

Precise Registration of 3D Models To Images by Swarming Particles

Joerg Liebelt
Computer Graphics & Visualization Group
Technical University Munich
D-80333 Munich, Germany
liebeltj@in.tum.de

Klaus Schertler
Dept. IW-SI
EADS Innovation Works Germany
D-81663 Munich, Germany
klaus.schertler@eads.net

Abstract

The precise alignment of a 3D model to 2D sensor images to recover the pose of an object in a scene is an important topic in computer vision. In this work, we outline a registration scheme to align arbitrary standard 3D models to optical and Synthetic Aperture Radar (SAR) images in order to recover the full 6 degrees of freedom of the object. We propose a novel similarity measure which combines perspective contour matching and an appearance-based Mutual Information (MI) measure. Unlike previous work, the resulting similarity measure is optimized using an evolutionary Particle Swarming strategy, parallelized to exploit the hardware acceleration potential of current generation graphics processors (GPUs). The performance of our registration scheme is systematically evaluated on an object tracking task using synthetic as well as real input images. We show that our approach leads to precise registration results, even for significant image noise, small object dimensions and partial occlusion where other methods would fail.

1. Introduction

Alignment of rigid 3D models to 2D scenes figures among the challenging problems in computer vision. Typically, a known 3D model of an object needs to be precisely aligned with the signature of the same object in sensor data as an important component for further scene analysis and object recognition tasks.

In the present work, we propose a comprehensive and flexible approach to precise 3D model registration, consisting of the generation of the model modalities corresponding to different sensor types, the computation of a robust and accurate probabilistic similarity measure and its efficient optimization based on the evolutionary concept of Particle Swarming to recover the six-dimensional object pose. Figure 1 illustrates the components of the proposed scheme. We present the results of extensive testing on synthetic and

real sensor input to analyze its precision and convergence properties and apply the system on an exemplary 3D tracking task.

In the following, we summarize previous literature on each of the topics that we address as part of our work.

Measure of Similarity

A vast number of approaches have been described in the literature to derive similarity measures between models and sensor data.

Some authors resort to computing the difference between 2D-projected edge models and edge-filtered sensor images, for example in [5], thereby increasing generality, but discarding other possibly discriminative information on both the model and the sensor side. Other researchers derive shape signatures or similar descriptors from the model and try to identify the corresponding cues on the sensor side by comparing the respective signatures, such as in [15]. [4] use a silhouette-based method to stitch textures from 2D images onto 3D models.

In their original work, [14] introduce the concept of *Mutual Information* (MI) as a similarity measure based on information theory and demonstrate its use for aligning untextured 3D objects to images using the interpolated surface normals as clues on the model side. In several publications, extensions to the classical MI formulation are proposed, notably by introducing normalization terms such as in [6, 12] to account for the amount of overlap or weights to account for spatial relationships, usually based on gradients or segmentation [11, 13].

Two fundamental problems of classical MI have been addressed by [10], i.e. the lack of spatial information in the similarity measure and the "curse of dimensionality" which prevented the use of classical density estimators for multivariate MI computation.

In the present work, we propose several extensions to the method of [10], notably the fusion of *Mutual Information* with an edge-based measure, thereby increasing robustness

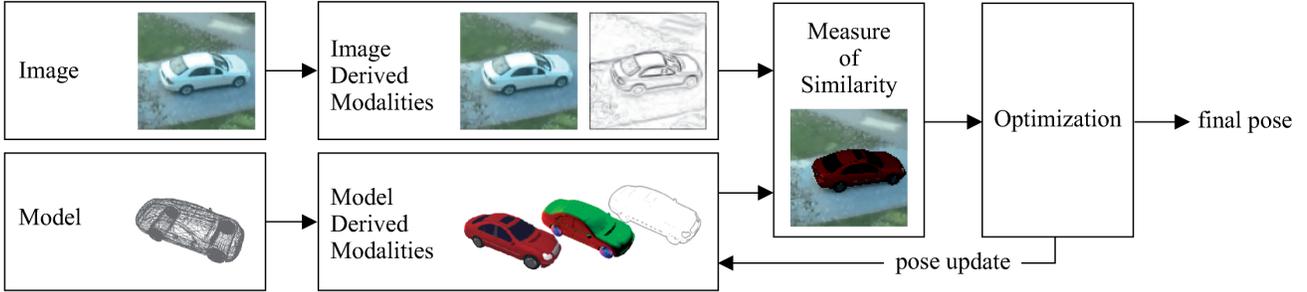


Figure 1. Outline of the proposed model registration scheme.

and registration precision.

