Transformed Polynomials for Global Registration of Point Clouds

Ruediger Schmedding* Barbara Frank† Wolfram Burgard‡ Matthias Teschner§

Abstract

In this paper, we introduce a novel approach for global registration of partially overlapping point clouds. The approach identifies feature points of matching objects based on surface-approximating polynomials and finds an initial transformation depending on these polynomials. We compute an extended set of rotationally-invariant features for polynomials. In contrast to purely feature-based approaches, we do not only compute transformations based on the invariant properties of polynomials, but actually transform the polynomials into a common coordinate system and compare the transformed coefficients. This results in an improved correspondence analysis of local surfaces. Hence, using transformed polynomials, we gain more discriminating information about different structures. Therefore, the approach can handle partial scans of different objects simultaneously. Each partial scan is assigned to one of the objects and registered accordingly. Moreover, the approach is robust against noise and can process real data.

CR Categories: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

Keywords: registration, correspondence, shape analysis, matching

1 Introduction

Registration, correspondence and matching algorithms are of high importance in the area of Computer Graphics. They are applied in various research fields such as object reconstruction, shape retrieval, and symmetry detection. The general task is to find matching parts of different objects with unknown relative orientations. These could be several scans of an object that should be aligned in a common coordinate system [Rusinkiewicz et al. 2002; Gelfand et al. 2005; Aiger et al. 2008], or a partial scan of an object which is used to identify a corresponding model in a given library [Kazhdan et al. 2004; Pauly et al. 2005; Gal and Cohen-Or 2006; Shalon et al. 2008]. Symmetry detection approaches look for similar sub-parts of objects [Mitra et al. 2006; Mitra et al. 2007; Pauly et al. 2008; Raviv et al. 2010]. A recent survey can be found in [van Kaick et al. 2010].

All these tasks can be tackled in a similar manner by looking for a transformation between pairs of objects to align them properly.

*schmedd@informatik.uni-freiburg.de
†bfrank@informatik.uni-freiburg.de
‡burgard@informatik.uni-freiburg.de
§teschner@informatik.uni-freiburg.de

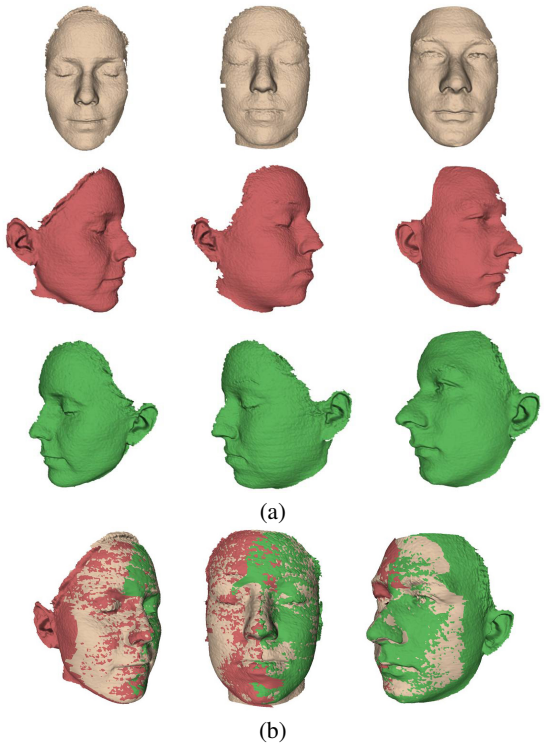

Figure 1: This example illustrates the application of our approach in simultaneous multi-scan registration. Three faces are registered simultaneously. Although, e. g., the left parts are very similar to each other and have large overlapping regions, the approach assigns them correctly to the corresponding front scan. (a) We start with nine arbitrarily oriented scans of different faces. The scans are not presorted in any way. (b) The algorithm computes a proper initial transformation. The result is obtained without local optimization.

Given at least three correspondent points on each object, such a transformation could be easily computed [Horn 1987]. Therefore, the task to establish an appropriate transformation is equivalent to finding appropriate correspondences between the query objects and computing a transformation based on these correspondences. In this paper, we focus on the registration problem, i. e. we want to align several partially overlapping scans of an object.

1.1 Rigid Registration

The Iterative Closest Points algorithm (ICP) [Besl and McKay 1992; Chen and Medioni 1991] is a well-known method for aligning overlapping point clouds. As outlined above, it looks for corresponding points on different objects, for example based on shortest distances, and computes a transformation using these correspondences. Much work has been done to improve the convergence rate and the results of the original algorithm, for example using better metrics and sophisticated criteria for corresponding points [Pulli 1999; Rusinkiewicz and Levoy 2001; Mitra et al. 2004]. However,

the ICP algorithm only optimizes locally and depends on a good initial guess of the relative orientations. This leads to the well-known problem of global registration, where a coarse initial alignment is computed.

In general, objects overlap only partially. Therefore, approaches using global quantities like principal component analysis are not applicable. A straightforward approach to get an initial alignment is to enumerate all three-point-pairs in different objects, compute the corresponding transformation and choose the one that leads to the smallest error with respect to some error metric. For a given triplet in one object, this leads to $O(n^3)$ possibly matching triplets. In case of outliers and noise, the RANSAC algorithm [Fischler and Bolles 1981] can be applied, and [Chen et al. 1998; Chen et al. 1999] build up on RANSAC to get an improved global registration algorithm. [Aiger et al. 2008] also build up on RANSAC, but choose coplanar four-point sets instead of triplets to establish an initial transformation. For a given four-point base, this method only requires to examine all pairs of points instead of all triplets by using affine invariant ratios, which reduces the complexity to $O(n^2)$.

Similarly, voting methods like the Hough transform also employ the fact that three correspondent points are sufficient to compute an initial alignment. The six-dimensional space of possible transformations is subdivided into cells and a transformation for each triplet of points on the different objects is computed. Each computed transformation results in a vote for the corresponding cell, and finally, the transformation represented by the cell with the largest number of votes is chosen as an initial alignment which could be refined using ICP.

