

Uniform distribution, distance and expectation problems for geometric features processing

XAVIER PENNEC, NICHOLAS AYACHE

INRIA, B.P. 93, 2004 route des Lucioles, 06902 SOPHIA ANTIPOLIS Cedex, FRANCE

Xavier.Pennec@sophia.inria.fr

Abstract. Complex geometric features such as oriented points, lines or 3D frames are increasingly used in image processing and computer vision. However, processing these geometric features is far more difficult than processing points, and a number of paradoxes can arise. We establish in this article the basic mathematical framework required to avoid them and analyze more specifically three basic problems: (1) what is a random distribution of features, (2) how to define a distance between features, (3) and what is the “mean feature” of a number of feature measurements?

We insist on the importance of an invariance hypothesis for these definitions relative to a group of transformations that models the different possible data acquisitions. We develop general methods to solve these three problems and illustrate them with 3D frame features under rigid transformations.

The first problem has a direct application in the computation of the prior probability of a false match in classical model-based object recognition algorithms. We also present experimental results of the two other problems for the statistical analysis of anatomical features automatically extracted from 24 three dimensional images of a single patient’s head. These experiments successfully confirm the importance of the rigorous requirements presented in this article.

keywords: geometric features, transformation groups, uniform distribution, invariant measure, invariant distance, expected features, mean features.

1. Introduction

Many algorithms in computer vision and object recognition deal with simple geometric features like points, for example the Iterative Closest Point [4, 31], the geometric hashing [18, 30, 26], and the alignment algorithm [3, 13]. On the other hand, models of the real world often lead to the consideration of more complex features: lines [9], planes, oriented points [6], frames [21, 22], etc. The handling of these features raises some problems, the first one being their representation, and can lead to paradoxes such as Bertrand’s paradox concerning geometric probabilities. We have previously shown [23] that additive noise is not suited for describing the uncertainty of frames and should be replaced by a “compositive” model of noise. Several other examples are presented in this article and demonstrate the need for particular attention when dealing with geometric features.

We investigate in this article three basic problems that often arise when processing geometric features or in the statistical analysis of these algorithms. The first one is the quantification of the probability of oc-

currence of an event when some geometric features are randomly distributed. A direct application is the quantification of the false positives rate in matching algorithms. The second problem concerns the distance between features. This is one of the operations mostly used in image processing algorithms and a change in its definition often leads to a different result. Last but not least, we analyze the notion of a mean feature, which turns out to be a difficult problem. For instance, if we want to obtain the mean 3D rotation, we can compute either the mean rotation matrix $\underline{R} = \frac{1}{n} \sum_i R_i$, the mean quaternion $\underline{q} = \frac{1}{n} \sum_i q_i$ by using the unit quaternion representation, or the mean rotation vector $\underline{r} = \frac{1}{n} \sum_i r_i$.

These three methods give different and incorrect results: the two first are not even rotations and none of them is stable with respect to a reference frame shift.

In order to give a meaningful solution to each of these problems, we have to consider them in a geometric framework. Indeed, in order to compare geometric objects in different locations (for instance, extracted from images with different view-points), we implicitly consider a set a space transformations that allows

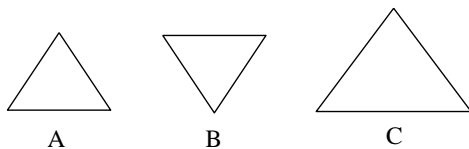


Fig. 1. Comparison of geometric objects: we can say that all three triangles are similar (relative to similarity transformations); or that only **A** and **B** are congruent (relative to rigid transformations); or that they are all different (relative to translations).

us to identify these objects. The choice of this set determines the properties of our objects, as in figure (1).

F. Klein formalized this idea in its “Erlangen program” [17]: let \mathcal{M} be a manifold of elements that we call points and \mathcal{G} a set of operations on these points forming a group. The geometry of \mathcal{M} for the group \mathcal{G} is the set of invariant properties of the manifold for the group action. A first application of this geometric framework is given with geometric probabilities (see [27] and section 3), where we need to define an invariant measure on random features (under the considered transformation group) in order to obtain a meaningful result. Similarly, for the distance and the mean feature, we have to design operators that are compatible with the action of a given transformation group.

The article is organized as follows. Section 2 focuses on the nature of geometric features, namely “points” on a manifold. We investigate transformation groups (rigid, affine...) that operate on this manifold. In Section 3, we investigate the standard geometric probabilities and, in particular, how to define an invariant measure on random features. This leads to the computation of the prior probability of a false match in recognition algorithms. Section 4 is devoted to invariant distances and section 5 to their use in the Fréchet expectation framework in order to provide a stable definition of the expected and average features. In the sixth and final section, we present an experimental application of the theory to the data fusion problem.

2. Sets of geometric features and sets of transforms

2.1. Geometric features: Manifolds and representations

Geometric features are generally defined as sets of points in the plane or 3D space, and the set of all geometric features of a given type can be described by a parameter p and a function $\beta(p, x)$ which associates the parameter p to the geometric feature (the set

$\{x \in \mathbb{R}^n / \beta(p, x) = 0\}$). The function β describes a particular type of geometric feature (lines, planes, curves, triangles...) with a specific *representation* p . For instance, 3D oriented planes can be represented by $p = (n, d)$ where n is a unit vector (the normal to the plane) and d the distance to the origin. The equation of “plane p ” is then:

$$\beta(p, x) = \langle n | x \rangle - d = 0$$

Usual sets of geometric features, such as lines, curves, surfaces... are regular and constitute differential manifolds. This means that the set is locally diffeomorphic to a vector space \mathbb{R}^m (i.e there exists, at each point of the manifold \mathcal{M} , a locally differentiable one-to-one mapping from \mathcal{M} to \mathbb{R}^m); m being the dimension of the manifold. In the above example, we can see that the parameter p is four dimensional with a quadratic constraint (which is differentiable), and planes are then a 3D-manifold equivalent to $S^2 \times \mathbb{R}^+$ (S^2 is the unit sphere in 3D). Despite the rather complex mathematical formulation, this simply means that manifolds are not traditional vector spaces, but that locally they may be treated as if they were. Spheres or smooth surfaces are such manifolds, as is the set of rotation matrices which is equivalent to \mathcal{P}^3 (the projective space of \mathbb{R}^4) by means of unit quaternions [22, 2]. Points trivially constitute a manifold since they already are a vector space. Another interesting type of features is oriented points, which are points associated with a vector. Such features can be extracted from a smooth surface, for instance, where the normal is attached at each point of the surface. A simple representation is given by $u = (x, n)$ where x is the position and n a unit vector. The manifold of oriented points is thus equivalent to $\mathbb{R}^3 \times S^2$ where S^2 is the unit sphere of the 3D space.

There are often numerous ways to represent a given manifold, with different properties. For instance, we can define a manifold as a subspace of \mathbb{R}^k with differentiable constraints and a one-to-one correspondence between features and parameters: this proves that the set of features is a differential manifold. For other purposes, in particular differentiation, it is necessary to have a minimal representation (where the dimension of the parameter is the dimension of the manifold), or more generally a set of charts forming an atlas of the manifold, exactly the same way we need several charts to represent the earth surface in a continuous way everywhere. Each chart is defined by a one-to-one dif-

ferentiable map $\varphi_i(p)$ from the representation into the manifold and an open definition domain \mathcal{D}_i . The set of charts must cover the manifold and must overlap each other so that it is possible to move from one chart to another. A study of different representations for 2D and 3D lines, planes and rotations is presented in [2]. We only assume for the moment that we have a one-to-one differentiable representation of the manifold, and we identify this representation with the manifold.

2.2. Transformations: Lie groups

There are a number of familiar transformations: translations, rotations, similarities, affine transformations. . . More generally, a transformation of a set X is a one-to-one map of X onto X . If g is a transformation, we denote by $g \star x = g(x)$ the action of the transformation on an element $x \in X$, and by $g^{(-1)}$ the inverse map. If g_1 and g_2 are two transformations, the map $g(x) = g_2(g_1(x))$ is also a transformation: the composition of g_1 and g_2 ($g = g_2 \circ g_1$). The set of all transformations with these operations is called the general transformation group of set X , and any subgroup \mathcal{G} is a transformation group of X .

The important class of Lie groups is obtained if \mathcal{G} has a separate topological structure (a Hausdorff space) and the composition and inversion maps are differentiable (\mathcal{G} is then a differentiable manifold). In fact, most usual transformation groups are Lie groups if they are continuous (in the non-discrete sense) and have reasonable operations. In this article, we use the 3D rigid motion group as an application example. An element of this group can be defined as the composition of a rotation with a translation. It can be represented by $f = (R, t)$, where the translation t belongs to \mathbb{R}^3 and R is a rotation matrix (a 3x3 matrix satisfying $R.R^\top = R^\top.R = Id$ and $\det(R) = +1$), and hence belongs to the special orthogonal group SO_3 . The inverse and compose maps are easily written (“” is the matrix multiplication):

$$f^{(-1)} = (R^\top, R^\top.t)$$

$$f = f_2 \circ f_1 = (R_2.R_1, R_2.t_1 + t_2)$$

2.3. From transformation of the Euclidean space to feature transformation

In the case of geometric objects, the transformation usually applies to the 2D plane or the 3D space (or more generally \mathbb{R}^n), but we wish to work directly on features and thus must take particular care that their nature is preserved during transformations. Consider, for instance, that two orthonormal axes are no longer orthonormal after a general affine transformation. The first constraint is then for the manifold \mathcal{M} to be globally invariant under the considered transformation group \mathcal{G} . We can then define the image of the feature p , satisfying $\beta(p, x) = 0$, by a transformation $g \in \mathcal{G}$ as being the feature $p' \in \mathcal{M}$ realizing $\beta(p', g \star x) = 0$. We write: $p' = g \star p$. With this definition, the group \mathcal{G} is also a transformation group of the manifold \mathcal{M} . It can be very tricky to make explicit the action on some geometric features with some representations and it can lead to highly non-linear transformations. However, usual cases are generally simple: in the case of oriented planes presented above under rigid motion, we have

$$\begin{aligned} \beta(p, f \star x) &= \langle n | R.x + t \rangle - d \\ &= \langle R^\top.n | x \rangle + \langle n | t \rangle - d \end{aligned}$$

which means that $f \star p = (R^\top n, d - \langle n | t \rangle)$. Similarly, the action of a rigid transform f on an oriented point $u = (x, n)$ is:

$$u' = f \star u = (R.x + t, R.n)$$

For the applications studied in this article, we are also interested in a third type of feature: frames. A frame is defined by a point with an orthonormal trihedron. We have already noted that we cannot use the affine group since orthonormal trihedra would not be conserved, but rigid motions are appropriate. A particularity of frames is that they are equivalent to rigid transformations. Indeed, any frame defines a basis for 3D space so that we can represent each frame by the rigid transformation which map the canonical basis to itself. It is easy then to verify that the composition and the action are equivalent.

