation of how an ideal ISAR image would be generated by a ground based radar under given conditions, and can be used to support a data interpretation process. Radar mappings have been used to detect non-nominal behavior and estimate the attitude states of the target by fitting them to observations. It has been demonstrated for the latter use case, that a coarse approximation of the target through a 3D-model is already sufficient to derive the attitude information from the generated mappings. The level of detail required for the 3D-model is limited due to the nature of the ISAR image, which is the superposition of point scatterers. We analyze the accuracy of an automated attitude determination method from ISAR images by varying the level of detail of the 3D-model. For a model with the required accuracy, the process of automated image and attitude fitting is compared with the results obtained by an operator. We show how the 3D-model can further be exploited to estimate the pose of different spacecraft components with respect to each other.

1. INTRODUCTION

In more than 50 years of space activities, more than 4800 launches have placed more than 5000 satellites in orbit, of which only a minor fraction of about 1000 are still operational today. Besides this large amount of intact space hardware, with a total mass of about 6000 tons, additional objects are known to orbit the Earth. They are regularly tracked by the US Space Surveillance Network and, today, more than 16000 of them are maintained in their public catalogue, which covers objects larger than approximately 5cm to 10cm in low Earth orbit (LEO) and 30cm to 1m at geostationary altitudes (GEO). Operational spacecraft make up only 6% of the catalogued orbit population, while 28% can be attributed to decommissioned satellites, spent upper stages, and mission related objects (launch adapters, lens covers, etc.). The remainder of about 66% is originating from more than 260 on-orbit fragmentations which have been recorded since 1961. A part of these fragmentations are due to on-orbit collisions. The high impact velocities of these collisions, which can reach 15km/s for most missions in LEO, are the reason for the destructive energy, even despite of the small object sizes. So far, there are four recorded examples of collisions, with the latest and most prominent one between the active Iridium-33 satellite and the decommissioned Cosmos-2251 satellite.

Today, there is little knowledge on the attitude state of decommissioned objects. Observational means have advanced in the past years, but are still limited w.r.t. accurately estimating the orientation of the rotational motion vector and its magnitude of most objects. In general, the attitude evolution of a decommissioned object is expected to be chaotic at first, and regularite slowly under the influence of external torques, e.g. due to the gravity gradient or Eddy currents. The actual attitude evolution depends strongly on the inertia tensor of the object and is influenced by the environment. Past observation, e.g. by the Fraunhofer Institute for High Frequency Physics and Radar Techniques (FHR) Tracking and Imaging Radar (TIRA) applying inverse synthetic aperture radar (ISAR) techniques, have mainly concentrated on objects which were about to undergo an uncontrolled re- entry in the following days. In a few cases attitude rates of about 10 deg/s have been estimated in this way. Generation of so-called light curves, i.e. evolution of the brightness of space objects in visible wavelengths, with the help of optical telescopes is a second promising measure to estimate attitude rates of targets in higher altitudes. Furthermore, laser ranging could be a potential tool to quantify varying offsets of a reflective surface from the center of mass of the parent body, if such a surface is present. In all cases, research is required to further improve the resolution of the data and/or to fit geometric models into the measurements for an estimation of the attitude.

The European Space Agency (ESA) Space Debris Office developed a software library which supports the interpretation and simulation of ISAR images, with as goal understanding and determining the attitude evolution of an object in orbit [1]. The library, MOWA (Models on Orbit With an Attitude), contains the required routines to generate the so called radar mappings. Such mappings are a 3D-model based simulation of how an ideal ISAR image would be generated by ground based radar under given processing conditions. This technology is a valuable tool for spacecraft operators when analyzing their spacecraft, e.g. during a contingency, where the identification of the attitude motion is key element in the process, or to confirm the deployment of structures. Moreover, knowledge on the rotational state, and its predictability, of an object is a driver for the preparation of an active debris removal mission. For example, attitude motions above a few degrees per second will impede the use of robotic capture mechanisms.

The focus of this paper is the collection of routines of MOWA dedicated to the automated processing of ISAR images to establish the attitude of an object during an observation pass. These routines include the pre-processing of ISAR images to binary representation, generation of distance maps and image fitting optimization. The theory is presented with the application to observations made from ESA's Envisat satellite on 2012-04-26 by FHR's TIRA installation. Different 3D models are compared to assess the level of a-priori knowledge required for the automated image fitting methods, and thus the attitude determination, to reach convergence. Moreover, given an ISAR image where the pose of the object is clear, the automated attitude fitting method can be applied to establish the pose of the 3D-model components w.r.t. each other. Finally, we will compare the automated attitude determination to the results obtained by an operator.