Using surface properties like geometry or color can improve the voting and RANSAC-based methods as well as the correspondence search in the ICP algorithm. While color, for example, is used for the matching of fresco fragments [Brown et al. 2008; Toler-Franklin et al. 2010], we concentrate on geometric features that can be used for arbitrary point clouds obtained by a range scanner. Various appropriate geometric descriptors are applied for surface alignment. Among these are differential quantities like curvature, which could be computed based on approximating polynomials [Gal and Cohen-Or 2006; Cazals and Pouget 2003], integral invariants [Pottmann et al. 2007] or volumetric descriptors [Gelfand et al. 2005; Huang et al. 2006]. Also, spin images [Johnson 1997] are used to establish a global initial alignment. In addition, spatial relations between points with similar features are used to enhance the registration [Gelfand et al. 2005; Aiger et al. 2008].

1.2 Non-rigid Registration

Most of the above-mentioned approaches are concerned with rigid registration problems. While articulated objects could be matched using rigid registration algorithms by dividing objects into parts [Gelfand et al. 2005], deformable objects have to be handled in a different manner. Therefore, there is a large research area dealing with non-rigid registration and similarity problems [Brown and Rusinkiewicz 2007; Huang et al. 2008; Raviv et al. 2010]. The possible transformations are extended allowing deformations instead of rigid transformations, and the goal is to minimize some deformation energy term [Huang et al. 2008]. For example, thin-plate splines are used for this purpose [Brown and Rusinkiewicz 2007]. Other methods are concerned with finding correspondences between non-rigid objects, which are possibly deformed, or pose-invariant correspondences [Ovsjanikov et al. 2008; Gal et al. 2007], e. g. between a standing and a sitting dog, without any knowledge about the relative positions. Many approaches are based on geodesic distances [Huang et al. 2008; Tevs et al. 2009] or surface metrics which are invariant under isometric deformations [Liu et al. 2009]. The so called Möbius voting algorithm [Lipman and Funkhouser 2009] makes use of the fact that isometries of simply-connected surfaces are contained in the Möbius group and performs a voting scheme similar to the Hough transform. Also, multi-dimensional scaling [Raviv et al. 2010], spectral correspondences [Jain et al. 2007] and skeleton-based approaches [Zheng et al. 2010; Kin-Chung Au et al. 2010] are used to perform non-rigid registration. As non-rigid registration is not the focus of this paper, we refer the reader to [van Kaick et al. 2010] for a thorough overview of the topic.

1.3 Contribution

In this paper, we introduce a novel method for the global registration of partially overlapping point clouds. We employ surface-approximating polynomials and compute an extended set of rotationally-invariant features for surface points. Further, we propose to transform polynomials with similar features into a common coordinate system and to compare the coefficients. This results in more discriminating information about different shapes than relying on features only. Further, it allows to optimize the transformation using a local optimization scheme such as Newton iteration. Moreover, we introduce a modified distance metric to account for the fact that feature values and polynomial coefficients can have different orders of magnitude. In summary, the contributions of this work are new rotational invariants for surface-approximating polynomials, a modified distance measure and transformed polynomials, which results in a more detailed comparison of the surfaces.

Finally, we show that the approach can be used in the simultaneous registration of different objects, each consisting of several partial scans. Fig. 1 shows the processing of nine partial scans. The three resulting faces are registered simultaneously, and the partial scans are assigned and registered correctly. Although there is significant overlap between the three left scans and among the three right scans, they are assigned to their corresponding front view of the face.

In the following section, we give an overview of our global registration approach. Details of the algorithm are given in Sec. 3 and Sec. 4. We outline the application of our approach in simultaneous multi-scan registration of several objects in Sec. 5. Finally, we present some experiments to illustrate the capabilities of the approach in Sec. 6.

2 Overview

Our approach belongs to the class of feature-based voting methods, where the basic idea is to compute local descriptors on surfaces which are invariant under rigid transformations, to compute candidate transformations for points with similar features on various objects, and to perform a voting scheme to locate the transformation with the largest number of votes. We aim at a coarse initial alignment for partially overlapping scans, which then could be refined using a local optimization algorithm such as ICP.

Alg. 1 gives an overview of our global registration algorithm. For two partially overlapping objects A and B, we first compute surface-approximating polynomials (Sec. 3.1) and a set of rotationally-invariant descriptors based on these polynomials (Sec. 3.2, 3.3). Then, we look for possible correspondences by comparing the invariant features. For a given point $\mathbf{y} \in B$, a point $\mathbf{x} \in A$ is considered to be a candidate if the descriptor values of the corresponding polynomials p_x and p_y differ less than some threshold ε. We then ex-

tract a rotation $\mathbf{R}_{xy}$ based on p_x and p_y that transforms p_x into the local coordinate system of p_y (Sec. 3.4). This allows to compare p_y with the transformed polynomial p_x^R. As different polynomials lead to similar descriptor values, the transformed polynomials return more discriminating information about the local neighborhood, which improves the correspondence search.

Let $\mathbf{x}$ be the candidate with the smallest distance of p_x^R and p_y with respect to the distance measure presented in Sec. 3.5. Then, $\mathbf{R}_{xy}$ is extended in a straightforward way to a transformation that aligns $\mathbf{x}$ and $\mathbf{y}$ and serves as a candidate for the initial alignment of A and B. We perform a vote for the corresponding transformation similar to the Hough transform and choose the transformation with the largest number of votes as the initial alignment. As we show in Sec. 6, using transformed polynomials, it is sufficient to insert only one vote per point, representing the candidate with the best-matching transformed polynomial, to get a global registration result. Indeed, this has some benefits in the simultaneous registration of partial scans of different similarly-shaped objects.

