

A Game-Theoretic Approach to Fine Surface Registration without Initial Motion Estimation

Andrea Albarelli, Emanuele Rodolà, and Andrea Torsello

Dipartimento di Informatica - Università Ca' Foscari - via Torino, 155 - 30172 Venice Italy

albarelli@unive.it rodola@dsi.unive.it torsello@dsi.unive.it

Abstract

Surface registration is a fundamental step in the reconstruction of three-dimensional objects. This is typically a two step process where an initial coarse motion estimation is followed by a refinement. Most coarse registration algorithms exploit some local point descriptor that is intrinsic to the shape and does not depend on the relative position of the surfaces. By contrast, refinement techniques iteratively minimize a distance function measured between pairs of selected neighboring points and are thus strongly dependent on initial alignment. In this paper we propose a novel technique that allows to obtain a fine surface registration in a single step, without the need of an initial motion estimation. The main idea of our approach is to cast the selection of correspondences between points on the surfaces in a game-theoretic framework, where a natural selection process allows mating points that satisfy a mutual rigidity constraint to thrive, eliminating all the other correspondences. This process yields a very robust inlier selection scheme that does not depend on any particular technique for selecting the initial strategies as it relies only on the global geometric compatibility between correspondences. The practical effectiveness of the proposed approach is confirmed by an extensive set of experiments and comparisons with state-of-the-art techniques.

1. Introduction

The distinction between coarse and fine surface registration techniques is mainly related to the different strategies adopted to find pairs of mating points to be used for the estimation of the rigid transformation. Almost invariably, fine registration algorithms exploit an initial guess in order to constrain the search area for compatible mates and minimize the risk of selecting outliers. On the other hand, coarse techniques, which cannot rely on any motion estimation, must adopt a mating strategy based on the similarity between surface-point descriptors or resort to random

selection schemes. The tension between the precision required for fine alignment versus the recall needed for an initial motion estimation stands as the main hurdle to the unification of such approaches.

The large majority of currently used fine alignment methods are modifications to the original ICP proposed by Zhang [23] and Besl and McKay [3]. These variants generally differ in the strategies used to sample points from the surfaces, reject incompatible pairs, or measure error. In general, the precision and convergence speed of these techniques is highly data-dependent and very sensitive to the fine-tuning of the model parameters. Several approaches that combine these variants have been proposed in the literature in order to overcome these limitations (see [16] for a comparative review). Some recent variants avoid hard culling by assigning a probability to each candidate pair by means of evolutionary techniques [14] or Expectation Maximization [10]. ICP variants, being iterative algorithms based on local, step-by-step decisions, are very susceptible to the presence of local minima. Other fine registration methods include the well-known method by Chen [6] and signed distance fields matching [15].

Coarse registration techniques can be roughly classified into methods that exploit some global property of the surface, such as PCA [8] or Algebraic Surface Model [18], and methods that use some 3D feature descriptor to find plausible candidates pairs over the model and data surfaces. Global techniques are generally very sensitive to occlusion. Feature-based approaches are more precise and can align surfaces that exhibit only partial overlap. Nevertheless, the unavoidable localization error of the feature points prevents them from obtaining accuracies on par with fine registration methods. Among the most successful descriptors are Point Signatures [7] and Spin Images [12]. A completely different coarse registration approach is the RANSAC-based DARCES [5], which is based on the random extraction of sets of mates from the surfaces and their validation based on the accuracy of the estimated transformation. Other recent methods include [1]; a recent and extensive review of all the different methods can be found in [17].

Regardless of the criteria used to obtain pairs of mating points, the subsequent step in surface registration is to search for the rigid transformation that minimizes the squared distance between them. Since many mature techniques are available to do this (for instance [11]), in this paper our effort is toward the matching step itself: specifically by proposing a novel game-theoretic approach that is able to deal equally well with both coarse and fine registration scenarios.

2. Game-Theoretic Surface Registration

We are looking for a robust set of inliers for correspondence selection from which we can estimate the rigid transformation. Most of the currently adopted matching schemes operate on a local level, and global information comes only as an afterthought by checking the quality of the candidate matches with respect to the registration error obtained. The approach we are proposing, on the other hand, brings global information into the matching process by favoring sets of point-associations that are mutually compatible with a single rigid transformation. Fundamental to our approach is the fact that requiring the compatibility to a single transformation is equivalent to requiring that there exists a compatible transformation for each pair of mates. Following [19, 2], we model the mating process in a game-theoretic framework, where two players extracted from a large population select a pair of corresponding points from two surfaces to be registered with one another. The player then receives a payoff from the other players proportional to how compatible his pairings are with respect to the other player's choice, where the compatibility derives from the existence of a common rigid transformation. More explicitly, if there exists a rigid transformation that moves both his point and the other player's point close to the corresponding mates, then both players receive a high payoff, otherwise the payoff will be low. Clearly, it is in each player's interest to pick correspondences that are compatible with the mates the other players are likely to choose. In general, as the game is repeated, players will adapt their behavior to prefer matings that yield larger payoffs, driving all inconsistent hypotheses to extinction, and settling for an equilibrium where the pool of mates from which the players are still actively selecting their associations forms a cohesive set with high mutual support. Within this formulation, the solutions of the matching problem correspond to evolutionary stable states (ESS's), a robust population-based generalization of the notion of a Nash equilibrium.

In a sense, this mating process can be seen as a contextual voting system, where each time the game is repeated the previous selections of the other players affect the future vote of each player in an attempt to reach consensus. This way the evolving context brings global information into the selection process.