2. ISAR IMAGES AND RADAR MAPPINGS

ISAR is the denomination of a technique where a fixed system is collecting data from a moving target, and by analyzing the Doppler histories of scattering centers creates a synthetic aperture. For the mathematical background of this process, the reader is referred to [2]. Being based on the Doppler effect, the image plane in ISAR imaging lies quite different from the optical one. The line of sight (LOS) is embedded in the image plane and not orthogonal to it as in optics. The other dimension of the image plane depends on the rotational motion of the object. In general, the set of LOS vectors at different times in the object-fixed coordinate system spans a surface. Usually, this surface can be approximated by a plane which can then be identified as the image plane. This is especially the case for comparatively small angles between the LOS vectors at different times, corresponding to a small time interval. The ISAR image then is the sum of the parallel projections to the image plane of the set of all scatterers of the object, weighted by their amplitudes and convolved by the point scatterer response of the system.

In certain observation situations, the surface spanned by the LOS vectors is exactly a plane, if Earth rotation is ignored. One such case is, for example, a zenith pass of a satellite in which the orbit vector, from Earth centre to the satellite, is fixed and is not rotating around an axis in the orbit plane. Another example is an object on a straight path relative to the Earth coordinate frame and not rotating in this frame. If the attitude of the satellite is always known, the surface spanned by the LOS vectors is also known and the image plane can be determined from that knowledge. However, in general Earth rotation has to be taken into account. For clarity, one can divide the rotational motion into two components:

1. The rotation of a coordinate system with one axis tied to the LOS vectors relative to an assumed stabilized attitude of the satellite.
2. The rotation of the assumed stabilized satellite relative to the real satellite attitude, called intrinsic rotation.

The intrinsic rotation is usually not known and difficult to determine. Therefore, in the imaging process it is often assumed to vanish. The composition of both rotational motions yields the rotation of the coordinate system with one axis tied to the LOS vectors relative to the real satellite attitude as required.

Whenever the attitude state of an object is known, it might be beneficial for an operator to simulate ISAR images and compare them to observations in order to detect anomalies, e.g. in the deployment of structures. For such use cases, it is not required to simulate the entire process of transforming Doppler-Range measurements of a scattering body into an image. Rather, we simulate how an ideal ISAR image would look like when the model is geometrically known, and only the attitude state of the object and the geometry between observer and observable are taken into account. These simulate images, approximating true ISAR images, we will call *radar mappings*.




Fig. 1. Left: ISAR image produced by TIRA. Centre: Radar mapping generated by MOWA for 3D model given the orbit state of the ISAR image and fitted attitude. Right: Representation of the attitude used to generate the radar mapping, i.e. representation of an optical observation given the attitude and orbit state of the ISAR image.

To generate radar mappings, the MOWA library requires a 3D-model of the object and orbit and attitude ephemerides. The Doppler shift which is used to generate an ISAR image comes from the change in aspect angle. For a stabilized motion, the Doppler axis, also called the cross axis, of the image plane is determined by the unit vector computed from the subtraction of the LOS vector relative to a target-fixed coordinate system from one ephemeris to the previous one. The rotation axis is thus perpendicular to the plane determined by the two LOS vectors. Here we assume an inertial stabilization, i.e. that the target doesn't rotate in the inertial system. If the image scaling orthogonal to the LOS in the ISAR images, generated due to the assumption of vanishing intrinsic rotation is obviously wrong, an additional intrinsic rotation has to be defined. Then the Doppler axis vector is the sum of two components. The first component is the direction determined by the LOS vectors relative to the stabilized, i.e. inertial, system, as in the previous case. The second component lies within the assumed plane of the intrinsic rotation and is proportional to the assumed intrinsic rotation velocity. In our case, the intrinsic rotation plane is assumed to coincide with the orbit plane. The second component thus lies in the orbit plane. The image plane for the radar mapping is then defined by using the normalized LOS as y-axis and the normalized cross axis as x-axis. By assuming an inertial reference system, these axes have an interpretation in the 3D model reference frame. In the latter frame, we define the body fixed axes to coincide with the coordinate axes for the default attitude, i.e. the unit quaternion.