Algorithm 1: Transformed Polynomials

1. Compute a surface-approximating polynomial for each surface point (Sec. 3.1);
2. Compute rotationally-invariant features and find candidate pairs with similar features (Sec. 3.3);
3. For each candidate pair, transform the polynomials into a common coordinate system to compare the coefficients (Sec. 3.4);
4. For matching polynomials, insert a vote into a subdivision scheme for the transformation space, and choose the transformation with the largest number of votes (Sec. 4);

Feature based approaches are widely used in rigid registration, symmetry and shape retrieval, e. g. [Gelfand et al. 2005; Li and Guskov 2005; Mitra et al. 2006; Gal and Cohen-Or 2006; Pauly et al. 2008; Aiger et al. 2008]. They employ differential features [Mitra et al. 2006; Pauly et al. 2008], integral features [Gelfand et al. 2005; Aiger et al. 2008] or polynomials to estimate surface descriptors [Li and Guskov 2005; Gal and Cohen-Or 2006; Pauly et al. 2008]. To the best of our knowledge, none of the approaches uses the transformed polynomials introduced in Sec. 3 or the invariants introduced in Sec. 3.3.

3 Transformed Polynomials

In this section, we describe the details of the transformed polynomials which are the essential step in the global registration approach. First, we introduce the surface-approximating polynomials that we use (Sec. 3.1), before we briefly review the transformation of polynomials (Sec. 3.2). Then, we describe the invariant descriptors (Sec. 3.3) and show how the rotation between polynomials in different coordinate systems is obtained (Sec. 3.4). Finally, we present a distance measure that accounts for the fact that the coefficients and invariants of one polynomial can have different orders of magnitude (Sec. 3.5).

3.1 Surface-approximating Polynomials

To establish correspondences between partial scans of an object, we first describe the surface at each point using a surface-approximating polynomial. In order to be robust against noise, we choose the Moving Least Squares approach [Levin 2004]. However, our algorithm does not depend on this choice and also works for other types like osculating jets [Cazals and Pouget 2003], which are used in [Pauly et al. 2008]. The Moving Least Squares approach first computes a best-fit plane for each point $\mathbf{x}$ such that the squared orthogonal distance of all points in a neighborhood is minimized. Then, the normal $\mathbf{n}_x$ of this best-fit plane is taken as the surface normal for $\mathbf{x}$ and is extended to a local orthonormal coordinate system $(\mathbf{l}_x^1, \mathbf{l}_x^2, \mathbf{n}_x)$. Within this coordinate system, a surface-approximating polynomial $p_x : \mathbb{R}^2 \to \mathbb{R}$ is calculated based on the neighborhood of $\mathbf{x}$. For the plane as well as for the polynomial, the influence of points is weighted depending on their distance to $\mathbf{x}$. The weight is controlled by the feature size h, which hereby implicitly defines the neighborhood of a point $\mathbf{x}$. Note that we actually do not use the Moving Least Squares approach for smoothing the input surface as this would lead to worse registration results [Aiger et al. 2008].

For our registration approach, we choose polynomials of degree 3, as they return good approximation results, while polynomials with higher degree tend to oscillate and result in a worse approximation [Alexa et al. 2001]. In general, the polynomial has the form

$$\begin{aligned} p(u,v) = \; & a_{30}u^3 + a_{20}u^2 + a_{10}u + a_{21}u^2v + a_{11}uv \\ & + a_{12}uv^2 + a_{01}v + a_{02}v^2 + a_{03}v^3 + a_{00}. \end{aligned} \tag{1}$$

In the following section, we introduce the transformation of polynomials, before the rotationally-invariant features are described.

3.2 Polynomial Transformation

In this section, we review the transformation of a polynomial $p : \mathbb{R}^2 \to \mathbb{R}$ into a rotated coordinate system and show how the coefficients are transformed. First, we need this transformation to determine polynomial invariants under rotation. Second, this is a necessary step for comparing different polynomials, as each surface-approximating polynomial p is defined in a local coordinate system. Thus, they have to be transformed into a common coordinate system in order to be comparable.

We assume that p is defined in the standard basis $\mathbf{B} = (\mathbf{e}_1, \mathbf{e}_2)$ and that $\mathbf{B}^R = (\mathbf{r}_1, \mathbf{r}_2)$ defines a coordinate system which is rotated by $\mathbf{R}$. Then, the basis transformation from $\mathbf{B}^R$ to $\mathbf{B}$ is also given by $\mathbf{R}$. We now look for the rotated polynomial p^R that is defined in $\mathbf{B}^R$ and equals p.

For a point given as (u,v) in $\mathbf{B}$ and (u^R, v^R) in $\mathbf{B}^R$, p^R has to fulfill $p^R(u^R, v^R) = p(u,v)$. Inserting the basis transformation leads to $p(u,v) = p(\mathbf{R}(u^R, v^R))$, which can be re-ordered in terms of u^R and v^R to get the coefficients of p^R. In order to compare two given polynomials p in $\mathbf{B}$ and q^R in $\mathbf{B}^R$, p has to be transformed to p^R or q^R to q. Then, the coefficients can be compared.

3.3 Invariants

Our invariants are partially based on differential properties of the polynomials. Differential invariants are also used in [Pauly et al. 2008] and [Mitra et al. 2006], for example. Therefore, we first discuss the relationship between the derivatives of rotated polynomials. Further, we combine both differential and integral invariants of a polynomial, which are developed in this section.

For the differential invariants of a polynomial $p(u,v)$, we consider the derivative at the point $(u,v) = (0,0)$. The partial derivatives are given as

$$\begin{aligned} &\partial_u p(0,0) = a_{10}, \quad \partial_v p(0,0) = a_{01} \\ &\partial_{uu} p(0,0) = 2a_{20},\ \partial_{vv} p(0,0) = 2a_{02},\ \partial_{uv} p(0,0) = a_{11} \\ &\partial_{uuu} p(0,0) = 6a_{30},\ \partial_{vvv} p(0,0) = 6a_{03} \\ &\partial_{uuv} p(0,0) = 2a_{21},\ \partial_{uvv} p(0,0) = 2a_{12}. \end{aligned} \tag{2}$$

The first two invariants are connected to the gradient of p and to the Hessian matrix. For a pair of transformed polynomials p and p^R, we can use the chain rule to see that the gradient of p^R is the rotated gradient of p. Hence, its length remains constant.

Further, the second partial derivatives are summarized in the Hessian matrix. Again, the chain rule leads to a relation between the Hessian matrices $\mathbf{H}$ of p and $\mathbf{H}^R$ of p^R, namely $\mathbf{H} = \mathbf{R}\mathbf{H}^R\mathbf{R}^T$. Hence, its determinant is constant.