2.4. Homogeneous features

A special kind of relation between the manifold and the group turns out to be very important: let $\mathcal{O} \in \mathcal{M}$

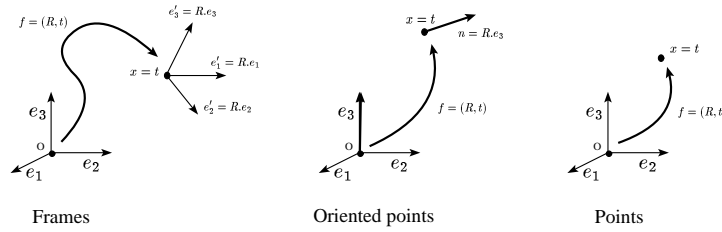


Fig. 2. Chosen origins for frames, oriented points and points, and their transformation by a rigid motion.

be an element called the origin. The manifold \mathcal{M} is *transitive* or *homogeneous* for the group \mathcal{G} if any other element of the manifold can be obtained by a transformation of \mathcal{G} , i.e. if $\mathcal{G} \star \mathcal{O} = \{g \star \mathcal{O} / g \in \mathcal{G}\} = \mathcal{M}$. This means that the features we consider have no invariants. In fact, we assume that we can split the features into an invariant part (which we do not consider here) and a variable part under the transformation group (it is not clear whether it is always possible but we did not have any problems with the features we studied).

In the case of homogeneous features, we identify the manifold with equivalence classes of group elements in the following way. Let \mathcal{H} be the subset of transforms that leave \mathcal{O} invariant:

$$\mathcal{H} = \{h \in \mathcal{G} / h \star \mathcal{O} = \mathcal{O}\} \quad (1)$$

\mathcal{H} is a group and is called the *isotropy* or *stability* group of \mathcal{G} at \mathcal{O} . The left cosets $g \circ \mathcal{H}$ can be identified with elements of \mathcal{M} . Indeed, if $g \star \mathcal{O} = x$, then $g \circ \mathcal{H}$ is the set of transformations which map \mathcal{O} to x . We write:

$$\mathcal{F}_x = \{g \in \mathcal{G} / g \star \mathcal{O} = x\} \quad (2)$$

For instance, if we consider point features with origin $\mathcal{O} = 0$ and rigid transformation, we have $\mathcal{H} = \{(R, 0) / R \in SO_3\}$ and $\mathcal{F}_x = \{(R, x) / R \in SO_3\}$ where R is any rotation. For frame, taking as the origin the canonical basis ($\mathcal{O} = (I_d, 0)$), then \mathcal{H} and all its cosets are reduced to a single point: $\mathcal{F}_f = \{f\}$. This special case where the manifold is equivalent to the group leads to important simplifications in the theory. As an intermediate example, we consider oriented points: we set the origin to $\mathcal{O} = (0, e_3)$. This origin is invariant by all rotations R_z around third axis e_3 . Thus the isotropy group is: $\mathcal{H} = \{R_z \in SO_3 / R_z \star e_3 = e_3\}$.

2.5. Back to the representation problem: 3D vectorial rotations, frames and points

It is well known that a 3D rotation matrix can be characterized by an angle θ around an axis n (unit vector), but since the coordinates of n are constrained, this couple is not minimal (the dimension of the representation is 4 instead of 3); moreover, the axis is not defined for the identity transformation. The rotation vector $r = \theta.n$ is always defined (in a multiple way since θ is modulo 2π) and differentiable (see [23] for the equations). In order to define an atlas of rotations, we need in fact four charts.

Chart 1: Non reflection rotations are represented by rotations vectors r from the open ball $\mathcal{B}^3(0, \pi)$.

Chart 2,3 and 4: Non identity rotations with axis not orthogonal to the x axis (respectively y, z) are represented by rotation vectors r from the open half ball $\mathcal{B}_{x^+}^3 = \{r \in \mathcal{B}^3(0, 2\pi) / r_x > 0\}$ (respectively $\mathcal{B}_{y^+}^3, \mathcal{B}_{z^+}^3$).

In theory we need to handle all four charts, but in practice only the principal chart (the first) is needed if we take care that, at the boundary of the domain, $r = \pi.n$ and $r' = -\pi.n$ are identical. Let $\mathcal{R}(r)$ and $\mathcal{r}(R)$ denote the mappings between rotation vectors and matrices, we can now write directly the composition and inversion laws on the representation:

$$r^{(-1)} = \mathcal{r}(\mathcal{R}(r)^\top)$$

$$r_2 \circ r_1 = \mathcal{r}(\mathcal{R}(r_2) \cdot \mathcal{R}(r_1))$$

Frames and motions are represented by a rotation vector and a translation: for convenient notation, we write $f = (r, t)$ and consider it as a column vector. In this framework, the representation of a point of the Euclidean space is denoted x (the standard coordinates). The canonical geometric operations can then be written:

- Composition: $f = f_2 \circ f_1 = (r_2 \circ r_1, r_2 \star t_1 + t_2)$

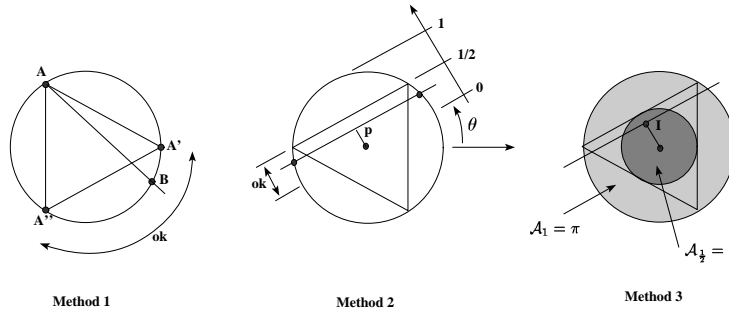


Fig. 3. Three methods for computing the probability that a “random” chord of a circle has a length greater than the side of an inscribed equilateral triangle. From left to right, the methods give a probability of $\frac{1}{3}$, $\frac{1}{2}$ and $\frac{1}{4}$.

- Inversion: $f^{(-1)} = (r^{(-1)}, r^{(-1)} \star (-t))$
- Action on a point: $x' = f \star x = r \star x + t$
- Cosets of points:
 $\mathcal{F}_x = \{f_x = (r, x) / \mathbb{R}(r) \in SO_3\}$
- Action on a frame: $f = f_2 \star f_1 = f_2 \circ f_1$
- Cosets of frames: $\mathcal{F}_f = \{f\}$.

3. Classical Geometric probabilities

The first class of problems in geometric probability is to quantify the probability of occurrence of an event when some geometric elements are randomly distributed. Bertrand’s paradox illustrates the need to consider invariance by a transformation group in order to obtain a single and well defined result. In fact, the problem lies in the notion of a *uniform* distribution (or measure). Some recent results in Lie group theory provide a mean for computing the left invariant measure on the group \mathcal{G} , which induces the invariant measure on homogeneous manifolds. An application is presented with the generalization of the false positives analysis.

3.1. Bertrand’s paradox

The problem raised by J. Bertrand in 1907 consists of finding the probability that a “random” chord of a circle has a length greater than the side of an inscribed equilateral triangle. Without loss of generality, we can fix the radius to 1 and the side length of the triangle is then $\sqrt{3}$. This problem can be tackled by at least three methods, which are illustrated in figure 3.

Method 1: By definition, a chord intersects the circle in two points, and we may assume that these two

points are equally and independently distributed on the circle. Assume that one of the points is A in figure 3. Then the second point has to lie between A' and A'' in the circle for the chord to be greater than the triangle side. This is just $\frac{1}{3}$ of the circumference and the searched probability is then $\frac{1}{3}$.

Method 2: A chord is characterized by its distance p to the center (between 0 and 1) and its orientation θ w.r.t. a fixed line (between 0 and 2π). If we draw the equilateral triangle with a side parallel to the chord, we can see that the distance d has to be less than $\frac{1}{2}$ in order to have a chord length greater than $\sqrt{3}$. By assuming a uniform orientation and distance to the origin, we find a normalized probability of $\frac{1}{2}$.

Method 3: A chord is uniquely defined by the orthogonal projection I of the circle center onto it. It has to lie inside the disc of radius $\frac{1}{2}$ in order to have a sufficient length. So, assuming I is uniformly distributed over the interior of the circle, the normalized probability is $\frac{1}{4}$.

The above three solutions are correct but they do not refer to the same notion of uniformity in the way we choose the chord. Using the (p, θ) representation (described in the second method), we can compute with [15] that the probability measures are respectively

$$d\sigma_1 = \frac{dp \cdot d\theta}{2\pi\sqrt{1-p^2}} \quad d\sigma_2 = \frac{dp \cdot d\theta}{2\pi} \quad d\sigma_3 = p \cdot \frac{dp \cdot d\theta}{\pi}$$

The solution to this problem is to impose an invariance constraint, or more precisely to *define* the notion of uniformity: for instance, uniform on \mathbb{R} means that the probability for a point to lie on an interval $]x; x+d[$ is the same for all x . This is basically an invariance by translation. In the same way, and since we can

only compare geometric objects with a transformation group, we define the uniform (or invariant) measure (the infinitesimal volume element) as the measure being invariant by the action of any fixed element f of the group. Let $d\mathcal{M}(x)$ be such a measure (\mathcal{M} stands for manifold); this means that $d\mathcal{M}(f \star x) = d\mathcal{M}(x)$ for any x . The invariant measure can be sought directly on a given representation [15], but a general formalism is developed in [27] to extract it from the left-invariant measure $d_L\mathcal{G}$ of the group \mathcal{G} . In the case of Bertrand's paradox, Poincaré showed in 1912 [25] that only the second measure is invariant under the action of rigid transformations (rotation and translation), although all three measures are invariant by rotations. Basically, this means that the computed probabilities for the first and third solutions would change if we translate the reference frame, whereas the second solution would stay 1/2, which is what was expected.

3.2. Left and right invariant measures on a Lie group (Haar measures)

We can require left invariance ($d_L\mathcal{G}(g \circ f) = d_L\mathcal{G}(f)$) for any fixed $g \in \mathcal{G}$, or right invariance ($d_R\mathcal{G}(f \circ g) = d_R\mathcal{G}(f)$). Since the group acts on the left (the action of the transformation f on the feature x is $f \star x$), we are mainly interested in left-invariance and, from now on, the left-invariant measure will be referred as the invariant measure.