2.1. Non-cooperative Games

Originated in the early 40's, Game Theory was an attempt to formalize a system characterized by the actions of entities with competing objectives, which is thus hard to characterize with a single objective function [21]. According to this view, the emphasis shifts from the search of a local optimum to the definition of equilibria between opposing forces. In this setting multiple players have at their disposal a set of strategies and their goal is to maximize a payoff that depends also on the strategies adopted by other players. Evolutionary game theory originated in the early 70's as an attempt to apply the principles and tools of game theory to biological contexts. Evolutionary game theory considers an idealized scenario where pairs of individuals are repeatedly drawn at random from a large population to play a two-player game. In contrast to traditional game-theoretic models, players are not supposed to behave rationally, but rather they act according to a pre-programmed behavior, or mixed strategy. It is supposed that some selection process operates over time on the distribution of behaviors favoring players that receive higher payoffs.

More formally, let $O = \{1, \dots, n\}$ be the set of available strategies (*pure strategies* in the language of game theory), and $C = (c_{ij})$ be a matrix specifying the payoff that an individual playing strategy i receives against someone playing strategy j . A *mixed strategy* is a probability distribution $\mathbf{x} = (x_1, \dots, x_n)^T$ over the available strategies O . Clearly, mixed strategies are constrained to lie in the n -dimensional standard simplex

$$\Delta^n = \left\{ \mathbf{x} \in \mathbb{R}^n : x_i \geq 0 \text{ for all } i \in 1 \dots n, \sum_{i=1}^n x_i = 1 \right\}.$$

The *support* of a mixed strategy $\mathbf{x} \in \Delta$, denoted by $\sigma(\mathbf{x})$, is defined as the set of elements chosen with non-zero probability: $\sigma(\mathbf{x}) = \{i \in O \mid x_i > 0\}$. The expected payoff received by a player choosing element i when playing against a player adopting a mixed strategy \mathbf{x} is $(C\mathbf{x})_i = \sum_j c_{ij}x_j$, hence the expected payoff received by adopting the mixed strategy \mathbf{y} against \mathbf{x} is $\mathbf{y}^T C\mathbf{x}$. The *best replies* against mixed strategy \mathbf{x} is the set of mixed strategies

$$\beta(\mathbf{x}) = \{ \mathbf{y} \in \Delta \mid \mathbf{y}^T C\mathbf{x} = \max_{\mathbf{z}} (\mathbf{z}^T C\mathbf{x}) \}.$$

A strategy \mathbf{x} is said to be a *Nash equilibrium* if it is the best reply to itself, i.e., $\forall \mathbf{y} \in \Delta, \mathbf{x}^T C\mathbf{x} \geq \mathbf{y}^T C\mathbf{x}$. This implies that $\forall i \in \sigma(\mathbf{x})$ we have $(C\mathbf{x})_i = \mathbf{x}^T C\mathbf{x}$; that is, the payoff of every strategy in the support of \mathbf{x} is constant.

A strategy \mathbf{x} is said to be an *evolutionary stable strategy* (ESS) if it is a Nash equilibrium and

$$\forall \mathbf{y} \in \Delta \quad \mathbf{x}^T C\mathbf{x} = \mathbf{y}^T C\mathbf{x} \Rightarrow \mathbf{x}^T C\mathbf{y} > \mathbf{y}^T C\mathbf{y}. \quad (1)$$

This condition guarantees that any deviation from the stable strategies does not pay.

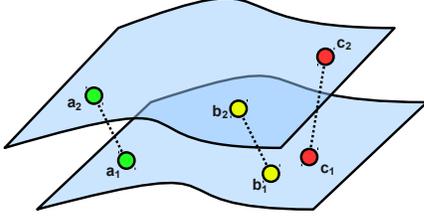


Figure 1. Example of mating strategies.

2.2. Mating Strategies and Payoffs

Central to this framework is the definition of a *mating game*, which implies the definition of the strategies available to the players and of the payoffs related to these strategies. Given a set of model points M and a set of data points D we call a *mating strategy* any pair (a_1, a_2) with $a_1 \in M$ and $a_2 \in D$. We call the set of all the mating strategies S . In principle, all the model and data points could be used to build the mating strategies set, thus giving $S = M \times D$. In practice, however, we adopt some heuristics that allow us to obtain good alignments with a much smaller set. Once S has been selected, our goal becomes to extract from it the largest subset that includes only correctly matched points: that is, strategies that associate a point in the model surface with the same point in the data surface. To enforce this we assign to each pair of mating strategies a payoff that is inversely proportional to a measure of violation of the rigidity constraint. This violation can be expressed in several ways, but since all the rigid transformations preserve Euclidean distances, we choose this property to express the coherence between mating strategies.

Definition 1. Given a function $\pi : S \times S \rightarrow \mathbb{R}^+$, we call it a rigidity-enforcing payoff function if for any $((a_1, a_2), (b_1, b_2))$ and $((c_1, c_2), (d_1, d_2)) \in S \times S$ we have that $\|a_1 - b_1\| - \|a_2 - b_2\| > \|c_1 - d_1\| - \|c_2 - d_2\|$ implies $\pi((a_1, a_2), (b_1, b_2)) < \pi((c_1, c_2), (d_1, d_2))$. In addition, if $\pi((a_1, a_2), (b_1, b_2)) = \pi((b_1, b_2), (a_1, a_2))$, π is said to be symmetric.