3. ENVISAT TIRA PASSAGE 2012-04-26

After 10 years of service, ESA's Earth observing satellite Envisat stopped sending data to ground from 8 April 2012 onwards. Following rigorous attempts to re-establish contact and the investigation of failure scenarios, the mission was declared over on 9 May 2012. The satellite is since then stranded in Sun-synchronous orbit at an altitude of around 760 km. On 2012-04-26 09:56 UTC, a circa 11 minutes pass Envisat was monitored with FHR's TIRA system, resulting in 68 ISAR images which are used for study in this paper. TIRA primarily serves as the central experimental facility for the development and investigation of radar techniques for the detection and reconnaissance of objects in space. It offers space agencies the possibility to measure the orbit of objects with high precision or produce a high resolution image of objects in LEO.

In the case of this observation sequence, an intrinsic rotation was not required for the generation of distortion free ISAR images. In its default Earth observing attitude, an intrinsic rotation of 0.06 deg/s with its axis of rotation perpendicular to the orbit plane would have been, theoretically, required. However, with integration times in the order of 10 sec for the generation of the ISAR images, the resulting rotation is low enough to be ignored. Based on 3D-model built from publically available data, see Section 5, the attitude of the satellite during each observation was determined by an operator by visually determining the best matching radar mapping. When looking to the observation sequence as a whole, a slow rotation becomes apparent, as can be seen later on in Fig. 9.

In Fig. 1, observation number 50 of the described sequence is given. The colors in the ISAR images are pseudo-colors indicating the strength of the reflected radar waves by the scatterers on the object. A mild remnant of

multipath reflection can be seen above the large structure, i.e. the solar array, in the top right of the image. In clear ISAR images such as the one above, shading effects can be seen, e.g. Envisat's Advanced Synthetic Aperture Radar (ASAR) is partially occluded by the main satellite bus. This ISAR image serves as the example for the methods described in the following section.

4. POSE ESTIMATION FROM RADAR MAPPING FIT OPTIMISATION

The direct interpretation of ISAR images is hampered due to the fact that they are generated based on the theory of scattering, which can lead to distorted images, the appearance of multipath reflections and a reduction in feature resolution scale. Moreover, to an average human, the range-Doppler plane in which these images are displayed is a non-intuitive geometrical environment and requires some a priori assumptions on the rotational state of the target. Thus, to automatically determine the attitude from ISAR images, the corresponding techniques from computer vision have to be identified. Well studied concepts such as structure from motion or feature tracking have been translated to the radar setting [3,4,5]. However, for the purpose of automation, they are limited in application due to the nature of the ISAR images.

The process of estimating which transformation is applied to a 3D object to generate a specific 2D image is called pose, i.e. position and orientation, estimation. One way of solving this problem, is by comparing the 2D image with a collection of projections from the 3D object and determining the best match. This avoids the limitations posed by feature tracking algorithms, e.g. iterative closest point requires a dense set of tracking points, at the cost of having to generate the collection of projections, which might be computationally intensive. We will apply this technique to determine the attitude from ISAR images, by automatically produce radar mappings and reformulate the problem as a numerical minimization problem. As a first step, the images are pre-processed to extract the shape of the object under scrutiny into a binary, i.e. black and white, image.

Some parts of the object reflect the radar waves only moderately, or display a certain regularity, which makes it easy for a human interpreter to find the edges of a shape. On the other hand, many scattering objects closely together make it hard to establish a clear boundary and to relate it back to its basic shape, and highly reflective parts can generate smeared-out streaks in the image. In most cases, a human interpreter outperforms a classical edge detection based algorithm when making educated guesses concerning the underlying shape defining, when posed with low-reflective parts. Therefore, it is still beneficial that an operator checks, and if necessary augments, the ISAR images.

The following procedure is applied to the pseudo-color ISAR images described in Section 3 to extract a primary binary image:

1. The input image is converted to a grey scale, between 0 and 1, based on the RGB luminance formula applied to each pixel.
2. The noise in the image is reduced by running a small, 3 to 5 pixels, median filter over it.
3. A gradient map of the image is computed from a Sobel filter.
4. A marker image is created by identifying the low, medium and high luminescence pixels.
5. The original image is segmented into two categories based on the gradient map and markers by applying the watershed method (As implemented in `skimage 0.8.0`).



Fig. 2. Application of the shape extraction routine to the ISAR image presented in Fig. 1. Left: Gradient based watershed. Centre: Morphological filling. Right: Contour extraction.

