

Automatic symmetry plane estimation of bilateral objects in point clouds

Benoît Combès, Robin Hennessy, John Waddington, Neil Roberts, Sylvain Prima

► **To cite this version:**

Benoît Combès, Robin Hennessy, John Waddington, Neil Roberts, Sylvain Prima. Automatic symmetry plane estimation of bilateral objects in point clouds. IEEE Conference on Computer Vision and Pattern Recognition (CVPR'2008), Jun 2008, Anchorage, United States. 10.1109/CVPR.2008.4587605 . inria-00331758

HAL Id: inria-00331758

<https://hal.inria.fr/inria-00331758>

Submitted on 17 Oct 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Automatic symmetry plane estimation of bilateral objects in point clouds

Benoît Combès
INRIA*

bcombes@irisa.fr

Robin Hennessy
RCSI†

John Waddington
RCSI†

Neil Roberts
MARIARC‡

Sylvain Prima
INRIA*

sprima@irisa.fr

Abstract

In this paper, the problem of estimating automatically the symmetry plane of bilateral objects (having perfect or imperfect mirror symmetry) in point clouds is reexamined. Classical methods, mostly based on the ICP algorithm, are shown to be limited and complicated by an inappropriate parameterization of the problem. First, we show how an adequate parameterization, used in an ICP-like scheme, can lead to a simpler, more accurate and faster algorithm. Then, using this parameterization, we reinterpret the problem in a probabilistic framework, and use the maximum likelihood principle to define the optimal symmetry plane. This problem can be solved efficiently using an EM algorithm. The resulting iterative scheme can be seen as an ICP-like algorithm with multiple matches between the two sides of the object. This new algorithm, implemented using a multiscale, multiresolution approach, is evaluated in terms of accuracy, robustness and speed on ground truth data, and some results on real data are presented.

1. Introduction

Most natural or manufactured objects exhibit some form of *symmetry*. In mathematical terms, it means that they remain invariant or almost invariant under certain transformations, typically: translations, rotations (with respect to a point or a line) and reflections (with respect to a point, a line or a plane). A single object can often display different types of symmetry. In computer vision, it is often crucial to use these symmetry features for common image processing tasks, such as recognition, denoising, registration, segmentation, *etc.* There is an extensive literature concerning the characterisation of symmetries of objects in 2D or 3D grey level images [23, 17]. Most of the works have been dedicated to the computation of the symmetry axis/plane

of bilateral objects, with a specific focus on magnetic resonance images (MRI) of the human brain [18, 13]. Interestingly, several methods have been developed separately for the computation of such a reflection axis/plane in point clouds, with a focus on the human face. Some applications include: shape completion/reconstruction [19], face authentication and recognition [22, 15, 17], surgical planning [14], intra-operative assistance [2], *etc.*

Two main approaches have been proposed to compute the symmetry axis/plane of bilateral objects in point clouds. In a first approach, the normals to the object under study are computed and stored in an orientation histogram to build the extended Gaussian image (EGI) associated with the object. If an object is symmetric, then so is its EGI. The position and orientation of the symmetry plane is then computed on the EGI (or the smoothed EGI) using eigen analysis of its inertia matrix and local optimisation techniques [15, 17]. However, practically this method is dedicated to structured point clouds (to be able to compute the normals) and is not robust to important occlusions.

In a second approach [24, 22, 14, 23] that we call the rigid-body transformation-based approach, the symmetry plane is estimated using a two-step algorithm. First, an initial plane K is chosen and the rigid-body transformation T best superposing the original data with its reflection with respect to K is computed. Secondly, the symmetry plane of the object under study is deduced from the resulting transformation T . The major advantage of this approach is to benefit from the ICP algorithm [3] relying on known closed-form solutions for the absolute orientation problem [10, 8, 1, 11, 21] to find T . However, computing the symmetry plane from the optimal rigid-body transformation is an ill-posed problem that needs *ad hoc* choices leading to different solutions.

In Section 2, we discuss the rationale behind the rigid-body transformation-based approach and show its limitations. In Section 3, we develop a simpler approach based on an explicit formulation of the searched plane. We derive a closed-form solution for the problem of finding the optimal symmetry plane when the matches between the two sides on the object are known, and then we propose an

* INRIA, VisAGeS Unité/Projet, F-35042 Rennes, France.

INSERM, U746, Faculté de Médecine CS 34317, F-35043 Rennes Cedex, France.

Univ. Rennes I, CNRS, UMR 6074, IRISA, F-35042 Rennes, France

† St Stephens Green, Dublin 2, Ireland

‡ University of Liverpool, Pembroke Place, Liverpool, L69 3GE, UK

iterative, ICP-like, algorithm to compute this plane when these matches are unknown. In Section 4, we experimentally show that this approach outperforms the rigid-body transformation-based approach in terms of computational time, accuracy and simplicity. Then, in Section 5, using this purely reflection-based approach, we reformulate the problem in probabilistic terms and estimate the optimal symmetry plane using the maximum likelihood principle and the expectation-maximisation (EM) algorithm. In Section 6, we evaluate the accuracy, robustness and computational time of this multiscale, multiresolution algorithm on synthetic, ground truth data, before proposing some useful applications on real data.

