SERIE B — INFORMATIK

A Simple and Robust Geometric Algorithm for Landmark Registration in Computer Assisted Neurosurgery

Frank Hoffmann[*]
Klaus Kriegel[*,◇]
Sven Schönherr[*,♯]
Carola Wenk[*,○]

B 99-21
December 1999

Abstract

We present a simple geometric and combinatorial algorithm for the approximate partial matching problem of small 3D point landmark patterns. It has been designed for and successfully applied in a computer aided neurosurgery navigation system. The algorithm relies mainly on a geometric voting and scoring technique.

* Institut für Informatik, Freie Universität Berlin, Takustr. 9, 14195 Berlin, Germany,
e-mail: `<name>@inf.fu-berlin.de`

◇ Deutsches Herzzentrum Berlin, Augustenburger Platz, 13353 Berlin, Germany.
This author was supported by Deutsche Forschungsgesellschaft, grant FL 165/4-2.

♯ Institut für Theoretische Informatik, ETH Zürich, ETH-Zentrum, 8092 Zürich, Switzerland.
This author was supported by the ESPRIT IV LTR Project No. 21855 (GALIA).

○ This author was supported by Deutsche Forschungsgemeinschaft, grant AL 253/4-2.

1 Introduction

1.1 The Application Background

We address a geometric point pattern matching question stemming from a problem that occurred in the development of a computer aided neurosurgery navigation system NEN [5] at Functional Imaging Technologies GmbH. The aim of that system is to enable a surgeon to utilize pre-operatively gathered information more effectively during the operation. This information could be for example a 3D–map of functional regions in human brain obtained by functional magnetic resonance imaging (MR). The surgeon of course wants to bypass functional regions. Now the overall idea consists in linking the information from the MR image directly to the operation environment for navigation purposes. Like in most such guidance systems (see for example [2] and the references there) the basic approach to solve this question relies on a technique called image registration. A set of markers is attached to the skin (or even bone implanted in other systems) and special points on the markers serve as orientation landmarks. These landmarks are visible in the pre-operative volumetric MR images and define therefore a 3D point set. Then, immediately prior to the operation, a 3D–tracking system can be used to localize (i.e. to 'register') the landmark points a second time in the operation environment while the patient's head is fixed. This, again, defines a 3D point set. Assuming the rigidity of the head, we are now able to relate the actual surgery position to the corresponding location in the MR image via the unique rigid transformation defined by the two point sets.

In practice however, we cannot assume such an ideal situation. Firstly, rigidity does not completely hold, since external forces are applied during the measurements or while fixing the patient's head. Secondly, we may be faced with distorted or missing landmarks in both point sets. This could be caused for example by noise/artifacts in the MR image, markers torn off unintentionally by the patient, or markers inaccessible for registration. Finally, in the application markers are intentionally not distinguishable by labels. The system designers wanted to keep both the procedure of attaching the markers to the skin and of registering them in the operation environment as simple as possible. Nevertheless, even under these weak assumptions the aim was to have a reliable algorithm that checks whether there is a unique nearly rigid transformation from world data to image points consistent with almost all landmarks within prescribed tolerances. The algorithm should compute an approximation of this mapping in case of an affirmative answer. Figure 1 shows cylindric markers as used in the system attached to a dummy skull. The landmark points are on top of the markers.

1.2 Mathematical Modeling and Algorithm Outline

Suppose that a number of landmark points mounted on a three-dimensional rigid body have been localized in two different coordinate systems representing world space and image space, respectively. This way we are given a set P of n world points and a set Q of m image points. We want to find an affine transformation A from world space into image space, which approximates a 1-to-1 correspondence between P and Q. Ideally, P and Q have equal size and A is the unique rigid transformation that maps each point of P onto the corresponding point of Q. Our goal is to find two large subsets $P' \subseteq P$ and $Q' \subseteq Q$ together with a 1-to-1 correspondence between them (called subsequently *combinatorial matching*) and an affine transformation that is nearly a rigid motion and maps each $p \in P'$