A rigidity-enforcing payoff function is a function that is monotonically decreasing with the absolute difference of the Euclidean distances between respectively the model and data points of the mating strategies compared. In other words, given two mating strategies, their payoff should be high if the distance between the model points is equal to the distance between the data points and it should decrease as the difference between such distances increases. In the example of Figure 1, mating strategies (a_1, a_2) and (b_1, b_2) are coherent with respect to the rigidity constraint, whereas (b_1, b_2) and (c_1, c_2) are not, thus it is expected that $\pi((a_1, a_2), (b_1, b_2)) > \pi((b_1, b_2), (c_1, c_2))$.

Further, if we want mating to be one-to-one, we must put an additional constraint on the payoffs, namely that mates sharing a point are incompatible.

Definition 2. A rigidity-enforcing payoff function π is said to be one-to-one if $a_1 = b_1$ or $a_2 = b_2$ implies $\pi((a_1, a_2), (b_1, b_2)) = 0$.

Given a set of mating strategies S and an enumeration $O = \{1, \dots, |S|\}$ over it, a *mating game* is a non-cooperative game where the population is defined as a vector $\mathbf{x} \in \Delta^{|S|}$ and the payoff matrix $C = (c_{ij})$ is defined as $c_{ij} = \pi(s_i, s_j)$, where $s_i, s_j \in S$ are enumerated by O and π is a symmetric one-to-one rigidity-enforcing payoff function. Intuitively, x_i accounts for the percentage of the population that plays the i -th mating strategy. By using a symmetric one-to-one payoff function in a mating game we are guaranteed that ESS's will not include mates sharing either model or data nodes. In fact, given a non-negative payoff function, a stable state cannot have in its support a pairs of strategies with payoff 0 [2]. Moreover, a mating game exhibits some additional interesting properties.

Theorem 1. Given a set of model points M , a set of data points $D = TM$ that are exact rigid transformations of the points in M , and a set of mating strategies $S \subseteq M \times D$ with $(m, Tm) \in S$ for all $m \in M$, and a mating game over them with a payoff function π , the vector $\hat{\mathbf{x}} \in \Delta^{|S|}$ defined as

$$\hat{x}_i = \begin{cases} 1/|M| & \text{if } s_i = (m, Tm) \text{ for some } m \in M; \\ 0 & \text{otherwise,} \end{cases}$$

is an ESS and obtains the global maximum average payoff.

Sketch of proof. Let $\hat{S} \subseteq S$ be the set of mates that match a point to its copy, clearly for all $s, q \in \hat{S}$, $s \neq q$ we have $\pi(s, q) = 1$, while for $s \in \hat{S}$ and $q \in S \setminus \hat{S}$, we have $\pi(s, q) < 1$. For all $s \in \hat{S}$ we have that $\pi(\hat{\mathbf{x}}, \hat{\mathbf{x}}) = \frac{|M|-1}{|M|}$ while, since π is one-to-one, for any $q \in S \setminus \hat{S}$ there must be at least one $s_q \in \hat{S}$ with $\pi(q, s_q) = 0$, thus $\pi(q, \hat{\mathbf{x}}) < \frac{|M|-1}{|M|}$, thus $\hat{\mathbf{x}}$ is a Nash equilibrium. Further, since the inequality is strict, it is an ESS. Finally, $\hat{\mathbf{x}}$ is a global maximizer of π since $\frac{|M|-1}{|M|}$ is the maximum value that a one-to-one normalized payoff function over $|M|$ points can attain. \square

This theorem states that when matching a surface with a rigidly transformed copy of itself the optimal solution (i.e., the population configuration that selects all the mating strategies assigning each point to its copy) is the stable state of maximum payoff. Since well established algorithms to evolve a population to such a state exist, this provides us with an effective mating approach. Clearly, aligning a surface to an identical copy is not very useful in practical scenarios, where occlusion and measurement noise come into play. While the quality of the solution in presence of noise will be assessed experimentally, we can give some theoretical results regarding occlusions.