To be mathematically correct, we require that for any continuous real function α on \mathcal{G} with compact support, we have:

$$\forall g \in \mathcal{G} \quad \int_{f \in \mathcal{G}} \alpha(g \circ f) \cdot d_L\mathcal{G}(g \circ f) = \int_{f \in \mathcal{G}} \alpha(f) \cdot d_L\mathcal{G}(f)$$

If the group is locally compact, then [12] proves that there exists only one left-invariant measure (up to a scale factor) that verifies the above properties. This measure is called the (left) Haar measure of the group. In a symmetric way, there is also a unique right Haar measure.

It is interesting to note that the left and right invariant measures are generally different; the group is called unimodular if they are equal. A compact group is always unimodular, but locally compact groups can have different left and right Haar measures. For instance, left and right Haar measures are identical on SO_3 , since the 3D rotation group is compact, and

the 3D rigid motion group is unimodular although the group is only locally compact due to the introduction of translations.

The left and right invariant measures can be generally computed from the Maurer-Cartan equations [27], but a very interesting theorem allows easy computation in the case of a *minimal* representation: assume that the definition domain of a chart almost covers the group (since we integrate function and not distributions, we can “forget” a subset of the group that has a null measure) and that the Jacobian of the left translation of the identity $J_L(f)$ exists and is continuous almost everywhere. Then the invariant measure can be written (see appendix A.1.1):

$$d_L\mathcal{G}(f) = \frac{df}{|J_L(f)|} \quad (3)$$

with

$$J_L(f) = \left. \frac{\partial(f \circ e)}{\partial e} \right|_{e=I_d} \quad \text{and} \quad |J| = |\det(J)|$$

The right-invariant measure can be derived in the same way using the Jacobian J_R of the right translation of the identity. Using this scheme, we can show [24] that the uniform measure for rotations using the rotation vector representation is

$$d_L\mathcal{G}(r) = d_R\mathcal{G}(r) = \frac{\sin^2(\theta/2)}{\theta^2} dr \quad (4)$$

where $\theta = \|r\|$. Thus, with our representation $f = (r, t)$, the invariant measure on rigid transformations is:

$$d\mathcal{G}(f) = d_R\mathcal{G}(f) = \frac{\sin^2(\theta/2)}{\theta^2} dr \cdot dt \quad (5)$$

3.3. Invariant measure on homogeneous manifolds

We saw in section 2.4 how to identify the homogeneous manifold \mathcal{M} with the quotient space \mathcal{G}/\mathcal{H} . We can find, thanks to the above section, the (left) invariant measures $d_L\mathcal{G}$ and $d_L\mathcal{H}$ on \mathcal{G} and \mathcal{H} , and write $d_L\mathcal{G} = d\mathcal{M} \cdot d_L\mathcal{H}$ where $d\mathcal{M}$ is a measure on the manifold \mathcal{M} (or the quotient space \mathcal{G}/\mathcal{H}). Santalo gives in [27] several forms of a necessary and sufficient condition for $d\mathcal{M}$ to be an invariant measure (i.e. $d\mathcal{M}(g \star x) = d\mathcal{M}(x)$). Let e be an element of the manifold \mathcal{M} . One of them can be stated as follows

(proof in appendix A.1.2):

$$\forall h \in \mathcal{H}, \quad \left| \frac{\partial(h \star e)}{\partial e} \Big|_{e=\mathcal{O}} \right| = 1 \quad (6)$$

This means that the measure of the infinitesimal volume element at the origin remains unchanged under any transformation that keeps the origin unchanged. If this condition is not satisfied, then there is no invariant measure on the manifold, otherwise we can compute it (with a minimal representation of the manifold) in a similar way to what we did for the group. Assume that the definition domain of a chart almost covers the manifold and that the Jacobian of the translation of the origin $J(f)$ exists and is continuous almost everywhere. Then the invariant measure is (the proof is obtained by replacing $(f \circ e)$ with $(f \star e)$ in appendix A.1.1):

$$d\mathcal{M}(x) = \frac{dx}{|J(f_x)|} \quad (7)$$

with

$$J(f_x) = \frac{\partial(f_x \star e)}{\partial e} \Big|_{e=\mathcal{O}} \quad \text{and} \quad f_x \in \mathcal{F}_x$$

3.4. Practical use: probability of a false match

Assume that x is a uniform random feature in the first image (characterized by a set of possible features I_1). What is the probability that this feature is accepted as a match with feature y in the second image, under a given global transformation f ?

If we characterize the possible matches for y by an “error volume” $Z(y)$ around y , we can write this probability as the conditional probability

$$P(f \star x \leftrightarrow y) = P((f \star x) \in Z(y) | x \in I_1)$$

$$\begin{aligned} P(f \star x \leftrightarrow y) &= \frac{\int_{(f \star I_1) \cap Z(y)} d\mathcal{M}(x)}{\int_{I_1} d\mathcal{M}(x)} \\ &= \frac{\mathcal{V}((f \star I_1) \cap Z(y))}{\mathcal{V}(I_1)} \end{aligned}$$

where $\mathcal{V}(X)$ is the “volume” of the set X . With the assumption that the volume $\mathcal{V}(Z(y))$ is sufficiently small with respect to the volume of the image, we can consider that the transformed image I_1 either contains the whole set $Z(y)$ or does not intersect it at all. This allows us to approximate the above probability by

$$P(f \star x \leftrightarrow y) = \varepsilon \frac{\mathcal{V}(Z(y))}{\mathcal{V}(I_1)}$$

with

$$\varepsilon = \begin{cases} 1 & \text{if } f^{(-1)} \star y \in I_1 \\ 0 & \text{otherwise} \end{cases}$$

A desirable property for our “error volume” $Z(y)$ is that it should be comparable at every point (as we usually fix the same bound for error on all the points): this means that, for any pair of points y and y' on the manifold, there exists a transformation f such that $f \star y = y'$ and $Z(y') = f \star Z(y)$ (the error volume is said to be homogeneous). A stronger hypothesis is that for every transformation f , the error volume on the transformed point is the transformation of the error volume: $Z(f \star y) = f \star Z(y)$. The volume is said to be isotropic in this case, and is completely determined by its shape around the origin (see [24] for an analysis of noise models). In both cases (homogeneity and isotropy), the volume of the error volume is invariant: $\mathcal{V}(Z(y)) = \mathcal{V}(Z(\mathcal{O})) = \mathcal{V}_0$. The basic probability of a false match $P(f \star x \leftrightarrow y) = \varepsilon \cdot \mathcal{V}_0 / \mathcal{V}(I_1)$ can now be applied as usual in an analysis of the frequency of false positives [10, 19, 11, 20].

As a practical example, we considered in [21] that two frames are matched if the distance between their origin is less than a threshold d_0 and if the rotation needed to adjust their trihedra has an angle less than a threshold θ_0 (this angle is $\theta = \|r_x^{(-1)} \circ r_y\|$), which, in fact, is a bound on an invariant distance (see section 4.6). Thus the volume is invariant and can be computed at the origin: a frame $f = (r, t)$ is in the error volume $Z(I_d)$ if $\theta = \|r\| < \theta_0$ and $\|t\| < d_0$. Using the invariant measure of equation (5), we can compute the volume of the error zone:

$$\begin{aligned} \mathcal{V}_0 &= \int_{\theta < \theta_0} \int_{\|t\| < d_0} d\mathcal{M}(r, t) \\ &= \left(\int_{\theta < \theta_0} \frac{\sin^2(\theta/2)}{\theta^2} dr \right) \cdot \left(\int_{\|t\| < d_0} dt \right) \\ \mathcal{V}_0 &= [2\pi(\theta_0 - \sin(\theta_0))] \cdot \left[\frac{4\pi}{3} d_0^3 \right] \end{aligned}$$

If we assume a cubic image of side l (256 for instance) this gives a Euclidean volume $V_I = l^3$ for points in

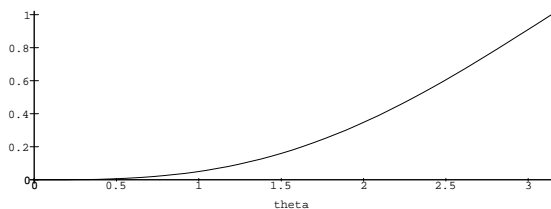


Fig. 4. Basic probability of a false match for trihedrons with a bound on the angle for the adjustment rotation of theta. Since the formula of the selectivity is multiplicative for frames, this curve is also the gain in selectivity when using frames instead of just points.

which trihedrons are not constrained: the rotation volume is $2\pi^2$. Finally, we obtain the basic probability of a false match:

$$P(f \star x \leftrightarrow y) = \left[\frac{\theta_0 - \sin(\theta_0)}{\pi} \right] \cdot \left[\frac{4\pi}{3} \left(\frac{d_0}{l} \right)^3 \right] \varepsilon$$

We have isolated in the first term the probability of a false match due to trihedra only, which reflects the gain in selectivity when using frames instead of points. This function is plotted in figure 4 and shows very interesting results: even for a bound of $\theta_0 = \pi/2 = 90$ deg, more than 80% of the random matches are rejected. For a bound of $\theta_0 = \pi/10 = 18$ deg, the probability of a false match drops to 0.0016: we would have to divide the bound on the position by 10 to obtain an equivalent selectivity using points only.

4. Invariant Distances

The distance between two points is often used in geometric algorithms: the Iterative Closest Point algorithm, developed in [4, 31] is the best example. Another classical example is the least-squares solution for registration between two sets of matched points (see section 4.2). All these algorithms can be extended to homogeneous features in a straightforward way using a distance between features. However, it is highly desirable that the results of these algorithms do not rely on the chosen representation nor on the reference frame of the physical space, as illustrated by the paradox of section (4.1).

Defining a distance directly on the manifold solves for the representation problem, but only the use of an invariant distance guarantees the stability of the results with respect to the action of the transformation group. We characterize in this section the properties of such invariant distances for the transformation group and the manifold. We give a general method to generate

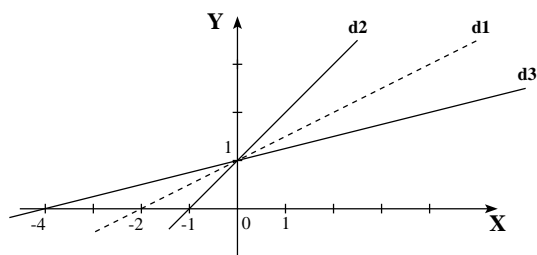


Fig. 5. Three lines in the plane.

an invariant distance on the manifold from a metric on the transformation space.

