

CORRELATION ICP ALGORITHM FOR POSE ESTIMATION BASED ON LOCAL AND GLOBAL FEATURES

Marco A. Chavarria and Gerald Sommer

Cognitive Systems Group. Christian-Albrechts-University of Kiel, D-24098 Kiel, Germany

Keywords: Pose estimation, ICP algorithm, monogenic signal, pre-alignment, global and local features.

Abstract: In this paper we present a new variant of ICP (iterative closest point) algorithm based on local feature correlation. Our approach combines global and local feature information to find better correspondence sets and to use them to compute the 3D pose of the object model even for the case of large displacements between model and image data. For such cases, we propose a 2D alignment in the image plane (rotation plus translation) before the feature extraction process. This has some advantages over the classical methods like better convergence and robustness. Furthermore, it avoids the need of a normal pre-alignment step in 3D. Our approach was tested on synthetical and real-world data to compare the convergence behavior and performance against other versions of the ICP algorithm combined with a classical pre-alignment approach.

1 INTRODUCTION

The estimation of the object position is crucial for an efficient robot-object interaction in research and industrial applications. In this context, geometric algebra has been introduced in computer vision as a problem adaptive algebraic language for modeling geometric related problems, see (Sommer, 2001). This mathematical framework was used by (Rosenhahn and Sommer, 2005) to formulate the monocular pose estimation problem and the model representation.

For every model-based pose estimation or 3D registration algorithm, correspondences must be found between model and acquired data. This is one of the most challenging problems for this kind of approaches. The most common and simple solution is the ICP algorithm introduced by (Besl and McKay, 1992), where the minimal Euclidean distance constraint is used to find correspondences. A comparison of several variants of the ICP algorithms is presented by (Rusinkiewicz and Levoy, 2001), where the original ICP is combined with different distance metrics and strategies to find correspondences and to align artificially generated 3D meshes. A tracking algorithm based on template and features data was presented by (Ladkos et al., 2007), where the systems changes adaptively between templates and features to deal with complex tracking scenarios. The above cited methods assume *tracking assumption* conditions. It means that the displacement between model

and scene is small enough to avoid convergence to a local minimum. The approach proposed by (Shang et al., 2007) uses known information about the limits of the object velocity and the image frame rate to reduce the space transformation between every frame. Then, the tracking is transformed into a classification problem. In the work of (Sharp et al., 2002) the ICP algorithm is combined with additional invariant features like curvature, moment invariants and spherical harmonics for registration of range images. In (Chavarria and Sommer, 2007), structural information from image and model (convexity, concavity and straightness of segments) is used as extra correspondence search constraints for monocular pose estimation. The combination of ICP with such structural features reduces the probability of being trapped in a local minima. That means the algorithm is robust against the tracking assumption up to certain limits.

For larger displacements between model and scene, a pre-alignment step is needed to get an initial rough estimation of the pose. Then, the ICP algorithm can be applied to obtain the final pose. In the work of (Brujic and Ristic, 1996), a pre-alignment based on the principal component analysis (PCA) is used. Once that the principal components of the sets of points are computed, the rough pose is obtained by finding the 3D pose that aligns these main components. Instead of extracting the main components by PCA, (Murino et al., 2001) align the extracted 3D skeletons of the model. A genetic algorithm is used

in the work of (Lomonosov et al., 2006) to compute the pre-alignment step. All these methods are based on the alignment of the main components or skeletons extracted from sets of 3D points. That means, the model points and the acquired data are defined in 3D. This is not the case for the monocular pose estimation, where the model is defined in 3D, but the acquired data (contour data) is defined in the 2D image plane. In order to perform a pre-alignment for the monocular pose estimation, the components of the data extracted from the image must be reconstructed in 3D. Another option is to project the 3D model components onto the image plane and apply a 2D version of the above cited methods.

In this paper we present a new correlation-based ICP algorithm for the monocular pose estimation of 3D free-form surfaces. It combines global and local orientation features in an approach that performs better even for cases where the classical versions of the ICP algorithm fail (e.g. when the tracking assumption is not met). Local orientation information in the image computed from the monogenic signal response (Felsberg and Sommer, 2004) and from the projected models are used to describe the orientation of contour segments for image and model points. Instead of the minimal distance criteria used by the normal variants of the ICP algorithm, correlation is used to measure the similarity (in terms of the local orientation) of the contour segments. Furthermore, global orientation is used to align the projected model data with the detected contour features in the image plane. That means, a simple 2D feature alignment is performed. This allows to find better conditioned correspondence sets even for larger displacements (rotations and translations) without the need of an extra pre-alignment step in 3D.