In the following, we introduce two new invariants that are related to the third derivative and to an integral, respectively. Integral invariants are also used in [Gelfand et al. 2005; Huang et al. 2006; Pottmann et al. 2007], for example.

To obtain an integral descriptor that is rotationally-invariant, the integration area has to be invariant under rotation. Thus, we use a disc with an arbitrary radius d, and integrate the polynomial using polar coordinates $u = r\cos(\phi), v = r\sin(\phi)$. Then, $p(u,v)$ corresponds to

$$\begin{aligned} p(r,\phi) =& a_{30}r^3\cos^3(\phi) + a_{20}r^2\cos^2(\phi) + a_{10}r\cos(\phi) \\ &+ a_{21}r^3\cos^2(\phi)\sin(\phi) + a_{11}r^2\cos(\phi)\sin(\phi) \\ &+ a_{12}r^3\cos(\phi)\sin^2(\phi) + a_{01}r\sin(\phi) \\ &+ a_{02}r^2\sin^2(\phi) + a_{03}r^3\sin^3(\phi) + a_{00}. \end{aligned} \tag{3}$$

We integrate this function for $\phi \in [0, 2\pi]$ and use

$$\int_0^{2\Pi} \cos(\phi)^k \sin(\phi)^l d\phi = 0 \tag{4}$$

with k or l being odd. It follows that most of the summands in (3) vanish as they contain a sin- or cos-term with an odd exponent. Only the purely quadratic terms $a_{20}r^2\cos^2(\phi)$ and $a_{02}r^2\sin^2(\phi)$ remain, and $\int_0^{2\pi} p(r,\phi)d\phi$ results in $r^2\pi(a_{20}+a_{02})$. Integrating this term for $r \in [0,d]$ with an arbitrary radius d results in $D\cdot(a_{20}+a_{02})$ with some constant D depending on d. It follows that the sum $a_{20}+a_{02}$ remains constant under rotations.

So far, a_{30}, a_{21}, a_{12} and a_{03} are not part of any invariant. However, as they represent the third partial derivatives, one could conclude that there is an invariant containing these quantities. We suggest the following expression that is invariant under rotation:

$$3a_{30}a_{12} + 3a_{03}a_{21} - a_{21}^2 - a_{12}^2 = const. \tag{5}$$

The invariance can be shown by transforming the coefficients a_{ij} to the corresponding rotated coefficients a_{ij}^R. For $\mathbf{R} = \begin{pmatrix} r_{00} & r_{01} \\ r_{10} & r_{11} \end{pmatrix}$, the transformed coefficients are

$$\begin{aligned} a_{30}^R =& a_{21}r_{00}^2r_{01} + a_{12}r_{00}r_{01}^2 + a_{30}r_{00}^3 + a_{03}r_{01}^3 \\ a_{03}^R =& a_{21}r_{10}^2r_{11} + a_{12}r_{10}r_{11}^2 + a_{30}r_{10}^3 + r_{11}^3a_{03} \\ a_{21}^R =& a_{12}r_{10}r_{01}^2 + 3a_{30}r_{00}^2r_{10} + 2a_{21}r_{00}r_{10}r_{01} \\ &+ 2a_{12}r_{00}r_{01}r_{11} + a_{21}r_{00}^2r_{11} + 3r_{11}r_{01}^2a_{03} \\ a_{12}^R =& 3a_{30}r_{00}r_{01}^2 + 2a_{21}r_{00}r_{01}r_{11} + a_{21}r_{01}^2r_{10} \\ &+ a_{12}r_{00}r_{11}^2 + 2a_{12}r_{01}r_{10}r_{11} + 3a_{03}r_{10}r_{11}^2. \end{aligned} \tag{6}$$

Inserting (6) into (5) and using the fact that $\mathbf{R}$ is a rotation, i. e. $r_{00}^2 + r_{01}^2 = 1$ and so on, leads to

$$\begin{aligned} &3a_{30}^Ra_{12}^R + 3a_{03}^Ra_{21}^R - {a_{21}^R}^2 - {a_{12}^R}^2 \\ &= 3a_{30}a_{12} + 3a_{03}a_{21} - a_{21}^2 - a_{12}^2, \end{aligned} \tag{7}$$

which means that (5) is invariant under rotations.

In summary, we use a set of four different invariants:

1. $a_{10}^2 + a_{01}^2 = const$, which is the length of the gradient.
2. $4a_{20}a_{02} - a_{11}^2 = const$, which is the determinant of the Hessian matrix.
3. $a_{20} + a_{02} = const$, which corresponds to an integral of the polynomial over a disc.
4. $3a_{30}a_{12} + 3a_{03}a_{21} - a_{21}^2 - a_{12}^2 = const$, which is related to the third derivative.

3.4 Polynomial Rotation Extraction

Having found a candidate pair $\mathbf{x}, \mathbf{y}$ with polynomials p_x, p_y by comparing the invariants described in the previous section, we compute the rotation $\mathbf{R}_{xy}$ between p_x and p_y that aligns the polynomials in one coordinate system in order to compare their coefficients. By first aligning the local coordinate systems of $\mathbf{x}$ and $\mathbf{y}$, $\mathbf{R}_{xy}$ is extended to a transformation $\mathbf{T}_{xy}$ aligning both points. The extraction of $\mathbf{R}_{xy}$ is based on the relationship between the Hessian matrices $\mathbf{H}_x$ and $\mathbf{H}_y$ of p_x and p_y. This is explained in this section. It is similar to [Mitra et al. 2006; Pauly et al. 2008], who use the principal curvatures to align different points.