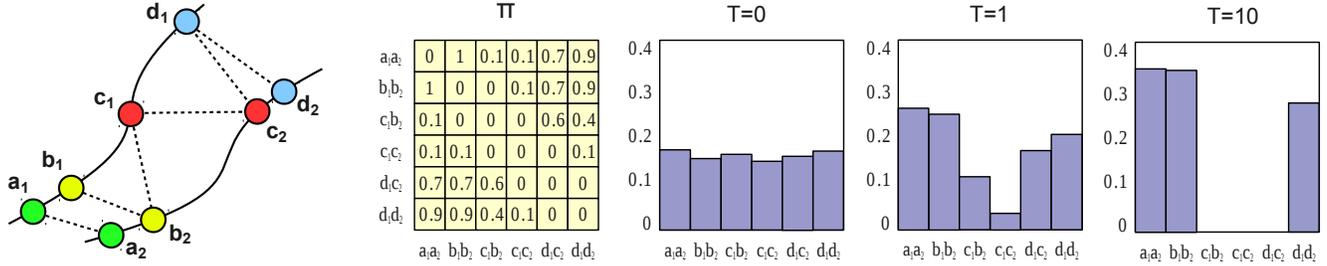


Figure 2. An example of the evolutionary process. Four points are sampled from the two surfaces and a total of six mating strategies are selected as initial hypotheses. The matrix Π shows the compatibilities between pairs of mating strategies according to a one-to-one rigidity-enforcing payoff function. Each mating strategy got zero payoff with itself and with strategies that share the same source or destination point (i.e., $\pi((b_1, b_2), (c_1, b_2)) = 0$). Strategies that are coherent with respect to rigid transformation exhibit high payoff values (i.e., $\pi((a_1, a_2), (b_1, b_2)) = 1$ and $\pi((a_1, a_2), (d_1, d_2)) = 0.9$), while less compatible pairs get lower scores (i.e., $\pi((a_1, a_2), (c_1, c_2)) = 0.1$). Initially (at $T=0$) the population is set to the barycenter of the simplex and slightly perturbed (3-5%). After just one iteration, (c_1, b_2) and (c_1, c_2) have lost a significant amount of support, while (d_1, c_2) and (d_1, d_2) are still played by a sizable amount of population. After ten iterations ($T=10$), (d_1, d_2) has finally prevailed over (d_1, c_2) (note that the two are mutually exclusive). Note that in the final population $((a_1, a_2), (b_1, b_2))$ have a larger support than (d_1, d_2) since they are a little more coherent with respect to rigidity.

Theorem 2. Let M be a set of points with $M_a \subseteq M$ and $D = TM_b$ a rigid transformation of $M_b \subseteq M$ such that $|M_a \cap M_b| \geq 3$, and $S \subseteq M_a \times D$ be a set of mating strategies over M_a and D with $(m, Tm) \in S$ for all $m \in M_a \cap M_b$. Further, assume that the points that are not in the overlap, that is the points in $E_a = M_a \setminus (M_a \cap M_b)$ and $E_b = M_b \setminus (M_a \cap M_b)$, are sufficiently far away such that for every $s \in S, s = (m, Tm)$ with $m \in M_a \cap M_b$ and every $q \in S, q = (m_a, Tm_b)$ with $m_a \in E_a$ and $m_b \in E_b$, we have $\pi(q, s) < \frac{|M_a \cap M_b| - 1}{|M_a \cap M_b|}$, then, the vector $\hat{x} \in \Delta^{|S|}$ defined as

$$\hat{x}_i = \begin{cases} 1/|M| & \text{if } s_i = (m, Tm) \text{ for some } m \in M_a \cap M_b; \\ 0 & \text{otherwise,} \end{cases}$$

is an ESS.

Sketch of proof. We have $\pi(\hat{x}, \hat{x}) = \frac{|M_a \cap M_b| - 1}{|M_a \cap M_b|}$. Let $q \in S$ be a strategy not in the support of \hat{x} , then, either it maps a point in M_a or M_b , thus receiving payoff $\pi(q, \hat{x}) < \frac{|M_a \cap M_b| - 1}{|M_a \cap M_b|}$ because of the one-to-one condition, or it maps a point in E_a to a point in E_b , receiving, by hypothesis, a payoff $\pi(q, \hat{x}) < \frac{|M_a \cap M_b| - 1}{|M_a \cap M_b|}$. Hence, \hat{x} is an ESS. \square

The result of theorem 2 is slightly weaker than theorem 1, as the face of the simplex corresponding to the “correct” overlap, while being an evolutionary stable state, is not guaranteed to obtain the overall highest average payoff. This is not a limitation of the framework as this weakening is actually due to the very nature of the alignment problem itself. The inability to guarantee the maximality of the average payoff is due to the fact that the original object (M) could contain large areas outside the overlapping subset that are perfectly identical. Further, objects that are able to slide (for instance a plane or a sphere) could allow to move between different mixed strategies without penalty. These situations cannot be addressed by any algorithm without re-

lying on supplementary information. However, in practice, they are quite unlikely, exceptional cases. In the experimental section we will show that our approach can effectively register even quasi-planar surfaces.

2.3. Building the Mating Strategies Set

From a theoretical point of view the total number of mating strategies in a registration problem is $|M \times D|$, which can be very large even with medium-sized surfaces. In practice, it is possible to apply several heuristics to select a lower number of candidates while still achieving good alignment results. Since the proposed approach is very selective it is not necessary to use all the model points: even a highly aggressive subsampling does not affect the registration quality, provided that some points in the overlapping region between model and data are retained. In fact, our approach does not try to find a good registration by means of a vote counting validation; instead it takes quite the opposite route, by self-validating the selected mixed strategy exploiting its internal coherence. Once the model points have been subsampled, the mating strategies set could be created by pairing each one of them with all the data points. Again, while this approach would work, it is somewhat wasteful since most of the mating strategies could be dropped on the basis of some local property of the surface surrounding the model and data point. For instance, the mean or Gaussian curvatures can be compared or some surface feature can be calculated in order to select only meaningful pairs. In the experimental section we will suggest an effective selection strategy. Once a proper set of mates has been chosen, a payoff function is needed. In principle, any proper one-to-one symmetric rigidity-enforcing payoff function could be used to capture the coherence between pairs of mating strategies. From a practical point of view it is often advisable to use bounded functions, usually in the interval $(0, 1]$. Very good candidates are the negative exponentiation of the difference