This paper is organized as follows, the image feature extraction based on the monogenic signal is briefly introduced in section 2. Section 3 describes the global feature extraction for image and projected model contours. The silhouette based pose estimation and the correlation based ICP algorithms are presented in section 4. The 2D feature alignment procedure follows in section 5. Finally, the result of several experiments made on artificial and real-world data to validate the efficiency and robustness of our algorithm are presented in section 6.

2 LOCAL IMAGE FEATURES

The monogenic scale-space representation and phase-based image processing techniques were introduced by (Felsberg and Sommer, 2004). If $p(\mathbf{x};s)$ and $q(\mathbf{x};s)$

are the filter responses of an image convolved with the Poisson and conjugate Poisson kernels respectively, local amplitude $a(\mathbf{x};s)$ and local phase $\mathbf{r}(\mathbf{x};s)$ are obtained for a scale s according to

$$\begin{aligned} a(\mathbf{x};s) &= \sqrt{|q(\mathbf{x};s)|^2 + |p(\mathbf{x};s)|^2} \\ \mathbf{r}(\mathbf{x};s) &= \frac{q(\mathbf{x};s)}{|q(\mathbf{x};s)|} \arctan\left(\frac{|q(\mathbf{x};s)|}{|p(\mathbf{x};s)|}\right). \end{aligned} \quad (1)$$

The amplitude is related to the local energy of the signal (presence of structure). The orientation and phase are combined in the phase vector. The phase gives information about the local symmetry of the signal (type of structure) and the orientation gives the direction of the highest signal variance. Once that the amplitude and phase are obtained for a scale factor s , a contour search algorithm based on the local phase and orientation is applied to extract the contour segments. By changing the scale factor, low contrast edges can also be detected. Thus, for every contour point we get as features its coordinates in the image and the local orientation $F_i^{im} = \{x_i^{im}, y_i^{im}, \alpha_i^{im}\}$.

3 GLOBAL FEATURE EXTRACTION

The Fourier transform of a closed contour delivers a set of complex coefficients that can be used to obtain low pass approximations of it. If only the first coefficient is used, the contour is approximated by a circle. As the number of the coefficients increases, a better approximation of the contour is obtained. We use a 2D real valued variant of this approach introduced by (Lin and Hwang, 1987). Then, a closed contour $\mathbf{c}(t) = f^1(t)\mathbf{e}_1 + f^2(t)\mathbf{e}_2$ in the image plane is approximated by

$$\begin{bmatrix} f^1(t) \\ f^2(t) \end{bmatrix} = \begin{bmatrix} a_0 \\ d_0 \end{bmatrix} + \sum_{k=0}^{N-1} \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} \begin{bmatrix} \sin(\frac{2k\pi t}{N}) \\ \cos(\frac{2k\pi t}{N}) \end{bmatrix}, \quad (2)$$

with

$$\begin{aligned} a_0 &= \frac{1}{N} \sum_{k=0}^{N-1} f^1(t) \\ d_0 &= \frac{1}{N} \sum_{k=0}^{N-1} f^2(t) \\ a_k &= \frac{1}{N} \sum_{k=0}^{N-1} f^1(t) \cos\left(\frac{2k\pi t}{N}\right) \\ b_k &= \frac{1}{N} \sum_{k=0}^{N-1} f^1(t) \sin\left(\frac{2k\pi t}{N}\right) \\ c_k &= \frac{1}{N} \sum_{k=0}^{N-1} f^2(t) \cos\left(\frac{2k\pi t}{N}\right) \\ d_k &= \frac{1}{N} \sum_{k=0}^{N-1} f^2(t) \sin\left(\frac{2k\pi t}{N}\right). \end{aligned} \quad (3)$$

The coefficients a_k, b_k, c_k, d_k are called *elliptic coefficients*. If only the first coefficient is used, the contour will be approximated by an ellipse. If the coefficients are arranged in a matrix, the parameters of such ellipse, which approximates the contour, are extracted




Figure 1: Examples of the extracted global orientation. Upper figures: projected surface models onto the image plane (triangle, motor part and power socket) and example of real-world images. Bottom figures: extracted contours and the corresponding major and minor axes.