The procedure above might still leave black spots where the original images has low reflective parts and white spots where noise reached maximum intensity for a few pixels. Thus a second procedure is applied to fill up the black holes and remove the white dots. It is important to remove the holes, i.e. connect the white regions, first in order not to remove true object features, e.g. edges.

The following procedure is applied to the watershed image:

1. The morphological closure of the image is compute a few consecutive times while reducing the radius of the disk used for the closing in each run [6].
2. The morphological opening of the resulting closed image is compute with a disk of radius 1 pixel [6].

The last procedure can also be applied to the radar mappings, as for these mappings it might happen that connected faces in the 3D-model result in separated streak in the generated radar mapping for certain attitudes. This is due to the discretization process in the radar mapping algorithm as described in [1]. The result of two pre-processing procedures applied to the ISAR image given in Fig. 1, is shown in Fig. 2. As can be seen from this example, when radar reflectivity is low, as for the back of the ASAR instrument, the shape extraction algorithm has difficulties in producing a faithful representation. This has to be taken into account for the optimization process.

From the generated binary images, which contain at least, but are not limited to, one white zone, the contour can easily be extracted with edge detection algorithms. When we want to translate the pose estimation problem into a minimization problem, candidates for cost function include quantifying the match between the contours of the binary ISAR images with the radar mappings or maximizing the overlapping area of both binary images. The later one was already preliminary explored in [1], and provides a viable option when the binary ISAR image is very close to reality, e.g. no artificial thickening of the borders by the morphological routines when some noise remains, and only one connected component is visible in the image. However, as in the example given, low reflectivity parts are problematic. Matching contours on the other hand, can be made more robust in the present of multiple project zones as contours can be close to each other without being penalized for have little overlapping area. Here, we will propose a cost function based on a harmonized contribution of both metrics.

When presented with two binary images, I_{ISAR} and I_{rmap} denoting respectively the binary ISAR image and binary radar mapping, the geometric center of both images is computed and used to center the images before the following operations. This is done to make sure that the focus of the radar installation for the ISAR images is aligned with the center of imaging used for the radar mapping. The projected area, pa , of a binary image is the total count of white pixels. Then we define the projected area metric α as:

$$\alpha = 1 - \frac{2pa_{I_{ISAR} \cap I_{rmap}}}{pa_{I_{ISAR}} + pa_{I_{rmap}}},$$

assuming that both I_{ISAR} and I_{rmap} have the same shape and are centered with respect to their geometric center before computing $I_{ISAR} \cap I_{rmap}$. Then α simply captures the ratio of overlapping area, with 0 indicating a perfect match and 1 no overlap at all. The graphical interpretation of α is given in Fig. 3.

To quantify the matching contours, the pixel boundaries from I_{ISAR} and I_{rmap} , denoted as b_{ISAR} and b_{rmap} respectively, are extracted and will be used in a Chamfer distance matching system [7]. In the first step, a new black image with the same shape as I_{ISAR} is generated and the pixels corresponding to b_{ISAR} are colored white. Secondly, for each pixel i,j , the distance $d_{i,j}$ to the nearest element of b_{ISAR} is computed, generating a so called distance map. The distance used for generating the map is the Chebyshev distance. Then our distance metric β is defined as

$$\beta = \frac{1}{N} \sum_{(i,j)}^N d_{i,j} \quad (i,j) \in b_{rmap},$$

where N is the total amount of pixels in the boundary of the radar mapping b_{rmap} . In summary, the metric β can be thought of as the pixel mismatch between the contours of the binary ISAR image and radar mapping per pixel. The graphical interpretation of β is given in Fig. 3.




Fig. 3. Graphical representations of the α and β metrics for the operator defined best fit. Left: Overlay of the ISAR image (white) and radar mapping (grey) shapes. Right: Chamfer distance map plot, with darker indicating closer to zero, overlaid with the contour of the radar mapping shape.