4.1. The paradox of the closest line

With the paradox of Bertrand, we saw different representations of 2D lines. In this section, we use another minimal representation based on the line equation [2]: the equation of a 2D line is $a.x + b.y + c = 0$. In order to obtain a minimal representation, we need to eliminate one parameter:

- **Chart 1:** lines that are not parallel to the X axis are represented by $d = (a, p) \in \mathbb{R}^2$ and have equation: $a.x + y + p = 0$.
- **Chart 2:** lines that are not parallel to the Y axis are represented by $d' = (a', p') \in \mathbb{R}^2$ and have equation: $x + a'.y + p' = 0$.

In the first chart, the line $d = (a, p)$ cuts the Y axes at the point $(0, -p)$ and has a director vector $(1, -a)$. This is symmetric in the second chart: the line $d' = (a', p')$ cuts the X axes at the point $(-p', 0)$ and has a director vector $(-a', 1)$.

We draw in figure 5 three lines. The problem is to choose which line d_2 or d_3 is the closest one to d_1 . A definition of the distance that seems to be reasonable is $\text{dist}(d_1, d_2) = \sqrt{(a_1 - a_2)^2 + (p_1 - p_2)^2}$. The coordinates of the three lines in the first chart are $d_1 = (-1/2; -1)$, $d_2 = (-1; -1)$, $d_3 = (-1/4; -1)$, and the distance between the lines d_1 and d_2 turns out to be greater than the distance between the lines d_1 and d_3 :

$$\text{dist}(d_1, d_2) = 1/2 \quad \text{and} \quad \text{dist}(d_1, d_3) = 1/4$$

Now if we consider the lines in the second chart, their coordinates are $d_1 = (-2; 2)$, $d_2 = (-1; 1)$, $d_3 = (-4; 4)$ and the distances are

$$\text{dist}(d_1, d_2) = \sqrt{2} \quad \text{and} \quad \text{dist}(d_1, d_3) = 2\sqrt{2}$$

We have then the following paradox: d_3 is the closest line to d_1 in one chart, and d_2 is the closest one in the second chart. Which chart gives the most reasonable result? In fact, visually, the line d_3 is the closest one: there is an angle of 12.5° between d_1 and d_3 versus an angle of 18.5° between d_1 and d_2 .

4.2. Using an invariant distance

Part of the above paradox can be raised by using a true distance on the feature manifold \mathcal{M} , and not a set of distances in different charts. This means that the distance verifies the following axioms:

- symmetry: $\forall(x, y) \in \mathcal{M}^2 \quad \text{dist}(x, y) = \text{dist}(y, x)$
- positivity: $\forall(x, y) \in \mathcal{M}^2 \quad \text{dist}(x, y) \geq 0$
- definite character: $\forall(x, y) \in \mathcal{M}^2 \quad (\text{dist}(x, y) = 0) \Leftrightarrow (x = y)$
- triangular inequality: $\forall(x, y, z) \in \mathcal{M}^3$
 $\text{dist}(x, y) + \text{dist}(y, z) \geq \text{dist}(x, z)$

However, if such a distance *on the manifold* is independent of the considered chart, this is not sufficient to ensure the stability of most results with respect to a change of the reference frame of the physical space (the action of a transformation on our features). The solution to this problem is to choose a distance which is invariant under the action of any transformation $g \in \mathcal{G}$: $\text{dist}(x, y) = \text{dist}(g \star x, g \star y)$.

Consider for instance the classical method of computing the transformation that maps a set of features x_i in one image to a set of features y_i in another image: this is the transformation that minimizes the sum of squared distances $C(f) = \sum_i \text{dist}^2(f \star x_i, y_i)$. Let F be the transformation minimizing this least square criterion (we can assume for simplicity that it is unique, but the same results hold for a set F of minima).

Assuming that the features x_i are transformed by a transformation g ($x'_i = g \star x_i$), the criterion becomes

$$\begin{aligned} C'(f) &= \sum_i \text{dist}^2(f \star x'_i, y_i) \\ &= \sum_i \text{dist}^2((f \circ g) \star x_i, y_i) \\ &= C(f \circ g) \end{aligned}$$

With or without the invariance constraint, the new result is $F' = F \circ g$. Now if we assume that the fea-

tures of the second image are globally transformed: $y'_i = g \star y_i$, then the criterion is

$$\begin{aligned} C''(f) &= \sum_i \text{dist}^2(f \star x_i, y'_i) \\ &= \sum_i \text{dist}^2(f \star x_i, g \star y_i) \end{aligned}$$

Here, we need the invariance property of the distance to conclude

$$\begin{aligned} C''(f) &= \sum_i \text{dist}^2((g^{(-1)} \circ f) \star x_i, y_i) \\ &= C(g^{(-1)} \circ f) \end{aligned}$$

This means that the new minimum is $F'' = g^{(-1)} \circ F$, which gives the expected result $F = g \circ F''$. The same experiment can be done if both images are transformed by the *same* transformation g (which means a global change of the reference frame), and the invariance of the distance is required to prove that the transformation found is $F''' = g^{(-1)} \circ F \circ g$, i.e. only the change of the reference frame.

4.3. Invariant distance on a manifold

Let $x, y \in \mathcal{M}$ and $g \in \mathcal{G}$. The distance is invariant if $\text{dist}(x, y) = \text{dist}(g \star x, g \star y)$. This means in particular that this distance is completely defined by the distance $N(x)$ of a feature x with the origin: if we use transformation $f_y^{(-1)}$ or $f_x^{(-1)}$, we have

$$\begin{aligned} \text{dist}(x, y) &= \text{dist}(f_y^{(-1)} \star x, \mathcal{O}) \\ &= N(f_y^{(-1)} \star x) = N(f_x^{(-1)} \star y) \end{aligned} \quad (8)$$

The axioms of the distance are translated, under the invariance assumption, into the three following properties:

- $\forall f_x \in \mathcal{F}_x \quad : \quad N(f_x^{(-1)} \star \mathcal{O}) = N(x)$
and thus $\forall h \in \mathcal{H} \quad : \quad N(h \star x) = N(x)$
- $N(x) \geq 0 \quad \text{and} \quad (N(x) = 0) \Leftrightarrow (x = \mathcal{O})$.
- $\forall f_x \in \mathcal{F}_x, f_y \in \mathcal{F}_y \quad :$
 $N(x) + N(y) \geq N(f_y^{(-1)} \star x) = N(f_x^{(-1)} \star y)$

These properties are very close to those required in order to define a norm on a vector space (without the positive linearity). Thus we call N the “norm” of the manifold. Note that we have so far defined the “norm” on the manifold and not in a particular chart. In practice, we use a “principal chart”, centered around the origin and covering almost the manifold. The “norm” N is defined in this chart, and when we have to use the distance $\text{dist}(x, y)$, we compute $N(f_y^{(-1)} \star x)$.

4.4. Invariant distances on a Lie group

Assume that we work now on the transformation group \mathcal{G} . We can require the distance to be either left or right invariant. As above, an invariant distance is determined by a “norm” N_g on the group, satisfying the following properties:

- $N_g(f^{(c)}) = N_g(f)$.
- $N_g(f) \geq 0$ and $(N_g(f) = 0) \Leftrightarrow (f = I_d)$.
- $N_g(f) + N_g(g) \geq N_g(g^{(c)} \circ f) = N_g(f^{(c)} \circ g)$.

The triangular inequality becomes

$$N_g(f) + N_g(g) \geq N_g(g \circ f^{(c)}) = N_g(f \circ g^{(c)})$$

for a right invariant distance. The corresponding left and right invariant distances are

$$\text{dist}_L(f, g) = N_g(g^{(c)} \circ f) = N_g(f^{(c)} \circ g)$$

$$\text{dist}_R(f, g) = N_g(g \circ f^{(c)}) = N_g(f \circ g^{(c)})$$

We are interested here only in the left-invariant distance since it induces an invariant distance on an homogeneous manifold.

4.5. Distance induced on the manifold by the group distance

Let N_g be a norm on the group \mathcal{G} . We define the induced “semi-norm” on the homogeneous manifold \mathcal{M} as

$$\begin{aligned} N(x) &= \inf_{(h \in \mathcal{H}, f_x \in \mathcal{F}_x)} (N_g(h \circ f_x)) \\ &= \inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1 \circ f_x \circ h_2)) \end{aligned} \quad (9)$$

If the infimum of $N_g(h_1 \circ f \circ h_2)$ is reached for every transformation f by $(h_1, h_2) \in \mathcal{H}^2$, then the semi-norm is separable and is thus a norm (see proofs in appendix A.2). This property is always true if the isotropy group \mathcal{H} is compact but is not automatically verified otherwise (for instance, there is no norm induced on points by the similarities or affine transformations).

Assuming that we have a norm, the distance associated with this norm is automatically invariant and satisfies

$$\begin{aligned} d(x, y) &= \inf_{\{f_x \in \mathcal{F}_x; f_y \in \mathcal{F}_y\}} (\text{dist}_L(f_x, f_y)) \\ &= \inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1 \circ f_x^{(c)} \circ f_y \circ h_2)) \end{aligned}$$

When a norm is chosen on the transformation space, we automatically have an invariant distance on the manifold. Since features are usually objects abstracted from an Euclidean space, a reasonable requirement to make is that the distance induced on points of the original space is the canonical distance of the space (possibly up to a scale factor). We have then the guarantee of reasonable invariant distances on all the features we consider.

4.6. Practical use on rigid transformations

The Euclidean distance on \mathbb{R}^3 is induced by the L_2 norm: $d_t(x, y) = \|x - y\|$. On the other hand, it can be shown [1] that the angle θ of a rotation is a “norm” that induces a left and right invariant distance on SO_3 , the rotation group. With the rotation vector representation, we have then $d_\theta(I_d, r) = \|r\| = \theta$ and thus $d_\theta(r_1, r_2) = \|r_2^{(c)} \circ r_1\| = \|r_1 \circ r_2^{(c)}\|$ (the last term of the equality comes from the right invariance).

We define the “norm” on the rigid motion group as (see appendix A.3):

$$N_\lambda(f) = N_\lambda((r, t)) = \|f\| = \sqrt{\lambda^2 \|r\|^2 + \|t\|^2}$$

where λ is a fixed parameter that allows to tune the importance of the trihedron (rotation part) with respect to the position (or translation part). Indeed, the angle of rotation θ is in radian (or degrees, ...) and the translation in millimeters, kilometers or inches... We usually scale each of the two terms by the inverse of their variation domain (π for θ and the diameter l_0 of the image or the interest object for the translation: $\lambda = l_0/\pi$). When we have information about the noise level (i.e. standard deviations σ_θ and σ_t), we can also use $\lambda = \sigma_t/\sigma_\theta$.