Optimization for Registration Tasks

Previous work on registration tasks has predominantly relied on gradient descent approaches in order to determine the optimum of the chosen similarity measure. The choice of the optimization scheme depends on the characteristics of the similarity measure as well as on its computation; [14] present several concepts for computing probabilistic similarity measures and outline their optimization using gradient descent.

Due to the high problem dimensionality and the lack of directly computable gradients, we propose an evolutionary optimization approach which uses Particle Swarming [3] to search for the global optimum. Moreover, this inherently parallel optimization approach is perfectly suited for an efficient implementation on dedicated hardware; in the present work, we employ the graphics processor (GPU) in a hybrid CPU-GPU implementation.

Registration to Multiple Modalities

When trying to extract information on objects visible in a scene, precision and robustness can be increased by resorting to different sensors working in different modalities and then combining the captured sensor responses.

Most of the existing methods are confined to working on optical imagery, sometimes exploiting different sensor characteristics such as color and polarization [8], as well as multiple channels from the model side, such as normals [9] or contours [13]. Most multimodal applications originate from medical 2D/2D and 3D/3D registration tasks, with the exception of [2, 8] who use MI for tracking 2D templates in images.

In our framework, we employ MI for registering fully textured complex 3D models to different sensor modalities by explicitly simulating and rendering the object appearance for a given sensor modality on the model side not only for optical sensors, but also for *Synthetic Aperture Radar* (SAR) devices.

Paper Outline

The rest of the paper is organized as follows: the problem setting is formulated in section 2.1, in section 2.2 we describe the type of 3D models used and in section 2.3 we present the similarity measure underlying the registration process. Section 2.4 focusses on the multimodal capabilities of our registration scheme, while section 2.5 details the optimization method and its efficient implementation. In section 3, we describe the results of testing our registration scheme on a 3D tracking task using synthetic as well as real sensor data.

2. Proposed Model Registration Scheme

Figure 1 gives an overview on the different parts of the proposed scheme for registering a 3D model to 2D sensor inputs. In the following, we give a formal description of the task and present our novel contributions to each of the components used to address the problem.

2.1. Formal Problem Setting

We suppose that a 3D CAD model M and a vector of sensor modalities $\mathbf{s} = s_1, \dots, s_n$ are given. The set of model modalities $\mathbf{m} = m_1, \dots, m_k$ can be described by

$$m_i(\mathbf{p}, \mathcal{P}_i, M) = \mathcal{P}_i \cdot (R \cdot M + \mathbf{t}), \quad (1)$$

where the 6 extrinsic parameters \mathbf{p} govern the camera-relative rotation R and translation \mathbf{t} of the 3D model M and \mathcal{P}_i specifies the projection rules for each type of modality m_i with known intrinsic parameters. For an optical camera, for example, \mathcal{P}_i is a perspective projection. The number of sensor modalities n and model modalities k can be different.

We now wish to determine for all modalities the common 6 extrinsic parameters \mathbf{p} such that the similarity between the model rendered in each of the k modalities and the n sensor modality inputs is maximized. More formally, we identify the optimal pose $\hat{\mathbf{p}}$ as

$$\hat{\mathbf{p}} = \underset{\mathbf{p}}{\operatorname{argmax}} S(\mathbf{s}, \mathbf{m}(\mathbf{p})) \quad (2)$$

where S represents the similarity measure which will be described in section 2.3.

2.2. Models

We aim at a flexible approach to model registration which does not require the time-consuming generation of models typical of most previous methods. Instead, we resort to using standard 3D CAD models. Given the choice of the similarity measure, our system can consistently register a 3D model even if its appearance does not exactly correspond to the object visible in the scene, since no direct equivalence between model and sensor object is required, as long as a functional relationship exists. This property is achieved by resorting to a probabilistic similarity measure as discussed in the next section.

2.3. Mutual Information as a Measure of Similarity

Mutual Information was first applied to computer vision problems by [14] and [6] and has since become a popular technique both for image registration and feature selection tasks. In the following, we summarize its theoretical basis and describe our own contribution.