by single value decomposition according to:

$$\text{svd} \begin{bmatrix} a_k & b_k \\ c_k & d_k \end{bmatrix} = \begin{bmatrix} \cos \theta_k & -\sin \theta_k \\ \sin \theta_k & \cos \theta_k \end{bmatrix} \begin{bmatrix} A_k & 0 \\ 0 & B_k \end{bmatrix} \mathbf{R}(\psi). \quad (4)$$

The indexes A_k and B_k are the lengths of the major and minor axes of the ellipse and θ_k is the angle of the major axis with respect to the image axis x . The phase matrix $\mathbf{R}(\psi)$, defines the angular distance from the main axis to the first point of the contour. Thus, for a projected model or a detected image contour, we obtain as global features $F_{img} = \{\mathbf{p}_{img} = [a_0, d_0], \theta_k\}$, which defines the global position and orientation of the contour in the image plane. Some examples for projected models and real images can be seen in figure 1. In the practice, this method of computing the mayor and minor axis is very similar than performing a PCA analysis on the point covariances. Despite of that, we apply the elliptical descriptors because of the ability to compute low pass approximations of the contours. This will eventually allows to avoid the local minima problem in some specific scenarios, see (Rosenhahn et al., 2003).

4 CORRELATION BASED ICP ALGORITHM

An algorithm for pose estimation of 3D surfaces models was proposed by (Rosenhahn et al., 2003), where the 3D silhouette of the model is extracted for every iteration of the minimization process. Originally, the classical ICP algorithm was applied to find the pose of the silhouette. The position of the complete surface model is updated and the process is repeated for a given number of iterations. We use a similar idea, but in our approach the 3D silhouette is projected onto the image plane and its global and local features are




Figure 2: Algorithm for the silhouette based pose estimation.

computed in 2D. The algorithm is summarized in figure 2.

4.1 Correlation as Similarity Criteria

Instead of the Euclidean metric used by the original ICP algorithm, we use the correlation as a similarity measure to find correspondences. For a contour segment or range n around a projected model point \mathbf{x}_i , $\{\mathbf{x}_{i-n}, \dots, \mathbf{x}_{i+n}\}$, we define a vector containing the local orientation values of this segment as $\mathbf{o}_i^{mod} = \{\alpha_{i-n}^{mod}, \dots, \alpha_{i+n}^{mod}\}$. We call it orientation profile vector and the local orientation is directly computed from the projected model points. Similarly, an orientation profile vector of an image contour segment (computed from the monogenic signal response) is defined as: $\mathbf{o}_i^{img} = \{\alpha_{i-n}^{img}, \dots, \alpha_{i+n}^{img}\}$. The similarity of these profile vectors can be computed by the correlation matrix as

$$\text{corr}(\mathbf{o}_i^{img}, \mathbf{o}_j^{mod}) = \frac{\text{cov}(\mathbf{o}_i^{img}, \mathbf{o}_j^{mod})}{\sqrt{V_{img} V_{mod}}}, \quad (5)$$

where $\text{cov}(\mathbf{o}_i^{mod}, \mathbf{o}_j^{img})$ is the covariance matrix and V_{mod} , V_{img} are the variances of image and model data. The correlation may vary in a range between -1 and 1, where -1 indicates perfect negative correspondence, 0 indicates no correspondence and 1 indicates perfect correspondence. For an image point with orientation profile \mathbf{o}_i^{img} , the corresponding model point (with orientation profile $\mathbf{o}_j^{mod} \in MOD$) will be the one with maximal correlation,

$$\text{corr}(\mathbf{o}_i^{img}, MOD) = \max_{j=1, \dots, n} \{\text{corr}(\mathbf{o}_i^{img}, \mathbf{o}_j^{mod})\}. \quad (6)$$

4.2 Outlier Elimination

Noise in the image and the presence of partial occlusions may cause not well conditioned correspondences that must be eliminated. To achieve that, the angular positions of image contour θ_i^{img} and model points θ_j^{mod} with respect to the major axis are used as an extra feature, see figure 3. Once that a candidate




Figure 3: Left: 2D feature alignment based on the difference of global orientation angles and position in the image plane. Angular position of a contour point with respect to the major axis (middle). Interval within the pair is considered to be a good conditioned correspondence (right).

correspondence pair has been found, it is rejected if the following criterion is fulfilled

$$\|\theta_i^{img} - \theta_j^{mod}\| > t, \quad (7)$$

where t is a given threshold value that defines the interval within which the correct correspondence must be. The introduction of this criterion in the algorithm reduces considerably the number of correspondence pairs as it can be seen in the examples of figure 7.

