Journal of Mathematical Imaging and Vision 11, 119–136 (1999)

Extension of Affine Shape

RIKARD BERTHILSSON
Centre for Mathematical Sciences, Lund University, Box 118, 221 00 Lund, Sweden
rikard@maths.lth.se

KALLE ÅSTRÖM
Centre for Mathematical Sciences, Lund University, Box 118, 221 00 Lund, Sweden
kalle@maths.lth.se

Abstract. In this paper, we extend the notion of affine shape, introduced by Sparr, from finite point sets to more general sets. It turns out to be possible to generalize most of the theory. The extension makes it possible to reconstruct, for example, 3D-curves up to projective transformations, from a number of their 2D-projections. An algorithm is presented, which is independent of choice of coordinates, is robust, does not rely on any preselected parameters and works for an arbitrary number of images. In particular this means that a solution is given to the aperture problem of finding point correspondences between curves.

Keywords: affine shape, projective reconstruction, 3D-curves

1. Introduction

Affine shape of finite point configurations, has been introduced and studied in a series of papers by Sparr, see for example [9, 10]. It has proved to be an important tool when analyzing the geometry of cameras and scene. Instead of working with camera matrices it all comes down to finding relations between linear subspaces, and the composition of two perspective transformations corresponds to an algebraically simple operation of multiplication of the depths. So far this analysis has been confined to finite point sets.

A reconstruction algorithm, by using affine shape, has been proposed in [10]. It works for arbitrary but finite numbers of points and camera views and it is based on aligning subspaces by using orthogonal projections and maximizing some of the largest eigenvalues of the sum of these projections.

In this paper, we extend the analysis of finite point configurations to more general sets, by introducing a new definition of affine shape. These configurations are here called shapeable and they include for example curves and surfaces in $\mathbb{R}^3$, and also finite point sets as a special case. As an application of the extension of shape, we propose an algorithm that is based on Sparr's but enables reconstructions of almost general curves once the curves have been extracted from the images. Furthermore, no point correspondences between the different images are needed beforehand, but are computed by the algorithm. It is independent of the choice of coordinates and works for an arbitrary number of images.

Another approach can be found in [5, 6], where a theory for computing the structure and motion of three-dimensional curves is developed. It is shown there that images of three-dimensional curves obey several constraints at each point. These constraints involve high order spatio-temporal derivatives of the image curve motion and camera motion and are thus difficult to handle in practice due to numerical problems.

Yet another type of reconstruction can be found in [1, 3, 7], where image-motion constraints that hold for certain points on the curve are developed. These constraints are more robust and can be applied to the silhouette of curved surfaces, but they do not exploit the full structure of the curve reconstruction problem.

Another application of shape is recognition of planar objects, by looking at their boundary contours. The algorithm for this is actually just a side effect of the reconstruction algorithm. The camera is modeled as an affine camera and all that has to be done is to compare the affine shapes of the curves.

Another way of doing this recognition task can be found in [2], where certain affine invariants of curves are compared.

2. Affine Shape of Finite Point Configurations

In this section, we will recapitulate what is meant by affine shape of an ordered finite point configuration. This will, in the next section, be generalized to more general configurations.

Let $\mathcal{C}_m^n$ be the set of ordered m-point configurations

$$\mathcal{X} = (p_1, p_2, \ldots, p_m) \in \mathbb{R}^{mn}$$

in $\mathbb{R}^n$, where $p_i \in \mathbb{R}^n$ is the coordinate vector of point number i in $\mathcal{X}$.

By an affine transformation, $a : \mathbb{R}^n \to \mathbb{R}^n$, is meant a map of the form,

$$a(x) = Mx + t, \tag{1}$$

where M is an $n \times n$ matrix and t and x are $n \times 1$ matrices. The matrix a is called nonsingular if $\det M \neq 0$. In a natural way, a can be extended to a transformation $\mathcal{C}_m^n \to \mathcal{C}_m^n$, by letting it act on all points of the configuration, i.e.

$$\begin{aligned} a \circ \mathcal{X} &= a \circ (p_1, p_2, \ldots, p_m) \\ &= (a(p_1), a(p_2), \ldots, a(p_m)). \end{aligned}$$

Now let A be the group of nonsingular affine transformations $\mathcal{C}_m^n \to \mathcal{C}_m^n$. By the A-orbit of $\mathcal{X}$ is meant the set

$$\{\mathcal{Y} \mid \mathcal{Y} = a(\mathcal{X}),\ a \in A\}.$$

We write $\mathcal{X} \sim \mathcal{Y}$, when $\mathcal{X}$ and $\mathcal{Y}$ are in the same orbit. The set of equivalence classes is denoted $\mathcal{C}_m^n/A$. Let

$$s : \mathcal{C}_m^n \to \mathcal{C}_m^n / A$$

be the natural projection.

It can be shown, [8], that $s(\mathcal{X})$ can be represented by the linear subspace

$$\left\{ \xi \;\middle|\; \begin{array}{l} \sum_i \xi_i p_i = 0, \sum_i \xi_i = 0, \\ \mathcal{X} = (p_1, p_2, \ldots, p_m) \end{array} \right\}, \tag{2}$$

i.e. a linear subspace of $\mathbb{R}^m$. In particular, this means that $s(\mathcal{X}) = s(\mathcal{Y})$ if and only if there exists $a \in A$ such that $\mathcal{Y} = a(\mathcal{X})$. The linear subspace $s(\mathcal{X})$ in (2) is called the affine shape of $s(\mathcal{X}) \in \mathcal{C}_m^n$. We will now extend affine shape to more general sets, as for example curves and surfaces in $\mathbb{R}^3$.

3. Extension of Shape

Let $\mathcal{X} = (p_1, \ldots, p_m) \subset \mathbb{R}^n$ be an m-point configuration in $\mathbb{R}^n$. Let $I = \{1, 2, \ldots, m\}$ and

$$\phi_{\mathcal{X}} : I \to \mathcal{X},$$

be defined by $\phi_{\mathcal{X}}(i) = p_i$, $i \in I$. The definition (2) can then be rewritten as,

$$s(\mathcal{X}) = \left\{ f : I \to \mathbb{R} \;\middle|\; \begin{array}{l} \sum_{x \in I} f(x)\phi_{\mathcal{X}}(x) = 0, \\ \sum_{x \in I} f(x) = 0 \end{array} \right\}.$$

An extension to more general sets is done by replacing the sums with integrals. First we need some basic facts and notions on integrable functions.

Let μ be a positive Radon measure on $\mathbb{R}^d$, with supp $\mu = I$. The measure μ is said to be finite if $\mu(\mathbb{R}^d) < \infty$. We say that a real valued function $f \in L_\mu^p$ if f is μ-measurable and

$$\begin{aligned} \|f\|_p &= \left(\int |f|^p \, d\mu \right)^{1/p} < \infty, \quad 1 \leq p < \infty, \\ \|f\|_\infty &= \inf\{t \mid \mu(|f(t)| > t) = 0\} < \infty. \end{aligned} \tag{3}$$

If $1/p + 1/q = 1$, $f \in L_\mu^p$ and $g \in L_\mu^q$, the integral

$$\langle f \mid g \rangle_\mu = \int f(x) g(x) \, d\mu,$$

is defined and finite. If $1 \leq p < \infty$, then L_μ^q can be identified with the dual of L_μ^p. More precisely, for every functional $T : L_\mu^p \to \mathbb{R}$ which is linear and continuous in the norm (3), there exists one and only one $u \in L_\mu^q$, such that $T(\cdot) = \langle \cdot \mid u \rangle_\mu$.

If $f \in L_\mu^p$ and $g = (g_1, g_2, \ldots, g_n)$, with $g_i \in L_\mu^q$, $i = 1, \ldots, n$, we set

$$\langle f \mid g \rangle_\mu = (\langle f \mid g_1 \rangle_\mu, \langle f \mid g_2 \rangle_\mu, \ldots, \langle f \mid g_n \rangle_\mu).$$

Also for vector valued functions we write $g \in L_\mu^q$, if each component of g belongs to L_μ^q.

When dealing with finite point configurations, as in Section 2, $\mathcal{X}$ denotes an ordered set of points. The ordering can be considered as a parametrisation of the set. To deal with more general point configurations, it is convenient to let $\mathcal{X}$ denote point sets, without ordering or parametrisation. This will be done below.

Definition 1. Let $\mathcal{X} \subseteq \mathbb{R}^n$. By a parametrisation of $\mathcal{X}$ is meant a surjective map

$$\phi_{\mathcal{X}} : I \to \mathcal{X}$$

where $I \subseteq \mathbb{R}^d$ is the support of a positive, finite Radon-measure μ, $\phi_{\mathcal{X}}$ is defined in a neighborhood of I, and $\phi_{\mathcal{X}} \in L^q_\mu$ for some q, $1 < q \leq \infty$. Then by the affine shape of $\phi_{\mathcal{X}}$ is meant the linear subspace of L^p_μ, where $1/p + 1/q = 1$,

$$s(\phi_{\mathcal{X}}) = \left\{ f \;\middle|\; \begin{array}{l} \langle f \mid \phi_{\mathcal{X}} \rangle_\mu = 0, \langle f \mid 1 \rangle_\mu = 0, \\ f \in L^p_\mu \end{array} \right\}.$$

If this construction is possible, $\mathcal{X}$ is called shapeable.

Below, we will restrict ourselves to shapeable sets. Often the subscript $\mathcal{X}$ of $\phi_{\mathcal{X}}$ will be dropped, when it is clear from the context which configuration is meant.

Note that the requirements for a configuration to be shapeable are most often satisfied. For example, any finite union of smooth surfaces and curves is shapeable.

It is sometimes convenient to invoke the constant function 1 in the parametrisation, writing $\phi_{\mathcal{X}} = (\phi_1, \ldots, \phi_n, \phi_{n+1})$, with $\phi_{n+1} \equiv 1$, instead of $\phi_{\mathcal{X}} = (\phi_1, \ldots, \phi_n)$. This will be called extended coordinates.

Recall that if $W \subseteq V$, where V is a linear space, then the linear hull, linhull(W), is the linear space of all finite linear combinations of elements in W. Furthermore, if V is a normed linear space, then the annihilator W^0 of W is the set of all continuous and linear functionals on V that vanish on W.

In Definition 1, observe that the constant function $1 \in L^q_\mu$, since μ is finite. The affine shape may thus be expressed as an annihilator. With terminology borrowed from the finite dimensional case, cf. [9], we define:

Definition 2. Let $\phi_{\mathcal{X}} = (\phi_1, \ldots, \phi_n, \phi_{n+1})$, with $\phi_{n+1} \equiv 1$, be a parametrisation of $\mathcal{X} \subseteq \mathbb{R}^n$ in extended coordinates. Then by the depth space of $\phi_{\mathcal{X}}$ is meant the linear subspace of L^q_μ,

$$d(\phi_{\mathcal{X}}) = \text{linhull}\big(\{\phi_i\}_1^{n+1}\big).$$

Notice that for an image curve $t \to (\psi_1(t), \psi_2(t))$, the depth space is the three-dimensional space spanned by the functions $(\psi_1, \psi_2, 1)$, where 1 is the constant function.

The name depth space will be justified below. As an immediate consequence of the Definitions 1 and 2, we obtain:

Proposition 1.

$$s(\phi_{\mathcal{X}})^0 = d(\phi_{\mathcal{X}})$$
$$\text{codim}\, s(\phi_{\mathcal{X}}) = \dim d(\phi_{\mathcal{X}}).$$

4. Examples

We now give some examples where $s(\phi_{\mathcal{X}})$ and $d(\phi_{\mathcal{X}})$ can be computed explicitly.