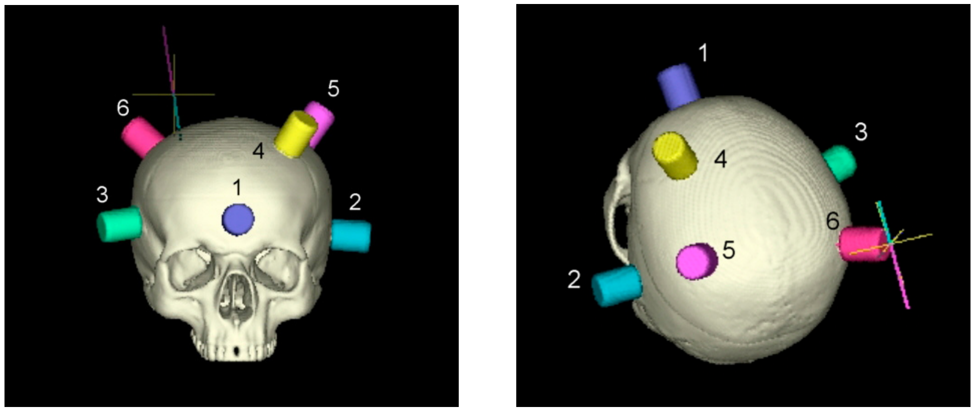

Figure 1: Two snapshots of a dummy skull with attached markers

approximately onto the corresponding $q \in Q'$. In [1] the reader finds a survey of geometric techniques developed for solving pattern matching problems.

Our approach is divided into three steps. The first one is a search for similar triangles in the point sets to build an $n \times m$ matrix S named *scoring table*. The value of an entry s_{ij}, called *score*, determines the relationship between $p_i \in P$ and $q_j \in Q$, i.e. a high (low) score represents a strong vote for matching (not matching) these two points. Based on the scoring table a *candidate table* is build in the second step. Each $q_j \in Q$ with a high score s_{ij} becomes a candidate for $p_i \in P$. From the candidate table we compute all possible combinatorial matchings. Finally, we determine for each combinatorial matching λ the best approximating affine transformation A_λ. If A_λ is nearly rigid (the term is defined later in dependence on a tolerance parameter, but should be intuitively clear), λ is a *geometric matching*. Usually a unique geometric matching will be found this way. Otherwise we return a list of all geometric matchings found.

2 Scoring Table

Let the two point sets $P = \{p^1, p^2, \dots, p^n\}$ and $Q = \{q^1, q^2, \dots, q^m\}$ and an error bound ε be given. We start with the key observation that a geometric matching maps world edges to image edges of nearly the same length and world triangles to nearly congruent image triangles. The idea of the scoring procedure is to turn this observation the other way round: each edge pair or triangle pair similarity gives a vote to match the corresponding points, such an approach was also used in [3]. The votes are collected in an $n \times m$ matrix S which is initialized with zero entries. The easiest way to fill S would be an edge comparison. Let E_P (E_Q) be the set of all edges in P (Q). For each pair of edges $e = (p^i, p^j) \in E_P$ and $e' = (q^u, q^v) \in E_Q$ check if they are ε-*similar*, i.e. whether their lengths differ by at most ε. A positive answer gives some evidence that q^u and q^v are candidates for p^i and p^j. Thus the scores s_{iu}, s_{jv}, s_{iv}, and s_{ju} are increased by

$$1 - \frac{|\,|e| - |e'|\,|}{\varepsilon}.$$

This *similarity value* is one for edges of equal length and tends to zero if edges exhaust the given error bound ε. If p^i corresponds to q^u in a geometric matching then s_{iu} will end up with a high score. We remark that the reverse implication is not always true for several reasons. Firstly, edges are symmetric objects and thus for any edge $e = (p^i, p^j)$ which is mapped to a similar edge $e' = (q^u, q^v)$ we have the same votes for s_{iu}, s_{iv}, s_{ju} and s_{jv}. Secondly, with increasing size of the point set the expected number of image edges similar to a given world edge also increases. Thus, we are faced with random noise in the scoring table which might cover the values corresponding to a geometric matching.

One can significantly reduce these difficulties replacing the role of edges by congruent triangles. Let T_P (T_Q) denote the set of all triangles in P (Q). Two triangles $t = (p^i, p^j, p^k) \in T_P$ and $t' = (q^u, q^v, q^w) \in T_Q$ are *ε-congruent* if the corresponding edges are ε-similar. In this case the scores of the corresponding point pairs are increased by the similarity value of the pairs of incident edges, e.g. score s_{iu} is increased by

$$\left(1 - \frac{\left||\overline{p^ip^j}| - |\overline{q^uq^v}|\right|}{\varepsilon}\right) + \left(1 - \frac{\left||\overline{p^ip^k}| - |\overline{q^uq^w}|\right|}{\varepsilon}\right).$$

Clearly, it is sufficient to consider only one representation for each triangle in T_P whereas all permutations of the vertices of each triangle in T_Q have to be considered.