5 FEATURE ALIGNMENT IN 2D

The principle of the 2D feature alignment is illustrated in figure 3. The dotted object represents the projected model and the solid object the detected object in the image with their corresponding main axes with global features $F_{mod} = \{\mathbf{p}_{mod}, \theta_{mod}\}$ and $F_{img} = \{\mathbf{p}_{img}, \theta_{img}\}$ respectively. First, the projected model points \mathbf{x}_i are aligned (rotated and translated in 2D) by the matrix $\mathbf{T}(\phi, \mathbf{t}_r)$. As can be seen in the figure, the angle ϕ is the orientation difference of the major axes and \mathbf{t}_r the translation vector between the respective centers of mass. Once that the projected model points have been aligned to the detected image points, the local features (orientation profiles) are computed and the correlation-based ICP algorithm is applied.

The effect of aligning the features is visualized in figure 4. The graphics show the orientation profiles of two corresponding points obtained with the simple correlation and with the aligned features. With aligned features the profiles are more similar and therefore better conditioned correspondences are found.

Let us notice that in contrast with the classical pre-alignment approaches where the rough pose is computed in 3D by a minimization approach, the rotation and translation in our approach are computed directly from the global orientation and position differences.




Figure 4: Orientation profiles of two corresponding points with simple correlation (upper graphics) and with additional feature alignment (lower graphics), and examples of correspondences between an object model and image contour.

6 EXPERIMENTS

The robustness against the tracking assumption (large rotations and translations) was tested and compared with the normal variant of ICP algorithm (see (Rusinkiewicz and Levoy, 2001)) and with the structural ICP algorithm (Chavarría and Sommer, 2007). In a second experiment, our ICP variant was compared with an approach based on the PCA pre-alignment used in (Bruzic and Ristic, 1996) and (Murino et al., 2001). In both cases the initial position of the model is known, then it is translated to a certain position (ground truth) and projected onto the image plane to generate an artificial image. On this artificial image the local and global features are extracted. The pose is calculated with the projective pose estimation algorithm of (Araujo et al., 1998) and compared with the ground truth. Several object models were used for our experiments (see figure 1). Finally, some examples for a real data scenario are presented.

6.1 Robustness Against Rotations

From the initial position of the model, its main orientation axes were extracted in 3D. Each axis defines a rotation axis α , β and γ , as it can be seen in figure 5. The model was rotated around these axes and in its new position the pose was computed and compared with the ground truth. The left graphics of figure 5, show a comparison of the convergence behavior for the motor and power socket models. In this case the model was rotated -30 degrees around the γ axis. The normal ICP algorithm does not converge to the ground truth pose as it can be seen in the graphics. In contrast to the structural ICP variant, the number of iterations needed to converge is significantly reduced with the correlation based ICP.

As the model rotates around the 3D axes, its appearance changes with respect of the image plane and therefore its local structure. Because of that, it is interesting to analyze the rotation ranges within which the algorithm is capable to converge to the

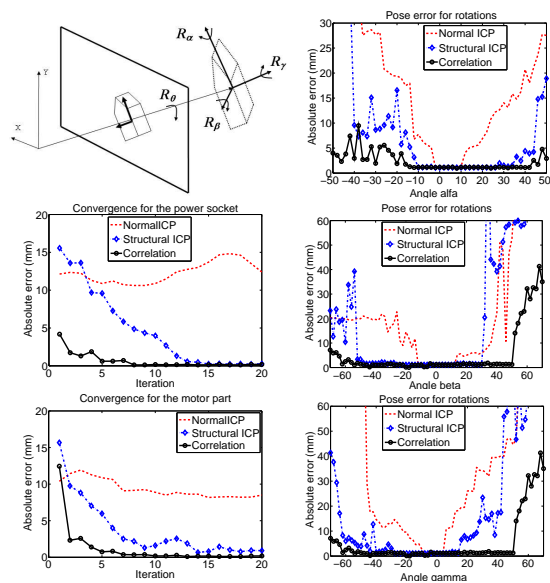


Figure 5: Left: Setup for the experiment for the rotation case (upper figure) and convergence behavior comparison for the power socket (middle figure) and motor part (bottom figure). Right: convergence ranges for rotations around the axes α (upper figure), β (middle figure) and γ (bottom figure).

ground truth pose. The figure 5 also shows a comparison of the convergence ranges for the power socket model. The normal variant of the ICP algorithm converges for relatively small rotations around all axes. Whereas the structural and correlation based ICP variants allow larger rotations. In the case of the rotation axes α and γ the extracted silhouette changes drastically with respect of the image plane as the angle increases, therefore the region where the algorithm converges is smaller. Despite of that, the correlation-based ICP variant shows larger convergence ranges than the structural and classical variants.