2. Rigid-body transformation-based methods

2.1. Formulation

Without loss of generality, we consider in the rest of this article that the object under study is represented by a cloud of points noted O , with $\text{card}(O) = N$. For an ideal bilateral object having a perfect symmetry, there exists a plane P superposing each point x with its counterpart in the other side of the object. The fact that P is a perfect symmetry plane simply writes:

$$O = S_P(O)$$

Based on the observation that for any plane K , $S_K \circ S_K = I_3$ (where I_3 is the 3×3 identity matrix) and noting $T = S_P \circ S_K$, the previous relation can be rewritten as:

$$\forall K, O = T \circ S_K(O)$$

Then one can note that:

- if $K \parallel P$ then $S_P \circ S_K$ is a translation
- if $K \not\parallel P$ then $S_P \circ S_K$ is a rotation

In other words, for any plane K , $T = S_P \circ S_K$ is a rigid-body transformation. Based on this observation, and on the need to deal with only grossly symmetrical objects, a first, intuitive approach to compute the symmetry plane of O is to fix a plane K and to find the rigid-body transformation \tilde{T} best superposing O and $S_K(O)$. In practice, this approach can be implemented efficiently using the classical ICP algorithm (or one of its numerous variants) between the two point clouds. Once K is fixed, \tilde{T} is then simply defined as:

$$\tilde{T} = \arg \min_{T, y_1, \dots, y_N} \sum_{x_i \in O} \|y_i - T \circ S_K(x_i)\|^2 = \arg \min_{T, y_1, \dots, y_N} \mathcal{E}_T,$$

where the points y_1, \dots, y_N belong to the cloud O . The ICP algorithm then iterates between computing the optimal matches $\{y_i\}$ while T is kept fixed (the solution is simply to choose y_i as the closest point of $T \circ S_K(x_i)$ in O) and then

computing a new estimate of T given these matches $\{y_i\}$. Different closed-form solutions exist for this last problem using, for instance the unit quaternions [10, 8] of the singular value decomposition [1, 11, 21].

2.2. Limitations of this approach

Unfortunately, this approach suffers from a major flaw: the optimal transformation \tilde{T} , composed with S_K , does not necessarily define a proper reflection, and thus the optimal plane \tilde{P} cannot be computed directly using $S_{\tilde{P}} = \tilde{T} \circ S_K$. For instance, if \tilde{T} is a pure translation that is not perpendicular to K then $\tilde{T} \circ S_K$ is not a reflection.

Three approaches have been proposed in the literature to estimate the optimal plane \tilde{P} from $\tilde{T} \circ S_K$, and are illustrated in Fig. 1:

- For each point $x \in O$, one can compute the perpendicular bisector plane P_x of the segment $[x, \tilde{T} \circ S_K(x)]$. The symmetry plane P_1 can then be estimated as the plane whose parameters (typically, unit normal and distance to the origin) are closest to that of the planes $P_x, x \in O$ in the least squares sense [22].
- Another solution is to compute the middle point m_x of the segment $[x, \tilde{T} \circ S_K(x)]$ for each point $x \in O$. The symmetry plane P_2 can then be estimated as the plane that best fits the points $m_x, x \in O$ in the least squares sense [24].
- A last approach is to consider the symmetry plane P_3 as the perpendicular bisector of the segment $[g(O), g \circ \tilde{T} \circ S_K(O)]$, where $g(\cdot)$ provides the centroid of a point cloud [14].

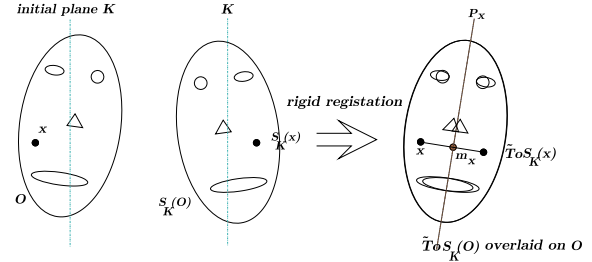


Figure 1. **Illustration of the rigid-body transformation-based method.** Example of computation of the plane P_x and the point m_x extracted from point x in O .

3. An adapted reflection-based method

3.1. Formulation

The main advantage of rigid-body transformation-based approaches is to be easy to implement using the ICP algorithm and to benefit from the numerous theoretical and

practical works performed in the literature to assess its behaviour. However, this rigid-body transformation-based formulation of the symmetry plane problem is ill-posed: the resulting optimal rigid-body transformation does not allow direct, unambiguous estimation of the optimal plane, as $\tilde{T} \circ S_K$ is usually not a reflection. In fact, it can be shown that $\tilde{T} \circ S_K$ defines a proper reflection if and only if:

- \tilde{T} is a pure rotation whose axis is in K , or
- \tilde{T} is a pure translation perpendicular to K .