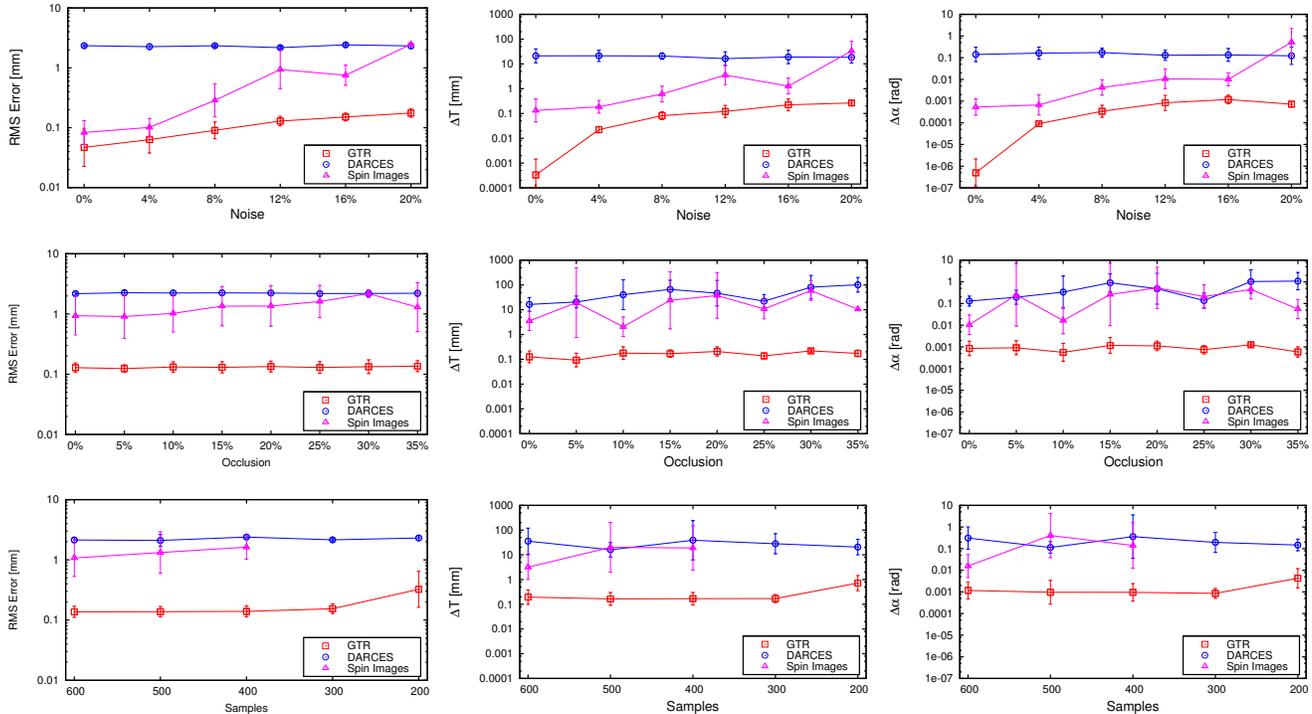


Figure 3. Comparison of coarse registration techniques using real range data, measuring ground RMS (column 1), translation (column 2) and rotation (column 3) errors as a function of noise (row 1), occlusion (row 2) and number of samples (row 3).

between the distances of the model and data points, or the ratio between the min and the max distance. In general, the steeper is the function, the more selective is the choice of the inlier mating strategies.

2.4. Evolving to an Optimal Solution

The search for a stable state is performed by simulating the evolution of a natural selection process. Under very loose conditions, any dynamics that respect the payoffs is guaranteed to converge to Nash equilibria [21] and (hopefully) to ESS's; for this reason, the choice of an actual selection process is not crucial and can be driven mostly by considerations of efficiency and simplicity. In this paper we chose to use the replicator dynamics, a well-known formalization of the selection process governed by the following equation

$$\mathbf{x}_i(t+1) = \mathbf{x}_i(t) \frac{(C\mathbf{x}(t))_i}{\mathbf{x}(t)^T C\mathbf{x}(t)} \quad (2)$$

where \mathbf{x}_i is the i -th element of the population and C the payoff matrix. A simple but complete example of the evolution process is shown in Figure 2.

Once the population has reached a local maximum, all the non-extincted mating strategies can be used to calculate the rigid transformation between data and model surfaces. A clear advantage of our approach is that in the final mixed strategy each pair of points is weighted proportionally to its degree of participation in the equilibrium (see Figure 2).

This is similar in spirit to the concept of compatibility between mates adopted by a number of fine registration algorithms, yet it does not depend at all on supplementary information such as surface normals or texture color. This compatibility can be used to weigh each pair when calculating the best surface alignment by using a weighted least squares fitting technique [11].

3. Experimental Results

Since the proposed technique can be used independently for coarse and fine registration, we evaluated its performance with respect to state-of-the-art algorithms of both fields. All the experiments have been executed on two sets of data: range images obtained from real-world scanners and synthetically-generated surfaces. For the first set of experiments we selected models from publicly available databases; specifically the Bunny [20], the Armadillo [13] and the Dragon [9] from the Stanford 3D scanning repository. To further assess the shortcomings of the various approaches, we used three synthetic surfaces representative of as many classes of objects: a wave surface, a fractal landscape and an incised plane (see Figure 4).