6.2 Pre-alignment Comparison

The PCA based pre-alignment algorithms ((Bruzic and Ristic, 1996) and (Murino et al., 2001)) align clouds of points in 3D by aligning its main axes. This implies that correspondences between the axes must be found. With these correspondences, the alignment is computed by a minimization process that takes in general several iterations. In contrast to that, the correlation algorithm combined with the 2D feature alignment needs only one iteration to compute a rough pose. Thus, in the next experiments we compare the first iteration of the correlation ICP algorithm against the pre-alignment based on the principal component analysis. Since we align model 3D data with



Figure 6: Pre-alignment comparison of PCA and 2D alignment. Initial position (upper row), PCA pre-alignment result (middle row) and computed pose after the first iteration of the correlation-based ICP (lower row).

Table 1: Error comparison for the pre-alignment step (motor part model).

Rotation (degree)	Absolute error (mm)	
	PCA	CORRELATION
0	14.9172	3.1682
10	15.8307	3.4555
20	17.0352	3.4002
30	18.6300	2.7695
40	20.5412	6.6615
50	22.6003	5.2424

image detected contours, we use a version of this approach which aligns the model in 3D by aligning the principal components in the image plane.

The model was translated and rotated around the γ axis from 0 to 50 degrees, see upper row of figure 6. The middle row shows the result of the pre-alignment with the PCA and the bottom row shows the result after the first iteration with the correlation-based ICP. Let us remember that the PCA alignment version used for the monocular pose estimation aligns the axes of the detected image contour and the projected model silhouette. In contrast to the 3D case, only two principal axes can be extracted and aligned in the image plane. This loss of information has the effect that, although the major and minor axes are roughly aligned, the error in 3D is significantly larger. On the other hand, the result of the correlation based ICP is considerably better than that of the PCA based algorithm. The table 1 shows the absolute error and the difference angle of the major orientation axis in 3D respectively.

6.3 Pre-alignment with Partial Occlusions

As described in the last experiments, an artificial image was generated where some partial occlusions were simulated. This can be seen in the figure 7. Correspondences were found with the correlation ICP algorithm with the outlier elimination criteria of equation (7). Additionally, the threshold value t was varied to see how it affects in the convergence behavior of the algorithm. The results of this experiment are shown in figure 8. The graphic of the left shows the convergence behavior for different threshold values (from 10 to 50 degrees). A better convergence is achieved for threshold values between 20 and 30 degrees. The right graphic shows the number of correspondences for every iteration of the algorithm after applying the outlier elimination criteria of equation (7). From initially 96 pairs, the number of correspondences is considerably reduced during the iterations.




Figure 7: From left to right: Artificial generated image with simulated occlusions. Initial position for the experiment. Correspondences with only the correlation ICP algorithm. Correspondences after eliminating the outliers.

The presence of partial occlusions affects the local structure of the contour and therefore the orientation profiles of each point. If only the correlation correspondence search is applied, a big number of ill conditioned correspondences are found (see figure 7). Once they are eliminated, a reduced set of correspondences is obtained. Despite of that, the remaining pairs are the best conditioned correspondences and the pose can be computed with them.




Figure 8: Left graphic: convergence behavior of the algorithm for the occlusion case. Right graphics: number of correspondences for each iteration.




Figure 9: Experiments in real images including the partial occlusions case. Initial position of the model (left column) and computed pose (right column).

6.4 Experiments on Real World Images

Finally, we applied our algorithm to real image data. A single calibrated camera system providing gray value images of 620×540 pixels is used. The algorithm was tested on a Linux based system with a 3 GHz Intel Pentium 4 processor. Figure 9 shows some examples taken from different test sequences and examples with partial occlusions in the image. The average computing time per frame for the image processing module (contour extraction, global and local feature extraction) was 224 milliseconds and the complete pose calculation process (image processing plus pose estimation) was 4.73 seconds. Additionally to the outlier elimination criterion of equation (7), the Euclidean distance criterion used by (Masuda et al., 1996) was combined with the outlier elimination of equation (7). Correspondence pairs are rejected if their point-to-point distance is larger than 2.5 times the standard deviation of the complete correspondence set.