Example 1. In Section 2 above, affine shape of finite point configurations $\mathcal{X} = \{p_1, \ldots, p_m\} \subset \mathcal{C}^n_m$ was defined. In the setting of Definition 1 it is obtained by using the measure $\mu = \sum_{i=1}^n \delta_i$, where δ_i is the Dirac measure at $I = \{1, \ldots, m\}$. Let $\phi_{\mathcal{X}}$ be a continuous parametrisation of $\mathcal{X}$. Then the affine shape is

$$s(\phi_{\mathcal{X}}) = \left\{ f \;\middle|\; \begin{array}{l} \langle f \mid \phi_{\mathcal{X}} \rangle_\mu = 0, \langle f \mid 1 \rangle_\mu = 0, \\ f \in C(\mathbb{R}) \end{array} \right\},$$

and if $\phi_{\mathcal{X}} = (\phi_1, \ldots, \phi_n, 1)$ in extended coordinates, then

$$d(\phi_{\mathcal{X}}) = \text{linhull}\big(\{\phi_i\}_1^{n+1}\big).$$

Example 2. Let $\mathcal{X}$ be a non degenerate ellipse in $\mathbb{R}^3$. By a nonsingular affine transformation, $\mathcal{X}$ can be transformed to a unit circle $\mathcal{X}'$ in a coordinate plane, say $z = 0$. Under this transformation shape unchanged by Corollary 1. One parametrisation $\phi_{\mathcal{X}'}$ of $\mathcal{X}'$ is $x = \cos(t)$, $y = \sin(t)$, $z = 0$, with $t \in [0, 2\pi]$. Since

$$1 \cup \{\cos(kt)\}_1^\infty \cup \{\sin(kt)\}_1^\infty$$

is an orthogonal basis of $L^2([0, 2\pi])$, we get

$$s(\phi_{\mathcal{X}}) = \overline{\text{linhull}\big(\{\cos(kt)\}_2^\infty \cup \{\sin(kt)\}_2^\infty\big)}$$

and

$$d(\phi_{\mathcal{X}}) = \text{linhull}(1, \cos(t), \sin(t)).$$

Example 3. Let $\mathcal{X}$ be a bounded piece of a straight line in $\mathbb{R}^2$. The set $\mathcal{X}$ can then be transformed to the line segment $\mathcal{X}' = I = [-1, 1]$, on a coordinate axis, say the x-axis, by a nonsingular affine transformation, leaving its shape unchanged. If $\phi_{\mathcal{X}}(x) = (x, 0, 0)$ for $x \in I$, then

$$s(\phi_{\mathcal{X}}) = \left\{ f \;\middle|\; \begin{array}{l} \int_{-1}^{1} f x \, dx = 0, \int_{-1}^{1} f \, dx = 0, \\ f \in L^2(I) \end{array} \right\}.$$

If we parametrise $L^2(I)$ by the orthogonal Legendre polynomials, defined on I by

$$P_n(x) = \frac{1}{2^n n!} \frac{d^n}{dx^n} (x^2 - 1)^n, \quad n \geq 0,$$

we thus have

$$s(\phi_{\mathcal{X}}) = \overline{\text{linhull}\big(\{P_i\}_2^\infty\big)}$$

and

$$d(\phi_{\mathcal{X}}) = \text{linhull}(x, 1).$$

For an infinite straight line, to get a finite measure one could use e.g. the measure $d\mu = e^{-x^2} dx$, and the Hermite polynomials.

Example 4. Example 3 can be extended to the situation when $\mathcal{X}$ is the union of a bounded piece of a straight line l and an m-point configuration $\mathcal{Y} = (p_1, \ldots, p_m) \subset \mathbb{R}^3$. As in Example 3, we can assume that the line is $-1 \leq x \leq 1$, $y = 0$, $z = 0$. Let $\mu = \mu_1 + \mu_2$, where μ_1 is the Lebesque measure on $I = [-1, 1]$ and $\mu_2 = \sum_{j \in J} \delta_j$, $J = \{2, \ldots, m + 1\}$, where δ_j is the Dirac measure at j. Let

$$\phi_{\mathcal{X}} : I \cup J \to \mathbb{R}^n$$

be a parametrisation that is continuous in a neighborhood of J and fulfills $\phi_{\mathcal{X}}(x) = (x, 0, 0)$, when $x \in I$ and $\phi_{\mathcal{X}}(j) = p_{j-1}$, when $j \in J$. The shape $s(\phi_{\mathcal{X}})$ is then given by

$$s(\phi_{\mathcal{X}}) = \left\{ (f, \alpha) \;\middle|\; \begin{array}{l} \int_{-1}^{1} f \begin{pmatrix} x \\ 0 \\ 0 \end{pmatrix} dx + \sum_1^m \alpha_i p_i = 0, \\ \int_{-1}^{1} f \, dx + \sum_1^m \alpha_i = 0, \\ f \in L^2(I) \end{array} \right\}.$$

Let $L^2(I)$ be parametrised by the Legendre polynomials P_i. For $f = P_0 \equiv 1$, the constraints imply that $\sum \alpha_i p_i = 0$ and $2 + \sum \alpha_i = 0$. For $f = P_1 = x$, they imply that $(2/3, 0, 0)^t + \sum \alpha_i p_i = 0$ and $\sum \alpha_i = 0$. For $f = P_i$, $i = 2, \ldots, m$, the constraints reduce to that of shape of finite point configurations, that is $s(\mathcal{Y})$.

To summarize, set

$$E_0(\mathcal{Y}) = \left\{ \alpha \;\middle|\; \sum \alpha_i p_i = 0, 2 + \sum \alpha_i = 0 \right\},$$

$$E_1(\mathcal{Y}) = \left\{ \alpha \;\middle|\; \begin{array}{l} (2/3, 0, 0)^t + \sum \alpha_i p_i = 0, \\ \sum \alpha_i = 0 \end{array} \right\},$$

and

$$E_i(\mathcal{Y}) = s(\mathcal{Y}), \quad i > 1.$$

Then

$$s(\phi_{\mathcal{X}}) = \overline{\text{linhull}\big(\{P_i \times E_i(\mathcal{Y})\}_0^\infty\big)},$$

where P_i are the Legendre polynomials. If $\phi_{\mathcal{X}} = (\phi_1, \phi_2, \phi_3, 1)$ in extended coordinates, then

$$d(\phi_{\mathcal{X}}) = \text{linhull}\{(x, \phi_1(J)), (0, \phi_2(J)), (0, \phi_3(J)), (1, \phi_4(J))\}.$$

Example 5. Let $\phi_{\mathcal{X}}$ be a parametrisation of a point configuration $\mathcal{X}$, with $I = \{x \mid 0 \leq x_i \leq 1\} \subset \mathbb{R}^d$, and let $T_k \subset L^2(I)$ be the set of step functions on I that are constant on $2^{-k} j_i < x_i < 2^{-k}(j_i + 1)$, $j_i = 0, 1, \ldots, k - 1$. Set $s_k(\phi_{\mathcal{X}}) = T_k \cap s(\phi_{\mathcal{X}})$, which is a closed linear space. If $f \in s_k(\phi_{\mathcal{X}})$, then $f = \sum \xi_l \chi_l$, where χ_l denotes characteristic functions, and the sum is finite. The defining property of $s(\phi_{\mathcal{X}})$ implies

$$\sum_l \xi_l \langle \chi_l \mid \phi_{\mathcal{X}} \rangle = 0, \quad \sum_l \xi_l \langle \chi_l \mid 1 \rangle = 0.$$

Since $\langle \chi_l \mid 1 \rangle = c \neq 0$, we get

$$\sum_l \xi_l \frac{\langle \chi_l \mid \phi_{\mathcal{X}} \rangle}{\langle \chi_l \mid 1 \rangle} = 0, \quad \sum_l \xi_l = 0.$$

Here $\langle \chi_l \mid \phi_{\mathcal{X}} \rangle / \langle \chi_l \mid 1 \rangle$ is the mass center of $\phi_{\mathcal{X}}$ on supp (χ_l). Hence $s_k(\phi_{\mathcal{X}})$ can be identified with the shape of the finite point configurations defined by the mass centers $\langle \chi_l \mid \phi_{\mathcal{X}} \rangle / \langle \chi_l \mid 1 \rangle$. Thus, as

$$s(\phi_{\mathcal{X}}) = \overline{\bigcup_0^\infty s_k(\phi_{\mathcal{X}})},$$

we can interpret $s_k(\phi_{\mathcal{X}})$ as the shape of $\phi_{\mathcal{X}}$ at scales k, $k = 1, 2, \ldots$.

5. Some Basic Theorems

We will now prove some basic theorems about affine shape and depth. For the finite dimensional versions of these, see [9].

Suppose that $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{X}' \subseteq \mathbb{R}^{n'}$, and that $\phi_{\mathcal{X}} : I \to \mathcal{X}$ and $\phi'_{\mathcal{X}'} : I \to \mathcal{X}'$ are parametrisations as in Definition 1, with the same μ and p. If $T : \mathbb{R}^n \to \mathbb{R}^{n'}$, we write $T(\phi_{\mathcal{X}}) = \phi'_{\mathcal{X}'}$ or $T : \phi_{\mathcal{X}} \to \phi'_{\mathcal{X}'}$ if $T(\phi_{\mathcal{X}}(x)) = \phi'_{\mathcal{X}'}(x)$ almost everywhere, with respect to the measure μ. Let Σ_0 be the subspace

$$\Sigma_0 = \left\{ f \mid \langle f \mid 1\rangle_\mu = 0,\ f \in L^p_\mu \right\}.$$

Theorem 1 (Affine Shape Theorem). *Let $\phi_{\mathcal{X}} : I \to \mathcal{X} \subseteq \mathbb{R}^n$ and $\phi'_{\mathcal{X}'} : I \to \mathcal{X} \subseteq \mathbb{R}^n$ be parametrisations as in Definition 1, with the same μ and p.*

(i) *If $a : \mathbb{R}^n \to \mathbb{R}^n$ is an affine transformation, then*

$$a : \phi_{\mathcal{X}} \to \phi'_{\mathcal{X}'} \iff s(\phi_{\mathcal{X}}) \subseteq s(\phi'_{\mathcal{X}'}) \iff d(\phi_{\mathcal{X}}) \supseteq d(\phi'_{\mathcal{X}'}).$$

(ii) *Let $S \subseteq \Sigma_0$ be a closed linear space, with $\operatorname{codim}(S) \le n+1$, where $\operatorname{codim}(S)$ is taken with respect to L^p_μ. Then there exists $\mathcal{X} \subseteq \mathbb{R}^n$ with a parametrisation $\phi_{\mathcal{X}} \in L^q_\mu$, such that $s(\phi_{\mathcal{X}}) = S$.*

Proof: In (i), the second equivalence is an immediate consequence of Proposition 1. To verify $\Rightarrow$, in the first, let $f \in s(\phi_{\mathcal{X}})$. Then by (1)

$$\begin{aligned}\langle f \mid \phi'_{\mathcal{X}'}\rangle_\mu &= \langle f \mid M\phi_{\mathcal{X}} + t\rangle_\mu \\ &= M\langle f \mid \phi_{\mathcal{X}}\rangle_\mu + t\langle f \mid 1\rangle_\mu = 0.\end{aligned}$$

To verify $\Leftarrow$, let $d(\phi'_{\mathcal{X}'}) \subseteq d(\phi_{\mathcal{X}})$. Then $\phi'_{\mathcal{X}'}$ is a linear combination of the components of $\phi_{\mathcal{X}} = (\phi_1, \ldots, \phi_n)$ and the constant function 1, which gives an affine transformation. This proves statement (i).