In practice, to our knowledge, these additional constraints on \tilde{T} have never been taken into account, probably because no closed-form solution is readily available in the literature for this problem.

We believe that a much more natural (but original) approach than searching for such a constrained rigid-body transformation is to look for the symmetry plane directly, *i.e.*, without relying on an intermediate rigid-body transformation. As a result, we can simply formulate the problem as:

$$\tilde{P} = \arg \min_{P, y_1, \dots, y_N} \sum_{x_i \in O} \|y_i - S_P(x_i)\|^2 = \arg \min_{P, y_1, \dots, y_N} \mathcal{E}_P,$$

where the points y_1, \dots, y_N belong to the cloud O . In a word, we propose to replace the two-step methods (computation of \tilde{T} and then \tilde{P}) of Section 2 by the one-step, purely *reflection-based*, minimisation of \mathcal{E}_P .

3.2. Algorithm

To our knowledge, there is no closed-form solution to minimise \mathcal{E}_P . Thus, we devise the following iterative scheme, inspired by the ICP algorithm, and that can be easily shown to converge to a (at least local) minimum of \mathcal{E}_P :

Step 0: Initialise \tilde{P}
Step 1: $\tilde{y}_1, \dots, \tilde{y}_N = \arg \min_{y_1, \dots, y_N} \sum_{x_i \in O} \|y_i - S_{\tilde{P}}(x_i)\|^2$
Step 2: $\tilde{P} = \arg \min_P \sum_{x_i \in O} \|\tilde{y}_i - S_P(x_i)\|^2$
Step 3: if \tilde{P} has changed go to Step 1 else finish

In Section 4, we call this optimal plane P_4 and compare it with the planes P_1, P_2 and P_3 obtained in Section 2.

3.3. Solving Step 1

One has to find N points \tilde{y}_i in O so as to minimise \mathcal{E}_P while P is kept fixed. The trivial solution to Step 1 is to simply choose \tilde{y}_i as the closest point of $S_P(x_i)$ in O .

3.4. Solving Step 2

To solve Step 2, we first characterise the plane P using a minimal parametrization: its unit normal n and its distance to the origin d . It can be easily shown that :

$$S_P(x) = (I_3 - 2nn^T)x + 2dn$$

where $x \in O$. We then reformulate the problem of Step 2 in a more general way and state the following theorem (proof in Appendix):

Theorem :

For a given square matrix A of size $N \times N$ (independent of P), the plane $P = (d, n)$ that minimises

$$\sum_{x_i \in O} \sum_{x_j \in O} A_{i,j} \|x_j - S_P(x_i)\|^2$$

is characterised by:

- n colinear with the eigenvector corresponding to the smallest eigenvalue of the 3×3 matrix B , where

$$B = \sum_{(x_i, x_j) \in O^2} A_{i,j} [(x_i - g_1 + x_j - g_2)(x_i - g_1 + x_j - g_2)^T - (x_i - x_j)(x_i - x_j)^T]$$

$$\text{and } \begin{cases} g_1 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in O^2} A_{i,j} x_i \\ g_2 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in O^2} A_{i,j} x_j \end{cases}$$

- $d = \frac{1}{2}(g_1 + g_2)^T n$.

Notice that the proof we give and the resulting formulas are linked to that obtained by Horn [10] and Faugeras and Hebert [8] for rigid-body registration of point clouds. There is an analogy between our parameters (n, d) (unit normal vector, distance to the origin) and their parameters (q, t) (unit quaternion, translation component).

In practice, in the problem of Step 2, the matrix A is a binary matrix built such that $A_{i,j} = 1 \Leftrightarrow x_j$ is the closest point of $S_P(x_i)$ in O .

4. Comparison of the approaches

In this section, we show experimentally that the rigid-body transformation-based approach is outperformed by the reflection-based solution. We use a dataset of about 80,000 points representing a human face (description in Section 6.3) and generate a perfectly symmetrical image from this raw data. Then we create three images with different levels of occlusions by removing a given number (1%, 3%, 7%) of adjacent points (example in Fig 3).

We then run the ICP algorithm and the ICP-like algorithm we devised in Section 3 on these data to find the optimal (rigid-body transformation-based) planes P_1, P_2, P_3 and the optimal (reflection-based) plane P_4 . At the end of each iteration of the two algorithms, we record the values $\mathcal{E}_T(T)$ and $\mathcal{E}_P(P_4)$. For each current estimate of T , we compute the three planes P_i ($i = 1, 2, 3$), for which we calculate the criterion values $\mathcal{E}_P(P_i)$. We plot these values, together with $\mathcal{E}_T(T)$ and $\mathcal{E}_P(P_4)$, for the image with 7%

of occlusions, on Fig. 2. We also display the final values $\mathcal{E}_P(P_i)$ ($i = 1, 2, 3, 4$) for the three images in Tab 1.