When describing the attitude of an object with respect to the J2000 inertial frame in intrinsic yaw-pitch-roll angles (z-y-x sequence, degrees), the operator identified the best match at the coordinates (-51, -15, 124) for our example ISAR image. Given this image and a coarse 3D-model for Envisat, as is described in Section 5, the metrics α and β are computed for all integer degree values within 60 degrees bins centered around the best match coordinates. When interpreting α as a function to be minimized over the search space, the phase space is convex and thus suitable for minimization, as can be seen in Fig. 4. However, the minimum value attainable is not the same as the best match found by an operator, as noise in the ISAR image between the main satellite bus and solar array gets transformed into an unphysical white zone when pre-processing the images. Moreover, the missing ASAR antenna skews the fit. When looking at the phase space defined by the metric β , the minimum value $\sim (-49, -15, 124)$ is close to the operator defined minimum, as the overall contours matching the main spacecraft bus and solar array outweigh the contribution from the missing ASAR antenna for all nearby coordinates. However, phase space seems convex on large scale but is highly irregular when looking at degree and sub-degrees scale, visible in Fig. 4.

Based on the nature of both metrics, they are combined into one metric γ as follows:

$$\gamma = \exp(a \alpha - b) \ln(\beta) .$$

Increasing the constant a , the phase space of γ will tend to be more like the one for α in form. The constant b allows us to set a value which we consider to be a good area match in terms of α . Taking the natural logarithm of β helps smoothing out the spiked behavior in the phase space and rewards matching contours when the area is not far away from the b threshold. The phase space of γ is still not convex but retains the desired global minimum. The phase space of γ is shown in Fig. 5.




Fig. 4. Sliced phase spaces of α (left) and β (right) for attitudes around the operator defined minimum with a fixed roll angle of 124 degree.

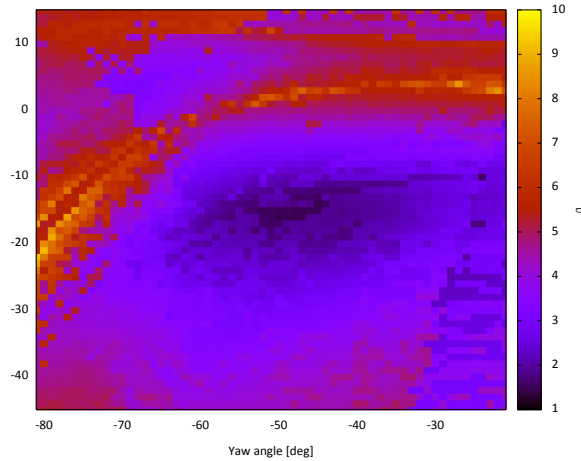


Fig. 5. Sliced phase spaces of γ for attitudes around the operator defined minimum with a fixed roll angle of 124 degree.

Different Jacobian and Hessian based constrained optimizers together with iterative least square fitting routines were tested to minimize γ interpreted as cost function, dependent on a and b . However, apart from large values of a , the non-convexity of γ impedes convergence even when starting close to the desired minimum, as line searching routines do not often find usable values along the path defined by the gradient. Simulated annealing was tested on the same phase spaces γ , with various different neighbor selection and cooling classes, but the run time is comparable to full brute force searching when the search space becomes larger than a few degrees, even though a point close to the minimum is often reached.

To combine information of the large scale gradient with the ability of simulated annealing methods to climb outside of local minima, a new search heuristic is defined. First the values of a and b are set to 2 and 0.35 respectively, based on empirical testing.

1. Starting from an attitude position x_i , phase space within one degree around x_i is sampled to find a lower value of γ . If it is found, x_i is updated to the new position.
2. If no lower value is found after 10 iterations, the gradient is numerically computed based on a large step size to avoid local irregularities. The negative gradient direction is sampled up to a few degrees for values of γ , and the lowest point x_l is retained.
3. Then x_i is updated to x_l whenever it is lower than x_i or when

$$\exp\left(-\frac{|\gamma_i - \gamma_l|}{c T_i}\right)$$

is smaller than a predefined acceptance threshold (< 1). Herein, c is a constant to increase acceptance for larger differences in the cost function $|\gamma_i - \gamma_l|$ and T_i is the systems' temperature. T_i is initially set to a value close to the maximum value of the cost function in the search space. The update schedule for T_i is dependent on the maximum number of search iterations and cools down exponentially [8].

By design, simulated annealing search heuristics require a lot of fine tuning and testing, making the final result difficult to transfer from one problem to another. In the routine above, step 1 can be replaced by an annealing like neighbor selection routine as well, and allowed to take positive step in the phase space. In the following sections, the heuristic as described above is used. However, it is important to note that further work is on-going to reduce the computational cost of the search step.




Fig. 6. Different 3D-Models used for automated attitude fitting on ISAR images. Left: Coarse model. Centre: Public model. Right: Expert model.