We know that $\mathbf{H}_x$ and $\mathbf{H}_y$ are symmetric. Hence, they are diagonalizable, and there are rotations $\mathbf{Q}_x$ and $\mathbf{Q}_y$ such that

$$\mathbf{D}_x = \mathbf{Q}_x^T\mathbf{H}_x\mathbf{Q}_x \tag{8}$$

$$\mathbf{D}_y = \mathbf{Q}_y^T\mathbf{H}_y\mathbf{Q}_y \tag{9}$$

have diagonal form. For a pair of transformed polynomials, the diagonal forms have to be equal. If this is not the case, the two points cannot correspond to each other. Otherwise, we conclude

$$\mathbf{Q}_x^T\mathbf{H}_x\mathbf{Q}_x = \mathbf{Q}_y^T\mathbf{H}_y\mathbf{Q}_y \tag{10}$$

$$\Rightarrow \mathbf{H}_x = \mathbf{Q}_x\mathbf{Q}_y^T\mathbf{H}_y\mathbf{Q}_y\mathbf{Q}_x^T. \tag{11}$$

Thus, $\mathbf{R}_{xy} = \mathbf{Q}_x\mathbf{Q}_y^T$ is the basis transformation matrix that transforms p_x into the coordinate system of p_y.

As we transform p_x into the coordinate frame of p_y, we can optimize $\mathbf{R}_{xy}$ using some optimization scheme like Newton iteration to minimize the distance between the transformed coefficients. This further improves the quality of the surface matching. The employed distance measure is explained in the following section.

3.5 Distance Measure

As the coefficients and invariant features of a polynomial can have different orders of magnitude, we suggest a modified distance measure to account for this fact. While for small coefficients, the squared distance is an appropriate measure, for large coefficients, a relative distance like the quotient is more appropriate. Obviously, the quotient should approximate 1, so $(1-a/b)^2$ would be a candidate for a distance measure for large coefficients a and b.

As the quotient is not symmetric and it would not be possible to use spatial subdivision techniques for the features, we suggest to compare the logarithms of large coefficients and invariants. Thus, we define $(\ln(a)-\ln(b))^2$ to be the distance of a and b. Taking the Taylor series of $\ln(x)$ around 1, $\ln(x) = \ln(1)+\ln'(1)(x-1)+O((x-1)^2) = x-1+O((x-1)^2)$, we see that $|\ln(a)-\ln(b)| = |\ln(a/b)|$ is a first-order approximation of $|1-a/b|$. As we demand that a/b approximates 1 for similar coefficients, $(\ln(a)-\ln(b))^2$ is a sufficient approximation for $(1-a/b)^2$.

There is a straightforward way to determine where to switch between the different measures. Assuming that it is reasonable to always take the minimum of both possible measures leads to the conclusion that $|1-a/b| < |a-b|$ holds iff $|b| > 1$. Hence, if $|a|$ and $|b|$ are greater than 1, we take the logarithmic difference.

4 Global Registration

In this section, we describe how the global registration of two objects A and B is computed.

For each point $\mathbf{x} \in A$ and $\mathbf{y} \in B$, we compute the surface-approximating polynomials p_x and p_y and determine their invariants. For each pair $\mathbf{x} \in A, \mathbf{y} \in B$ with similar features, we transform the polynomial p_x into the local coordinate system of p_y and compare the coefficients. Then, for each $\mathbf{x} \in A$, we take the point $\mathbf{y} \in B$ with the best-matching polynomial and compute the aligning transformation $\mathbf{T}_{xy}$. Taking only the best-matching point is possible due to the transformed polynomials, as we get a thorough comparison of the local neighborhood.

Finally, we perform a voting scheme similar to the Hough transform and take the transformation $\mathbf{T}_{AB}$ with the largest number of votes as the initial alignment of A and B. Further, for each point pair $(\mathbf{x},\mathbf{y})$ whose transformation $\mathbf{T}_{xy}$ voted for $\mathbf{T}_{AB}$, $\mathbf{y}$ is stored as a corresponding point to $\mathbf{x}$. As an alternative to $\mathbf{T}_{AB}$, the stored correspondences can be used to compute an initial alignment using the method presented in [Horn 1987].

Of course, it is possible to take more than only the best-matching polynomial, but as we show in Sec. 6, in our application the best-matching polynomial is sufficient. Especially for similarly-shaped objects, we avoid mismatches that influence the global alignment.

4.1 Runtime

The theoretical runtime is dominated by the search for the local neighborhood to establish the surface polynomials and by the search for point pairs with similar features. Using kd-trees, we have to perform $O(n)$ range queries which take $O(\log n + k)$ time if n is the number of object points and k is the number of points returned by the range query. In the worst case, all points have similar features, which results in $k = n$ and a runtime of $O(n^2)$.

Indeed, the largest part of the runtime is spent in the computation of the polynomials and the search for candidate points, while the remaining parts - including the voting scheme - typically take only a few seconds. However, the computation of the polynomials and candidates for two points $\mathbf{x}_1$ and $\mathbf{x}_2$ is completely independent of each other, which means that the time-consuming steps can be perfectly parallelized.

5 Multi-scan Registration

In this section, we outline an application of the transformed polynomials approach in multi-scan registration and show that it returns discriminating information when registering different objects simultaneously.

For a given set of partial scans $A_1,\ldots,A_k$, we register all scans pairwise to obtain a set of transformations $\mathbf{T}_{ij}$ and a set of correspondences between all objects. Like [Li and Guskov 2005; Huang et al. 2006], we establish a connectivity graph with nodes $A_1,\ldots,A_k$ and weighted edges $e(A_i,A_j)$ if we find an aligning transformation $\mathbf{T}_{ij}$. The weight $c(e(A_i,A_j))$ reflects the registration quality from A_i to A_j. We perform two independent registration steps for each pair of object (A_i,A_j). Therefore, we can control the edge weights in a twofold way. First, we check if $\mathbf{T}_{ij}\mathbf{T}_{ji}$ approximates the identity. If not, the corresponding edges are rejected. This is similar to [Huang et al. 2006] who check circles for consistency. Second, if $\mathbf{y} \in A_j$ is correspondent to $\mathbf{x} \in A_i$, we check if the reverse relation is also true. Therefore, we get a bijection between the corresponding points of A_i and A_j and count the number of bijective pairs, which then serves as the weight for the edge $e(A_i,A_j)$.