In all the experiments the set of mating strategies was obtained using the same selection technique. We used the MeshDOG [22] 3D feature detector to find interesting points in both the model and the data range images. A descriptor was associated to each point of interest; after considering both the MeshHOG and the Spin Image descrip-

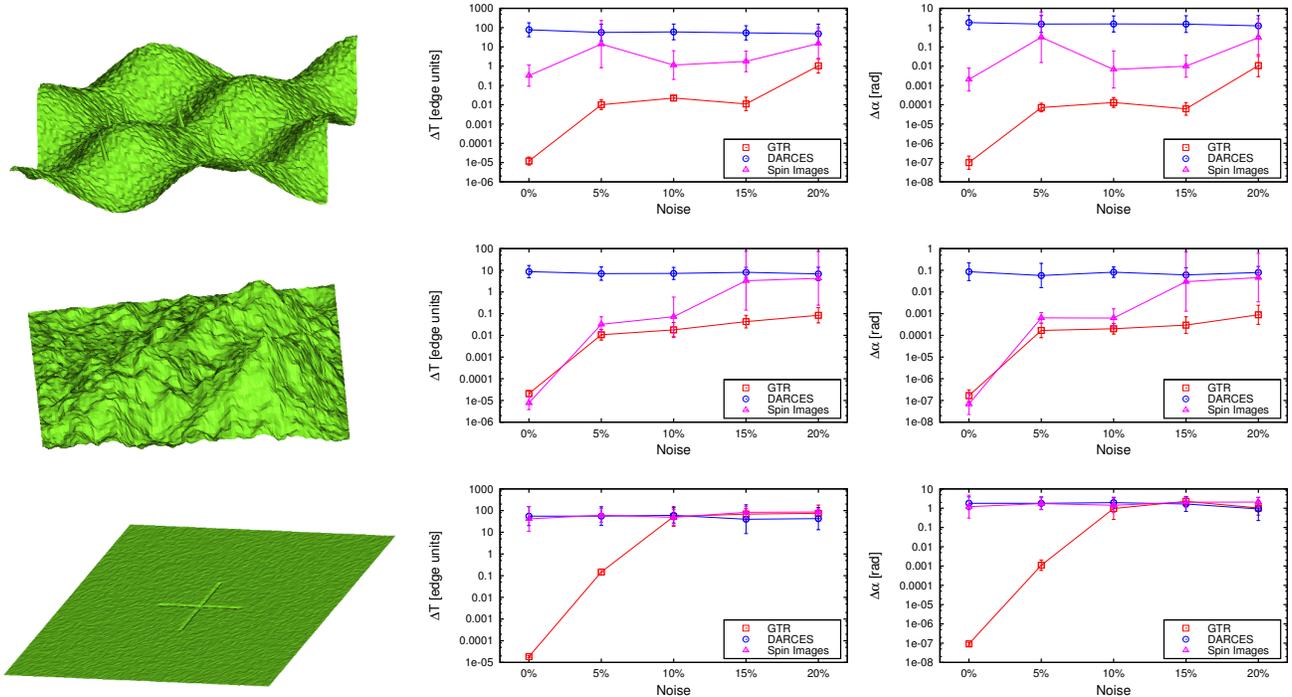


Figure 4. Comparison of coarse registration techniques using synthetic objects, measuring motion error as a function of noise.

tors, we preferred the latter as we found it to be more distinctive. Then a set of candidate source points was subsampled from the model and for each source point we created 5 mating strategies by connecting it to the 5 points with the most compatible descriptors. The rigidity-enforcing payoff function chosen was

$$\pi((a_1, b_1), (a_2, b_2)) = \frac{\min(|a_1 - a_2|, |b_1 - b_2|)}{\max(|a_1 - a_2|, |b_1 - b_2|)} \quad (3)$$

where a_1, a_2, b_1 and b_2 are respectively the two model (source) and data (destination) points in the compared mating strategies.

3.1. Coarse Registration

We compared our method with two coarse registration methods: RANSAC-based DARCES [5] and Spin Images [12]. DARCES has been implemented according to the original paper, while we used the Spin Images variant suggested in [4] to obtain a higher accuracy. In Figure 3 we show the results obtained using the set of surfaces from the Stanford repository. Each test was made under different conditions of noise, occlusion and subsampling and was run for a total of 12 times over the set of range images. For each set of experiments we plot the RMS distance for the actual point correspondences in the two meshes, and the estimation errors of the translation vector and rotation angle. In order to obtain a ground-truth for precise error measurement we generated the data points by adding Gaussian

noise, random occlusion and motion to the model points. In these experiments the surfaces were obtained from laser scans of objects of hundreds of millimeters in size, with a resolution of about one tenth of millimeter. The first row of Figure 3 plots the sensitivity to Gaussian noise exhibited by the different techniques. The noise level is expressed as the ratio between the standard deviation of the noise and the average edge length. While DARCES is not very sensitive to noise, it delivers by far the worst overall results. By contrast, Spin Images give fairly good results at low noise levels, but their performance worsens quickly as noise is increased. The proposed approach (GTR), on the other hand, exhibits errors that are consistently an order of magnitude below Spin Images. In the second row we show the effect of occlusion under a constant level of Gaussian noise with standard deviation equal to 12% of the average edge length. The results show that the tested techniques are substantially insensitive to occlusions, our technique constantly outperforming the other approaches. Finally, the third row shows the effect of subsampling. Our game-theoretic method outperforms the other approaches. Note that the Spin Images based technique was never able to find a correct transformation when provided with less than 300 samples.