Statement (ii) follows from the fact that L^q_μ is the dual to L^p_μ. In fact, S can be written

$$S = \bigcap_1^{n+1} H_i,$$

where H_i are closed hyperplanes, not necessarily all different, and $H_{n+1} = \Sigma_0$. Define linear functionals $L_i : L^p_\mu \to \mathbb{R}$, by $L_i(H_i) = 0$ and $L_i(u_i) = 1$, when $u_i \in L^p_\mu \backslash H_i$, where $A\backslash B = \{x \mid x \in A,\ x \notin B\}$. Then L_i is continuous and there exists, by the duality, $\phi_i \in L^q_\mu$, so that $L_i(\cdot) = \langle \cdot \mid \phi_i\rangle_\mu$, $i = 1, \ldots, n+1$. Let $\phi_{\mathcal{X}} : I \to \mathbb{R}^n$ be defined by $\phi_{\mathcal{X}}(x) = (\phi_1(x), \ldots, \phi_n(x))$. Then $\phi_{\mathcal{X}}(I) = \mathcal{X}$ is a point configuration fulfilling (ii). □

Corollary 1. *Under the assumptions of Theorem 1,*

$$a : \phi_{\mathcal{X}} \to \phi'_{\mathcal{X}'},$$

with a nonsingular affine transformation

$$\iff s(\phi_{\mathcal{X}}) = s(\phi'_{\mathcal{X}'}).$$

Proof: Apply Theorem 1 (i) to a and a^{-1}. □

It follows that $s(\phi_{\mathcal{X}})$ is a complete affine invariant.

Let Π be an affine hyperplane in $\mathbb{R}^n$. The corresponding projective hyperplane Π_* is obtained by adjoining points at infinity, which can be identified with the direction vectors in Π.

Definition 3. Let Π_*, Π'_* be projective hyperplanes in $\mathbb{P}^n$. If $c \notin \Pi'_*$, then $P : \Pi_* \to \Pi'_*$ is called a perspective transformation with center c, if for $x \in \Pi_*$ there exists $y \in \Pi'_*$ and $\alpha \in \mathbb{R}$, such that $\overrightarrow{cx} = \alpha\,\overrightarrow{cy}$, where $\overrightarrow{ab}$ is the vector from a to b in $\mathbb{R}^n$. α is called the depth of x. We also allow c to be a point at infinity in which case P is a parallel projection along the direction corresponding to c, and the depth α is set to 1.

Remark. For perspective transformations only two cases can appear: either $c \in \Pi_*$, in which case P is called singular, or $c \notin \Pi_*$, in which case P is bijective, and is called nonsingular.

We will now use affine shape to characterize perspective transformations. Following Definition 3, we assume that, when $\mathcal{Y}$ is the perspective image of $\mathcal{X}$, the parametrisations $\psi_{\mathcal{Y}} : I \to \mathcal{Y}$ and $\phi_{\mathcal{X}} : I \to \mathcal{X}$ are related by

$$\overrightarrow{c\phi_{\mathcal{X}}(x)} = \alpha(x)\overrightarrow{c\psi_{\mathcal{Y}}(x)}, \quad \text{a.e.,} \tag{4}$$

for some μ-measurable function $\alpha : I \to \mathbb{R}$. Then α is called the depth of $\phi_{\mathcal{X}}$ with respect to $\psi_{\mathcal{Y}}$. We also write $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$. The depth of $\phi_{\mathcal{X}}$ gives how much the vector from the camera center to the image point $\psi_{\mathcal{Y}}(x)$ should be enlonged to coincide with the vector from the camera center to the object point $\phi_{\mathcal{X}}(x)$ almost everywhere.

We say that α is non degenerate if α and $1/\alpha$ belong to L^{∞}_{μ}. The following lemma will be useful.

Lemma 1. *If $\alpha \in d(\phi_{\mathcal{X}})$ is non degenerate, then*

(i) $\alpha L^p_\mu = L^p_\mu$, $\alpha L^q_\mu = L^q_\mu$,
(ii) $\operatorname{codim}(s(\phi_{\mathcal{X}})) = \operatorname{codim}(\alpha s(\phi_{\mathcal{X}}))$ *and* $\alpha s(\phi_{\mathcal{X}})$ *is closed.*

Proof: (i) follows from Hölder's inequality.

The linear operator

$$T : L^p_\mu \ni f \to \alpha f \in L^p_\mu$$

is seen to be a homeomorphism and the restriction of T to the closed space $s(\phi_{\mathcal{X}})$ gives (ii). □

The following theorem motivates the name depth space in Definition 2.

Theorem 2 (Depth Theorem). *Let $\phi_{\mathcal{X}} : I \to \mathcal{X} \subseteq \mathbb{R}^n$ and $\psi_{\mathcal{Y}} : I \to \mathcal{Y} \subseteq \mathbb{R}^n$ be parametrisations as in Definition 1, with the same μ and p.*

(i) *If α is the depth of a perspective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$, then $\alpha \in d(\phi_{\mathcal{X}})$ and*

$$\alpha s(\phi_{\mathcal{X}}) \subseteq s(\psi_{\mathcal{Y}}).$$

(ii) *Let $\phi_{\mathcal{X}}(I) \subseteq \Pi_*$. If $\alpha \in d(\phi_{\mathcal{X}})$ is non degenerate and $c \notin \Pi_*$, then there exists $\psi_{\mathcal{Y}}$, and a nonsingular perspective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$ with depth α and center c.*

(iii) *If $\alpha \in d(\phi_{\mathcal{X}})$ is non degenerate and $\lambda \in \Sigma_0 \backslash \alpha \Sigma_0$, then there exists $\psi_{\mathcal{Y}}$ and a singular perspective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$ with depth α, such that*

$$s(\psi_{\mathcal{Y}}) = \operatorname{linhull}(\{\alpha s(\phi_{\mathcal{X}})\}, \lambda).$$

The perspective center is given by

$$c = \left\langle \frac{\lambda}{\alpha} \,\middle|\, \phi_{\mathcal{X}} \right\rangle_\mu \Big/ \left\langle \frac{\lambda}{\alpha} \,\middle|\, 1 \right\rangle_\mu .$$

If $\alpha \equiv C$, where $C \neq 0$ is a constant, and $\lambda \in \Sigma_0 \backslash s(\phi_{\mathcal{X}})$, then there exists $\psi_{\mathcal{Y}}$ and a parallel perspective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$ along the direction $\langle \lambda \mid \phi_{\mathcal{X}} \rangle$, such that

$$s(\psi_{\mathcal{Y}}) = \operatorname{linhull}(\{s(\phi_{\mathcal{X}})\}, \lambda).$$

Remark. Note that if α is not constant modulo sets of measure zero then $\Sigma_0 \backslash \alpha \Sigma_0 \neq \emptyset$.

Theorem 2 states that the only depths that can occur when applying a perspective transformation on $\phi_{\mathcal{X}}$ are given by $d(\phi_{\mathcal{X}})$.

Before giving the proof, recall that if $\mathcal{X} \subseteq \mathbb{R}^n$, then the affine hull, affhull $(\mathcal{X})$, is the affine space in $\mathbb{R}^n$ of smallest dimension that contains $\mathcal{X}$.

Proof: (i) The assumption means that there exists $c \in \mathbb{R}^n$, such that $\alpha(x)\overrightarrow{c\psi}(x) = \overrightarrow{c\phi}(x)$, $x \in I$. Let $f \in s(\phi)$. Then

$$0 = \langle f \mid \overrightarrow{c\phi} \rangle_\mu = \langle f \mid \alpha \overrightarrow{c\psi} \rangle_\mu = \langle f \mid \alpha\psi \rangle_\mu - c\langle f \mid \alpha \rangle_\mu .$$

If $\langle f \mid \alpha \rangle_\mu \neq 0$, then

$$c = \frac{\langle f \mid \alpha\psi \rangle_\mu}{\langle f \mid \alpha \rangle_\mu} = \frac{\langle f\alpha \mid \psi \rangle_\mu}{\langle f\alpha \mid 1 \rangle_\mu} = \langle g \mid \psi \rangle_\mu,$$

with

$$g = \frac{f\alpha}{\langle f\alpha \mid 1 \rangle_\mu}.$$

Since $\langle g \mid 1 \rangle_\mu = 1$, it follows that

$$c = \langle g \mid \psi \rangle_\mu \in \operatorname{affhull}(\psi(I)),$$

i.e., $c \in \Pi'_*$, which is a contradiction. Hence $\langle f \mid \alpha \rangle_\mu = 0$, i.e., $\alpha \in s(\phi)^0 = d(\phi)$, and $\langle f\alpha \mid \psi \rangle_\mu = 0$. This proves that $\alpha s(\phi) \subseteq s(\psi)$.

(ii) Define ψ by $\overrightarrow{c\psi} = \overrightarrow{c\phi}/\alpha$. Then, by Lemma 1, $\psi \in L^q_\mu$ and thus is a parametrisation for some $\mathcal{Y}$. Since $\alpha s(\phi) \subseteq s(\psi)$, by (i), this implies that $\operatorname{codim}(s(\psi)) \leq \operatorname{codim}(s(\phi))$, by Lemma 1. Thus,

$$\dim \operatorname{affhull} \psi(I) \leq \dim \operatorname{affhull} \phi(I)$$

and $P : \phi \to \psi$ is a perspective transformation with depth α. Since $P : \phi \to \psi$ is bijective, we actually have $\dim \operatorname{affhull} \psi(I) = \dim \operatorname{affhull} \phi(I)$.

(iii) First assume that α is not constant and define ψ as in (ii). Then

$$\langle \lambda \mid \overrightarrow{c\psi} \rangle_\mu = \left\langle \frac{\lambda}{\alpha} \,\middle|\, \overrightarrow{c\phi} \right\rangle_\mu = \left\langle \frac{\lambda}{\alpha} \,\middle|\, \phi \right\rangle_\mu - \left\langle \frac{\lambda}{\alpha} \,\middle|\, c \right\rangle_\mu .$$

Let

$$c = \left\langle \frac{\lambda}{\alpha} \,\middle|\, \phi \right\rangle_{\mu} \Big/ \left\langle \frac{\lambda}{\alpha} \,\middle|\, 1 \right\rangle_{\mu}.$$

Then $\langle \lambda \mid \psi \rangle_\mu = 0$, so that $\lambda \in s(\psi)$ and linhull$(\alpha s(\phi), \lambda) \subseteq s(\psi)$.

To verify the reverse inclusion, let $f \in s(\psi)$. Then

$$\langle f \mid \psi \rangle_\mu = \left\langle f \,\middle|\, \frac{\phi}{\alpha} - \frac{c}{\alpha} + c \right\rangle_{\mu} = 0.$$

Since $\langle f \mid 1 \rangle_\mu = 0$ and $c = \langle \frac{\lambda}{\alpha} \mid \phi \rangle_\mu / \langle \frac{\lambda}{\alpha} \mid 1 \rangle_\mu$,

$$\left\langle f \,\middle|\, \frac{\phi}{\alpha} - \frac{\langle \lambda \mid \alpha^{-1}\phi \rangle_\mu}{\langle \lambda \mid \alpha^{-1} \rangle_\mu} \frac{1}{\alpha} \right\rangle_{\mu} = 0. \tag{5}$$

If $f \notin \alpha s(\phi)$, (5) implies $\langle f \mid \frac{1}{\alpha} \rangle_\mu \neq 0$ and since $\lambda \notin \alpha \Sigma_0$,

$$\left\langle f \frac{\langle \lambda \mid \alpha^{-1} \rangle_\mu}{\langle f \mid \alpha^{-1} \rangle_\mu} - \lambda \,\middle|\, \alpha^{-1}\phi \right\rangle_{\mu} = 0.$$

Thus

$$f \frac{\langle \lambda \mid \alpha^{-1} \rangle_\mu}{\langle f \mid \alpha^{-1} \rangle_\mu} - \lambda \in \alpha s(\phi),$$

and $f \in$ linhull$(\alpha s(\phi), \lambda)$.