5. 3D-MODEL COMPARISON

Observation number 50 from the sequence described in Section 3 was selected as it provides a favorable attitude for the assessments of the object's dimension, if no prior knowledge would be available. The short range dimension of the solar array, together with its span in cross track, means that almost the full length would be visible when observed optically. The same reasoning applies to the interpretation of the shape coming from the main satellite bus. The presence of the fixed mounted ASAR antenna in the image can help to determine the angle between the main bus and the solar array. Moreover, a communications antenna is clearly distinguishable. These four parts are thus the minimal components required to make a 3D model of the object.

Three models are built with the purposed of automatically determine their attitude with the procedure described in the previous section. The models will exhibit three levels of detail:

1. The coarse model. The dimensions of the four compounds are estimated form the ISAR image and rounded to the nearest integer meter. This procedure slightly underestimates the dimensions of the solar array and main bus, and overestimates the communications and ASAR antennas lengths. The connection between the solar array and main bus is model as a simple stick of 30 cm width. All components are modeled as boxes.
2. The public model. The dimensions of the coarse mode are updated with publically available to increase the accuracy to the centimeter level. The position of the communication antenna is set to its correct location and the ASAR antenna is inclined w.r.t. the main satellite bus. The lower part of the main bus is an adapter ring which we model as a box. Its dimensions are taken into account to correct the location of the solar array w.r.t. the main bus.
3. The expert model. For the dimensions of the third model, actual design drawings are used, which is consistent the public model. Instead of modeling the main bus as one box, the large scientific instruments on top of the main bus are modeled separately and cones and cylinders are used to replace the boxes where required. There is a difference in solar array pose of one degree w.r.t the public and coarse model.

The coarse and public model are constructed directly in MOWA, the expert model was constructed in ESA's DRAMA software and afterwards exported to MOWA's TDO format [9]. For all models, the solar array is assumed to be in a position consistent with the observation sequence. The three models are shown in Fig. 6.

To assess which level of detail is required for automated attitude fitting, one hundred different attitude states per model are randomly selected from the yaw-pitch-roll attitude bins $[-61, -41]$, $[-25, -5]$ and $[114, 134]$. These bins include a worst case instantaneous rotation of 2 deg/s, which is far higher than the rotations rates observed in this sequence. The optimization scheme as described in Section 5 is run for maximally 50 iterations, and the last and lowest cost function value state are returned. The difference between attitude states can conveniently be described in terms of the quaternion distance, or equivalently as the angle of the quaternion rotation to transform one state into the other. We will use the later one to assess the quality of the automated fitting. When two attitude state, $q1$ and $q2$, are given as quaternions, the angle δ between them is given by

$$\delta = \arccos(2 \langle q1, q2 \rangle^2 - 1),$$

where \langle, \rangle is the standard inner-product on the quaternion space.