Information theory provides a concept for quantizing the amount of new information contained within a signal when interpreted as the possible states of a discrete random variable X , the entropy H

$$H(X) = - \sum_x p_X(x) \log(p_X(x)). \quad (3)$$

Analogously, the joint entropy of two signals can be computed as

$$H(X, Y) = - \sum_x \sum_y p_{XY}(x, y) \log(p_{XY}(x, y)). \quad (4)$$

A strong statistical relationship between the two signals reduces their joint entropy, while a weaker relationship will increase this value. Entropy can thus be used as a means of determining how well one signal describes the other by minimizing the joint entropy of two signals given their separate marginal entropies,

$$MI(X, Y) = H(X) + H(Y) - H(X, Y). \quad (5)$$

This term is called *Mutual Information* (MI); it will tend towards zero for completely unrelated signals and assume its maximum, the sum of the marginal entropies, for statistically identical signals. Its main advantage over other similarity measures such as SSD is that no direct per-value equivalence is required as long as some form of statistical dependency between the signals can be determined. This concept does not, however, take into account the spatial distribution of a signal, but relies exclusively on a per-element

relationship. In our work, we use the normalized version MI_{ECC} of this term as suggested by [6, 12].

[10] assume an underlying normal distribution to be able to directly compute the joint entropy of an arbitrary number d of inputs simultaneously from their covariance matrix Cov_d using

$$H(Cov_d) = \log((2\pi e)^{d/2} \det(Cov_d)^{1/2}). \quad (6)$$

The above term allows to perform the similarity computation on multiple input modalities without requiring costly multidimensional histograms. Moreover, spatial relationships can be incorporated into the entropy estimation by considering neighborhood information instead of single pixels as input [10]. Despite the simplifying assumptions, we have found this approach to produce by far the best results and we have therefore decided to build our registration scheme on this method.

Unlike the registration of entire images, we wish to compute the similarity of a 3D model and a number of images containing input from various sensors. Using the projection of the model into a sensor image, we can determine a precise region of interest which will serve as a mask containing the area of the sensor image relevant for similarity computation. This mask is not precomputed or limited to certain variation models as is the case for most template-based methods, but instead it varies depending on the current pose of the 3D model. We can thus much more effectively limit the support of the similarity computation to the relevant region in the sensor image and reduce computation time.

Due to the neighborhood size chosen in the MI computation, the increased robustness comes with a small tendency for an underfitting of the model. In order to increase alignment precision, we introduce another supporting similarity measure based on edges. We compute precise perspective contours from the available 3D model and determine the fraction of the model contour which matches the edges derived from the sensor input image. This edge weight E is then multiplied with the MI value using

$$S = (MI_{ECC})^w \cdot (E)^{(1-w)} \quad (7)$$

where $w \in [0, 1]$ allows to vary the influence of either of the two components. In section 3, we analyze the gain in precision which can thus be obtained.

2.4. Image Multimodality

In the following, we describe two sensor input modalities which can be used in our multimodal registration scheme.

2.4.1 Optical Intensity Images

Based on a 3D CAD model, a perspective projection camera model with known intrinsic parameters and a set of light

sources, intensity images are rendered. For testing purposes, a background and various noise levels can be added and the resulting intensity images can be used as synthetic input with known ground truth parameters for systematic testing.

2.4.2 Synthetic Aperture Radar (SAR)

Synthetic Aperture Radar (SAR) is a remote sensing technique which exploits post-processing of a sequence of radar signals over time to fusion the echos from points at different distances. As a consequence, an object remains longer inside the beam of the radar and its resolution can be increased in the resulting radar image by merging the separate echos, given the sensor movement is known.

In the present work, SAR images are generated on the model side using a simple algorithm to simulate the SAR specific projection properties. Although we do not simulate speckles and highlights which would require raytracing precision, the quality of the simulated SAR model modalities is sufficient to register 3D models to real SAR sensor input and to recover the 3D pose of the objects in the scene.

To simulate the SAR modality for a given 2D optical model intensity image, we save the original 3D coordinates (x, y, z) for each pixel together with its intensity. Using these 3D coordinates, we then map each pixel to its new 2D SAR screen coordinates (x, z) . If several pixels map to the same screen coordinates, their intensity values are accumulated at this position.