Now assume $\alpha \equiv C$, where $0 \neq C \in \mathbb{R}$. Since $\Sigma_0 \backslash s(\phi) \neq \emptyset$, there exists $1 \not\equiv \alpha' \in d(\phi)$ such that $\langle \lambda \mid \alpha' \rangle \neq 0$. Let

$$\psi = \phi - \frac{\langle \lambda \mid \phi \rangle}{\langle \lambda \mid \alpha' \rangle} \alpha'.$$

It is easily verified that $s(\psi) =$ linhull$\{s(\phi), \lambda\}$ and since linhull$(\{\alpha s(\phi)\}, \lambda)$ is closed, the theorem is proved. □

Before stating the next theorems we must exclude a highly degenerate case for perspective transformations where all depth information is lost. This occurs when a perspective transformation is the composition of two perspective transformations where at least one is a parallel transformation. We say that a perspective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$, with depth α is flat if $s(\psi_{\mathcal{Y}}) \subseteq \alpha \Sigma_0$. In the following we assume that all perspective transformations are non flat.

Theorem 3 (Shape transform theorem). *Let $\phi_{\mathcal{X}} : I \to \mathcal{X} \subseteq \mathbb{R}^n$ and $\phi'_{\mathcal{X}'} : I \to \mathcal{X}' \subseteq \mathbb{R}^n$ be parametrisations as in Definition* 1, *with the same μ and p. Then*

$$\alpha s(\phi_{\mathcal{X}}) \subseteq s(\psi_{\mathcal{Y}}),\ \alpha \in d(\phi_{\mathcal{X}}) \text{ is non degenerate}$$
$$\iff$$
$$s(P(\phi_{\mathcal{X}})) = s(\psi_{\mathcal{Y}})$$
$$\text{for some perspective transformation } P,$$
$$\text{with depth } \alpha.$$

Proof: To prove $\Rightarrow$, first assume that

$$\alpha s(\phi) \subset s(\psi),$$

with strict inclusion. Let $\lambda \in s(\psi) \backslash \alpha s(\phi)$, so that Theorem 2(iii) is fulfilled. Then there exists a point configuration $\mathcal{Y}$, with a parametrisation φ, and a perspective transformation $P : \phi \to \varphi$, with depth α, such that $s(\varphi) =$ linhull$(\{\alpha s(\phi)\}, \lambda) = s(\psi)$, by a dimensionality argument. Thus, $s(P(\phi)) = s(\psi)$.

If instead $\alpha s(\phi) = s(\psi)$, then $\alpha s(\phi - c) = s(\psi)$, where $c \in \mathbb{R}^n$ and $s((\phi - c)/\alpha) = s(\psi)$. Let $c \notin$ affhull$(\phi(I))$ and define φ by $\overrightarrow{c\varphi} = \overrightarrow{c\phi}/\alpha$. Then $c \notin$ affhull$(\varphi(I))$ and

$$s(\varphi) = s\left(\frac{\phi - c}{\alpha}\right) = s(\psi).$$

Thus, $P : \phi \to \varphi$ is a perspective transformation and $s(P(\phi)) = s(\psi)$.

The left implication $\Leftarrow$, is an immediate consequence of Theorem 2(i). □

Theorem 4. *Let $P : \phi_{\mathcal{X}} \to \phi_{\mathcal{Y}}$ and $Q : \phi_{\mathcal{Y}} \to \phi_{\mathcal{Z}}$ be perspective transformations, with depths α and β respectively. Then there exists a perspective transformation*

$$R : \phi_{\mathcal{X}} \to \phi_{\mathcal{Z}'},$$

with depth $\alpha\beta \in d(\phi_{\mathcal{Z}})$, such that $s((Q \circ P)(\phi_{\mathcal{X}})) = s(R(\phi_{\mathcal{X}}))$.

Proof: From Theorem 3 follows that $\alpha s(\phi_{\mathcal{X}}) \subseteq s(\phi_{\mathcal{Y}})$, $\beta s(\phi_{\mathcal{Y}}) \subseteq s(\phi_{\mathcal{Z}})$. Hence $\alpha\beta s(\phi_{\mathcal{X}}) \subseteq s(\phi_{\mathcal{Z}})$ and $\alpha\beta \in d(\phi_{\mathcal{Z}})$. By applying Theorem 3 once more the theorem is proved. □

We call a composition of perspective transformations a projective transformation. By Theorem 4, the depth of a projective transformation equals the product of the depths of the corresponding perspective transformation by which it is built up. This holds independently of the factorization.

Since any affine transformation of an image can be realized by a series of parallel perspective projections ($\alpha \equiv 1$), Theorem 3 expresses the fact that $\psi_{\mathcal{Y}}$ is the projective image of $\phi_{\mathcal{X}}$ if and only if $\alpha s(\phi_{\mathcal{X}}) \subseteq s(\psi_{\mathcal{Y}})$. As all the theorems concern affine shape and everything is just known up to an affine transformation, in the sequel there is no reason to distinguish between projective and perspective transformations.

6. Applications

We will here present two applications of shape. The first will deal with reconstruction of curves up to nonsingular projective transformations, and the second with recognition of boundary curves of planar objects.

In [10], a reconstruction algorithm for arbitrary numbers of points and images, was proposed, based on affine shape. It works by aligning subspaces, using orthogonal projections and maximising some of the largest eigenvalues of the sum of these projections. That algorithm is here extended and modified to handle curves as well. As an additional issue, no point correspondences between the different images are needed. First, it is assumed that the curves are non-closed, i.e. they have different start and end points and we assume that these are known in the different images. By a simple modification, the method is then extended to closed curves as well. The algorithms are independent of the choice of coordinates and work for an arbitrary number of images.

In a second application, recognition of boundary curves of planar objects, we use a number of different curves, stored in a data base. An image of an unknown object is taken and then compared with these of the data base. The algorithm for this is actually just a variant of the algorithm for reconstruction.

7. Projective Reconstruction of Three-Dimensional Curves

Our aim is to extend the algorithm of [10] for projective reconstruction of finite point configurations to general 3D-curves.

In the following, we will assume the pinhole camera model, which means that the image is formed by a perspective transformation. Let $\mathcal{X} \subset \mathbb{R}^3$ be a 3D-curve of finite extent and $\mathcal{Y}$ a projective 2D-image of $\mathcal{X}$. Furthermore, let

$$L^2(I) \ni \psi_{\mathcal{Y}} : I \to \mathcal{Y}$$

and

$$L^2(I) \ni \phi_{\mathcal{X}} : I \to \mathcal{X},$$

be their respective parametrisations. Then, by Theorem 3 there exists a projective transformation $P : \phi_{\mathcal{X}} \to \psi_{\mathcal{Y}}$ with depth α if and only if $\alpha s(\phi_{\mathcal{X}}) \subseteq s(\psi_{\mathcal{Y}})$. However, to use this theorem, a correspondence between the parametrisations of $\mathcal{X}$ and $\mathcal{Y}$ must have been established. This is a difficult problem, referred to as the correspondence problem or aperture problem. The solving of this problem is one of the contributions of this paper.

For the moment assume that the point correspondences are known. Let $\mathcal{X}$ be a fixed but unknown object and $\{\mathcal{Y}_i\}_{i=0}^{m-1}$ a sequence of projective images of $\mathcal{X}$, taken by uncalibrated pinhole cameras. Uncalibrated means that nothing is known about the orientations of the cameras or their internal parameters. The only thing known is that $\{\mathcal{Y}_i\}_0^{m-1}$ are projective images of a fixed object, thus nothing about the projective transformations. Denote by $\phi : I \to \mathcal{X}$ and $\psi_i : I \to \mathcal{Y}_i$, parametrisations so that for some P_i with non degenerate depth $\alpha_i : I \to \mathbb{R}$,

$$\psi_i(x) = P_i \phi(x), \quad i = 0, 1, \ldots, m-1$$

holds for all $x \in I$. Then, by Theorem 3

$$\alpha_i s(\phi) \subseteq s(\psi_i), \quad i = 0, 1, \ldots, m-1,$$

and thus

$$s(\phi) \subseteq \frac{1}{\alpha_i} s(\psi_i), \quad i = 0, 1, \ldots, m-1.$$

Together this implies

$$s(\phi) \subseteq \bigcap_0^{m-1} \frac{1}{\alpha_i} s(\psi_i),$$

or equivalently, by multiplying both sides by α_0,

$$\alpha_0 s(\phi) \subseteq s(\psi_0) \bigcap_1^{m-1} q_i s(\psi_i),$$

where $q_i = \alpha_0/\alpha_i$ are called kinetic depths. Since $s(\phi)$ is usually unknown, we replace $\alpha_0 s(\phi)$ by $s(\phi)$. That

$\alpha_0 s(\phi)$ is a shape space follows from $\alpha_0 \in d(\phi)$, Theorem 1 (ii) and Lemma 1. We get the relation

$$s(\phi) \subseteq s(\psi_0) \bigcap_1^{m-1} q_i s(\psi_i). \tag{6}$$

Remark. By Theorem 3, this freedom in replacing $\alpha_0 s(\phi)$ by $s(\phi)$ corresponds to a projective ambiguity of the reconstructed configuration $\mathcal{X}$.

In the following, assume that $\mathcal{X} \subset \mathbb{R}^3$ is a bounded 3D-curve, such that $\mathcal{X}$ does not belong to any affine plane. Assume also that the focal points corresponding to the different images do not belong to an affine plane. We then say that the cameras and object are in general position.

For parametrisations, we use the Hilbert space $L^2(I)$, with $I = [0, 1]$. Then, because of the generality of the curve and camera positions, a dimensionality (or rather codimensionality) argument implies that (6) holds with equality,

$$s(\phi) = s(\psi_0) \bigcap_1^{m-1} q_i s(\psi_i). \tag{7}$$

By considering the annihilator, or in this case the orthogonal complement, we obtain

$$d(\phi) = d(\psi_0) + \sum_1^{m-1} q_i^{-1} d(\psi_i), \tag{8}$$

where $q_i^{-1} = 1/q_i$, and where the sum of two linear subspaces is defined by

$$A + B = \{a + b \mid a \in A, \ b \in B\}.$$

In fact, since $\mathcal{X}$ is a general 3D-curve, codim $s(\phi) = 4$, and in the same way codim $q_i s(\psi_i) = 3$, $i = 0, \ldots, m-1$. Since the spaces $q_i s(\psi_i)$, $i = 0, \ldots, m-1$, do not coincide, the left hand side of (6) has codimension four.