Given a world triangle t we denote by $C_\varepsilon(t)$ the set of all image triangles which are ε-congruent to t. We remark that the expected number $E[|C_\varepsilon(t)|]$ of image triangles ε-congruent with a randomly chosen world triangle t is significantly smaller than the expected number of image edges ε-similar to a random world edge. Moreover, only few triangles are (nearly) symmetric (i.e. ε-isosceles or ε-equilateral). Thus, in contrast to scoring by edge similarities, now in the similarity scores will be updated for only one permutation of the vertices with the before mentioned exceptions. This shows that the random noise in the scoring table is essentially smaller if ε-congruent triangles are used.

The only drawback caused by the triangle method is an increase of the running time. The naive approach requires $\Omega(n^3 m^3)$ triangle comparisons. We introduced a more efficient data structure sorting triangles lexicographically according to the lengths of the longest, the middle, and, the shortest edge. Thus, for any world triangle t with maximal edge length l the set $C_\varepsilon(t)$ can be limited performing a binary search for the first and last triangle t' with maximal edge length $l' \geq l - \varepsilon$ and $l' \leq l + \varepsilon$, respectively.

3 Combinatorial Matchings

A scoring table S provides a lot of useful information about geometric matchings. A small score s_{ij} testifies against matching p^i with q^j, whereas a high score $s_{i'j'}$ is an advise that there should be a geometric matching mapping $p^{i'}$ onto $q^{j'}$. However, if there are several (relatively) high scores in the ith row we have several candidates which can be matched with p^i and the highest one may not always be the best one. Thus, we have to check several combinatorial matchings. In the next section we will see that many such matchings can be excluded because they are not geometric matchings or because they are reflections. The problem is to find a balance ensuring that the desired geometric matching will be found but at the same time keeping the total number of matching to be checked small.

In practice the following approach proved to be very useful. Let α_i, β_j, and γ denote the maximal scores in the i-th row of S, in the j-th column of S, and in the whole matrix S, respectively. An image point q_j is a candidate of p_i if the following condition holds:

$$s_{ij} \geq \max\left\{\frac{\alpha_i}{2}, \frac{\beta_i}{2}, \frac{\gamma}{4}\right\}.$$

The parameters 1/2, 1/2, and 1/4 have been fixed by numerous experiments. If $\alpha_i < \gamma/2$ for some p_i, i.e. there is no candidate with very high score for this world point, a formal symbol `unmatched` is inserted in the list of candidates. The collection of all candidate lists forms the candidate table C. A combinatorial matching generated from C selects for each world point one of its candidates such that each image point is used at most once. This guarantees that world points with high scored candidates are always matched. The list of all possible combinatorial matchings is computed with a back-tracking method.

4 Least-Squares Approximation

Given a combinatorial matching, i.e. two large subsets $P' \subseteq P$ and $Q' \subseteq Q$ together with a 1-to-1 correspondence between them, assume for simplicity of notation that $P' = \{p^1, p^2, \ldots, p^l\}$ and $Q' = \{q^1, q^2, \ldots, q^l\}$ and that p^ν corresponds to q^ν for all $\nu \in \{1, \ldots, l\}$. In the sequel a homogeneous representation of the points is used, i.e. $\tilde{p}^\nu = (p_1^\nu, p_2^\nu, p_3^\nu, 1)^T$ and $\tilde{q}^\nu = (q_1^\nu, q_2^\nu, q_3^\nu, 1)^T$. The goal is to find an affine transformation A with

$$A = \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{1}$$

such that

$$\sum_{\nu=1}^{l} w^\nu ||A\,\tilde{p}^\nu - \tilde{q}^\nu||^2 \tag{2}$$

is minimal. Here, the symbol w^ν denotes an additional weight function. The standard approach would work with $w^\nu = 1$ for all $\nu \in \{1, \ldots, l\}$. Taking into account that points p^ν which are far from the center of gravity are more essential for the alignment of the transformation than closer points, the weight w^ν is defined as $||p^\nu - \overline{p}||$, i.e. as the distance between the world point p^ν and the center of gravity $\overline{p}$ of all matched world points $p^1, \ldots, p^l$.