Figure 4 plots the alignment results on the three synthetic surfaces. Each set of experiments was conducted over a single type of surface (displayed at the beginning of the row) with 12 runs for each technique and noise level. Since these objects are synthetic, errors on translation are expressed

in edge units. The “wave” test object (first row) offers a regular surface with few outstanding features and high redundancy of the pattern; in this scenario the Spin Images technique is affected by the inability to discern among a large amount of similar descriptors, thus it performs poorly at all noise levels. Conversely, the geometric-based consensus exploited by our registration approach allows for a more precise selection and thus a more accurate registration. The “fractal landscape” test object (second row) is an irregular surface that allows to produce very distinctive feature descriptors. In fact, with low levels of noise both Spin Images and our technique perform very well, albeit as noise increases we achieve better results. Finally, the “incised plane” object (third row) is a big flat domain with a small cross just half an edge deep. This represents a very difficult target for most registration techniques, since very few and faint features are available, while a large planar surface dominates the landscape. Despite the lack of good detectable points, our technique is able to register the surface as long as noise is minimal. With higher noise levels the bumped cross fades and becomes almost indistinguishable from the plane itself. Note that DARCES achieves mediocre results under all tested conditions.

3.2. Fine Registration

The performance of our approach with respect to fine registration has been studied in a separate batch of experi-

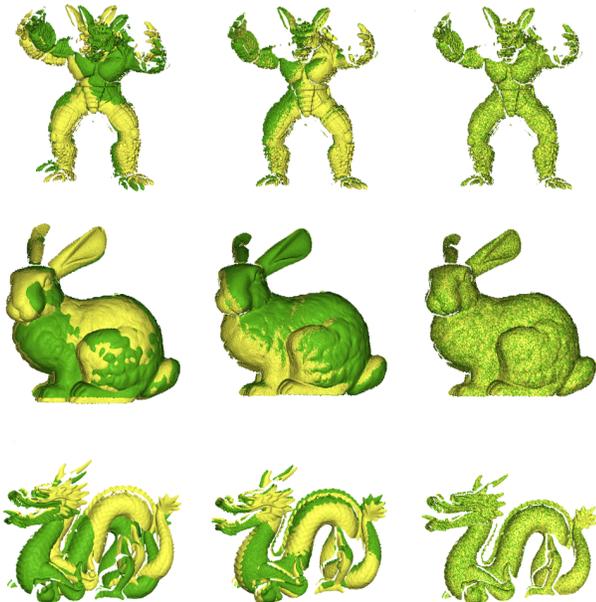


Figure 5. Examples of surface registration obtained respectively with RANSAC-based DARCES (first column), Spin Images (second column), and our game-theoretic registration technique (third column).

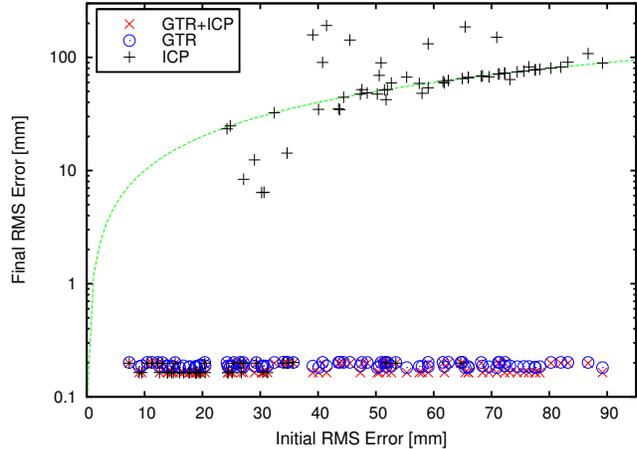


Figure 6. Comparison of fine registration accuracies (the green dashed line represents $y=x$). Graph best viewed in color.

ments. The goal of this test is two-fold: we want to evaluate our quality as a complete alignment tool and, at the same time, find the breaking point for traditional fine registration techniques. The method we used for comparison is a best-of-breed ICP variant, similar to the one proposed in [20]. Point selection is based on Normal Space Sampling [16], and point-surface normal shooting is adopted for finding correspondences; distant mates or candidates with back-facing normals are rejected. To minimize the influence of incorrect normal estimates, matings established on the boundary of the mesh are also removed. The resulting pairings are weighted with a coefficient based on compatibility of normals, and finally a 5%-trimming is used. Each test was performed by applying a random rotation and translation to different range images selected from the Stanford 3D scanning repository. Additionally, each range image was perturbed with a constant level of Gaussian noise with standard deviation equal to 12% of the average edge length. We completed 100 independent tests and for each of them we measured the initial RMS error between the ground-truth corresponding points and the resulting error after performing a full round of ICP (ICP) and a single run of our registration method (GTR). In addition, we applied a step of ICP to the registration obtained with our method (GTR + ICP) in order to assess how much the solution extracted using our approach was further refinable. A scatter plot of the obtained errors before and after registration is shown in Figure 6. We observe that ICP reaches its breaking point quite early; in fact with an initial error above the threshold of about 20mm it is unable to find a correct registration. By contrast, GTR is able to obtain excellent alignment regardless of the initial motion perturbation. Finally, applying ICP to GTR decreases the RMS only by a very small amount.

While we did not carry out any formal benchmark of the execution time required by our technique, we always observed a very fast convergence of the replicator dynamics, even with several thousands of mating strategies. In the

worst scenarios our unoptimized C++ implementation¹ of the framework required less than 2 seconds (on a typical desktop PC) to evolve a population of 4000 to a stable state.