We can check that the distance induced by this norm on the original space is the Euclidean distance (see appendix A.3.1). Thus the left-invariant distance is

$$\begin{aligned} \text{dist}_L(f_1, f_2)^2 &= \|f_2^{(c)} \circ f_1\|^2 \\ &= \lambda^2 \|r_2^{(c)} \circ r_1\|^2 + \|t_1 - t_2\|^2 \end{aligned}$$

whereas the right invariant distance is

$$\begin{aligned} \text{dist}_R(f_1, f_2)^2 &= \|f_1 \circ f_2^{(c)}\|^2 \\ &= \lambda^2 \|r_1 \circ r_2^{(c)}\|^2 + \|t_1 - (r_1 \circ r_2^{(c)}) \star t_2\|^2 \end{aligned}$$

Although the rigid motion group is unimodular (and thus left and right invariant measures are identical), the

left and right invariant distances are obviously different.

4.7. Discussion

From a (left) invariant distance on the transformation group, we can determine an induced invariant distance on the feature manifold. However, if we have a sufficient condition for its existence, this condition does not seem to be necessary and is moreover difficult to handle. The second problem in this approach is that we need to know a left invariant distance on the transformation group, and nothing helps us to construct it. We reserve a forthcoming article to a Riemannian approach of this problem that gives necessary and sufficient conditions for the existence of such distances and a way to construct them via geodesics (see however [24]).

5. Expectation and mean of random features

Uncertainty on geometric features (and more generally on measurements) is usually characterized by a probability density function for which the expected value corresponds to the exact value. From a computational point of view, however, we need to keep only a few number of parameters characterizing this pdf. The usual way ([5, 2, 32]) is to consider the representation of the random feature as a random vector and, assuming that the pdf is quasi-Gaussian, approximate it up to the second order by its expectation and covariance matrix. We focus in the sequel on the expectation \bar{x} and its statistical measurement: the empirical mean (in the following, the term expectation refers to the expectation of a random feature of pdf p_x , whereas a mean feature denotes the empirical mean of a set of measured features). The classical definition is, for a pdf p_x (in the parameter space) and a set of measured features $\{x_i\}$:

$$\bar{x} = E(x) = \int_{\mathcal{D}} y \cdot p_x(y) \cdot dy$$

$$M(\{x_i\}) = \frac{1}{n} \sum_i x_i$$

We claim that these operators are not properly defined. In particular, the result of the integral or the sum is *not* ensured to be in the definition domain and defines not necessarily a feature: for instance, the arith-

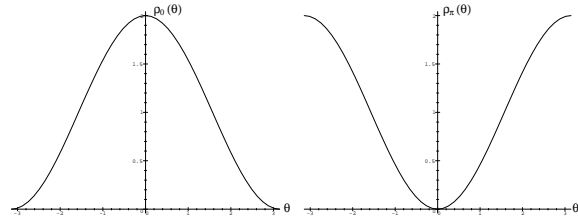


Fig. 6. Left: Original pdf $\rho_0(\theta) = 2 \cos(\theta/2)^2$. Expectation is $\bar{\theta}(0) = 0$. Middle: pdf after a rotation of angle $\lambda = \pi$. The expectation $\bar{\theta}(\pi)$ is also 0, whereas the rotated expectation is $\underline{\theta}_\pi = \pi$.

metic mean of several rotation matrices is generally not an orthogonal matrix and is therefore not a rotation itself (particularly for large deviations). The second reason, is that the expectation does not commute in general with the action of a fixed transformation (see example below). This means that the mean value of a pdf depends on the chosen reference frame, which is unacceptable.

5.1. Standard expectation of a 2D random line

We consider for this example 2D oriented vector lines, which can be represented by a point on the unit circle, and therefore an angle θ with a given axis. We fix the domain of θ to be $\mathcal{D} =]-\pi, \pi]$. The action of a rotation of angle λ is simply the addition (modulo 2π). We can define an uncertain line by its probability density function $\rho(\theta)$, and the classical way to obtain the expected value is to integrate in the parameter space:

$$\bar{\theta} = E(\theta) = \int_{\mathcal{D}} \theta \cdot \rho(\theta) \frac{d\theta}{2\pi}$$

where the term 2π is a normalization factor. We note that $d\theta$ is the uniform measure for lines under rota-

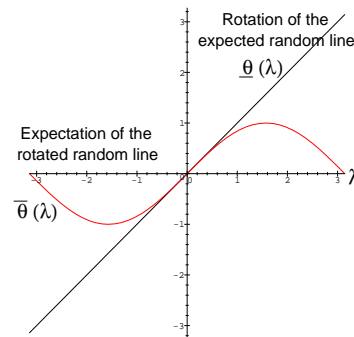


Fig. 7. Comparison of the expectation of the rotated line and the rotation of the expected line.

tion. Let $\rho_0(\theta) = 2 \cos(\theta/2)^2$ be such a pdf (drawn in figure 6). It is normalized and its expectation is $\bar{\theta}_0 = 0$. If we now change the reference frame, i.e. apply a rotation of angle λ , we obtain the new pdf $\rho_\lambda(\theta) = \rho_0(\theta - \lambda)$, and the expected value is now

$$\bar{\theta}(\lambda) = \int_{\mathcal{D}} \theta \cdot \rho_0(\theta - \lambda) \frac{d\theta}{2\pi} = 2 \sin(\lambda/2) \cos(\lambda/2)$$

which is different from the rotated expectation $\bar{\theta}(\lambda) = \lambda$. In particular, for a rotation of $\lambda = \pi$, we find that $\bar{\theta}_\pi = \pi$ and $\bar{\theta}_\pi = 0$.

To avoid this kind of problem, the first idea is to “center” the definition domain of the chart around the expected feature. If this can be easy in the case of a circle, as above, it may be more problematical for some other features such as 3D lines or frames, where the manifold is more complex. The question is similar for the mean value, especially with scattered measurements. Last, but not least, if the chart is centered around the expected (or the mean) feature, then the problem is already solved. The Fréchet expectation is a well-posed formalism to implement this idea: the “centrality” of a feature is based on its distance with other measurements, and the mean or expected values are the features that optimize the “centrality”.

5.2. Fréchet expectation of a random feature

Let v be a random vector in \mathbb{R}^n . Fréchet [8] observed that the variance $\sigma_v^2(x) = E(\text{dist}(v, x)^2)$ is minimized at the expected value \bar{v} . The second point is that if the expectation of the representation of a feature (a vectorial integral) is not well defined (because features are not vectors), the expectation of a real mapping (a positive function of the random feature, for instance) is properly defined. In particular, the expectation of the squared distance between features is properly defined.

Let dist be an invariant distance on the manifold \mathcal{M} under the group \mathcal{G} , and x a random feature of pdf p_x (in the parameter space). The expected square distance of a deterministic feature y with the random feature x is defined by

$$\begin{aligned} \sigma_x^2(y) &= E(\text{dist}(y, x)^2) \\ &= \int_{\mathcal{D}} \text{dist}(y, x')^2 \cdot p_x(x') \cdot dx \end{aligned} \quad (10)$$

If $\sigma_x^2(y)$ is finite for all y , we call every feature \bar{x} minimizing σ_x^2 an expected feature, and we denote by $\mathbb{E}(x)$ the set of all expected features of the random

feature x . Thus we have

$$\mathbb{E}(x) = \arg \min_{y \in \mathcal{M}} (E(\text{dist}(y, x)^2)) \quad (11)$$

If there is at least one expected feature \bar{x} , then $\sigma_x = \sigma_x(\bar{x})$ is called the standard deviation of x and σ_x^2 the variance. In general, there can be several expected values. However, Karcher [14] and Kendall [16] show, under some strong conditions, the existence and uniqueness of the expected value.

In a very similar way, we can define the set of empirical means of features $\{x_i\}$ by

$$M(\{x_i\}) = \arg \min_{y \in \mathcal{M}} \left(\sum_i \text{dist}(y, x_i)^2 \right) \quad (12)$$

and if there is at least one mean \bar{y} , we call $s = \sqrt{\frac{1}{n} \sum_i \text{dist}(y, x_i)^2}$ the empirical standard deviation and s^2 the empirical variance.

Other types of central values can be defined in this framework: we define more generally the “mean deviation” at order α by

$$\begin{aligned} \sigma_\alpha(y) &= (E(\text{dist}(y, x)^\alpha))^{1/\alpha} \\ &= \left(\int_{\mathcal{M}} \text{dist}(y, x')^\alpha \cdot p_x(x') \cdot dx' \right)^{1/\alpha} \end{aligned} \quad (13)$$

If this function is finite over \mathcal{M} , the features \bar{x}_α minimizing it are called central features of order α . To be more practical, we obtain the modes of the pdf for $\alpha = 0$, the median features for $\alpha = 1$, of course the mean or expected features for $\alpha = 2$, and the “barycenter” of the support of the pdf (which is a compact set) if the mean deviation is finite for $\alpha = \infty$. This can also be applied to define the empirical central features at any order.

The nice properties of the Fréchet expectation and mean features are, in our case, due to the invariant distance (see appendix A.4 for proofs); these sets are stable under the transformation group:

$$\begin{aligned} \mathbb{E}(g \star x) &= g \star \mathbb{E}(x) \\ M(\{g \star x_i\}) &= g \star M(\{x_i\}) \end{aligned} \quad (14)$$

Since the distance we use does not depend on the representation, the results of all minimization are ensured to be also independent of the representation. Thus we have obtained a stable definition (and a mean of computation via optimization) for the expected features of a pdf and for the empirical mean features.

5.3. Application: the mean frame

Assume that we have a set of frame measurements $f_i = (r_i, x_i)$. We are looking for the mean frame $f = (r, x)$ in the Fréchet sense. Since the distance from f to f_i is

$$\text{dist}(f, f_i) = N_\lambda(f^{(-1)} \circ f_i) = N_\lambda(f_i^{(-1)} \circ f)$$

with $f_i^{(-1)} \circ f = (r_i^{(-1)} \circ r ; r_i^{(-1)} \star (x - x_i))$, the least squares criterion reduces to

$$C_\lambda(f) = \lambda^2 \sum_i \|r_i^{(-1)} \circ r\|^2 + \sum_i \|x - x_i\|^2$$

We can thus minimize independently for the position and the orientation. Moreover, the solution is independent of the parameter λ .

The position is given by the barycenter of the frame positions and the orientation is obtained by an iterative gradient descent (equations can be found in [22]). The gradient descent can be repeated for several starting points to verify that the global minimum is obtained and to test its uniqueness. Another method following the same principles but incorporating second order informations (covariance matrices) was also proposed in [22].