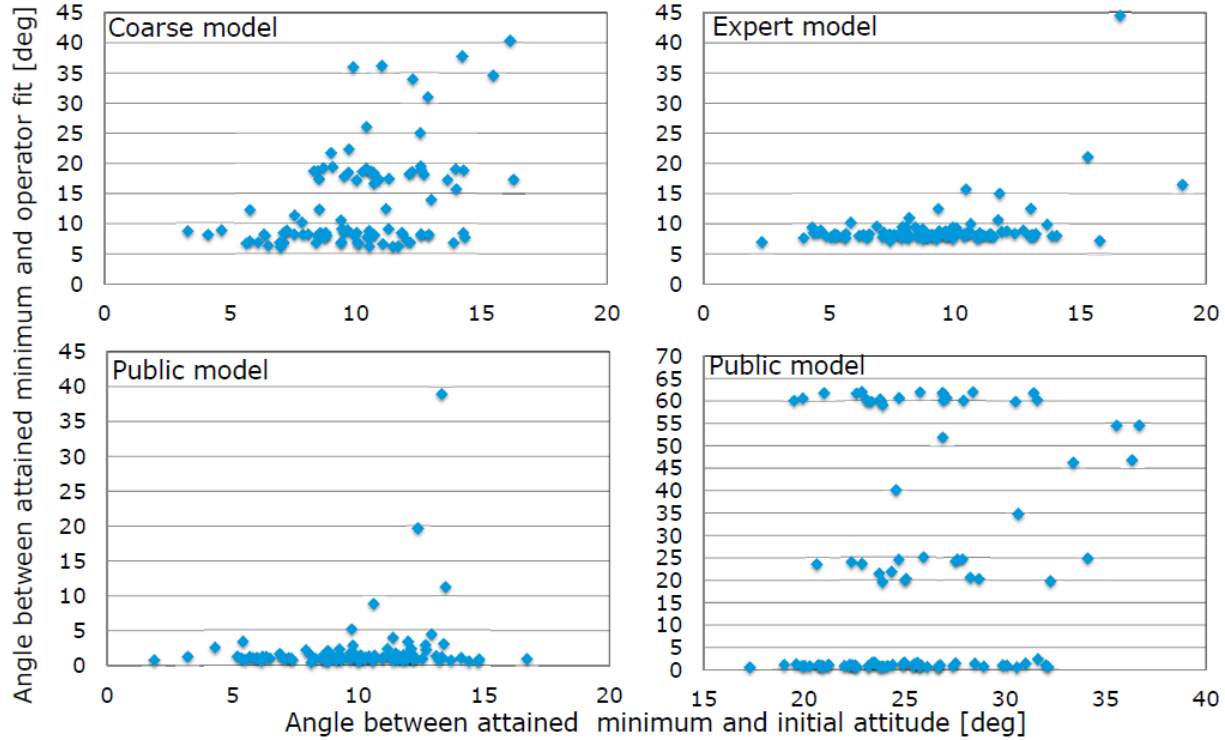


Fig. 7. Scatterplots of the distance from the operator defined minimum of the starting position versus the minimum attained for the different 3D-models.

We consider a run as diverged when the angle between the attained minimum given by the optimizer and the operator defined best fit attitude is larger than 10 degrees. For respectively the coarse, public and expert model, 49, 4 and 23 of the 100 runs are divergent. Excluding the divergent cases, the median of the remaining angle for respectively the coarse, public and expert model are 8.04, 1.92 and 8.50. Therefore, in both perspectives, the public model significantly outperforms the coarse and expert model when it comes to mimicking the behavior of an operator. As can be seen from the scatter plots for the different cases given in Fig. 7, the divergence is not correlated with the initial distance from the operator defined attitude fit.

In the case of the expert model, the optimizer converges to a minimum in phase space, but not the desired one for the current set up. The cause is the tilt offset in the pose of the solar array w.r.t. the main body, which was corrected for the coarse and public model. For the coarse model, no consistent fit is found. One must thus conclude that even though the ISAR images are noise, the 3D model must be sufficiently accurate in order not to misinterpret the results of an automated fit.

The experiment is repeated with the public model but pushing the initial starting condition further away from the predefined best fit. In this case, shown in Fig. 7, the convergence is worse. 54 of the 100 trials reach the desired minimum, with a median of 1.34 degrees, but two other minima are found farther away. Again, the initial distance from the desired minimum does not influence the attained minimum per se, but the non-convex nature of the phase space leads to different basins of attraction, which can be found via complex paths.

6. SOLAR ARRAY ORIENTATION

The models in Section 5 were all tested under the same assumption w.r.t. the pose of the solar array. However, as was clear from the difference in fitting between the public and expert model, this pose can have a significant impact on the quality of the fit. Therefore, we repeat the test of Section 5, but now by varying the solar array pose on the public model. Fifty different attitude states are randomly selected from the yaw-pitch-roll attitude bins $[-61, -41, [-25, -5], [114, 134]$. The solar array of the original public model is rotated with an angle, in degrees, from the following list: -20, -15, -11, -8, -6, -5, -4, -3, -2, -1, 0, 1, 2, 3, 4, 5, 6, 8, 11, 15, and 20.




Fig. 8. Boxplots summarizing the quality of the automated fit for given solar array offset angles.

The results of the 21 trials poses are summarized in Fig. 8, where the data per trials is displayed as a box plot. The outer whiskers for the boxplots are the maximum and minimum differences between attained minimum and operator defined best fit for the 50 runs. The boundaries of the boxes are defined by the 75th and 25th percentile, and the blue star is the median, determined for the data from each individual trial. Based on this data, solar array poses which require an offset angle outside the interval $[-3, 3]$ are excluded. The observation geometry of the original ISAR image is more favorable for discriminating between the lengths and tilt of solar array, as was done in the previous section, rather than the cross-track component.