The quality of these edge weights is illustrated in Fig. 1. The left parts of the three faces are similar to each other and have more overlap among themselves than to their corresponding front scan, and the same holds for the right parts. Nevertheless, the approach finds about an order of magnitude more correspondences from the left and right scans to their respective front scan than among each other and to the not corresponding front scans (Sec. 6). This is due to the transformed polynomials which allow to take only one candidate pair per point. Thus, lots of matches are excluded which would occur over the cheeks, for example, and would lead to higher weights between the left and between the right scans. Therefore, the edge weights obtained by the transformed polynomials approach return information about the connectivity of the objects, although the structures are quite similar and are likely to be mixed up.

Similar to [Li and Guskov 2005; Huang et al. 2006], we register objects (A_i,A_j) that share the edge with the highest weight. Then, if an object A_k has correspondences to A_i and A_j, the corresponding transformations are checked for consistency as in [Huang et al. 2006], and the edge weights are summed up. This is done up to some user-defined threshold. In Fig. 1, we end up with separated objects.

6 Results

In this section, we illustrate the capabilities of our developed global registration approach. All results are obtained using the trans-

formed polynomials approach without any local improvement. The experiments have been performed on an AMD Opteron 8435 with 24 cores at 2.6 GHz and 64 GB memory.

In our implementation, the feature size for the Moving Least Squares approach (Sec. 3.1) is set to 3% of the bounding box diagonal. Two points $\mathbf{x}, \mathbf{y}$ are considered as candidates if the distance of their polynomial invariants (Sec. 3.5) is smaller than 10^{-5}. As we transform the polynomials afterwards and take only the best-matching candidate, this threshold rather affects the runtime than the registration result. The subdivision scheme for the aligning transformations has a cell size of 6° for the rotation angle and of 0.05 for the components of the rotation axis, which is represented as a unit vector. For the translational components, the cell size corresponds to the feature size.

In order to illustrate the quality of the global alignments, we applied the ICP algorithm [Besl and McKay 1992; Chen and Medioni 1991] after the transformed polynomials approach. For an optimal global alignment, ICP should end up with no additional rotation. Therefore, we denote the additional rotation angle obtained by ICP to show the quality of the transformed polynomials approach. However, the figures show the global registration results without any local improvement.

The first experiment shows an artificial test setup with two models of the Stanford Buddha. We used a resampled mesh consisting of 135634 points, where one model is rotated by an arbitrary angle. Although there are lots of similar surface parts which lead to similar surface polynomials, the approach locates 135631 correspondences correctly. Note that we allow exactly one possible correspondence per point. Using the transformed polynomials, the best-matching polynomial leads to the correct correspondence for most of the points. The runtime for this experiment was 211 s. Here, applying ICP resulted in no additional rotation angle, so the optimal alignment was reached.

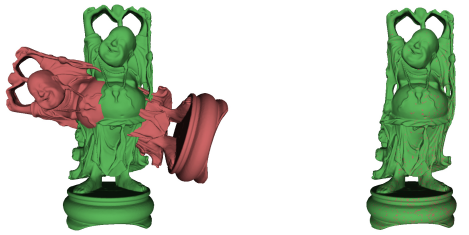

(a) (b)

Figure 2: The Stanford Buddha consisting of 135634 points. (a) Relative orientation before registration. (b) Transformation found by the global registration approach. 135631 points are assigned correctly by the transformed polynomials approach.

All of the following experiments are performed with real data collected by different range scanners.

Fig. 3 shows the registration of ten range scans of the Stanford bunny. A similar experiment has also been performed in [Gelfand et al. 2005]. In contrast to [Gelfand et al. 2005], where overlapping scans are pre-assigned and registered pairwise in single registration steps, all scans are reliably assigned and registered simultaneously within 583 s using the transformed polynomials approach. The scans are not pre-sorted in any way, and no pair of scans is excluded a priori. Compared to [Gelfand et al. 2005], where only few correspondences are found between overlapping objects, we locate 5376 correspondences on average. Applying ICP, we obtain an average additional rotation angle of 0.44°.

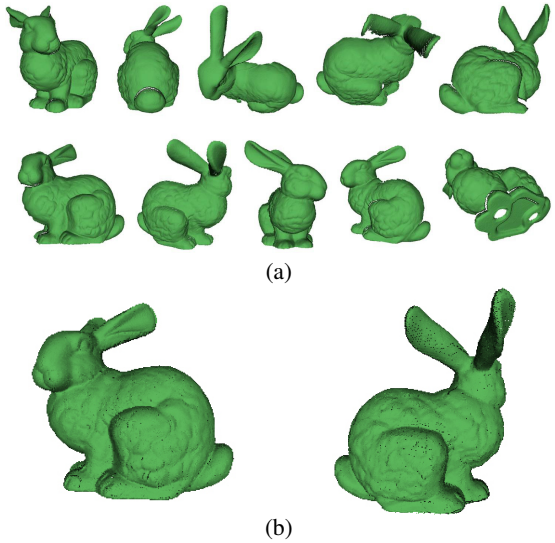

(a) (b)

Figure 3: (a) 10 range scans of the Stanford Bunny with arbitrary orientations. (b) Point clouds of the registered scans. The result is obtained without any local refinement.

As shown in Fig. 1, the algorithm is also able to assign partial scans of different faces correctly. The face scans are obtained by a faceScan-III system from Breuckmann and have an average size of 79800 points. The approach locates 9500 correspondences on average between correct pairs of scans, while there are typically 200-300 correspondences between false left-left or right-right pairs, which serve as weights for the simultaneous registration approach (Sec. 5). As in Fig. 3, the scans are not pre-sorted in any way and no pairs are excluded a priori. We emphasize that this is a special strength of our approach. For example, algorithms looking for a largest common point set would be likely to assign a higher weight between two left scans than between the correct left and front scan. Registering nine partial scans simultaneously required 72 single registration steps, which took 100 s on average. Applying ICP resulted in an additional rotation angle of 0.3° on average.

Fig. 4 and Fig. 5 illustrate that the approach is able to handle noisy data. In Fig. 4, we added some random noise to the face scans to illustrate the robustness of the approach. Although the noise is in the same order of magnitude as the feature size, the global registration finds a proper initial alignment. The objects in Fig. 5, a teddy bear besides a mobile robot, were scanned with a Microsoft® Kinect™-camera and are rather noisy (Fig. 5 (b)). Nevertheless, the scans are assigned and registered simultaneously within 1766 s. Applying ICP, we obtain an additional rotation angle of 1° on average.