4. Conclusions

In this paper we introduced a novel game-theoretic technique that solves both the coarse and fine surface registration problems at once. Our approach has several advantages over the state-of-the-art: it does not require any kind of initial motion estimation, as it does not rely on spatial relationships between model and data points, it does not need any threshold as it produces a continuous compatibility weight on the selected point matches that can be used directly for alignment estimation, and, differently from most inlier selection techniques, it is not affected by a large number of outliers since it operates an explicit search for good inliers rather than using random selection or vote counting for validation. From a theoretical point of view, a sound correspondence between optimal alignments and evolutionary equilibria has been presented and a wide range of experiments validated both the robustness of the approach with respect to noise and its performance in comparison with other well-known techniques.

Acknowledgments

We acknowledge the financial support of the Future and Emerging Technology (FET) Programme within the Seventh Framework Programme for Research of the European Commission, under FET-Open project SIMBAD grant no. 213250.

References

- [1] D. Aiger, N. J. Mitra, and D. Cohen-Or. 4-points congruent sets for robust surface registration. *ACM Transactions on Graphics*, 27(3):#85, 1–10, 2008. **1**
- [2] A. Albarelli, S. R. Bulò, A. Torsello, and M. Pelillo. Matching as a non-cooperative game. In *ICCV 2009: Proc. of the 2009 IEEE Intl. Conf. on Comput. Vis.* IEEE Computer Society, 2009. **2, 3**
- [3] P. J. Besl and N. D. McKay. A method for registration of 3-d shapes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(2):239–256, 1992. **1**
- [4] O. T. Carmichael, D. F. Huber, and M. Hebert. Large data sets and confusing scenes in 3-d surface matching and recognition. In *3DIM*, pages 358–367, 1999. **6**
- [5] C.-S. Chen, Y.-P. Hung, and J.-B. Cheng. Ransac-based darcs: A new approach to fast automatic registration of partially overlapping range images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(11):1229–1234, 1999. **1, 6**
- [6] Y. Chen and G. Medioni. Object modeling by registration of multiple range images. In *Proc. 1991 IEEE Intl. Conf. on Robotics and Automat.*, pages 2724–2729 vol.3, 1991. **1**
- [7] C. S. Chua and R. Jarvis. Point signatures: A new representation for 3d object recognition. *Intl. J. of Comput. Vis.*, 25(1):63–85, October 1997. **1**
- [8] D. H. Chung, I. D. Yun, and S. U. Lee. Registration of multiple-range views using the reverse-calibration technique. *Pattern Recognition*, 31(4):457–464, 1998. **1**
- [9] B. Curless and M. Levoy. A volumetric method for building complex models from range images. In *Proc. of SIGGRAPH 96*, pages 303–312, New York, NY, USA, 1996. ACM. **5**
- [10] S. Granger, X. Pennec, and A. Roche. Rigid point-surface registration using an em variant of icp for computer guided oral implantology. In *Proc. of the 4th Intl. Conf. on Medical Image Comput. and Computer-Assisted Interv.*, pages 752–761, London, UK, 2001. Springer-Verlag. **1**
- [11] B. K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. of the Optical Society of America. A*, 4(4):629–642, Apr 1987. **2, 5**
- [12] A. E. Johnson and M. Hebert. Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(5):433–449, 1999. **1, 6**
- [13] V. Krishnamurthy and M. Levoy. Fitting smooth surfaces to dense polygon meshes. In *Proc. of SIGGRAPH 96*, pages 313–324, 1996. **5**
- [14] Y. Liu. Replicator dynamics in the iterative process for accurate range image matching. *Intl. J. Comput. Vis.*, 83(1):30–56, 2009. **1**
- [15] T. Masuda. Registration and integration of multiple range images by matching signed distance fields for object shape modeling. *Comput. Vis. Image Underst.*, 87(1-3):51–65, 2002. **1**
- [16] S. Rusinkiewicz and M. Levoy. Efficient variants of the icp algorithm. In *Proc. of the Third Intl. Conf. on 3D Digital Imaging and Modeling*, pages 145–152, 2001. **1, 7**
- [17] J. Salvi, C. Matabosch, D. Fofi, and J. Forest. A review of recent range image registration methods with accuracy evaluation. *Image Vis. Comput.*, 25(5):578–596, 2007. **1**
- [18] J.-P. Tarel, H. Civi, and D. B. Cooper. Pose estimation of free-form 3d objects without point matching using algebraic surface models. In *Proc. of IEEE Workshop Model Based 3D Image Analysis*, pages 13–21, Mumbai, India, 1998. **1**
- [19] A. Torsello, S. R. Bulò, and M. Pelillo. Grouping with asymmetric affinities: A game-theoretic perspective. In *CVPR '06: Proc. of the 2006 IEEE Conf. on Comput. Vis. and Pattern Recognit.*, pages 292–299, Washington, DC, USA, 2006. IEEE Computer Society. **2**
- [20] G. Turk and M. Levoy. Zipped polygon meshes from range images. In *Proc. of SIGGRAPH 94*, pages 311–318, New York, NY, USA, 1994. ACM. **5, 7**
- [21] J. Weibull. *Evolutionary Game Theory*. MIT Press, 1995. **2, 5**
- [22] A. Zaharescu, E. Boyer, K. Varanasi, and R. P. Horaud. Surface feature detection and description with applications to mesh matching. In *Proc. of the IEEE Conf. on Comput. Vis. and Pattern Recognit.*, Miami Beach, Florida, June 2009. **5**
- [23] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Intl. J. Comput. Vis.*, 13(2):119–152, 1994. **1**

¹A reference implementation of the framework will be made available for downloading at <http://www.dsi.unive.it/~rodola/sw.html>