Let $\mathbf{P}_\phi$, $\mathbf{P}_i$, $\mathbf{Q}_\phi$ and $\mathbf{Q}_i$ be the orthogonal projections from $L^2(I)$ onto $s(\phi)$, $q_i s(\psi_i)$, $d(\phi)$ and $q_i^{-1} d(\psi_i)$, respectively. These projection operators can be explicitly written using orthonormal bases. For example, using extended coordinates, $\phi = (\phi_1, \phi_2, \phi_3, \phi_4)$, where $\phi_4 \equiv 1$, let $\{\tilde{\phi}_1, \tilde{\phi}_2, \tilde{\phi}_3, \tilde{\phi}_4\}$ be an orthonormal basis for linhull$(\phi_1, \phi_2, \phi_3, \phi_4)$. Then

$$\mathbf{Q}_\phi(f) = \sum_{k=1}^{4} \langle \tilde{\phi}_k \mid f \rangle \tilde{\phi}_k,$$

and

$$\mathbf{P}_\phi(f) = (I - \mathbf{Q}_\phi)(f) = f - \sum_{k=1}^{4} \langle \tilde{\phi}_k \mid f \rangle \tilde{\phi}_k.$$

The conditions (7) and (8) can be rewritten in either of the following four ways:

1. the operator $\mathbf{P}_\phi \mathbf{Q}_i$ is the zero operator for every i,
2. the restriction of $\frac{1}{m} \sum_{i=0}^{m-1} \mathbf{P}_i$ to $s(\phi_{\mathcal{X}})$ equals the identity operator,
3. the restriction of $\frac{1}{m} \sum_{i=0}^{m-1} \mathbf{Q}_i$ to $d(\phi_{\mathcal{X}})$ equals the identity operator,
4. the operator $\frac{1}{m} \sum_{i=0}^{m-1} \mathbf{Q}_i$ has only four non-zero eigenvalues.

For real images, these equalities will never hold exactly, due to noise and other errors. It is of interest to introduce an error criteria to minimize.

Any criterion that is based on the linear spaces or the projection operators above is invariant to the choice of affine coordinate system in the images. Any such criterion also has the property that all images are treated in a symmetrical fashion and works for an arbitrary number of images. Such an invariant criterion, will be called a proximity measure.

There are several possibilities. Using the fact that $\mathbf{P}_\phi \mathbf{Q}_i = 0$ for all i, one proximity measure is

$$\mu_1 = \sum_{i=0}^{m-1} \|\mathbf{P}_\phi \mathbf{Q}_i\|_{HS}^2.$$

Here *HS* stands for the Hilbert-Schmidt norm, see [4], defined by

$$\|A\|_{HS}^2 = \sum_k \|Ae_k\|^2,$$

where $\{e_i\}_1^\infty$ is an orthonormal basis for $L^2(I)$. For finite dimensional spaces it is the same as the Frobenius norm. The *HS*-norm is independent of the choice of orthonormal basis. By choosing it so that the first three basis vectors $\{e_1, e_2, e_3\}$ span $q_i^{-1} d(\psi_i)$ (and consequently $\mathbf{P}_\phi \mathbf{Q}_i e_k = 0$ for all $k > 3$), it is seen that

$$\|\mathbf{P}_\phi \mathbf{Q}_i\|_{HS}^2 = \sum_{k=1}^{3} \|\mathbf{P}_\phi e_k\|^2.$$

Thus, if $\{\tilde{\psi}_{i,1}, \tilde{\psi}_{i,2}, \tilde{\psi}_{i,3}\}$ is an orthonormal basis of $q_i^{-1} d(\psi_i)$, $i = 0, \ldots, m-1$, it follows that

$$\mu_1 = \sum_{i=0}^{m-1} \sum_{k=1}^{3} \|\mathbf{P}_\phi \tilde{\psi}_{i,k}\|^2.$$

Another proximity measure can be based on the fact that the operator

$$\frac{1}{m}\sum_{i=0}^{m-1}\mathbf{P}_i$$

equals the identity operator on $s(\phi)$. Let U be a finite dimensional subspace of $L^2(I)$, and let $u \in U \bigcap L^2(I)$. Then (7) yields

$$\frac{1}{m}\sum_{i=0}^{m-1}\|\mathbf{P}_i u\|_2^2 = \|u\|_2^2 \iff u \in s(\phi).$$

If we define M to be the restriction to U of operator

$$\frac{1}{m}\sum_{i=0}^{m-1}\mathbf{P}_i,$$

this can be rewritten as

$$\langle u, Mu\rangle = \langle u, u\rangle \iff u \in s(\phi).$$

By properties of quadratic forms, a necessary and sufficient condition for

$$\dim(U \cap s(\phi)) \geq k$$

is that k eigenvalues of M are equal to 1. We introduce the proximity measure

$$\mu_2 = \sum_{i=5}^{\dim U}\left(1 - \lambda_i^2\right),$$

where λ_i are the eigenvalues of M sorted in an increasing sequence. Here the interpretation of M by means of projections implies that all terms are non-negative. This is the proximity measure used in [10].

For small values of μ_2, the space spanned by the eigenvectors corresponding to the $\dim U - r$ largest eigenvalues, can be taken as an approximation of $U \cap s(\phi)$.

8. Algorithm for Non-Closed Curves

In the following, let $I = [0, 1]$ and $\{\mathcal{Y}_i\}_0^{m-1}$ be a sequence of projective images of an unknown 3D-curve $\mathcal{X}$. Furthermore, let all the parametrisations $\phi : I \to \mathcal{X}$ and $\psi_i : I \to \mathcal{Y}_i$, $i = 0, \ldots, m-1$, be expressed in extended coordinates.

We propose the following algorithm, which is based on repeatedly finding $s(\phi)$, adjusting the kinetic depths q_i, and the parametrisations ψ_i. It is assumed that the curve $\mathcal{X}$ has two distinct end points, which can be identified in each image curve $\mathcal{Y}_i$, $i = 0, \ldots, m-1$. To obtain a reconstruction, the problem is to find the parametrisations ψ_i, and the kinetic depths q_i in (8).

- I *Initialization*: Choose one parametrisation ψ_i in each image curve, for example by using the image based arc-length. Set $q_i(x) \equiv 1$ and $d(\psi_i) = \text{linhull}(\psi_i)$, $i = 0, \ldots, m-1$.
- II *Update $d(\phi)$*: Keeping $d(\psi_i)$ fixed for all i, find $\mathbf{P}_\phi$ that minimizes μ.
- III *Update q_i*: Keeping $d(\phi)$ and $d(\psi_i)$ fixed, find q_i such that $q_i^{-1}d(\psi_i)$ minimizes μ. Set $d(\psi_i) := q_i^{-1}d(\psi_i)$, $i = 0, \ldots, m-1$.
- IV *Update parametrisation*: Keeping $d(\phi)$ and $d(\psi_i)$ fixed, find a continuous bijection $\gamma_i : I \to I$, such that $d(\psi_i) \circ \gamma_i$ minimizes μ. Set $d(\psi_i) := d(\psi_i) \circ \gamma_i$, $i = 0, \ldots, m-1$ and go to II.

It is difficult to minimize μ with respect to all parameters simultaneously. It is, however, reasonably fast to solve each of the three steps II to IV approximately, as will be demonstrated. Since the procedure is iterated, we do not have to be very precise in each step. Below, we use the proximity measure μ_1.

8.1. *Step I—Initialization*

Let each image curve

$$\psi_i(t) = (\psi_{i1}(t), \psi_{i2}(t), \psi_{i3}(t))$$

be parametrised using scaled image arclength $t \in I$, so that $\psi_i(0)$ and $\psi_i(1)$ are the endpoints, and such that $(\psi'_{i1})^2 + (\psi'_{i2})^2$, $i = 0, \ldots, m-1$, is constant. Initially, let the depths be $q_i(t) \equiv 1$ for all points in all curves, and let $d(\psi_i) = \text{linhull}(\psi_{i,1}, \psi_{i,2}, \psi_{i,3})$.

8.2. *Step II—Computation of $d(\phi)$ given $d(\psi_i)$*

Let $\{\psi_{i,1}, \psi_{i,2}, \psi_{i,3}\}$ be an orthonormal basis for the three-dimensional linear space $d(\psi_i)$. By (8), the four-dimensional linear space $d(\phi)$, corresponding to the 3D-curve to be reconstructed, is then the linear span of all basis functions $\psi_{i,k}$, $i = 0, \ldots, m-1$, $k = 1, 2, 3$, i.e.,

$$d(\phi) = \text{linhull}\{\psi_{i,k}, i = 0, \ldots, m-1, k = 1, 2, 3\}.$$

An estimate of $d(\phi)$ is obtained by solving

$$\min_{\dim d(\phi)=4} \mu_1 = \min_{\dim d(\phi)=4} \sum_{i=0}^{m-1} \sum_{k=1}^{3} \|\mathbf{P}_\phi \psi_{i,k}\|^2. \qquad (9)$$

This optimization problem can be solved using singular value decomposition. Form the symmetric matrix

$$M_1 = \begin{pmatrix} \langle \psi_{0,1} \mid \psi_{0,1} \rangle & \dots & \langle \psi_{0,1} \mid \psi_{m-1,3} \rangle \\ \langle \psi_{0,2} \mid \psi_{0,1} \rangle & \dots & \langle \psi_{0,2} \mid \psi_{m-1,3} \rangle \\ \langle \psi_{0,3} \mid \psi_{0,1} \rangle & \dots & \langle \psi_{0,3} \mid \psi_{m-1,3} \rangle \\ \langle \psi_{1,1} \mid \psi_{0,1} \rangle & \dots & \langle \psi_{1,1} \mid \psi_{m-1,3} \rangle \\ \vdots & \ddots & \vdots \\ \langle \psi_{m-1,3} \mid \psi_{0,1} \rangle & \dots & \langle \psi_{m-1,3} \mid \psi_{m-1,3} \rangle \end{pmatrix}.$$

Compute a singular valued decomposition $M_1 = USV^T$, where U and V are orthogonal matrices and S is a non-negative diagonal matrix. In the case of exact data, the matrix M_1 has rank 4. In the case of measured data, the matrix which is closest in Frobenius norm to a matrix of rank 4, is $\hat{M} = US_4V^T$, where S_4 is obtained by setting all but the four largest diagonal elements in S to zero. An orthonormal basis for $d(\phi)$ is then

$$\phi_k = \frac{1}{\sqrt{S_{k,k}}}(V_{1,k}\psi_{0,1} + V_{2,k}\psi_{0,2} + V_{3,k}\psi_{0,3} + V_{4,k}\psi_{1,1} + \dots + V_{3m,k}\psi_{m-1,3}), \quad k = 1, 2, 3, 4.$$

This $d(\phi)$ solves the optimization problem (9).

8.3. *Step III—Computation of Kinetic Depths*

Let $\{\psi_{i1}, \psi_{i2}, \psi_{i3}\}$ be an orthonormal basis for $d(\psi_i)$ and let $\{\phi_1, \phi_2, \phi_3, \phi_4\}$ be an orthonormal basis for the four-dimensional linear space $d(\phi)$, corresponding to the curve to be reconstructed. The projection operator $\mathbf{P}_\phi : L^2(I) \to s(\phi)$ is given by

$$\mathbf{P}_\phi f = f - \sum_{j=1}^{4} \phi_j \langle \phi_j \mid f \rangle.$$

We want to find q_i^{-1}, $i = 0, \dots, m-1$ so that

$$\mathbf{P}_\phi \psi_{ik} q_i^{-1}$$

are small in some sense, e.g. by minimizing

$$\sum_{k=1}^{3} \left\| \mathbf{P}_\phi \psi_{ik} q_i^{-1} \right\|^2$$

over all q_i^{-1} with $\|q_i^{-1}\| = 1$. For convenience, the index i is dropped below, since each image can be treated separately. Parametrise q^{-1} using a finite basis f_j according to

$$I \times \mathbb{R}^n \ni (t, x) \to q^{-1}(x) = \sum_{j=1}^{n} x_j f_j(t). \qquad (10)$$

In the case of exact data, the projection $\mathbf{P}_\phi \psi_k q^{-1}(x) = \sum(\mathbf{P}_\phi(\psi_k f_j))x_j = \sum_j \vartheta_{k,j} x_j$ vanishes, where

$$\vartheta_{kj} = \mathbf{P}_\phi(\psi_k f_j).$$

In the case of non-exact data, the least squares solution can be found by singular value decomposition USV^T of the matrix

$$M_2 = \sum_{k=1}^{3} \begin{pmatrix} \langle \vartheta_{k,1} \mid \vartheta_{k,1} \rangle & \dots & \langle \vartheta_{k,1} \mid \vartheta_{k,n} \rangle \\ \vdots & \ddots & \vdots \\ \langle \vartheta_{k,n} \mid \vartheta_{k,1} \rangle & \dots & \langle \vartheta_{k,n} \mid \vartheta_{k,n} \rangle \end{pmatrix}.$$

By taking x as the last column of V we obtain the vector x of unit length which inserted in (10) gives q_i^{-1}, and minimizes

$$\min_{\|x_i\|=1} \sum_{k=1}^{3} \left\| \mathbf{P}_\phi \psi_{ik} q_i^{-1}(x_i) \right\|^2.$$

Set $d(\psi_i) = \text{linhull}\{q_i^{-1}\psi_{i1}, q_i^{-1}\psi_{i2}, q_i^{-1}\psi_{i3}\}$.