For a fair comparison, the two algorithms are initialised in the same way:

- ▷ the initial plane P for the ICP-like algorithm of Section 3 and the initial plane K for the ICP algorithm are computed using the principal axes and the centre of mass of an uniformly resampled version of O [17].

- ▷ the initial transformation T for the ICP algorithm is the identity.

Note that the comparisons of the different algorithms on the perfectly symmetrical image is of little interest, as in this case this initialisation already gives the optimal symmetry plane. Our conclusions are threefold, and suggest the superiority of the reflection-based approach over the others:

- The final value of $\mathcal{E}_T(T)$ is lower than that of $\mathcal{E}_P(P_4)$. This observation was expected; considering that any reflection S_P can be rewritten as the composition of a rigid-body transformation and of a reflection with respect to a fixed plane K , we have necessarily $\min \mathcal{E}_T(T) \leq \min \mathcal{E}_P(P)$. However, as mentioned earlier, the optimal rigid-body transformation T , combined with S_K , does not always define a proper reflection.

- The convergence to T is two times slower than for P_4 in terms of computational speed and number of iterations, because six parameters have to be estimated instead of three.

- The three planes P_i , $i = 1, 2, 3$ perform worse than P_4 . For instance, for 1% of occlusions, the final \mathcal{E}_P values for P_1 , P_2 and P_3 are 2%, 7% and 6% larger than for P_4 . Interestingly, these percentages drastically increase with the level of occlusions, and become equal to respectively 40%, 70% and 66% when there are 7% of occlusions.




Figure 2. Values of the different criteria through the iterations for a perfectly symmetrical face with 7% of points removed.

	$\mathcal{E}_P(P4)$	$\mathcal{E}_P(P1)$	$\mathcal{E}_P(P2)$	$\mathcal{E}_P(P3)$
1%	4.67	4.80	5.04	4.98
3%	17.09	19.11	24.09	22.01
7%	52.62	74.33	93.33	87.02

Table 1. Final values of the criterion $\mathcal{E}_P(P_i)$, $i = 1, 2, 3, 4$, for the three images with different levels of occlusions (1%, 3%, 7%).

5. A fast, robust and accurate algorithm

In Section 3, we have proposed to define the symmetry plane as the minimiser of the criterion \mathcal{E}_P , and we have shown how to perform this minimisation using an ICP-like scheme with iterative, closed-form formulas. We have also shown this approach to be superior than two-step methods based on the intermediate minimisation of \mathcal{E}_T . However, as the classical ICP algorithm, which is essentially a least squares minimisation, our ICP-like algorithm is not robust to outliers. Classical techniques can be adapted to deal with this issue [4, 12] and we have recently proposed a solution based on M-estimators [5]. However, another critical problem is that, as the criterion \mathcal{E}_T in the classical ICP algorithm, \mathcal{E}_P (and our modified criterion based on M-estimators) is based on discrete, one-to-one matches between the two sides of the object under study, relying on the closest point principle. In practice, this causes the criterion \mathcal{E}_P to exhibit lots of local minima, which are prone to make the ICP-like algorithm (which has a monotonic, local convergence) fail. To deal with this limitation, in the context of rigid-body registration of point clouds, Rangarajan and colleagues proposed to define a new criterion using multiple, weighted matches [16]. They showed this criterion to be well-behaved and smoother, with a reduced number of local minima, than \mathcal{E}_T . A simple way to implement this idea, and to include such fuzzy matches between points, is to cast the problem in a probabilistic framework, and then to use classical techniques of statistical inference to estimate the unknown transformation.

Following Chui & Rangarajan [6] and Granger & Penneç [9], who introduced these techniques in the context of rigid-body registration of point clouds, we propose an adaptation of the criterion \mathcal{E}_P , based on mixture models, to compute the symmetry plane of the object.