We can now come back to the mean rotation problem raised in introduction : the mean rotation matrix $\underline{R} = \frac{1}{n} \sum_i R_i$ is generally not a rotation matrix, nor is the mean unit quaternion $\underline{q} = \frac{1}{n} \sum_i q_i$ a unit quaternion. These two solutions have thus to be discarded. The standard mean rotation vector $\underline{r} = \frac{1}{n} \sum_i r_i$ is always a rotation vector, but the results is not coherent with the action of rotations. On the contrary, the Fréchet expectation is always correctly (but not always uniquely) defined, as illustrated in figure (8).

6. Experiments: a data fusion problem

We presented in [22] an algorithm for the registration based on 3D frames which also quantifies the uncertainty on both the data and the transformation. We used it to register medical images and showed that the accuracy of the registration is far below both the voxel size and the uncertainty of the individual features. In this method, only the most stable features are used to compute the registration, and a lot of matches are discarded due to their large uncertainty.

The aim of the present experiments is to fuse together several registered images of a single patient in order to construct an average model based on extremal points. We are interested here in both the ‘‘topological’’ stability of the model features (their probability of observation) and their geometric stability, i.e. their deviation from the model in different observation. Thus the selectivity of the features (section 3.4) is of the upper importance. Such a study on several patients will eventually lead to identify the most stable anatomical features (landmarks), and will allow to reduce better the complexity while increasing the robustness of the registration task.

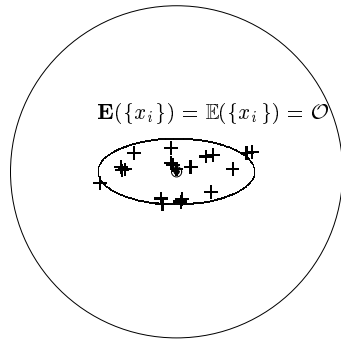
The key point is to be able to compute the ‘‘mean feature’’ even in the presence of large deviations. We saw for instance in section 3.4 that the selectivity of a trihedron match remains high even for a large error bound: it is most interesting to keep in our model the mean frames and not only the mean points. In such a case, the Fréchet expectation we defined in section 5.2 is particularly well suited. We use more precisely the mean feature (equation 12), defined with the invariant distance on rigid transformations (see 4.6 and section 5.3 for the algorithm).

6.1. 3D medical images

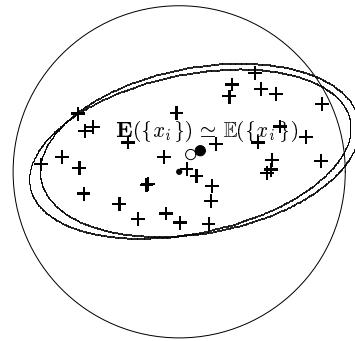
We present results from an experiment performed using 3D Magnetic Resonance images (MRI) in collaboration with Dr. R. Kikinis and Dr C. Guttmann from the Brigham and Woman’s Hospital. These images are part of an extensive study of the evolution of the Multiple Sclerosis (MS) disease. The same patient gets a complete 3D MR examination several times during one year (typically 24 different 3D acquisitions). The aim is to register precisely in 3D all those images in order to segment the lesions and evaluate very accurately their evolution. The images are first echo, $256 \times 256 \times 54$ voxels, the voxel size is $1mm \times 1mm \times 3mm$.

6.2. Extremal points and frames

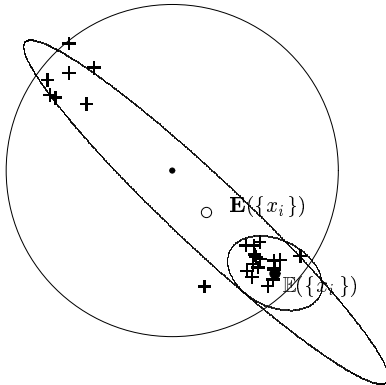
Our registration algorithm relies on the extraction of feature points in 3D medical images, defined with differential geometry criteria: the *Extremal Points* (defined in [29]). These are points on the object surface for which both principal curvatures are extremal. The interesting thing is that not only do we get some



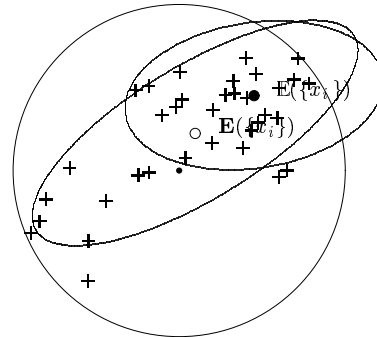
(a) Standard and Fréchet expectations are identical at the origin.



(b) When the Fréchet expectation is near the origin, the standard expectation can be considered as a first order approximation.



(c) Near the Domain boundary, standard and Fréchet expectations greatly differ. We note that with this relatively small noise, it would be possible to detect the boundary effect.



(d) With a greater amount of noise, it is no longer possible to guess the correct clustering to avoid the boundary effect.

Fig. 8. Behavior of expected rotations and their covariance matrices: projection of the measured rotation vectors, their standard and Fréchet expectation and the corresponding uncertainty ellipsoid at $\chi^2 = 15$ onto the (r_x, r_y) plane. The circle represent the domain boundary for the rotation vector ($\theta = \|r\| = \pi$). Remember that when we cross this boundary on one side at point $r = \theta.n$, we reenter the domain at the symmetric point $r' = -\theta.n$.

invariant measurements associated with those points (the principal curvatures), which are used to reduce the complexity of the matching process, but we get also the principal directions, which form, with the normal to the surface and the extremal point itself, an orthonormal basis, that is, a frame.

Typically, we extract about 3000 extremal points from a 3.5 million voxels image. Our matching algorithm produces about 600 pairs of associated extremal points between two images with a residual mean square error (RMS) of about $1mm$, and about 1000 additional matches with a RMS around $5mm$.

6.3. Building a model

Among the 24 images of a patient, one is considered as the reference for registration: the algorithm presented in [22] registers the 23 others using the 600 most ge-

ometrically stable frames. The accuracy obtained for the transformation is sufficiently small (compared to the large deviations on frames that we want to handle) to consider that the transformations are exact.

Then we regroup the frames that we can match without ambiguity in several images. We consider that two frames $f_1 = (r_1, x_1)$ and $f_2 = (r_2, x_2)$ are matched without ambiguity if the distance between them is less than a given threshold (more exactly a threshold $\theta_0 \simeq 20^\circ$ on $\|r_1^{(-1)} \circ r_2\|$ for the orientation and a threshold $d_0 = 0.8mm$ on $\|x_1 - x_2\|$ for the position), and if there is no other frame that can be matched with any of the two frames within these bounds. To find matched frames among multiple images, we compute matches between each pair of images, and look for maximal cliques of correspondences between the whole set of images. This process is rather time consuming and can certainly be improved. However, the high selectiv-

ity of frames keeps the algorithm complexity at a reasonable level, which would not be the case with only points (see section 3.4).

For each group of matched frames, we compute the Fréchet mean frame (section 5.3) using the invariant distance of section 4.6. To build a interesting model, we add to this mean frame its variance and more precisely the variance σ_θ on the orientation and the variance σ_x on the position. Let $\bar{f} = (\bar{r}, \bar{x})$ be the mean frame; these variances are computed with:

$$\sigma_\theta^2 = \sum_i \|\bar{r}^{(i)} \circ r_i\|^2 \quad \text{and} \quad \sigma_x^2 = \sum_i \|x_i - \bar{x}\|^2$$

This characterizes the geometric stability of this feature. To characterize its “topological” stability, we add also its probability of observation, i.e. the number of images where it is observed divided by the total number of images (here 24).

6.4. Results

In figure 9 we present the surface of the brain and the crest lines extracted from the first image and the most stable frames from the model. Remember that those frames do not exist in any of the 24 images: they are mean frames. We observed that about 30 frames are extremely well preserved, both geometrically and topologically, and 70 others are observed in more than 80% of the images.

An interesting result is that the most stable extremal points are located on the surface of the brain and not on the skin nor on the skull. Since the images come from the magnetic resonance modality, the skull is indeed not very visible. This points out the fact that the registration is mainly done on the surface of the brain, as expected.

The probability of observation of frames is linked with the choice of the error bound for the multiple matching step. The above figures were obtained with

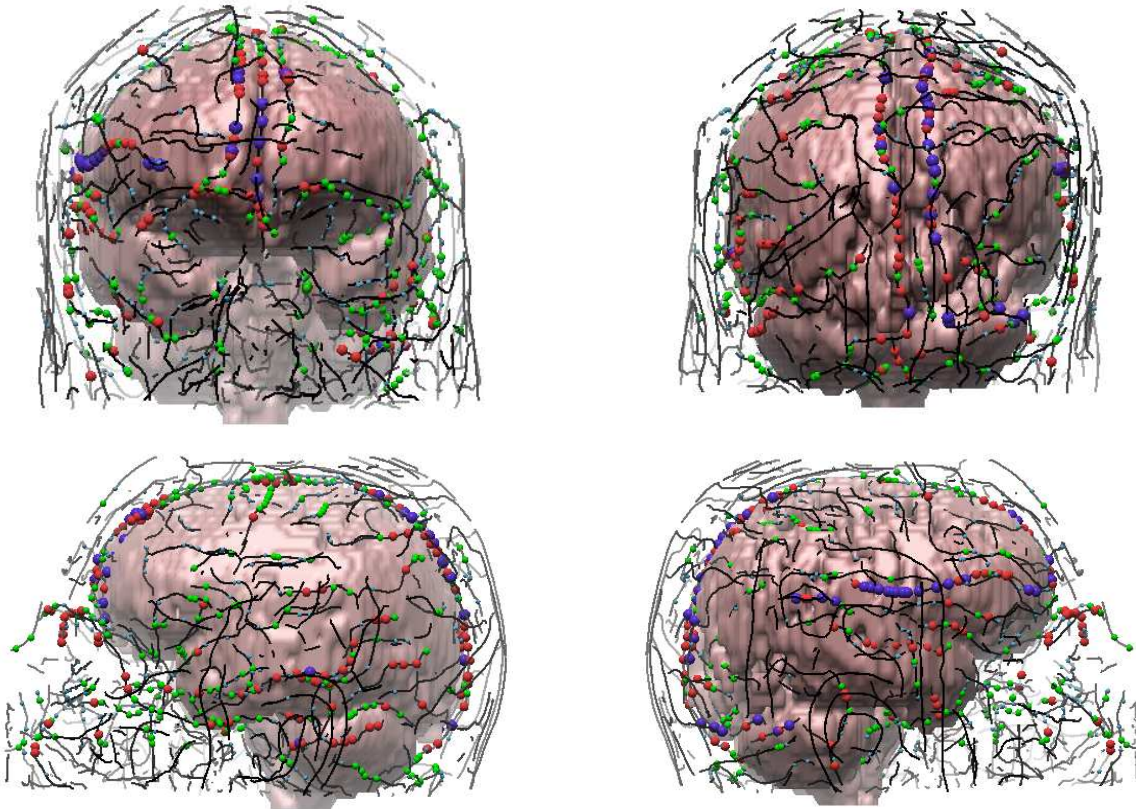


Fig. 9. The surface of the brain is displayed along with the crest lines extracted from the first image. Spheres represent the model extremal points computed as the mean over the 24 images. Their size is inversely proportional to their stability. Top line from left to right: front and rear views of the head. Bottom line: left and right views.