8.4. *Step IV—Reparametrisation of the Image Curves*

Let $\{\psi_{i1}, \psi_{i2}, \psi_{i3}\}$ be an orthonormal basis for $d(\psi_i)$, and let $\{\phi_1, \phi_2, \phi_3, \phi_4\}$ be an orthonormal basis for the four-dimensional linear space $d(\phi)$, corresponding to the curve to be reconstructed. We want to find a reparametrisation γ_i in each image $i = 0, \dots, m-1$, such that

$$\sum_{k=1}^{3} \|\mathbf{P}_\phi(\psi_{ik} \circ \gamma)\|^2$$

is minimized over some set of reparametrisations. Again, we drop the index i for convenience.

Parametrise γ by using a finite basis g_j, according to

$$I \times \mathbb{R}^n \ni (t, x) \to \gamma(t, x) = t + \sum_j x_j g_j(t), \qquad (11)$$

where the basis function fulfill $g_j(0) = 0$ and $g_j(1) = 0$. g_j can for example be a translated and dilated

Gaussian function multiplied by $\sin(2\pi x)$ in order to fulfill $g_j(0) = g_j(1) = 0$. The function $\gamma(t, x)$ is monotonic for small x, that is

$$\frac{\partial \gamma}{\partial t} = 1 + \sum_j x_j g_j'(t) > 0, \quad t \in I,$$

if $|x|$ is sufficiently small. This is guaranteed by

$$|x|^2 < \min_{t \in I} \frac{1}{\sum_j |g_j'(t)|^2}. \tag{12}$$

Now study the linearisation of $\Theta_k(x) = \mathbf{P}_\phi(\psi_k \circ \gamma(x))$ around $x = 0$, i.e.,

$$\Theta_k(x) \approx \Theta_k(0) + \nabla_x \Theta_k(0)x,$$

where ∇_x is the gradient operator in the x-variables. The derivatives are given by

$$\theta_{k,j} = \left.\frac{\partial \Theta_k}{\partial x_j}\right|_{x=0} = \mathbf{P}_\phi(\psi_k' g_j).$$

The Gauss-Newton iteration for the minimization problem

$$F(x) = \min_x \sum_{k=1}^{3} \|\Theta_k(x)\|^2,$$

is obtained from the normal equations

$$-Ax = b, \tag{13}$$

where

$$A = \sum_{k=1}^{3} \begin{pmatrix} \langle \theta_{k,1} \mid \theta_{k,1} \rangle & \dots & \langle \theta_{k,1} \mid \theta_{k,n} \rangle \\ \vdots & \ddots & \vdots \\ \langle \theta_{k,n} \mid \theta_{k,1} \rangle & \dots & \langle \theta_{k,n} \mid \theta_{k,n} \rangle \end{pmatrix}$$

and

$$b = \sum_{k=1}^{3} \begin{pmatrix} \langle \theta_{k,1} \mid \Theta_k \rangle \\ \vdots \\ \langle \theta_{k,n} \mid \Theta_k \rangle \end{pmatrix}.$$

If the solution of (13) gives an x not fulfilling (12), or if $F(x)$ is larger than $F(0)$ due to the non-linearities of the function F, since A is positive definite and therefore x is a descent direction, it is always possible to decrease the error function by restricting the step length.

After having solved (13) for x, (11) defines a reparametrisation of the basis for $d(\psi_i)$, and we set

$$d(\psi_i) = \text{linhull}\{\psi_{i1} \circ \gamma_i(x), \psi_{i2} \circ \gamma_i(x), \psi_{i3} \circ \gamma_i(x)\}.$$

Observe that in this step we differentiate ψ_{ij} in order to use the Gauss-Newton iteration. We therefore have to assume that also $\psi_{ij}' \in L^2(I)$.

Remark. The functions q_i^{-1} and γ_i are computed by expanding them as finite linear combinations of some basis functions (cf. formulae (10) and (11) respectively). This rises the question of how these bases should be chosen or what bases are best to use in the algorithm. A straight forward and natural way to represent the curves ψ_i is by sampling at some fixed number of equidistant positions along the curve and approximating ψ_i by step functions. When doing this it is also natural to let q^{-1} be a step function. This is done by letting

$$f_0(t) = \begin{cases} 1, & 0 < t \le \tau, \\ 0, & \text{otherwise}, \end{cases}$$

for some $\tau > 0$ and setting $f_j(t) = f_0(t - j\tau)$. If the curves ψ_j are smooth it might be worth while to use a Fourier basis for ψ_j and q^{-1}, as this will increase the computational efficiency.

For the basis g_j of γ, there is no really obvious choice, aside from that it should fulfill the boundary conditions $g_j(0) = 0$ and $g_j(1) = 0$ and be rich enough for solving Step IV. In the experiments below we have chosen

$$g_j = \sin(\pi t)e^{-a(t-\tau_j)^2},$$

for some τ_j and $a > 0$. These g_j obviously fulfill the boundary conditions. Another choice is to use polynomials fulfilling the boundary conditions.

8.5. Experimental Validation

We will here give a demonstration of how the algorithm performs in an experimental setup. The experiment will be on simulated data. A simulation was made resulting in 6 images of a common 3D-curve as shown in Fig. 1. The 3D-curve $\mathcal{X}$ is illustrated in Fig. 2.

Initially, each image curve was parametrised by arclength, $(\psi_{i1}(s), \psi_{i2}(s), 1)$, $i = 0, \dots m - 1$. Note that the endpoints of these curves were assumed to be known, but that points with the same curve parameter s in different images are not necessarily in correspondence. Twenty iterations of the algorithm were performed. After each step, the proximity measure μ_1 was stored, and the reconstructed curve was compared with the true curve $\mathcal{X}$. Figure 2 shows the reconstructed

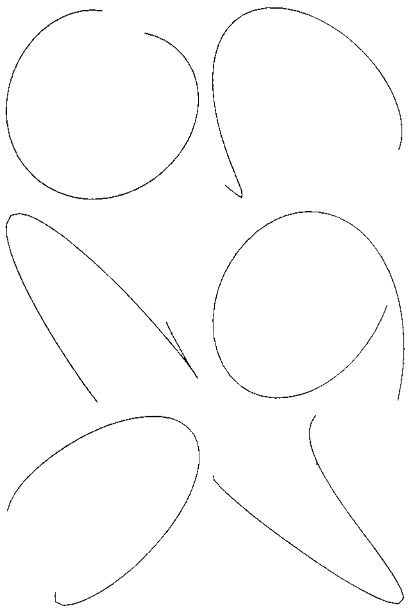

Figure 1. Six images of a curve.

curve after the first and after the 20th iteration. Notice the relatively good alignment already after the first iteration.

To show the necessity of the reparametrisation Step IV, computations were made with this step omitted. Figure 3 shows the proximity measure as a function of the number of iterations, with and without reparametrisation. Figure 4 shows the RMS residual as functions of the number of iterations with and without reparametrisation. The RMS is defined by

$$\left(\int_0^1 |\phi - \phi_{\text{rec}}|^2 \, dx\right)^{1/2}$$

where ϕ is the true curve and ϕ_{rec} is the reconstructed curves which has been chosen so as to minimize RMS. Recall that the reconstruction ϕ_{rec} is only given up to nonsingular projective transformations.

9. Algorithm for Closed Curves

It is straight forward to extend the algorithm described above to handle closed curves as well. It is then convenient to use the torus $\mathbb{R}/\mathbb{Z}$ as the domain of the functions. Let $\tau_t : L^2(\mathbb{R}/\mathbb{Z}) \to L^2(\mathbb{R}/\mathbb{Z})$ be a cyclic translation operator defined by $\tau_t \circ f(x) = f(x-t)$. All we have to do in order to extend the previous algorithm to closed curves is to divide Step IV into two substeps IV/a and IV/b according to:

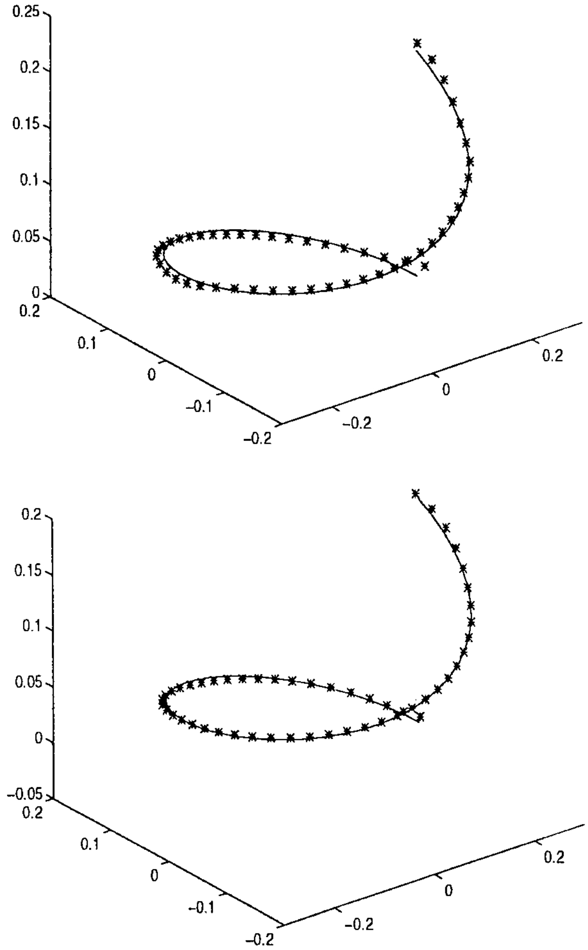

Figure 2. The reconstructed curve (∗) and the true curve (-) after the first and the 20th iteration.

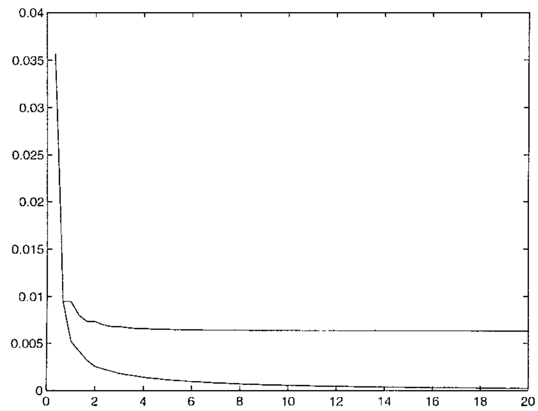

Figure 3. The proximity measure as a function of the number of iterations, with (lower) and without (upper) reparametrisation.