The quadratic function (2) in the variables a_{ij} assumes a unique minimum if and only if the gradient vector vanishes, i.e.

$$\frac{\partial \sum_{\nu=1}^{l} w^\nu ||A\,\tilde{p}^\nu - \tilde{q}^\nu||^2}{\partial a_{ij}} = 0 \tag{3}$$

for all $i \in \{1, 2, 3\}$, $j \in \{1, \ldots, 4\}$.

Let $p \in P$ and $q \in Q$ be corresponding points. Using (1) gives

$$\begin{aligned} ||A\,\tilde{p} - \tilde{q}||^2 &= \sum_{i=1}^{3}\Big(\Big(\sum_{j=1}^{4} a_{ij}\tilde{p}_j\Big) - \tilde{q}_i\Big)^2 \\ &= \sum_{i=1}^{3}\Big(\Big(\sum_{j=1}^{4} a_{ij}\tilde{p}_j\Big)^2 - 2\tilde{q}_i \sum_{j=1}^{4} a_{ij}\tilde{p}_j + \tilde{q}_i^2\Big). \end{aligned}$$

Note that the upper border of index i is three (not four) because the product of the last row of A with $\tilde{p}$ is one, and so is the last coordinate of $\tilde{q}$. Consider the partial derivative for some a_{ij},

$$\frac{\partial||A\tilde{p}-\tilde{q}||^2}{\partial a_{ij}} = 2a_{ij}\tilde{p}_j^2 + 2\tilde{p}_j \sum_{k=1,k\neq j}^{4} a_{ik}\tilde{p}_k - 2\tilde{q}_i\tilde{p}_j \tag{4}$$

$$= \sum_{k=1}^{4} 2\tilde{p}_j\tilde{p}_k a_{ik} - 2\tilde{p}_j\tilde{q}_i. \tag{5}$$

Substituting (5) into (3) results in

$$\sum_{k=1}^{4}\Big(\sum_{\nu=1}^{l} w^\nu \tilde{p}_j^\nu \tilde{p}_k^\nu\Big) a_{ik} = \sum_{\nu=1}^{l} w^\nu \tilde{p}_j^\nu \tilde{q}_i^\nu, \tag{6}$$

for all $i \in \{1,2,3\}$, $j \in \{1,\dots,4\}$. Note that a factor of two has been cancelled from each equation. The coefficients of A which minimize (2) are obtained by solving this system of 12 linear equations in 12 variables as follows:

A more careful analysis of (6) shows that the variables and the equations can be partitioned into three independent groups. Each group consists of four variables and four equations and corresponds to one row in A. This leads to the following three linear equation systems

$$\left(\sum_{\nu=1}^{l}\begin{pmatrix} w^\nu p_1^\nu p_1^\nu & w^\nu p_1^\nu p_2^\nu & w^\nu p_1^\nu p_3^\nu & w^\nu p_1^\nu \\ w^\nu p_2^\nu p_1^\nu & w^\nu p_2^\nu p_2^\nu & w^\nu p_2^\nu p_3^\nu & w^\nu p_2^\nu \\ w^\nu p_3^\nu p_1^\nu & w^\nu p_3^\nu p_2^\nu & w^\nu p_3^\nu p_3^\nu & w^\nu p_3^\nu \\ w^\nu p_1^\nu & w^\nu p_2^\nu & w^\nu p_3^\nu & w^\nu \end{pmatrix}\right)\begin{pmatrix} a_{i1} \\ a_{i2} \\ a_{i3} \\ a_{i4}\end{pmatrix} = \sum_{\nu=1}^{l}\begin{pmatrix} w^\nu p_1^\nu q_i^\nu \\ w^\nu p_2^\nu q_i^\nu \\ w^\nu p_3^\nu q_i^\nu \\ w^\nu q_i^\nu \end{pmatrix}$$

with $i \in \{1,2,3\}$. Note that the matrix does not depend on i, so it is stored only once. To solve the systems of linear equations an LU-decomposition of this matrix is computed.