7. PASS ATTITUDE FITTING

After identifying a suitable model from one ISAR image in the previous sections, the selected model, i.e. the public model with zero offset angle for the solar array pose, is used to run an automatic attitude fit on the entire series. The operator defined fit at the 50th observation is used as initial guess. Once a radar mapping has been fitted to an ISAR image, the computed attitude is used as initial guess for the next step in the sequence. The results of the automated fit in comparison with the operator fit are given in Fig. 9.

The automated fit follows to within a few degrees the operator defined best fit attitude until halfway the sequence a strong divergence appears. This point coincides with an observation where the solar array vanishes from the ISAR image, as it has its long side aligned with the range direction and short side at a fixed cross-track value. The noise on the main bus allows for multiple attitudes to fit the image, and a diverging sequence starts. An operator can swiftly identify such cases, and the process can be restarted after a few observations, when the aspect angle changed sufficiently.

In general, this makes the automated attitude fitting procedure a useful tool for pre-processing a sequence of ISAR images once a model has been identified and an operator makes a first fit. The result has to be manually verified by an operator to assess the correctness.




Fig. 9. Comparison between operator defined best fit and automated attitude fit in yaw-pitch-roll angles. The black line indicates the starting point of the automated fitting process.

8. CONCLUSIONS

In this paper, a new search heuristic to solve the problem of automated attitude determination from ISAR images via interpreting the problem in a shape matching context, was presented. The phase space search strategy is based on mixing simulated annealing and gradient descent paradigms. This optimization problem avoids the automated feature identification difficulties at the cost of increased computation time. Further research is required in order to reduce the computational effort involved.

The automated attitude fitting procedure, as tested on a representative ISAR image, can be used to identify suitable 3D- models of the observed object in different test scenarios. Given favorable attitude observation geometries, pose and dimension hypotheses of the observed object can be tested as convergence or divergence of the automated fit in comparison to an operator defined match. To obtain an operator fit, some experience is required but not an accurate model. However, the identification of an accurate model is of significant use when many observation sequences for the same object have to be processed.

In the case where a full observation sequence has to be processed, the automated fit can approximate the fit by an operator to within a few degrees. However, but problematic geometries can occur, which trigger divergence of the solution. These divergences are easily identified upon manual inspection. This implies that the automated fitting can be used as an initial guess before an operator does the fine tuning, which still reduces the effort with respect to a full operator fitting or feature tracking algorithms.

9. ACKNOWLEDGMENTS

The authors would like to thank Quirin Funke and Benjamin Bastida Virgili for many fruitful discussions during the development of the attitude fitting method. Furthermore, we would like to thank ESA's General Studies Programme for their support.

10. REFERENCES

1. Lemmens, S. and Krag, H., *Radar Mappings for Attitude Analysis of Objects in Orbit*, Proceedings of the 6th European Conference on Space Debris, April 20-24, ESOC, Darmstadt, Germany, ESA SP-723, 2013.
2. Borden, B., *Radar Imaging of Airborne Targets*, Institute of Physics Publishing, ed. 1, 1999.
3. Rosebrock, J. , *3-D point scatterer estimation for 3-D ISAR imaging and attitude determination*, Proceedings of the 6th European Conference on Synthetic Aperture Radar, 2006.
4. Rosebrock, J., *Absolute Attitude From Monostatic Radar Measurements of Rotating Objects*, IEEE Transactions On Geoscience And Remote Sensing, vol. 49, no. 10, pp. 3737-3744, 2011.
5. Ferrara, M., Arnold, G., and Stuff, M. A., *Shape and Motion Reconstruction from 3D-to-1D Orthographically Projected Data via Object-Image Relations*, IEEE Transactions On Pattern Analysis and Machine Intelligence, vol. 31, no. 10, pp. 1906-1912, 2009.
6. Soille, P., *Morphological Image Analysis: Principles and Applications*, Springer, ed. 2, 2004.
7. D. M. Gavrilu and P. Vasanth, *Real-time object detection for "smart" vehicles*, The Proceedings of the Seventh IEEE International Conference on Computer Vision, IEEE, 1999.
8. A. L. Ingber, *Very Fast Simulated Re-Annealing*, Journal of Mathematical and Computer Modelling, vol. 12, pp. 967-973, 1989.
9. N. Sanchez-Ortiz, J. Correira de Oliveira, J. Gelhaus, H. Krag, *Computation of Cross Section of Complex Bodies in ESA DRAMA tool*, Proceedings of the 6th European Conference on Space Debris, April 20-24, ESOC, Darmstadt, Germany, ESA SP-723, 2013.