5.1. Algorithm

The cloud $S_P(O)$ is now considered as a noised version of O . In practice, we define the probability function of the data points $y_i \in S_P(O)$ as a mixture density, using the points $x_j \in O$ as follows:

$$p(y_i|P) = \sum_{x_j \in O} p(y_i, x_j|P) = \sum_{x_j \in O} A_{i,j} p(y_i|x_j, P)$$

There are as many mixture components as there are

points in O . The unknown mixture component $A_{i,j}$ represents the unknown probability that $y_i \in S_P(O)$ is a measure of a given point $x_j \in O$, knowing nothing else. Intuitively, $A_{i,j}$ conveys the affinity between the points y_i and x_j . Note that $\sum_j A_{i,j} = 1$. If we consider all the data points y_i to be independent, then the likelihood of the cloud $S_P(O)$ can be written as (noting $x_i = S_P(y_i) \in O$):

$$L(P) = \prod_{x_i \in O} \sum_{x_j \in O} A_{i,j} p(S_P(x_i)|x_j, P)$$

We then define the optimal plane P using the maximum likelihood principle. Considering the noise as isotropic Gaussian with variance σ , the EM algorithm can be used for this purpose [7] and can be shown to yield the very simple following iterative scheme:

<p>Step 0: Initialise \tilde{P}</p> <p>E-step: $\tilde{A}_{i,j} = \frac{\exp(-\ x_j - S_{\tilde{P}}(x_i)\ ^2 / (2\sigma^2))}{\sum_{x_k \in O} \exp(-\ x_k - S_{\tilde{P}}(x_i)\ ^2 / (2\sigma^2))}$</p> <p>M-step: $\tilde{P} = \arg \min_P \sum_{(x_i, x_j) \in O^2} \tilde{A}_{i,j} \ x_j - S_P(x_i)\ ^2$</p> <p>Step 3: if \tilde{P} has changed go to E-step else finish</p>
--

The M-step can be solved efficiently using the general theorem we stated in Section 3.4. The ICP-like algorithm we presented in Section 3 can actually be seen as a simplified version of this EM algorithm. The point x_j with the highest value $A_{i,j}$ is the closest point of $S_P(x_i)$ in O , so the E-step is very similar to Step 1. If the matrix A is made binary, by keeping only this closest point into account, then the M-step is identical to Step 2.

5.2. Multiscale scheme

The convergence of this new algorithm is very dependent on the σ parameter. A small σ allows to be very selective on the pairs of points to be significantly taken into account for the estimation of P . Thus, if the algorithm is given a good initial plane, a small σ allows good accuracy. If this is not the case, there is little chance that two mismatched points will ever be matched properly during the iterations. On the contrary, a large value of σ allows these symmetrical, but initially mismatched points, to have growing influence during the iterations. Based on this, to allow both accuracy and robustness, we follow a multiscale approach, by running successive EM algorithms with decreasing σ values. In this scheme, σ acts as a scale parameters, and allows to progressively refine the plane estimation. In practice, the initial scale factor σ_0 is successively divided by a constant value $f > 1$ until it is equal to, or lower than a predefined final value σ_f . Note that this approach is slightly different from that of Granger & Pennec, who update the scale factor at each iteration of the EM algorithm. Experimentally, we observed that letting the EM algorithms converge and choosing a higher f value led to better results in our case,

in terms of both robustness and accuracy, without significantly increasing the computational time. This multiscale scheme can be seen as a deterministic annealing procedure where σ is analogous to the temperature parameter.

5.3. Coarse-to-fine approach

At the beginning of this multiscale scheme, large values of σ only lead to a gross estimation of the unknown plane. Consequently, it is useless to take the entire point set O into account at these stages. As a result, we propose a coarse-to-fine approach, where O is decimated at large σ values, and then refined progressively when σ decreases. We use the same technique as Granger & Pennec to compute the successive decimated clouds, named O' , in which are chosen the points x_i . At a given σ value, we iteratively merge each point of O included in spheres of radius $k\sigma$ (where k is a constant to define). Each of these spheres finally contains N_m points (with $\sum_m N_m = N = \text{card}(O)$), which are then replaced by their centroids. The set of centroids, each one being given a weight N_m (later used in the estimation of P), constitutes the decimated cloud used at scale σ . Note that the original set O is entirely taken into account when $k\sigma$ is lower than or equal to the minimal resolution of O .

5.4. Rejection of outliers

The E-step involves the computation of the distance between each point $S_{\tilde{P}}(x_i), x_i \in O'$ and all the points x_j of O . Actually, when the point $S_{\tilde{P}}(x_i)$ is far away from O (which can happen for instance when x_i has no satisfying counterpart on the other side of the object), the values $A_{i,j}$ can be not negligible because of the normalisation term (see E-step) and thus influence negatively the criterion to maximise. A simple way to tackle this problem is to reject the points x_j in O that are farther from $S_{\tilde{P}}(x_i)$ than a predefined threshold, which amounts to give them a null weight $\tilde{A}_{i,j}$. This last process can be seen as a threshold on a Mahalanobis distance and one can easily show that the algorithm converges to a local maximum of a modified likelihood function. In practice, the threshold must depend on σ : there is a high uncertainty on the matches when σ is large, while there is increased confidence on these matches when σ decreases. Thus we choose to set this threshold at the value $l\sigma$. The E-step then becomes:

<p>i) Initialise \tilde{A} as the null matrix</p> <p>ii) $\forall x_i \in O'$</p> <p>$L_i = \{x_j \in O \text{ such that } \ x_j - S_{\tilde{P}}(x_i)\ < l\sigma\}$</p> <p>$\forall x_j \in L_i, \tilde{A}_{i,j} = \exp(-\ x_j - S_{\tilde{P}}(x_i)\ ^2 / 2\sigma^2)$</p> <p>Normalise the i^{th} line of \tilde{A}</p>
--

The set L_i is built by a search in a kd -tree space-partitioning built at the beginning of the algorithm.