Figure 7 shows the results of the above algorithm (center) as opposed to an industry-grade SAR simulation ([7], left) and the recovered 3D model (right).

2.5. Optimization

Once the sensor inputs and the model renderings for the given modalities are available, we wish to determine the parameter $\hat{\mathbf{p}}$ which maximizes the probabilistic similarity measure S described in 2.3.

The high dimensionality of the parameter vector to estimate, the lack of a closed-form gradient computation, the costly function evaluations and the necessary alignment precision require a careful choice of the optimization method. Another aspect we took into consideration was the potential of the optimization method for an efficient implementation on acceleration hardware. Tests with several traditional optimization approaches, notably a Simplex method, have shown to work reliably only for up to three parameter dimensions.

As a consequence, we propose an evolutionary optimization approach which has consistently displayed an excellent convergence behavior.

2.5.1 Initialization

The initialization of the registration process consists of choosing an approximate location and orientation of the object in the sensor input. In our registration scheme, initialization can be performed manually, alternatively one can resort to using local features for a rough pose estimation or a colour-based region of interest selection. In section 3 we show that our registration scheme is capable of dealing with poor initializations and a relatively large search space.

2.5.2 Particle Swarming

In order to overcome some of the difficulties observed with other optimization methods, an evolutionary optimization strategy was chosen. *Particle Swarm Optimization* (PSO) has been introduced by [3] as a concept adopted from nature, where swarms of animals converge towards the richest feeding grounds by communicating between each other. Classical PSO as implemented in this work boils down to randomly placing S particles as search agents in a parameter space R^n . Each particle can be described by

1. its position in R^n , $x = (x_1, \dots, x_n)$, which corresponds to a pose parameter vector \mathbf{p} ,
2. its velocity $v = (v_1, \dots, v_n)$
3. and the position vector which has yielded the currently best function value, $b = (b_1, \dots, b_n)$.

In addition, randomly chosen subgroups of the particles are linked amongst each other in order to be able to communicate the best position currently found by the members of this subgroup, $g = (g_1, \dots, g_n)$; the link topology changes if no improvement has been made during one time step. The movement of a particle in each dimension d is governed by

$$v_d^{t+1} = \alpha(v_d^t + \text{rand}(0, \beta_1)(b_d^t - x_d^t) + \text{rand}(0, \beta_2)(g_d^t - x_d^t)) \quad (8)$$

$$x_d^{t+1} = x_d^t + v_d^t, \quad (9)$$

where the parameters β_1, β_2 allow varying the influence of the currently known best function values on the particle movement, α is the so-called *constriction factor* to prevent swarm divergence and $\text{rand}(0, \beta)$ adds a random component to the particle behaviour.

A review of the properties of PSO is given in [1]. In our setting, we obtained excellent results with a satisfactory tradeoff between runtime and convergence precision which are detailed in section 3.

2.6. Implementation

Since we are working with 3D CAD models and require a rendering pipeline for the simulation of different sensor

modalities, we have chosen to implement our framework as a hybrid system on the CPU and on the graphics processor (GPU). More specifically, the rendering of optical and SAR modalities on the model side, the contour filtering and the edge distance are computed entirely on the GPU, and the similarity measure described in 2.3 is calculated on the CPU.

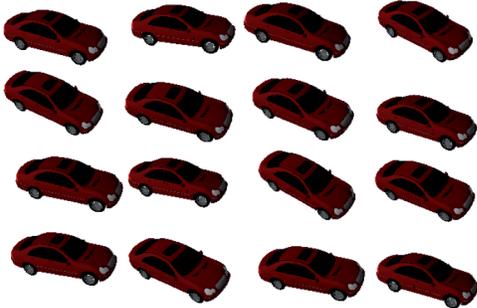


Figure 2. Simultaneously rendered poses of a particle swarm.

Figure 2 shows a set of rendered particles of a swarm, each of which is associated with a distinct model pose. Since there are no interdependencies between the particles in one iteration, their associated model poses can be rendered simultaneously on the GPU and their respective similarity to the sensor input can be computed in parallel. For the experiment shown in figure 3, our framework registers a 3D model with 6 degrees of freedom to a single 2D sensor image of size 256x256 on average in about 2 seconds on a 2GHz Pentium with an Nvidia GeForce 7900 GPU. To obtain these results, convergence is declared when the initial kinetic energy of the swarm has been reduced by more than 90%, approaching the purely stochastically induced level which depends on the chosen PSO parameters. The parameters for the optimization follow the suggestions in [1].