7 CONCLUSIONS AND FUTURE WORK

An ICP algorithm based on feature correlation for pose estimation of 3D surfaces was presented. The experimental results show that our approach performs more efficiently than the normal and structural ICP variants. It also shows better convergence behavior, which reduces the probability of being trapped in a local minimum during the minimization process. An important feature of our approach is that in the first iteration of the process, the pose error is smaller than that of the PCA based pre-alignment step. The experiments show the convergence limits of the algorithm when only one camera is available. The integration of an additional camera would increase the view range over the object and therefore the convergence ranges. The computation of local and global features in every iteration and the 3D silhouette extraction step increase the computation time of the algorithm. Real time is not reached with our approach, but the reported computation time is a good tradeoff considering the robustness of the algorithm. A natural extension for our approach is to adapt the correlation ICP algorithm and combine it with the structural ICP variant in a system which deals with more complex scenarios like more general object occlusions, local model deformations, illumination changes or similar.

REFERENCES

- Araujo, H., Carceroni, R., and Brown, C. (1998). A fully projective formulation to improve the accuracy of Lowe's pose-estimation algorithm. *Comput. Vis. Image Underst.*, 70(2):227–238.
- Besl, P. and McKay, N. (1992). A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256.
- Brujic, D. and Ristic, M. (1996). Analysis of free form surface registration. In *International Conference of Image Processing. Lausanne, Switzerland*, pages 393–396.
- Chavarria, M. and Sommer, G. (2007). Structural icp algorithm for pose estimation. In *2nd International Conference on Computer Vision Theory and Applications, VISAPP 2007, March 8-11, Barcelona, Spain*, pages 341–346, Portugal. INSTICC.
- Felsberg, M. and Sommer, G. (2004). The monogenic scale-space: A unifying approach to phase-based image processing in scale-space. *Journal of Mathematical Imaging and Vision*, 21(1):5–26.
- Ladkos, A., Behimane, S., and Navab, N. (2007). A real-time tracking system combining template-based and feature-based approaches. In *2nd International Conference on Computer Vision Theory and Applications, VISAPP 2007, March 8-11, Barcelona, Spain*, pages 325–332, Portugal. INSTICC.
- Lin, C.-S. and Hwang, C.-L. (1987). New forms of shape invariants from elliptic fourier descriptors. *Pattern Recognition*, 20(5):535–545.
- Lomonosov, E., Chetverikov, D., and Ekárt, A. (2006). Pre-registration of arbitrarily oriented 3d surfaces using a genetic algorithm. *Pattern Recognition Letters*, 27(11):1201–1208.
- Masuda, T., Sakaue, K., and Yokoya, N. (1996). Registration and integration of multiple range images for 3-d model construction. In *ICPR '96: Proceedings of the 1996 International Conference on Pattern Recognition (ICPR '96) Volume I*, page 879, Washington, DC, USA. IEEE Computer Society.
- Murino, V., Ronchetti, L., Castellani, U., and Fusiello, A. (2001). Reconstruction of complex environments by robust pre-aligned ICP. In *Proc. Third Int. Conf. on 3-D Digital Imaging and Modeling*, pages 187–194.
- Rosenhahn, B., Perwass, C., and Sommer, G. (2003). Pose estimation of free-form surface models. In Michaelis, B. and Krell, G., editors, 25. *Symposium für Mustererkennung, DAGM 2003, Magdeburg*, volume 2781 of LNCS, pages 574–581. Springer-Verlag, Berlin.
- Rosenhahn, B. and Sommer, G. (2005). Pose estimation in conformal geometric algebra, part II: Real-time pose estimation using extended feature concepts. *Journal of Mathematical Imaging and Vision*, 22:49–70.
- Rusinkiewicz, S. and Levoy, M. (2001). Efficient variants of the ICP algorithm. In *Proceedings of the Third Intl. Conf. on 3D Digital Imaging and Modeling*, pages 145–152, Quebec City, Canada.
- Shang, L., Jasiobedzki, P., and Greenspan, M. (2007). Model-based tracking by classification in a tiny discrete pose space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(6):976–989.
- Sharp, G., Lee, S., and Wehe, D. (2002). ICP registration using invariant features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(1):90–102.
- Sommer, G., editor (2001). *Geometric Computing with Clifford Algebras*. Springer-Verlag, Heidelberg.