Note that Granger & Pennec [9] have experimentally

shown (in case of rigid-body registration) that the transformation yielding the global maximum likelihood is different from that globally minimising the ICP criterion. Tsin & Kanade [20] have actually shown that in case of two identical, aligned point clouds, the identity transformation is usually not a stationary point of the likelihood function. The same can be said about our problem, where we look for a reflection best superposing the cloud with itself. However, the bias between the correct solution and the ML estimate is virtually null when σ is very small, as at the end of the multiscale scheme.

5.5. Algorithm and implementation details

Finally, we have the algorithm (called **mEM**, as *modified EM*):

Initialise \tilde{P} and σ
Repeat
Decimation O' = decimation of O with a radius $k\sigma$
Repeat
E-step: see Section 5.4
M-step: $\tilde{P} = \arg \min_P \sum_{x_i \in O', x_j \in O} N_j \tilde{A}_{i,j} \|x_j - S_P(x_i)\|^2$
Until \tilde{P} does not change
Decrease scale factor $\sigma = \max(\sigma/f, \sigma_f)$
Until \tilde{P} does not change and $\sigma = \sigma_f$

The initialisation of P

Each iteration of this new algorithm is about three times longer than that of the ICP-like algorithm we devised in Section 3. As a consequence, we propose to initialise it with the plane given by the ICP-like algorithm rather than the plane deduced from the analysis of the inertia matrix, to allow faster convergence. In practice, we propose a slight modification of the ICP-like algorithm to make it more robust to outliers and asymmetries, based on a least trimmed squares technique [4]. The rejection rate of this LTS is arbitrarily set to 40% and our implementation includes also an acceleration step, analogous to that proposed by Besl & McKay [3], within a multiresolution process. In the following this algorithm will be called **tICP**. This tICP algorithm is preferred to the modification based on M-estimators we proposed before [5] (termed **mICP** here) because of its lower run time.

The stopping criterion

We choose as *ad hoc* stopping criterion $\|P^t - P^{t-1}\| \leq \epsilon$ where ϵ is a parameter to define and $P^t = (n, d)^t$ is the parameter of the searched plane P at the end of iteration t .

Values of the parameters

In the following, we will take as default parameters:

σ_0	σ_f	f	ϵ	l	k
5	0.5	1.5	0.01	3	1

6. Evaluation and results

6.1. Evaluation on symmetrical data

In this section, we investigate the accuracy and capture range of the mEM algorithm on perfectly symmetrical data. For this, we work on face data (description in Section 6.3). We apply angular offsets between 0 and 40 degrees and linear offsets between 0 and 60 mm to the ground truth symmetry plane and use it to initialise our algorithm. After convergence, we compute the angular and linear errors (called respectively θ and τ) of the estimated plane compared to the ground truth solution. With a step of 1 mm and 1 degree, this yields over 3200 experiments. For large linear offsets (below 60 mm) and large angular offsets (below 31 degrees), the algorithm always converges to a plane for which θ and τ are less than 10^{-15} . The tICP and mICP algorithms have a similar capture range expect that they always provide angular and linear errors greater than 0.1 degree and 0.05 mm.

6.2. Evaluation on asymmetrical data

In this section, we evaluate the robustness and accuracy properties on mEM on asymmetrical data. For this purpose, we add artefacts to the perfectly symmetrical face point cloud of the previous section whose symmetry plane is still considered as the ground truth. We generate artefacts as follows:

- ▷ Noise is modeled by an isotropic Gaussian white noise of variance δ^2 and is added to each point.
- ▷ Occlusions are generated by removing a given quantity of adjacent points. In the following, we term *outliers* the points with no symmetrical counterpart resulting from this removal.
- ▷ Asymmetries are generated by choosing a point D close to the surface and deforming each point P of the cloud to a new position P' according to:

$$\triangleright P' = P + K \times G_v(P - D) \frac{\overrightarrow{PD}}{\|\overrightarrow{PD}\|}$$

where G_v is a 3D non-normalised Gaussian function of variance v^2 and K is the deformation strength.

By randomly combining these artefacts, we generate a set of 150 images with varying levels of artefacts. The parameters are chosen such that: $(K, v^2) \in [0, 20] \times [0, 25]$ (one deformation on the right cheek and another on the right forehead), 0 to 20% of outliers and a noise of variance $\delta^2 = 0.3$. Example of data built this way are shown on Fig 3 (middle) and the resulting estimated plane on Fig 3 (right).