6.1 Limitations

The transformed polynomials approach leads to more discriminating information than using surface features only and therefore, it is able to distinguish different objects. However, the approach naturally has some limitations.

Obviously, our approach runs into problems if we cannot compute distinct features like it is the case for objects with large featureless parts, e. g. planes. For such situations, the approach described in [Aiger et al. 2008] is more appropriate.

Like all registration approaches, the approach has problems to establish an initial transformation if there is only a small overlap.

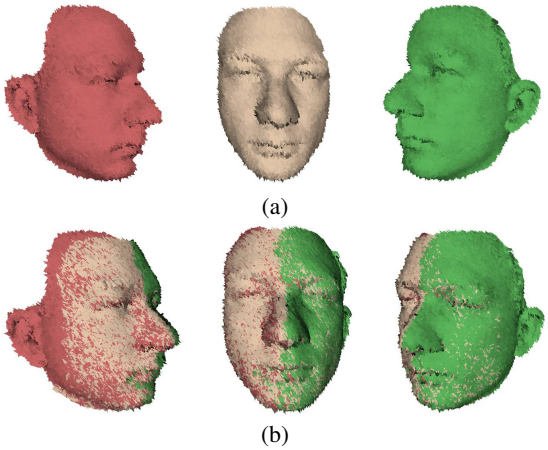

Figure 4: (a) Although the scanned data naturally contains noise, we added some random noise to a face scan to illustrate the robustness of our approach. (b) The simultaneous global registration algorithm is able to obtain a proper alignment within 1118 s.

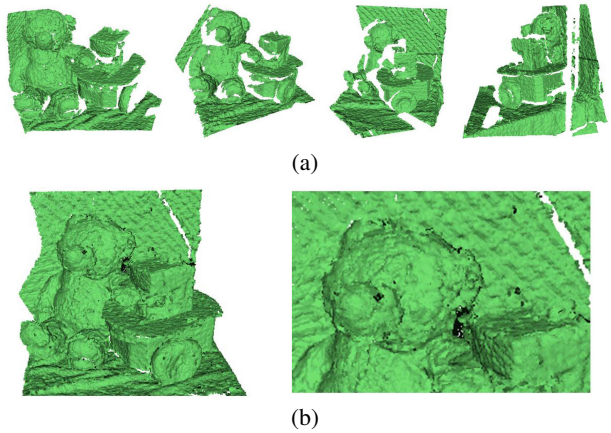

Figure 5: (a) Range scans of a teddy bear and a mobile robot. (b) The transformed polynomials approach returns a good global registration, although the scans are quite noisy.

The multi-scan registration is based on a heuristic, which does not work in every case. However, due to the transformed polynomials, the approach at least returns few false correspondences between different objects, although they have a similar shape. Consequently, a lot of partial scans are registered correctly before the first false match would occur. Therefore, similar to [Li and Guskov 2005], the number of non-registered objects decreases, which simplifies the manual separation of the objects, if the heuristic does not completely work.

7 Conclusion

In this paper, we introduced a novel approach for the global rigid registration of partially overlapping point clouds. We introduced new invariants for the search of possible candidates and the transformed polynomials which lead to an improved correspondence search. We outlined that the established correspondences can be applied for the simultaneous registration of several similarly shaped objects. In the results, we showed the capabilities of the approach for global registration tasks.

Future work will be concerned with the simultaneous registration and scaling of objects that have different sizes using ratios of the rotational invariants.

8 Acknowledgements

This project is supported by the German Research Foundation (DFG) under contract number SFB TR8. The Bunny and the Buddha models are taken from the Stanford 3D Scanning Repository of the Stanford University Computer Graphics Laboratory.

References

AIGER, D., MITRA, N. J., AND COHEN-OR, D. 2008. 4-points congruent sets for robust pairwise surface registration. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, ACM, New York, NY, USA, 1–10.

ALEXA, M., BEHR, J., COHEN-OR, D., FLEISHMAN, S., LEVIN, D., AND SILVA, C. T. 2001. Point set surfaces. In *VIS '01: Proceedings of the conference on Visualization '01*, IEEE Computer Society, Washington, DC, USA, 21–28.

BESL, P., AND MCKAY, N. 1992. A method for registration of 3-d shapes. *IEEE Trans. PAMI 14*, 2, 239–256.

BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2009. Topology-invariant similarity of nonrigid shapes. *Int. J. Comput. Vision 81*, 3, 281–301.

BROWN, B. J., AND RUSINKIEWICZ, S. 2007. Global non-rigid alignment of 3-d scans. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers*, ACM, New York, NY, USA, 21.

BROWN, B. J., TOLER-FRANKLIN, C., NEHAB, D., BURNS, M., DOBKIN, D., VLACHOPOULOS, A., DOUMAS, C., RUSINKIEWICZ, S., AND WEYRICH, T. 2008. A system for high-volume acquisition and matching of fresco fragments: reassembling theran wall paintings. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, ACM, New York, NY, USA, 1–9.

CAZALS, F., AND POUGET, M. 2003. Estimating differential quantities using polynomial fitting of osculating jets. In *Proceedings of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, SGP '03, 177–187.

CHEN, Y., AND MEDIONI, G. 1991. Object modeling by registration of multiple range images. In *Proc. of IEEE International Conference on Robotics and Automation*, 2724–2729.

CHEN, C.-S., HUNG, Y.-P., AND CHENG, J.-B. 1998. A fast automatic method for registration of partially-overlapping range images. In *Proc. ICCV*, 242–248.

CHEN, C.-S., HUNG, Y.-P., AND CHENG, J.-B. 1999. RANSAC-based DARCES: a new approach to fast automatic registration of partially overlapping range images. *IEEE Transactions on Pattern Analysis and Machine Intelligence 21*, 11 (November), 1229 –1234.

FISCHLER, M. A., AND BOLLES, R. C. 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM 24*, 6, 381–395.

GAL, R., AND COHEN-OR, D. 2006. Salient geometric features for partial shape matching and similarity. *ACM Trans. Graph. 25*, 1, 130–150.