$\theta_0 \simeq 20^\circ$ and $d_0 = 0.8 \text{ mm}$. Since only 70 frames are observed in more than 80% of the cases, these bounds seem to be quite restrictive. Thus we have remade the experiment with larger bounds ($\theta_0 \simeq 30^\circ$ and $d_0 = 1.4 \text{ mm}$) and 200 frames are now observed at 80% (figure 10). However, the multiple matching step is in this case much longer and less reliable.

In order to obtain more accurate statistics, we plan to adapt automatically the error bound for each group of matched feature in order to maximize for each model feature the number of non ambiguous matches. This should allow to compute even more robustly the stable features. At the current stage, we did not observe correlation between the probability of observation and the geometrical stability of the frames (variances σ_θ and σ_x). We believe that a more thorough analysis with an automatic choice of the error bounds should allow to detect interesting correlations.

This kind of statistical model comprising second order informations characterizes the features that are geometrically and topologically stable in one object (here the head of a patient). It can thus be considered as a reference for the evaluation of feature selection criterions. For instance, we plan to compare several multi-scale criterions for selecting extremal points [7] with this model in order to select the best ones and validate them. Another interesting experiment would be to compute the models of several patients and compare them to characterize the anatomically stable extremal points. The results could then be incorporated into an anatomical atlas (see for instance [28]). The problem is more complex since transformations are not any more rigid.

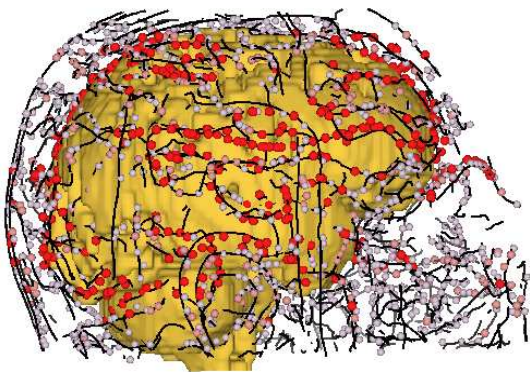


Fig. 10. Extremal points matched in more than 80% of the 24 images with a larger error bound.

7. Conclusion

We show in this paper that a wide range of paradoxes arise when we try to generalize to geometric features the classical algorithms used for points, and we demonstrate, in the case of homogeneous features, that they can be avoided by the careful definition of basic operators which respect the following rules:

- Independence of the representation.
- Invariance or “commutativity” with respect to the action of the associated transformation group.

We develop general methods that allow to define three basic operators following these rules: the invariant measure, invariant distances and stable expectation and mean features. These methods are illustrated with frame features under the rigid transformation group, and an application to the data fusion problem is presented.

We believe that the application of these two basic rules to a large number of geometrical problems can lead to a proper mathematical framework that will give reasonable and robust results in any situation without the need to design ad-hoc heuristics. We are currently working on a theory of uncertainty on geometric features continuing the formalism introduced in this paper. A further interesting development will concern the invariants and their relationship with the statistical theory of shapes.

Appendix

A.1. Proofs for classical geometric probabilities

A.1.1. Invariant measures on a group

Let $d_L \mathcal{G}(f)$ be the following measure:

$$d_L \mathcal{G}(f) = \frac{df}{|J_L(f)|} \quad \text{with} \quad J_L(f) = \left. \frac{\partial(f \circ e)}{\partial e} \right|_{e=I_d}$$

where $|J| = |\det(J)|$. We want to show that it is left invariant. We have:

$$d_L \mathcal{G}(g \circ f) = \frac{d(g \circ f)}{\left| \frac{\partial((g \circ f) \circ e)}{\partial e} \Big|_{e=I_d} \right|}$$

$$d(g \circ f) = \det \left(\frac{\partial(g \circ f)}{\partial f} \right) df$$

and we get by the chain rule:

$$\frac{\partial((g \circ f) \circ e)}{\partial e} \Big|_{e=I_d} = \frac{\partial(g \circ f)}{\partial f} \cdot \frac{(f \circ e)}{e} \Big|_{e=I_d}$$

Since $\det(A.B) = \det(A) \cdot \det(B)$ for square matrices, we can conclude that $d_L \mathcal{G}(g \circ f) = d_L \mathcal{G}(f)$. The proof for the right invariance of the proposed measure $d_R \mathcal{G}(f)$ is analogous.

A.1.2. Invariant measure on an homogeneous manifold

Let $d\mathcal{M}(x)$ be the following measure:

$$d\mathcal{M}(x) = \frac{dx}{|J(f_x)|}$$

with:

$$J(f_x) = \frac{\partial(f_x \star e)}{\partial e} \Big|_{e=\mathcal{O}} \quad \text{and} \quad f_x \in \mathcal{F}_x$$

The condition (6):

$$\forall h \in \mathcal{H}, \quad \left| \frac{\partial(h \star e)}{\partial e} \Big|_{e=\mathcal{O}} \right| = 1$$

is indeed required for $d\mathcal{M}(x)$ to be invariant with respect to the choice of $f_x \in \mathcal{F}_x$. Let f_x and $f'_x = f_x \circ h$ (with $h \in \mathcal{H}$) be two transformations of \mathcal{F}_x . By the chain rule, we can write:

$$\begin{aligned} J(f'_x) &= \frac{\partial(f_x \star (h \star e))}{\partial e} \Big|_{e=\mathcal{O}} \\ &= \frac{\partial(f_x \star e')}{\partial e'} \Big|_{e'=\mathcal{O}} \cdot \frac{\partial(h \star e)}{\partial e} \Big|_{e=\mathcal{O}} \\ &= J(f_x) \cdot \frac{\partial(h \star e)}{\partial e} \Big|_{e=\mathcal{O}} \end{aligned}$$

and we have $|J(f_x)| = |J(f'_x)|$ if and only if $\left| \frac{\partial(h \star e)}{\partial e} \Big|_{e=\mathcal{O}} \right| = 1$. The proof of the invariance of the measure $d\mathcal{M}(x)$ is then very similar to the proof for the group measure (and can be obtained by replacing $(f \circ e)$ by $(f \star e)$).

A.2. Norm induced on the manifold by the group

Let N_g be a norm on the group \mathcal{G} . We define the induced semi-norm on the homogeneous manifold \mathcal{M} as

$$\begin{aligned} N(x) &= \inf_{(h \in \mathcal{H}, f_x \in \mathcal{F}_x)} (N_g(h \circ f_x)) \\ &= \inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1 \circ f_x \circ h_2)) \end{aligned} \quad (\text{A1})$$

The positivity of N follows immediately from the positivity of N_g . The symmetry comes from the symmetry of the norm on the transformation space and the symmetry of the norm definition:

$$\begin{aligned} N(f_x^{(-1)} \circ \mathcal{O}) &= \inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1 \circ f_x^{(-1)} \circ h_2)) \\ &= \inf_{(h'_1, h'_2) \in \mathcal{H}^2} (N_g(h'_1 \circ f_x^{(-1)} \circ h'_2)) \\ &= N(x) \end{aligned}$$

The triangular inequality is preserved by the infimum:

$$\begin{aligned} &\inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1^{(-1)} \circ f_x^{(-1)} \circ f_y \circ h_2)) \\ &\leq \inf_{h_1 \in \mathcal{H}} \{N_g(f_x \circ h_1)\} + \inf_{h_1 \in \mathcal{H}} \{N_g(f_y \circ h_2)\} \\ &\leq \inf_{h_1, h_2} (N_g(h_1 \circ f_x \circ h_2)) + \inf_{h_1, h_2} (N_g(h_1 \circ f_y \circ h_2)) \end{aligned}$$

and we eventually get $N(f_x^{(-1)} \star y) \leq N(x) + N(y)$.

Now the definiteness: if the infimum of $N_g(h_1 \circ f \circ h_2)$ is reached for every transformation f by $(h_1, h_2) \in \mathcal{H}^2$, then the semi-norm is separable and is thus a norm (this is in particular always true if the isotropy group \mathcal{H} is compact):

$$\begin{aligned} N(x) = 0 &\Leftrightarrow \exists (h_1, h_2) \in \mathcal{H}^2 / N_g(h_1 \circ f_x \circ h_2) = 0 \\ &\Leftrightarrow \exists (h_1, h_2) \in \mathcal{H}^2 / f_x = h_1 \circ h_2 \\ &\Leftrightarrow f_x \in \mathcal{H} \Leftrightarrow \mathcal{F}_x = \mathcal{H} = \mathcal{F}_{\mathcal{O}} \\ &\Leftrightarrow x = \mathcal{O} \end{aligned}$$

A.3. Norm on rigid transformations

The ‘‘norm’’ definition is (λ is a fixed parameter):

$$N_\lambda(f)^2 = N_\lambda((r, t))^2 = \|f\|^2 = \lambda^2 \|r\|^2 + \|t\|^2$$

This ‘‘norm’’ is positive and null only for $\|r\| = \|t\| = 0$, that is for identity. If $f = (r, t)$, we have

$f^{(-1)} = (r^{(-1)}, r^{(-1)} \star (-t))$ (see [23] for equations on rotations and frames) and since $\theta = \|r\|$ is a “norm” on rotations:

$$N_\lambda(f^{(-1)})^2 = \lambda^2 \|r^{(-1)}\|^2 + \|\ -R^\top \cdot t\|^2 = N_\lambda(f)^2$$