Tab. 2 and 3 show statistics over the 150 experiments for mICP, tICP and mEM. We observe that over the 150 experiments, mEM has a very low mean and maximum error for both θ and τ , with very low variance. These values




Figure 3. **Estimation of the symmetry plane on ground truth, asymmetrical data.** Left: Original point cloud, acquired with a laser scanner. Middle: The original point cloud is made symmetrical, and then noise, occlusions and deformations (color-coded) are added. Right: Estimation of the plane by mEM.

	mICP	tICP	mEM
max τ	0.61	5.14	0.12
max θ	1.2	2.82	0.15
mean (θ, τ)	(0.12,0.65)	(0.76,1.14)	(0.06,0.04)
var (θ, τ)	(0.02,0.01)	(0.75,0.16)	(10^{-3} ,0.01)

Table 2. Statistics on (θ, τ) for mICP [5], tICP and mEM initialised with tICP.

	tICP	mEM init by PA	mEM init by tICP
mean t	30s	600s	160s
max t	80s	1360s	380s

Table 3. Statistics on run time t for tICP, mEM initialised with PA and mEM initialised with tICP performed on a standard PC with an Intel Core Duo T7700 at 2.4GHz with 2GB Ram.

are higher for tICP and mICP, but we also demonstrate the relevance of using tICP rather than PA (principal axes) to initialise mEM in terms of computational time.

6.3. Results on real data

6.3.1 Results on difficult data

A population of 131 healthy subjects, 49 males and 82 females has been face scanned with a portable handheld laser scanner with resolution and accuracy below 1 mm (Polhemus FastScan, Polhemus Inc, VT, USA, <http://www.fastscan3d.com>). Each of these point clouds contains about 80,000 points. In Fig. 4, we display the symmetry plane obtained with mEM for one of these point clouds, representing a woman’s face with a large area of missing points on the left cheek. We also display a point cloud of about 26,000 points representing a chair with a missing leg, acquired with the same device, and the Stanford bunny, represented by a cloud of about 36,000 points.

6.3.2 Asymmetry mapping

Once mEM has converged and a symmetry plane P has been estimated, we quantify the local asymmetry at $x_i \in O$ by computing the distance between $S_P(x_i)$ and O .

The symmetry plane can also be conveniently aligned with the centre of the image coordinate system, allowing




Figure 4. **Estimation of the symmetry plane using mEM on different objects.**

improved display of the point cloud or easier further post-processing. If we name C the plane at the centre of the coordinate system, it can be shown that rigid-body transformation $(S_C \circ S_P)^{-1/2}$ aligns P with C [18] (where S_C and S_P are respectively the reflection symmetry with respect to planes C and P). Once P has been aligned with C , there are three remaining degrees of freedom (two translations and one rotation), and additional constraints can be specified to completely define the position of the object on the grid. We show mappings of face asymmetry for four subjects with this improved display on Fig. 5.




Figure 5. **Asymmetry maps.** Four healthy males. Blue (resp. red) corresponds to symmetrical (resp. asymmetrical) areas.

6.3.3 Profile extraction

Once the symmetry plane is estimated on human face data, one can easily extract the vertical profile. This can be highly useful for face recognition and identification [22]. To compute this profile, we project each point of the cloud on the plane, provided its distance to the plane is lower than a predefined threshold (typically 1mm for the laser data). Two such profiles are shown on Figure 6.




Figure 6. **Two vertical profiles and corresponding faces.** Left: a male. Right: a female.

7. Conclusion

In this paper, we have presented a new general formulation for the problem of estimating the symmetry plane of

bilateral objects in point clouds. We have shown this formulation to outperform the classical one. Then we have encompassed the problem in a probabilistic framework, and devised a multiscale, multiresolution algorithm shown to be fast, robust and accurate. A possible improvement could be to automate the choice of σ through the iterations of the multiscale scheme. Alternative approaches, based on kernel correlation for instance [20], could also be implemented and compared to our ML-based method. We expect our algorithm to have numerous applications in many domains where point clouds of symmetrical objects are obtained.

Appendix: Proof of theorem 2

The Lagrangian $\mathcal{L}(d, n, \lambda)$ associated to \mathcal{E}_P is equal to:

$$\sum_{(x_i, x_j) \in O^2} A_{i,j} \left\| x_j - (I_3 - 2nn^T)x_i - 2dn \right\|^2 - \lambda (\|n\|^2 - 1)$$