3. Results

The proposed model registration scheme has been systematically tested for different object tracking tasks. We use synthetic scenes with known ground truth trajectories as well as real sensor input using a large number of different models.

A test run typically consists of choosing an initial parameter search space and a model and then starting the optimization process. Since we do not use any explicit movement model, it is necessary to specify the center and the range of the parameter search space; the particles of our optimization will initially be placed arbitrarily inside this search space. The pose for the first frame is initialized manually; for each consecutive frame the last estimated pose is used as initialization.

3.1. Synthetic Sequences

In order to test the registration precision exhaustively, synthetic input sequences for each sensor modality have been generated by defining trajectories for an object visible in the scene as well as for the camera and the light sources. Gaussian blur and speckle noise with different characteristics can then be added to the image data. Figure 3 shows an example image of a synthetically generated scene and its noised version.

In the following, we analyze the influence of noise, initial search space size and edge distance weight on the registration precision. Moreover, we present the results of the model registration to SAR sensor input. In the result plots, we use the standard errors in position and tangential and normal orientation of the model. The position error is given in multiples of the model radius, while the rotation errors are given in degrees.

3.1.1 Noise

A circle was chosen as the ground truth trajectory of the object. The registration process was then initialized on the first of a sequence of synthetic optical images with added noise. Figure 3 shows a part of a recovered trajectory (red) as opposed to the ground truth (blue). We systematically increased the noise which was added to the synthetic sequence to analyze the robustness of the registration process and for each noise level used, we plotted the standard errors in figure 4. The registration process remains stable up to a significant noise level, while the recovery of the rotation parameters appears to be more sensitive towards strong noise than that of the position parameters. Still, registration begins to fail only for noise levels where even the human eye can no longer reliably distinguish the model.

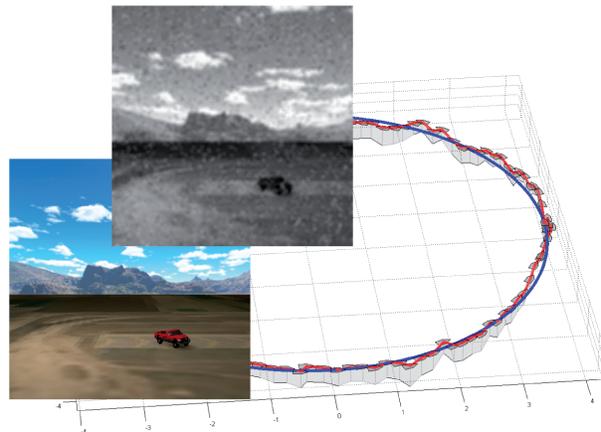


Figure 3. Part of the recovered 125-frame trajectory (red) and ground truth trajectory (blue) for a synthetic optical test sequence with added noise.

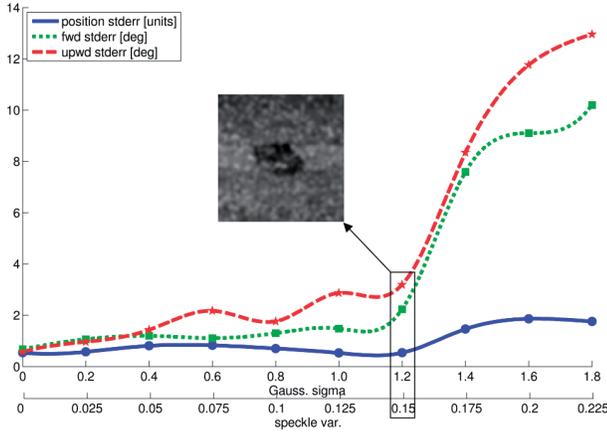


Figure 4. Standard errors (as defined in section 3.1) for a registration sequence of 125 frames (deviation from ground truth position (blue) and tangential (green) and normal vectors (red)) for increasing noise; the optical image of one noise level is shown for visualization.