GAL, R., SHAMIR, A., AND COHEN-OR, D. 2007. Pose-oblivious shape signature. *IEEE Transactions on Visualization and Computer Graphics 13*, 2, 261–271.

GELFAND, N., MITRA, N. J., GUIBAS, L. J., AND POTTMANN, H. 2005. Robust global registration. In *Proceedings of the third Eurographics symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, 197:1–197:10.

HORN, B. K. P. 1987. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Amer. A 4*, 4, 629–642.

HUANG, Q.-X., FLÖRY, S., GELFAND, N., HOFER, M., AND POTTMANN, H. 2006. Reassembling fractured objects by geometric matching. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, ACM, New York, NY, USA, 569–578.

HUANG, Q.-X., ADAMS, B., WICKE, M., AND GUIBAS, L. J. 2008. Non-rigid registration under isometric deformations. *Computer Graphics Forum 27*, 5, 1449–1457.

JAIN, V., ZHANG, H., AND VAN KAICK, O. 2007. Non-rigid spectral correspondence of triangle meshes. *International Journal on Shape Modeling 13*, 1, 101–124.

JOHNSON, A. 1997. *Spin-Images: A Representation for 3-D Surface Matching*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA.

KAZHDAN, M., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2004. Symmetry descriptors and 3D shape matching. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, ACM, New York, NY, USA, SGP '04, 115–123.

KIN-CHUNG AU, O., TAI, C.-L., COHEN-OR, D., ZHENG, Y., AND FU, H. 2010. Electors voting for fast automatic shape correspondence. *Computer Graphics Forum 29*, 2, 645–654.

LEVIN, D. 2004. Mesh-independent surface interpolation. In *Geometric Modeling for Scientific Visualization*, G. Brunnett, B. Hamann, K. Mueller, and L. Linsen, Eds. Springer-Verlag, 37–50.

LI, X., AND GUSKOV, I. 2005. Multi-scale features for approximate alignment of point-based surfaces. In *Proceedings of the third Eurographics symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, 217:1–217:11.

LIPMAN, Y., AND FUNKHOUSER, T. 2009. Möbius voting for surface correspondence. *ACM Trans. Graph. 28*, 3, 1–12.

LIU, R., ZHANG, H., SHAMIR, A., AND COHEN-OR, D. 2009. A part-aware surface metric for shape analysis. *Computer Graphics Forum (Special Issue of Eurographics 2009) 28*, 2, 397–406.

MITRA, N. J., GELFAND, N., POTTMANN, H., AND GUIBAS, L. 2004. Registration of point cloud data from a geometric optimization perspective. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, ACM, New York, NY, USA, SGP '04, 22–31.

MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2006. Partial and approximate symmetry detection for 3D geometry. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, ACM, New York, NY, USA, 560–568.

MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2007. Symmetrization. *ACM Trans. Graph. 26*, 3 (July), 63:1–63:8.

OVSJANIKOV, M., SUN, J., AND GUIBAS, L. 2008. Global intrinsic symmetries of shapes. In *Proceedings of the Symposium on Geometry Processing*, Eurographics Association, Aire-la-Ville, Switzerland, SGP '08, 1341–1348.

PAULY, M., MITRA, N. J., GIESEN, J., GROSS, M., AND GUIBAS, L. J. 2005. Example-based 3D scan completion. In *Proceedings of the third Eurographics symposium on Geometry processing*, Eurographics Association, Aire-la-Ville, Switzerland, 23:1–23:10.

PAULY, M., MITRA, N. J., WALLNER, J., POTTMANN, H., AND GUIBAS, L. J. 2008. Discovering structural regularity in 3D geometry. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, ACM, New York, NY, USA, 1–11.

PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. In *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, ACM, New York, NY, USA, 549–559.

POTTMANN, H., WALLNER, J., YANG, Y.-L., LAI, Y.-K., AND HU, S.-M. 2007. Principal curvatures from the integral invariant viewpoint. *Computer Aided Geometric Design 24*, 8-9, 428–442.

PULLI, K. 1999. Multiview registration for large data sets. In *3-D Digital Imaging and Modeling, 1999. Proceedings. Second International Conference on*, 160–168.

RAVIV, D., BRONSTEIN, A. M., BRONSTEIN, M. M., AND KIMMEL, R. 2010. Full and partial symmetries of non-rigid shapes. *Int. J. Comput. Vision 89*, 1, 18–39.

RUSINKIEWICZ, S., AND LEVOY, M. 2001. Efficient variants of the ICP algorithm. *3D Digital Imaging and Modeling, International Conference on*, 145–152.

RUSINKIEWICZ, S., HALL-HOLT, O., AND LEVOY, M. 2002. Real-time 3D model acquisition. *ACM Trans. Graph. 21*, 3, 438–446.

RUSTAMOV, R. M. 2008. Augmented planar reflective symmetry transform. *Vis. Comput. 24*, 6 (May), 423–433.

SHALON, S., SHAPIRA, L., SHAMIR, A., AND COHEN-OR, D. 2008. Part analogies in sets of objects. In *Proc. of Eurographics Symposium on 3D Object Retrieval*, 33–40.

TEVS, A., BOKELOH, M., WAND, M., SCHILLING, A., AND SEIDEL, H.-P. 2009. Isometric registration of ambiguous and partial data. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 1185–1192.

TOLER-FRANKLIN, C., BROWN, B., WEYRICH, T., FUNKHOUSER, T., AND RUSINKIEWICZ, S. 2010. Multi-feature matching of fresco fragments. *ACM Trans. Graph. 29*, 6 (December), 185:1–185:12.

VAN KAICK, O., ZHANG, H., HAMARNEH, G., AND COHEN-OR, D. 2010. A survey on shape correspondence. In *Proc. of Eurographics State-of-the-art Report*.

ZHENG, Q., SHARF, A., TAGLIASACCHI, A., CHEN, B., ZHANG, H., SHEFFER, A., AND COHEN-OR, D. 2010. Consensus skeleton for non-rigid space-time registration. *Computer Graphcis Forum (Special Issue of Eurographics) 29*, 2, 635–644.