\mathcal{L} must be stationary at the optimal values d and n . We must have $\frac{\partial \mathcal{L}}{\partial d} = 0$, which writes:

$$\sum_{(x_i, x_j) \in O^2} A_{i,j} [-2n^T (x_j - (I_3 - 2nn^T)x_i - 2dn)] = 0$$

$$\Leftrightarrow d = \frac{1}{2}(g_1 + g_2)^T n$$

$$\text{with } g_1 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in O^2} A_{i,j} x_i,$$

$$g_2 = \frac{1}{\sum_i \sum_j A_{i,j}} \sum_{(x_i, x_j) \in O^2} A_{i,j} x_j$$

Substituting this expression of d in \mathcal{L} , we obtain:

$$\mathcal{L}(n, \lambda) = \sum_{(x_i, x_j) \in O^2} A_{i,j} (x_i - x_j)^2 - n^T (B - \lambda I_3)n + \lambda$$

$$\text{with } B = \sum_{(x_i, x_j) \in O^2} A_{i,j} [(x_i - g_1 + x_j - g_2)(x_i - g_1 + x_j - g_2)^T - (x_i - x_j)(x_i - x_j)^T]$$

$$\frac{\partial \mathcal{L}}{\partial n} = 0 \Leftrightarrow (B - \lambda I_3)n = 0$$

$$\text{so } \mathcal{L}(\lambda) = \sum_{(x_i, x_j) \in O^2} A_{i,j} (x_i - x_j)^2 + \lambda \text{ at the optimum}$$

and \mathcal{L} is minimised by the (unit) eigenvector of B with the smallest eigenvalue λ .

References

- [1] K. S. Arun *et al.* Least-squares fitting of two 3-D point sets. *IEEE PAMI*, 9(5):698–700, 1987.
- [2] M. Benz *et al.* The symmetry of faces. *Vision Modeling and Visualization*, 43–50, 2002.
- [3] P. Besl and N. Mckay. A method for registration of 3-D shapes. *IEEE PAMI*, 14(2):239–256, 1992.
- [4] D. Chetverikov *et al.* Robust Euclidean alignment of 3D point sets: the trimmed iterative closest point algorithm. *Image Vision Computing*, 23(3):299–309, 2005.
- [5] B. Combès *et al.* An algorithm to map asymmetries of bilateral objects in point clouds. *IEEE ISBI*, 2008.
- [6] H. Chui and A. Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3):114–141, 2003.
- [7] A. Dempster *et al.* Maximum likelihood from incomplete data via the EM algorithm. *JRSS, Series B*, 39(1):1–38, 1977.
- [8] O. Faugeras and M. Hebert. The representation, recognition, and locating of 3-D objects. *International Journal of Robotics Research*, 5(3):27–52, 1986.
- [9] S. Granger and X. Pennec. Multi-scale EM-ICP: A fast and robust approach for surface registration. *ECCV*, 2353:418–432, 2002.
- [10] B. Horn. Closed-form solution of absolute orientation using unit quaternions. *J of the Optical Society of America*, 4(4):629–642, 1987.
- [11] B. Horn *et al.* Closed-form solution of absolute orientation using orthonormal matrices. *Journal of the Optical Society of America A*, 5(7):1127–1135, 1988.
- [12] S. Kaneko *et al.* Robust matching of 3D contours using iterative closest point algorithm improved by M-estimation. *Pattern Recognition*, 36(9):2041–2047, 2003.
- [13] Y. Liu *et al.* Robust midsagittal plane extraction from normal and pathological 3-D neuroradiology images. *IEEE TMI*, 20(3):175–192, 2001.
- [14] E. D. Momi *et al.* Automatic extraction of the mid-facial plane for crano-maxillofacial surgery planning. *Int J Oral Maxillofac Surg*, 35(7):636–642, 2006.
- [15] G. Pan *et al.* Finding symmetry plane of 3D face shape. *IEEE ICPR*, 1143–1146, 2006.
- [16] A. Rangarajan *et al.* A robust point matching algorithm for autoradiograph alignment. *Visualization in Biomedical Computing*, 277–286, 1996.
- [17] C. Sun and J. Sherrah. 3D symmetry detection using the extended gaussian image. *IEEE PAMI*, 19(2):164–168, 1997.
- [18] J.-P. Thirion *et al.* Statistical analysis of normal and abnormal dissymmetry in volumetric medical images. *MedIA*, 4(2):111–121, 2000.
- [19] S. Thrun and B. Wegbreit. Shape from symmetry. *IEEE ICCV*, 1824–1831, 2005.
- [20] Y. Tsin and T. Kanade. A correlation-based approach to robust point set registration. *IEEE ICCV*, 558–569, 2005.
- [21] S. Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE PAMI*, 13(4):376–380, 1991.
- [22] Y. Wu *et al.* Face authentication based on multiple profiles extracted from range data. *Audio- and Video-Based Biometric Person Authentication*, 2688:515–522, 2003.
- [23] H. Zabrodsky *et al.* Symmetry as a continuous feature. *IEEE PAMI*, 17(12):1154–1166, 1995.
- [24] L. Zhang *et al.* 3D face authentication and recognition based on bilateral symmetry analysis. *The Visual Computer*, 22(1):43–55, 2006.