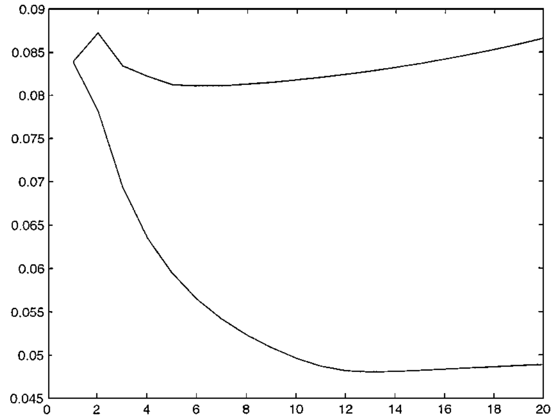

Figure 4. The RMS residual between reconstructed curve and true curve as a function of the number of iterations, with (lower) and without (upper) reparametrisation.

IV/a *Update translation*: Keeping $d(\phi)$ and $d(\psi_i)$ fixed, find cyclic translations t_i, such that $\tau_{t_i} \circ d(\psi_i)$ minimizes μ. Set $d(\psi_i) = \tau_{t_i} \circ d(\psi_i)$, $i = 0, \ldots, m-1$.

IV/b *Update parametrisation*: Keeping $d(\phi)$ and $d(\psi_i)$ fixed, find a continuous bijection $\gamma_i : I \to I$, such that $d(\psi_i) \circ \gamma_i$ minimizes μ. Set $d(\psi_i) := d(\psi_i) \circ \gamma_i$, $i = 0, \ldots, m-1$, and go to II.

9.1. Step IV/a

The only new thing compared to the previous case of non-closed curves is Step IV/a, which will be studied in more detail here. Let $\{\psi_{i1}, \psi_{i2}, \psi_{i3}\}$ be an orthonormal basis for $d(\psi_i)$ and $\{\phi_1, \phi_2, \phi_3, \phi_4\}$ an orthonormal basis for $d(\phi)$. The objective is to find t_i such that

$$f_i(t_i) = \sum_j \left\| \mathbf{P}_\phi \tau_{t_i} \psi_{i,j} \right\|^2$$

is minimized for each i. Since $\mathbf{P}_\phi = I - \mathbf{Q}_\phi$ and $\mathbf{Q}_\phi$ is a projection, we have

$$\begin{aligned} f_i &= \sum_j \left\| \tau_{t_i} \psi_{i,j} - \mathbf{Q}_\phi \tau_{t_i} \psi_{i,j} \right\|^2 \\ &= \sum_j \left(\|\psi_{i,j}\|^2 - \left\| \mathbf{Q}_\phi \tau_{t_i} \psi_{i,j} \right\|^2 \right) \\ &= 3 - \sum_{j,k} \langle \phi_k \mid \tau_{t_i} \psi_{i,j} \rangle^2. \end{aligned}$$

Let $\check{\psi}(x) = \psi(-x)$, then $\langle \phi_k \mid \tau_{t_i} \psi_{i,j} \rangle = \phi_k * \check{\psi}_{i,j}(t_i)$, where $*$ denotes cyclic convolution. If ψ_{ij} and ϕ_k are sampled, with the number of samples being 2^n for some integer n, then f_i can be computed very fast by using the fast Fourier transform. There are several ways of finding a minimum for f_i, like for example the Gauss-Newton method, when the minimum is known approximately, or some more robust method when less is known.

Step IV/a can also be used in Step I in order to get a better initial parametrisation of the curves.

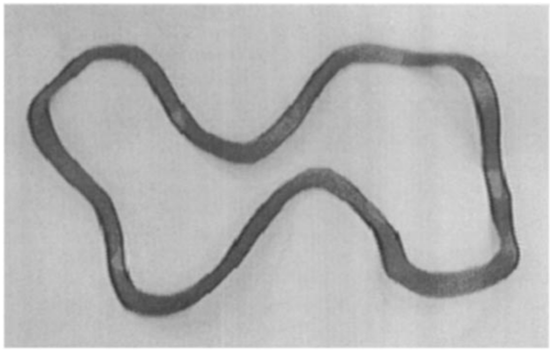

Figure 5. Image of the object.

9.2. Experimental Validation

We will here give a demonstration of how this algorithm performs in an experimental setup. The experiment will be on real data.

Four images of a closed 3D-curve, see Fig. 5, were taken in a laboratory environment.

The cameras were not calibrated and no two consecutive view points were close to each other. The inner curve was extracted from the images by an edge-detector and parametrised by arc-length in anti-clockwise direction. Note that this does not give any point correspondences as the parametrisations might start on different places on the curves. Furthermore, as was noted previously, arc-length is not a projective invariant parametrisation of a curve and so we would not have point correspondences even if the parametrisations started at a corresponding point in all images. Twenty iterations of the algorithm were performed. After each step, the proximity measure μ was stored. Step IV/a was also included as a routine in the initialization process (Step I), to improve the initial correspondences of the parametrisations ψ_j, $j = 0, \ldots, m-1$. For the experimental results, Fig. 6 shows the proximity measure as a function of the iteration number. Recall that the reconstruction $\hat{\phi}$ is only given up to nonsingular projective transformations, which makes it difficult visualize. A back projection of the reconstruction $\hat{\phi}$ to the images can always be computed though and Fig. 7 shows an image curve together with its back projected reconstruction.

A third experiment was performed on the reclining chair of one of the authors, see Fig. 8.

In this object there are no really well defined edges, since these are rounded off for comfort reasons. Furthermore, the boundary curve was extracted by hand, which brings in a high degree of uncertainty. Otherwise, the experiment was performed in the same manner as the preceding, and with the same number of images. Figure 9 shows the extracted curve (solid curve) of Fig. 8 (middle) together with the back projected reconstruction (dashed curve).

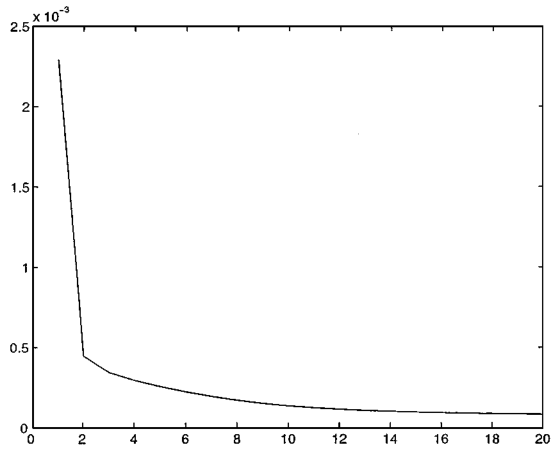

Figure 6. The proximity measure as a function of the number of iterations.

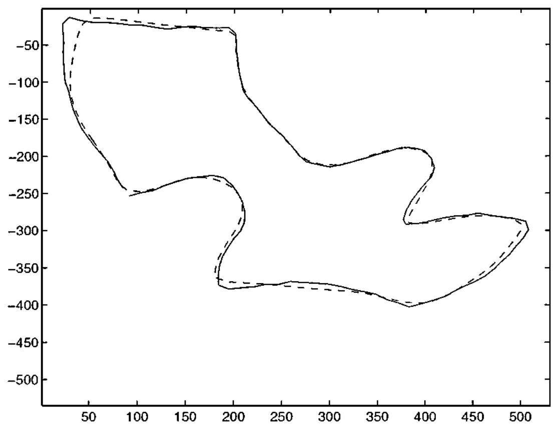

Figure 9. The extracted contour (solid line) together with the back projected contour (dashed line).

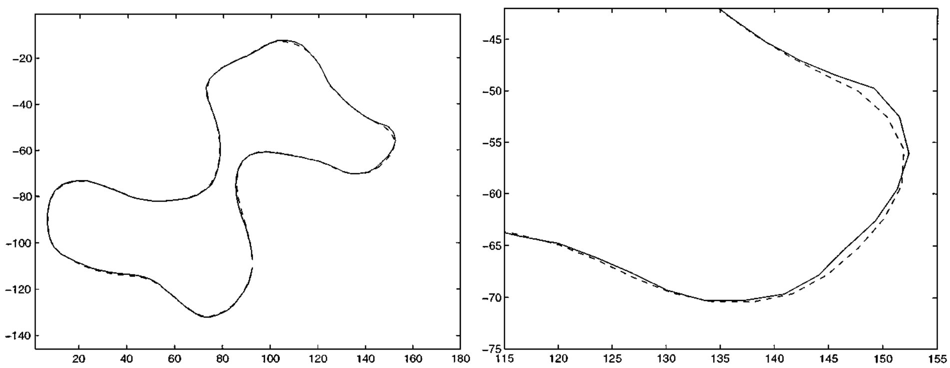

Figure 7. Backprojected image of the reconstruction with close up image.

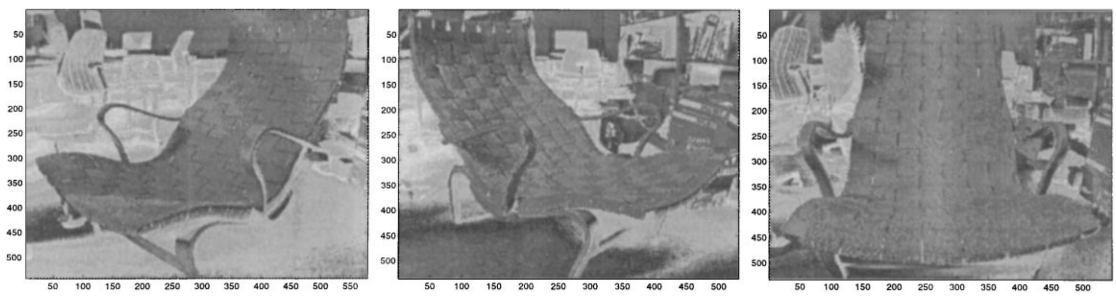

Figure 8. One of the authors reclining chair.

10. Convergence of Algorithm

In each step the algorithm performs a minimization of the same goal function (the proximity measure) over (i) the space curve (ii) the depths (iii) the parametrisations of the image curves. Thus the algorithms is really a descent method. As for all types of descent methods, e.g. the steepest descent method [6] one can under some condition prove global convergence to a stationary point, but not guarantee that the global optimum is found. The algorithms has proved to give reasonable estimates to the problem within the first iterations, but then the convergence is slow. This is due to the fact that each iteration involves three optimizations in turn. For faster convergence one might consider switching to another optimization method after a few iterations.

The problem of showing when the algorithm converges to a global minimum is hard and not solved yet. However, experiments indicate that, in practice, the algorithm, converges to the right solution, and it does so in only a few iterations.

11. Recognition of Closed Planar Boundary Curves

In this section we will develop an algorithm for recognition of closed planar boundary curves.

The experimental setup is a pinhole camera in a fixed position over a flat surface, whose normal is parallel to the optical axis. Images are taken of a number of planar objects, and their boundary curves are extracted. The affine shape of these curves are compared with shapes of a number of model objects, whose shapes are stored in a data base. The goal is to find the item that, in some sense, is closest to the measured item. Because of the experimental setup, the depths are constant functions. We therefore can model the camera as an affine camera, and we need only iterate Step IV/a and IV/b, with one image, in the algorithm above for curves. This means that for each item in the data base, we reparametrise the image curve ψ to minimize the proximity distance to $s(\psi)$. This is then repeated for all items in the data base, in order to find the best fit.