The triangular inequality follows from triangular inequalities on the rotation and translation “norms”: let $f_1 = (r_1, t_1)$ and $f_2 = (r_2, t_2)$ be two frames. We have

$$f_1^{(-1)} \circ f_2 = (r_1^{(-1)} \circ r_2 ; r_1^{(-1)} \star (t_2 - t_1))$$

and thus $N_\lambda(f_1^{(-1)} \circ f_2)^2 = \lambda^2 \|r_1^{(-1)} \circ r_2\|^2 + \|R_1^\top \cdot (t_2 - t_1)\|^2$. The triangular inequality on rotations ensures that $\theta(r_1^{(-1)} \circ r_2) \leq \theta_1 + \theta_2$ (where $\theta_i = \|r_i\|$) and we have on vectors $\|t_2 - t_1\|^2 = \|t_1\|^2 + \|t_2\|^2 - 2 \langle t_1 | t_2 \rangle$. Hence:

$$N_\lambda(f_1^{(-1)} \circ f_2)^2 \leq \theta_1^2 + \theta_2^2 + 2\theta_1\theta_2 + \|t_1\|^2 + \|t_2\|^2$$

Since $\theta_1^2\theta_2^2 \leq (\theta_1^2 + \|t_1\|^2)(\theta_2^2 + \|t_2\|^2)$, we obtain that

$$\begin{aligned} N_\lambda(f_1^{(-1)} \circ f_2)^2 &\leq \frac{(\theta_1^2 + \|t_1\|^2) + (\theta_2^2 + \|t_2\|^2)}{+2 \cdot \sqrt{(\theta_1^2 + \|t_1\|^2)(\theta_2^2 + \|t_2\|^2)}} \\ &\leq \left(\sqrt{\theta_1^2 + \|t_1\|^2} + \sqrt{\theta_2^2 + \|t_2\|^2} \right)^2 \end{aligned}$$

Taking the root-square, we obtain the requested inequality

$$N_\lambda(f_1^{(-1)} \circ f_2) \leq N_\lambda(f_1) + N_\lambda(f_2)$$

A.3.1. Metric induced on points

The norm induced on points is defined by

$$\begin{aligned} N(x) &= \inf_{(h \in \mathcal{H}, f_x \in \mathcal{F}_x)} (N_g(h \circ f_x)) \\ &= \inf_{(h_1, h_2) \in \mathcal{H}^2} (N_g(h_1 \circ f_x \circ h_2)) \end{aligned}$$

Let $f_x = (0, x) \in \mathcal{F}_x$ and $h_1 = (r_1, 0)$, $h_2 = (r_2, 0) \in \mathcal{H}$. We have thus $h_1 \circ f_x \circ h_2 = (r_1 \circ r_2; r_1 \star x)$ and the (squared) norm of this transformation is simply

$$N_\lambda^2(h_1 \circ f_x \circ h_2) = \lambda \|r_1 \circ r_2\|^2 + \|x\|^2$$

The infimum is reached for $r_1 = r_2^{(-1)}$ and we have $N(x) = \|x\|$. We note that in this case the isotropy group \mathcal{H} is reduced to the rotation group SO_3 , which is compact.

A.4. Stability of the expected and mean features

Assume that $z = g \star x$ is the random feature obtained by the transformation of the random feature x by g :

$$\begin{aligned} \sigma_z^2(y) &= E(\text{dist}(g \star x, y)^2) \\ &= E(\text{dist}(x, g^{(-1)} \star y)^2) \\ &= \sigma_x^2(g^{(-1)} \star y) \end{aligned}$$

thanks to the invariance of the distance. $\bar{x} \in \mathbb{E}(x)$ implies that $\bar{z} = g \star \bar{x} \in \mathbb{E}(z)$. Eventually, we get

$$\mathbb{E}(z) = g \star \mathbb{E}(x) \quad \text{and} \quad \sigma_z = \sigma_x \quad (\text{A2})$$

The same argument holds for the stability of central features of any order and in particular the mean features of a set x_i . If $z_i = g \star x_i$, we have

$$\mathbb{M}(\{z_i\}) = g \star \mathbb{M}(\{x_i\}) \quad \text{and} \quad s_z = s_x$$

Acknowledgments

Part of this work was supported by a fellowship from D.R.E.T./D.C.N. (France). We want to thank the whole Epidaure Research team for the help provided during this work.

References

1. S. L. Altmann. *Rotations, Quaternions, and Double Groups*. Clarendon Press - Oxford, 1986.
2. N. Ayache. *Artificial Vision for Mobile robots - Stereo-vision and Multisensor Perception*. MIT-Press, 1991.
3. N. Ayache and O.D. Faugeras. Hyper: A new approach for the recognition and positioning of two-dimensional objects. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 8(1):44–54, 1986.
4. P.J. Besl and N. McKay. A method for registration of 3D shapes. *PAMI*, 14(2):239–256, 1992.
5. H.F. Durrant-Whyte. *Integration, Coordination and Control of Multi-Sensor Robot Systems*. Kluwer Academic Publishers, 1988.
6. J. Feldmar and N. Ayache. Rigid, affine and locally affine registration of free-form surfaces. *IJCV*, 18(2):99–119, May 1996.
7. M. Fidrich and J.P. Thirion. Multiscale representation and analysis of features from medical images. In N. Ayache (ed.), editor, *Computer Vision, Virtual Reality and Robotics in Medicine (proc. first CVRMed'95)*, number 905 in LNCS, pages 358–364. INRIA, Springer, 1995.

8. M. Fréchet. Les éléments aléatoires de nature quelconque dans un espace distancié. *Ann. Inst. H. Poincaré*, 10:215–310, 1948.
9. W.E.L. Grimson. *Object Recognition by Computer - The role of Geometric Constraints*. MIT Press, 1990.
10. W.E.L. Grimson and D.P. Huttenlocher. On the sensitivity of geometric hashing. In *Proc. third ICCV*, pages 334–338, 1990.
11. W.E.L. Grimson, D.P. Huttenlocher, and D.W. Jacobs. A study of affine matching with bounded sensor error. *Int. Journ. of Comput. Vision*, 13(1):7–32, 1994.
12. G. Hochschild. *The structure of Lie groups*. Holden-Day, 1965.
13. D.P. Huttenlocher and S. Ullman. Object recognition using alignment. In *Proc. of ICCV*, pages 72–78, 1987.
14. H. Karcher. Riemannian center of mass and mollifier smoothing. *Comm. Pure Appl. Math.*, 30:509–541, 1977.
15. M.G. Kendall and P.A.P. Moran. *Geometrical probability*. Number 10 in Griffin’s statistical monographs and courses. Charles Griffin & Co. Ltd., 1963.
16. W.S. Kendall. Probability, convexity, and harmonic maps with small image i: uniqueness and fine existence. *Proc. London Math. Soc.*, 61(2):371–406, 1990.
17. F. Klein. Erlangen program. Inaugural address at the University of Erlangen, 1872.
18. Y. Lamdan and H.J. Wolfson. Geometric hashing: A general and efficient model-based recognition scheme. In *Proc. of Second ICCV*, pages 238–289, 1988.
19. Y. Lamdan and H.J. Wolfson. On the error analysis of geometric hashing. In *IEEE Int. Conf. on Comput. Vis. and Patt. Recog.*, pages 22–27, 1991.
20. X. Pennec. Correctness and robustness of 3D rigid matching with bounded sensor error. Research report 2111, INRIA, November 1993.
21. X. Pennec and N. Ayache. An $\mathcal{O}(n^2)$ algorithm for 3D substructure matching of proteins. In A. Califano, I. Rigoutsos, and H.J. Wolfson, editors, *Shape and Pattern Matching in Computational Biology*, Proc. First Int. Workshop, Seattle, Wash, June 20, 1994. Plenum Publishing, 1996.
22. X. Pennec and J.P. Thirion. Validation of 3D registration methods based on points and frames. In *Proceedings of the 5th Int. Conf on Comp. Vision (ICCV’95)*, pages 557–562, Cambridge, Ma, June 1995. Part of the INRIA Research Report n° 2470.
23. X. Pennec and J.P. Thirion. A framework for uncertainty and validation of 3D registration methods based on points and frames. *Int. Journal of Computer Vision*, in press, 1997.
24. Xavier Pennec. *L’incertitude dans les Problèmes de Reconnaissance et de Recalage – Applications en Imagerie Médicale et Biologie Moléculaire*. PhD thesis, Ecole Polytechnique, Palaiseau (France), December 1996. <http://www.inria.fr/epidaure/personnel/pennec/These.html>
25. H. Poincaré. *Calcul des probabilités*. 2nd edition, Paris, 1912.
26. I. Rigoutsos and R. Hummel. Distributed Bayesian object recognition. In *Proceedings of Int. Conf on Comput. Vis. and Pat. Recog.*, pages 180–186. IEEE Computer Society Press, June 1993.
27. L.A. Santalo. *Integral Geometry and Geometric Probability*, volume 1 of *Encyclopedia of Mathematics and its Applications*. Addison-Wesley Publishing Company, 1976.
28. G. Subsol, J.Ph. Thirion, and N. Ayache. A general scheme for automatically building 3D morphometric anatomical atlases: application to a skull atlas. In *MRCAS’95*, November 1995.
29. J.P. Thirion. New feature points based on geometric invariants for 3D image registration. *Int. Journ. Comp. Vis.*, 18(2):121–137, 1996.
30. H.J. Wolfson. Model-based recognition by geometric hashing. In O. Faugeras, editor, *Proc. of 1st Europ. Conf. on Comput. Vision (ECCV’90)*, pages 526–536. Springer-Verlag, April 1990. Lecture Notes in Computer Science 427.
31. Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Int. Journ. Comp. Vis.*, 13(2):119–152, 1994. Also Research Report No.1658, INRIA Sophia-Antipolis, 1992.
32. Z. Zhang and O. Faugeras. *3D Dynamic Scene Analysis: a stereo based approach*, volume 27 of *Springer series in information science*, chapter 2: Uncertainty Manipulation and Parameter Estimation, pages 9–27. Springer Verlag, 1992.



Xavier Pennec graduated from the Ecole Polytechnique of Paris (1992) and received his PhD from the same institution (1996) while working in the Epidaure group at INRIA (French Research Institute on Computer Science and Automatic Control), Sophia-Antipolis (France).

He is currently a postdoctoral associate at MIT, Artificial Intelligence Lab. Dr Pennec’s research interests include geometric data processing, geometry-based object recognition, geometric uncertainty and statistics, medical image analysis and protein structure analysis.



Nicholas Ayache is a Research Director at INRIA Sophia-Antipolis, France, where he leads the EPIDAURE research group on medical image analysis and robotics since 1989. He is also teaching graduate courses on computer vision at Univ. Paris XI, Ecole Centrale and consulting for a number of private companies, including currently Matra-Cap-Systèmes and Focus Medical.

He received his Ph.D. in 1983, and his “Thèse d’Etat” in 1988 both in computer science from the University of Paris XI, on topics related to model based object recognition, passive stereovision and multisensor fusion. His current research interests are in medical image processing and analysis, shape and motion representation, rigid and nonrigid registration, tracking and analysis of deformable objects, image guided and simulated therapy.