3.1.2 Search Space Size

In section 2.5.2, we outlined the optimization scheme used. The initial search space size is of crucial importance for a successful registration since the particles will tend to get dispersed inside a search space which is chosen too large. As a result, the global optimum might not be found reliably any more. In figure 5, we show that, while keeping the swarm size constant, the search space can be safely increased to $\pm 15^\circ$ for rotation parameters and $\pm 40\%$ of the model size for translational parameters without sacrificing an inadmissible amount of registration precision. The result shows that the chosen evolutionary optimization scheme is surprisingly stable even for imprecise initializations.

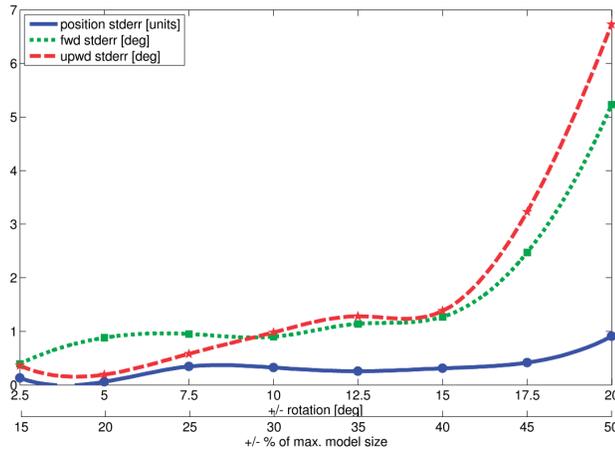


Figure 5. Standard errors (as defined in section 3.1) for a registration sequence of 125 frames while increasing the initial search space size.

3.1.3 Varying the Contribution of the Edge Distance

In section 2.3, we outlined a contour-based extension to the MI similarity measure, allowing to vary the influence of either of the two components as a function of the parameter w . In systematic tests, we have found an optimal choice to be about $w = 0.5$, which corresponds to the geometric mean. For this choice, the three standard errors are minimal in the majority of test settings, while for $w \rightarrow 0.0$ signifying an exclusive use of the edge distance component and likewise for an exclusive use of the MI measure with $w \rightarrow 1.0$, stability and registration precision deteriorate significantly. To illustrate the gain in stability, we registered a model to 605 frames of an input sequence using different w values. In figure 6, the frames are shown for which the tracking failed for the first time. When using only the edge distance ($w \rightarrow 0.0$), registration failed after 24 frames, when using only the MI measure ($w \rightarrow 1.0$), precision was inadmissible after 117 frames showing a characteristic underfitting of the model. The combined measure with $w = 0.5$ completed the sequence successfully and produced precise results.

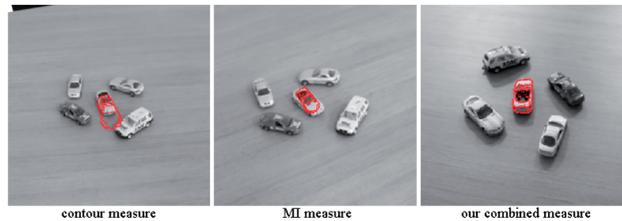


Figure 6. Tracking with different w values: using only contour matching, tracking fails after a few frames (left), the MI measure fails after 117 frames (center), the weighted combination handles the full sequence successfully (right).

3.1.4 SAR

Using the simulated SAR response, we performed registration on SAR input data. Figure 7 shows two frames of a SAR sequence, with the input images from an industry-grade SAR simulator ([7]) on the left, the registered SAR signature in the center and the recovered optical 3D model corresponding to the SAR signature on the right. Despite the simple SAR modality generation used, the results show a precision comparable to the optical sequences.

3.2. Video Sequences

In order to verify the suitability of the proposed registration scheme for realistic sensor input, we registered different models to videos taken with a standard interlaced Sony camera. We did not perform any preprocessing on the input data and did not remove lens distortion, motion blur and interlacing artifacts. The tested sequences displayed fully perspective changes and fast movements. Since we recover for each input frame the full 6 pose parameters of the object, we can reconstruct the camera trajectories for the input