Finally, for testing whether the resulting transformation A is nearly rigid, consider the matrix $R = (a_{ij})_{1\le i,j\le 3}$, i.e. the linear part of the affine transformation A. The following necessary conditions for R are checked:

$$|\det(R) - 1| \le \delta, \tag{7}$$

$$||R(e_i)| - 1| \le \delta, \quad i \in \{1,2,3\}, \tag{8}$$

$$||R((1,1,1) - e_i)| - \sqrt{2}| \le \delta, \quad i \in \{1,2,3\}, \tag{9}$$

$$||R((1,1,1))| - \sqrt{3}| \le \delta, \tag{10}$$

with e_i being the unit vectors and for δ the value 0.1 is chosen in the implementation. Condition (7) guarantees $\det(R) > 0$ (i.e. A contains no reflection) and $\det(R) \approx 1$ (the determinant is ± 1 for rigid transformations). The remaining conditions are relaxations of the fact, that rigid transformations do not scale, i.e. they are length preserving.

5 Implementation

Our algorithm has been implemented in C++. We followed the generic programming paradigm as it is realized in the standard template library (STL) [4]. The whole computation is encapsulated in a class which provides an interface to easily adjust the parameters.

We tested the implementation on several different input sets. Besides computer generated point sets to check degenerated configurations, we used data stemming from dummy skulls as well as patients' data from real world surgeries. A selection of these data sets is shown in Table 1.

name	**#points**		**description**
	world	**image**	
dummy1	5	5	dummy skull (MRI)*
dummy2	5	4	dummy skull (MRI)* one marker not detected in preprocessing
dummy3	14	5	dummy skull (MRI)* nine "virtual" markers measured at random
dummy4	8	8	dummy skull (CT)**
dummy5	6	8	dummy skull (CT)** two markers not measured
phantom1	4	10	second dummy skull (MRI)* only four markers measured
phantom2	6	10	second dummy skull (MRI)* only six markers measured
patient1	5	5	patient's head (MRI)*
patient2	5	6	patient's head (MRI)* one marker lost during surgery preparation
equal	6	6	two equal sets (generated)
cube	6	6	six vertices of a 3-cube (slightly perturbed)

Table 1: Data sets.

The algorithm computed the right geometric matching in all cases. Table 2 gives the results, i.e. the number of matched point pairs, the determinant of the resulting transformation, and the mean and maximal point deviation. Note that the missing point pair of input sets 'dummy4' and 'phantom2' is due to one badly measured marker. All running times are far below one second, which is completely satisfactory for the designated application.

Our implementation has been integrated in the computer aided surgery system NEN [5] and gives very good results in practice.

Acknowledgements

We would like to thank the team of Functional Imaging Technologies GmbH for a stimulating and fruitful cooperation. Moreover we are indebted to the Universitätsklinikum Benjamin Franklin of Freie Universität Berlin where the navigation system has been successfully tested.

*Magnetic Resonance Imaging.

**Computer Tomography.

name	**#matched pairs**	**deter-minant**	**point deviation mean**	**max.**
dummy1	5 of 5	1.04	1.48	2.60
dummy2	4 of 4	0.95	0.00	0.00
dummy3	5 of 5	1.04	0.64	1.10
dummy4	7 of 8	0.98	2.74	6.46
dummy5	6 of 6	1.01	0.41	0.98
phantom1	4 of 4	1.00	0.00	0.00
phantom2	5 of 6	1.01	0.13	0.21
patient1	5 of 5	0.97	0.51	1.27
patient2	5 of 5	0.95	0.45	1.15
equal	6 of 6	1.00	0.00	0.00
cube	6 of 6	1.01	0.55	0.80

Table 2: Test Results.

References

[1] H. Alt and L. Guibas. Discrete geometric shapes: Matching, interpolation, and approximation. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 121 – 153. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 1999.

[2] J. M. Fitzpatrick, J. B. West, and C. R. Maurer. Predicting error in rigid-body point-based registration. *IEEE Transactions on Medical Imaging*, 17(5):694–702, 1998.

[3] F. Hoffmann, K. Kriegel, and C. Wenk. An applied point pattern matching problem: comparing 2d patterns of protein spots. *Discrete Applied Mathematics*, 93:75–88, 1999.

[4] D. Musser and A. Saini. *STL Tutorial and Reference Guide: C++ Programming with the Standard Template Library*. Addison-Wesley, 1996.

[5] A. Schleich, S. Jovanovic, B. Sedlmaier, U. Warschewske, F. Oltmanns, S. Schönherr, W. Oganovsky, B. Ohnesorge, T. Hoell, and H. Scherer. Electromagnetic ENT navigation: the NEN system. In *Proc. 4th Annu. Conf. International Society for Computer Aided Surgery*, 2000.