To be more precise let

$$\{s(\phi_i)\}_1^m$$

be a data base of shapes of some items, where $\phi_i : [0, 1] \to \mathcal{X}_i \subset \mathbb{R}^2$ are parametrisations of $\mathcal{X}_i$. As the camera is affine and shape is a complete affine invariant, $s(\phi_i)$ can be measured from the image curve $P(\phi_i)$, which we also call ϕ_i. We propose the following algorithm, with $d(\phi_i) = \text{linhull}\{\phi_{1i}, \phi_{2i}, \phi_{3i}\}$, and $d(\psi) = \text{linhull}\{\psi_1, \psi_2, \psi_3\}$ being the depth spaces for the curves of the data base and a boundary curve of an unknown item, respectively. As proximity measure for ϕ_i and ψ we chose

$$\mu_i = \left\| \mathbf{P}_{\phi_i} \mathbf{Q} \right\|_{HS}^2. \tag{14}$$

1. Initialize: Parametrise all ϕ_i and ψ by arc length and set $i = 0$.
2. Set $i := i + 1$ and $\widetilde{d(\psi)} = d(\psi)$.
3. Update translation: Keeping $d(\phi_i)$ and $\widetilde{d(\psi)}$ fixed, find a translation t, such that $\tau_t \circ \widetilde{d(\psi)}$ minimizes μ_i. Set $\widetilde{d(\psi)} := \tau_t \circ \widetilde{d(\psi)}$.
4. Update parametrisation: Keeping $d(\phi_i)$ and $\widetilde{d(\psi)}$ fixed, find a continuous bijection $\gamma : I \to I$, such that $\widetilde{d(\psi)} \circ \gamma$ minimizes μ_i. Set $\widetilde{d(\psi)} := \widetilde{d(\psi)} \circ \gamma$. If finished, store μ_i and go to 2 else go to 3.

These steps are just Step IV/a and IV/b in the algorithm of closed curves and need no further comments. A stop criterion in step 4 could for example be that a predefined number of iterations have been performed. When all μ_i have been computed, choose the i that minimizes μ_i.

11.1. Experimental Validation

For the experiment, we had a set of black and white drawings. Three of these were chosen to construct a data base. One image of each was taken. The respective boundary curves were extracted and the affine shapes were computed and stored. Images were also taken of all the drawings from varying camera positions, but with the optical axis parallel with the normal of the drawings. The boundary curves were extracted and the affine shape was computed from each boundary curve. In Table 1 the images of the drawings are listed in the column to the left. At the top, the images of the drawings in the data base are listed. The indices, 1 and 2, indicate that the drawings are different, while a and b indicate that the camera has been moved between the imaging instants.

Figure 10 shows the boundary curves of the drawings in the data base and Figs. 11 and 12 some of those that were not. The results can be seen in Table 1, where μ is the proximity measure (14). Notice that if the drawing corresponding to the smallest proximity measure μ_i is chosen, we will always get the right decision. Moreover, if a suitable threshold is set, the stranger drawings, such as bone and fish, will not misinterpreted to be an apple, a banana or a pear. Note also that the proximity measure is robust against small changes. The proximity measures within, for example, the apple group are very similar, when comparing with the proximity measures across groups. On the other hand, it does not seem possible to distinguish between items within a group, when noise is present.

Table 1. Proximity measure after four iterations.

Image	$\mu(\text{Apple}_{1a}, \cdot)$	$\mu(\text{Banana}_{1a}, \cdot)$	$\mu(\text{Pear}_{1a}, \cdot)$
Apple_{1a}	0.0000	0.0440	0.0151
Apple_{1b}	0.0001	0.0634	0.0180
Apple_{2a}	0.0003	0.0393	0.0093
Apple_{2b}	0.0002	0.0372	0.0079
Banana_{1a}	0.0188	0.0000	0.0065
Banana_{1b}	0.0203	0.0003	0.0452
Banana_{2a}	0.0104	0.0005	0.0138
Banana_{2b}	0.0199	0.0008	0.0353
Pear_{1a}	0.0084	0.0285	0.0000
Pear_{1b}	0.0083	0.0251	0.0001
Pear_{2a}	0.0164	0.0197	0.0012
Pear_{2b}	0.0101	0.0185	0.0005
Bone_a	0.0174	0.0184	0.0119
Bone_b	0.0319	0.0688	0.0217
Fish_a	0.0111	0.0446	0.0118
Fish_b	0.0117	0.0483	0.0261

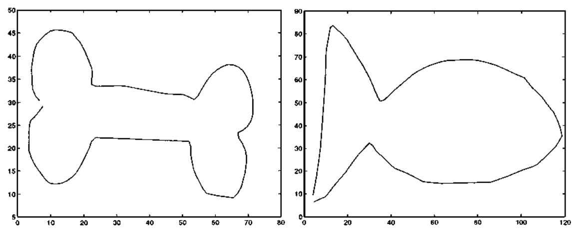

Figure 12. Boundary curves of bone_a and fish_a.

12. Conclusions

In this paper, affine shape has been extended from finite point sets to very general sets, so called shapeable sets. The only requirement of a set to be shapeable, is that it can be parametrised by a measurable map, according to some positive Radon measure. All theorems that hold for affine shape of finite point sets still hold in the new setting. Shape of some curves has been computed analytically, e.g. circles and lines. Furthermore, shape of finite point sets can be obtained as a special case from the extension of shape, by choosing the proper measure.

The extension makes it possible to, projectively, reconstruct three dimensional curves from a number of

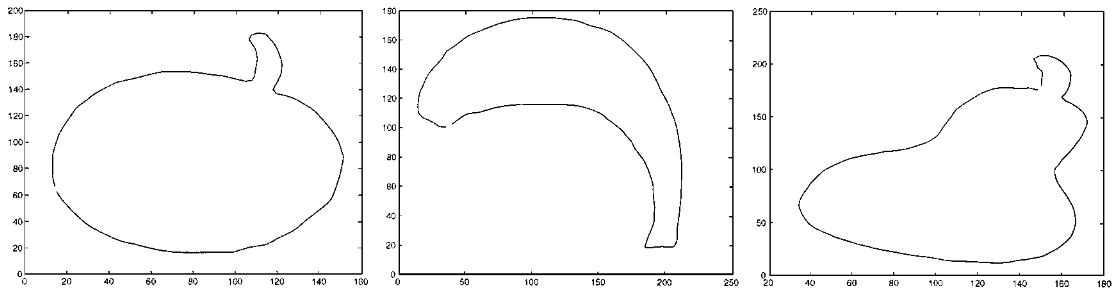

Figure 10. Boundary curves of the drawings in the data base, i.e. apple_{1a}, banana_{1a} and pear_{1a}.

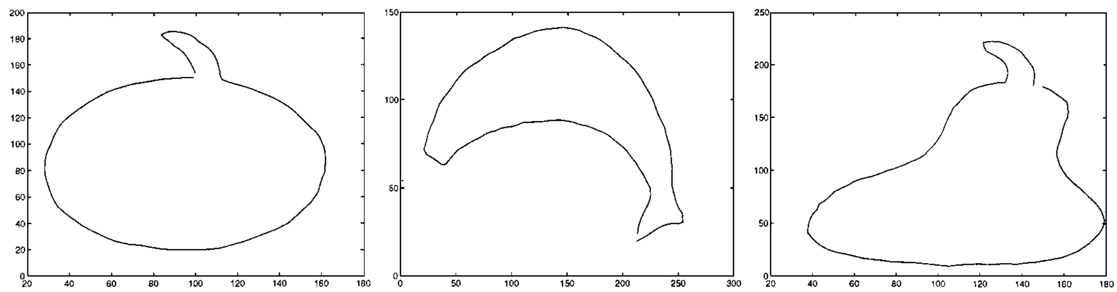

Figure 11. Boundary curves of apple_{2a}, banana_{2a} and pear_{2a}.

uncalibrated cameras. The algorithm does not rely on any derivatives of the image curves, but rather on integrals. This makes the algorithm very robust and insensitive to noise, which is indicated by the experiments. More experiments has to done here to be able to draw more exact conclusions about the impact of noise. This will be done in the future. A drawback, of the algorithm, is that the entire curve has to be visible in all images. Since occlusion is common in practice, it is important to be able to reconstruct even in this situation. Furthermore, the algorithm does not make any assumptions about the measurement errors and all images are treated on an equal footing. When nothing is known it seems best to have such a least committed algorithm, but when something can be said about the errors it would be advantageous to incorporate such knowledge into the treatment. All this will be the subject for further studies.

In the paper it is also shown that it is possible to use affine shape of curves to recognize planar objects, by extracting their boundary curve and comparing its shape with a data base. Further, another application of this technique is for recognition of hand written letters. Further studies on this important topic will be done in the future too.

Acknowledgments

The authors thank Gunnar Sparr and Anders Heyden for valuable discussions. This work has been done within the ESPRIT Reactive LTR 21914, CUMULI and the Swedish Research Council for Engineering Sciences (TFR), project 95-64-222.

References

1. K. Åström, R. Cipolla, and P.J. Giblin, "Generalised Epipolar Constraints," in *Proc. 4th European Conf. on Computer Vision*, Cambridge, England, 1996, pp. 97–108.
2. K. Åström, "Invariancy methods for points, curves and surfaces in computer vision," Lund University, Lund Institute of Technology, Department of Mathematics, 1996.
3. S. Carlsson, "Sufficient image structure for 3D motion and shape estimation," in *Proc. 3rd European Conf. on Computer Vision*, Stockholm, Sweden, 1993, pp. 83–91.
4. N. Dunford and J.T. Schwarz, "Linear operators," *Spectral theory*, Vol. 2, Interscience, 1963.
5. O.D. Faugeras and T. Papadopoulo, "A theory of the motion fields of curves," *Int. Journal of Computer Vision*, Vol. 10, pp. 125–156, 1993.
6. R. Fletcher, "Practical Methods of Optimization", John Wiley and sons, 1987.
7. T. Papadopoulo and O. Faugeras, "Computing structure and motion of general 3D curves from monocular sequences of perspective images," in *Proc. 4th European Conf. on Computer Vision*, Cambridge, England, 1996, pp. 696–708.
8. J. Porrill and S.B. Pollard, "Curve matching and stereo calibration," *Image and Vision Computing*, Vol. 9, pp. 45–50, 1991.
9. G. Sparr, "A common framework for kinetic depth, reconstruction and motion for deformable objects," in *Proc. 4th European Conf. on Computer Vision*, Cambridge, England, 1994, pp. 471–482.
10. G. Sparr, "Structure and motion from kinetic depth," in *Proc. of the Sophus Lie International Workshop on Computer Vision and Applied Geometry*, Nordfjordeid, Norway, 1995, to appear.
11. G. Sparr, "Simultaneous reconstruction of scene structure and camera locations from uncalibrated image sequences," in *Proc. 13th International Conference on Pattern Recognition*, Vienna, Austria, 1996, pp. 328–333.

Rikard Berthilsson received his M.Sc. degree in Engineering Physics in 1994 and his Licentiate degree in 1997. He is currently pursuing a Ph.D. in mathematics at the University of Lund, Sweden. His interests in research include reconstruction problems in computer vision, statistical theory of geometric invariants, correlation techniques in image analysis, and handwriting recognition. The latter has resulted in a filed patent application in 1998.

Kalle Åström received his B.Sc. in Mathematics in 1990, M.Sc. degree in Engineering Physics in 1991 and Ph.D. in Mathematics in 1996 from Lund University, Sweden. His thesis was awarded Best Nordic Ph.D. Thesis in pattern recognition and image analysis 1995–1996 at the Scandinavian Conference in Image Analysis, 1997. He has been a postdoctoral research fellow and subsequently research associate at the Department of Mathematics, Lund University. He is the president of the Swedish Society for Automated Image Analysis. His teaching include undergraduate and graduate courses in mathematics, image processing and computer vision. His current research interests include stochastic analysis of low level vision, computer vision for autonomous guided vehicles, geometry and algebra of multiple views of points, curves and surfaces.