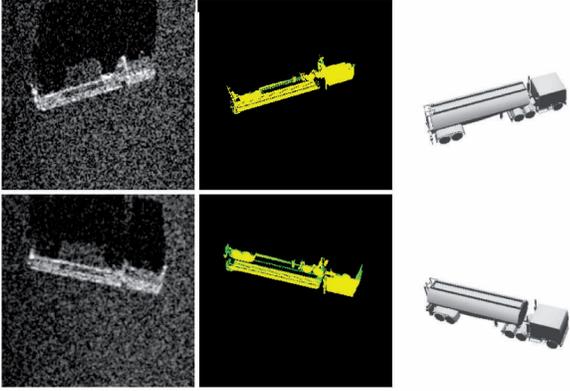


Figure 7. 2 frames from a tracking sequence on SAR input images (left, [7]), the registered SAR model modality (center) and the recovered optical 3D model (right). Images are scaled for better visualization.

sequences when supposing that the only movement in the scene stems from the camera.

Figure 8 shows some images taken from a 654-frame video of a set of small toy cars on an indoor office table with artificial lighting. For each frame, we show the input image with the reprojected model edges on the left, and the recovered 3D model pose on the right. In figure 9, the spline-interpolated camera trajectory for the same scene is plotted together with the model position and the camera position and orientation for each of the 4 images shown in figure 8. The average size of the toys in this sequence varies between 60x40 and 120x80 pixels.

In figure 10, we present more results taken from different input sequences. The topmost frame is taken from another indoor sequence of a different toy car, while the next two images show real cars in an outdoor setting. Due to the chosen robust similarity measure, significant occlusions of the object can be sustained as illustrated in the third image of figure 10. The registration is also suitable for textured objects with shiny surfaces, even in complex settings, as can be seen on the last image of figure 10 for an indoor recording of a soda can. Model sizes in the tested sequences vary from 80x60 to 160x120 pixels with video resolutions of 256x256 and 512x512 pixels. For non-matching objects, the value of the similarity measure after convergence ranges an order of a magnitude below the result for a correct match, thus allowing to determine whether the choice of the model for registration to a given scene ought to be reconsidered.

4. Conclusion

We have proposed a simple but efficient similarity measure derived from information theory and perspective contour matching. In conjunction with a robust and precise evolutionary optimization strategy, we have been able to achieve excellent results for the registration of 3D models



Figure 8. 4 frames from a 654-frame sequence with toy cars; input image with recovered object (left), recovered 3D model (right).

to 2D sensor inputs of different modalities. By exploiting the hardware acceleration potential of the GPU and by using standard 3D CAD models, we outperform many of the existing registration approaches as regards precision, speed and universality. Further work will be dedicated to integrating this method into an extensive framework for object recognition tasks.

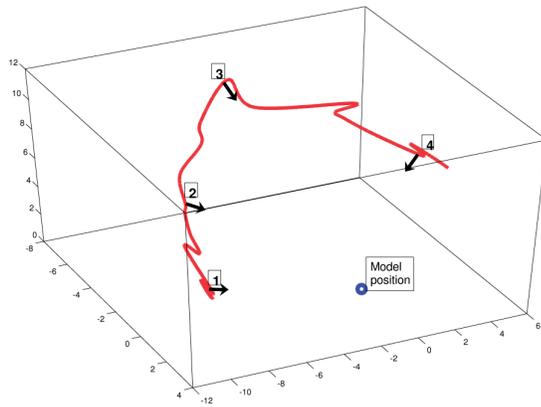


Figure 9. Recovered camera trajectory of the tracking sequence 8 with the camera positions of the 4 frames shown in the figure.

Acknowledgements We would like to thank Christoph Stahl, Peter Knappe (EADS Military Air Systems) and Christoph Neumann (EADS Defence Electronics) for their valuable support.

References

- [1] J. Clerc, M. Kennedy. The particle swarm - explosion, stability, and convergence in a multidimensional complex space. *Trans. on Evolutionary Computation*, 6:58–73, 2002.
- [2] N. Dowson and R. Bowden. Metric mixtures for mutual information (m3i) tracking. In *Proc. 17th Int. Conf. on Pattern Recognition*, pages 1051–1061, 2004.
- [3] J. Kennedy and R. Eberhart. Particle swarm optimization. In *Proc. Int. Conf. on Neural Networks*, volume IV, pages 1942–1948, 1995.
- [4] H. Lensch, W. Heidrich, and H.-P. Seidel. Automated texture registration and stitching for real world models. *ACM Graphical Models*, 63:245–252, 2001.
- [5] D. G. Lowe. Robust model-based motion tracking through the integration of search and estimation. *Int. Journal of Computer Vision*, 8:113 – 122, 1992.
- [6] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens. Multimodality image registration by maximization of mutual information. *Trans. on Medical Imaging*, 16:187–198, 1997.
- [7] J. Meyer-Hilberg. Pirdis: A new versatile tool for sar/mti systems simulation. In *6. Europ. Conf. on Synthetic Aperture Radar (EUSAR2006)*, 2006.
- [8] J. L. Mundy and C.-F. Chang. Fusion of intensity, texture, and color in video tracking based on mutual information. In *Proc. 33rd Appl. Imagery Patt. Recog. Workshop*, 2004.
- [9] H.-K. Pong and T.-J. Cham. Alignment of 3d models to images using region-based mutual information and neighborhood extended gaussian images. In *Proc. Asian Conf. for Computer Vision*, pages 60–69, 2006.
- [10] D. Russakoff, C. Tomasi, T. Rohlfing, and C. Maurer. Image similarity using mutual information of regions. In *Proc. Europ. Conf. on Comp. Vision*, pages 596–607, 2004.

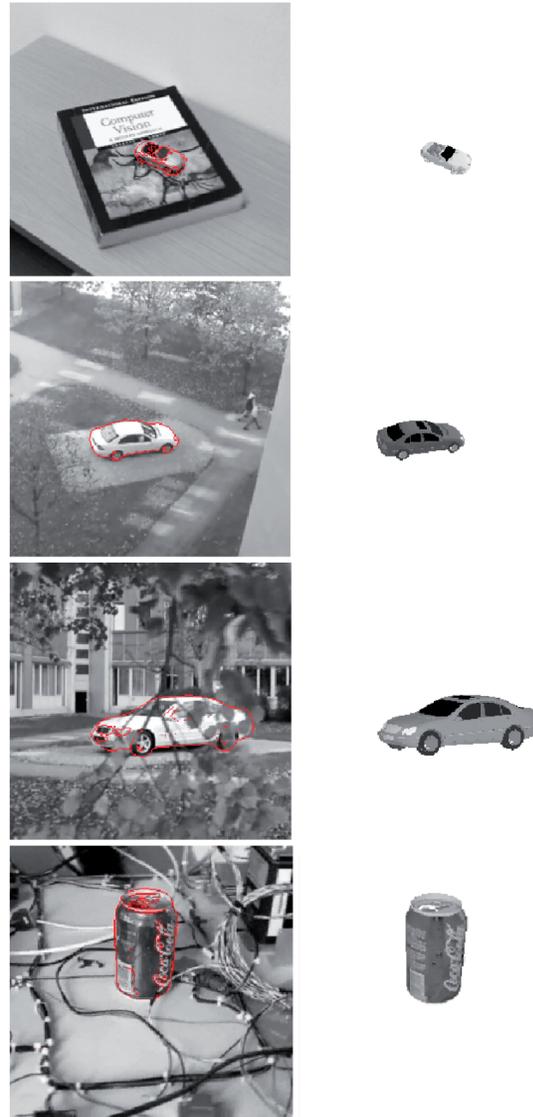


Figure 10. Selected frames from four different registration sequences and the recovered 3D models.

- [11] C. Studholme, D.L.G.Hill, and D.J.Hawkes. Incorporating connected region labelling into automated image registration using mutual information. In *Proc. Workshop on Math. Methods in Biomed. Image Analysis*, pages 23–31, 1996.
- [12] C. Studholme, D. Hill, and D. J. Hawkes. Overlap invariant entropy measure of 3d medical image alignment. *Pattern Recognition*, 32(1):71–86, 1999.
- [13] I. Suveg and G. Vosselman. Mutual information based evaluation of 3d building models. In *Int. Conf. Pattern Recognition*, volume 3, pages 557–560, 2002.
- [14] P. Viola and W. M. Wells. Alignment by maximization of mutual information. *Int. Journal of Computer Vision*, 24(2):137–154, 1997.
- [15] S. Yamany and A. Farag. Free-form surface registration using surface signatures. In *Proc. 7. Int. Conf. on Computer Vision*, volume 2, pages 1098–1104, 1999.