

An affine invariant deformable shape representation for general curves

A. Ericsson and K. Åström
Mathematics, Center for Mathematical Sciences,
Institute of Technology, Lund University,
Lund, Sweden

Abstract

Automatic construction of Shape Models from examples has been the focus of intense research during the last couple of years. These methods have proved to be useful for shape segmentation, tracking and shape understanding. In this paper novel theory to automate shape modelling is described. The theory is intrinsically defined for curves although curves are infinite dimensional objects. The theory is independent of parameterisation and affine transformations. We suggest a method for implementing the ideas and compare it to minimising the Description Length of the model (MDL). It turns out that the accuracy of the two methods is comparable. Both the MDL and our approach can get stuck at local minima. Our algorithm is less computational expensive and relatively good solutions are obtained after a few iterations. The MDL is, however, better suited at fine-tuning the parameters given good initial estimates to the problem. It is shown that a combination of the two methods outperforms either on its own.

1. Introduction

Statistical models of shape [7, 10] have turned out to be a very effective tool in image segmentation and image interpretation. Such models are particularly effective in modelling objects with limited variability, such as medical organs.

The basic idea behind statistical models of shape is that from a given training set of known shapes be able to describe new formerly unseen shapes, which still are representative. The shape is traditionally described using landmarks on the shape boundary. A major drawback of this approach is that during training a dense correspondence between the boundaries of the shapes must be known. In practice this has been done by hand. A process that commonly is both time consuming and error prone.

The goal of this paper is to formulate a new model for shape variation of parameterised curves that is invariant to affine transformations, curve parameterisation and also invariant in the ordering of the curves. The model is a generalisation of earlier work on affine shape [17] and its ex-

tensions to curves [4]. The work is inspired by the work on active shape [6] and work on 3D-reconstruction of a deformable object [1].

This paper is organised as follows. In Section 2 the necessary background on affine shape and shape variation is given. In Section 3, a new model for shape variation is presented and a novel algorithm for learning and using shape variation are explained in Section 4. Experiments with this new model are given in Section 5. The model is evaluated and compared to models built using the minimum description length.

1.1. Previous Work

Consider m parameterised curves $c_i : [0, 1] \mapsto \mathbb{R}^2$. The problem of finding a dense correspondence among the shape boundaries is equivalent of reparameterising the shape boundary curves (to obtain $\psi_i = c_i \circ \gamma_i$), so that $\psi_i(t)$ is the point that corresponds to $\psi_j(t)$ for all $(i, j = 0, \dots, m - 1)$ and $t \in [0, 1]$. Here $\gamma_i : [0, 1] \mapsto [0, 1]$ represents the reparameterisation of curve i . The same formulation can be used for, e.g. closed curves by changing the interval $[0, 1]$ to the circle \mathbb{S}^1 .

There have been many suggestions on how to automate the process of building shape models or more precise finding a dense correspondence among a set of shapes [2, 3, 13, 14, 20, 22]. Many approaches that correspond to human intuition attempt to locate landmarks on curves using shape features. [3, 13, 20]. The located features have been used to establish point correspondences, using equal distance interpolation in between. Local geometric properties, such as geodesics, have been tested for surfaces [22]. Different ways of parameterising the training shape boundaries have been proposed [2, 14]. The above cited are not clearly optimal in any sense. Many have stated the correspondence problem as an optimisation problem [5, 9, 12, 15, 16].

In [16] a measure is proposed and dynamic programming is applied to find the reparameterisation functions. A problem with this method is that it can only handle contours, for which the shape not changes too much, correctly.

In [8, 9, 15] the description length of the model is min-

imised. A very successful algorithm but a problem with this method is that the objective function is not stated explicitly. Besides the optimisation process is often tedious and can easily get stuck at a local minima.

2 Preliminaries

The standard shape model used in e.g. [6] describes the variation of a set of points as linear combinations of a set of basis vectors

$$\psi = \bar{\psi} + \Phi v.$$

The goal of this paper is to derive a shape theory that is intrinsically defined for curves. Similar to the point theory it is reasonable to use a linear model, i.e.

$$\psi_i(t) = \phi_{n+1}(t) + \sum_{k=1}^n v_k \phi_k(t).$$

This implies that the curve ψ_i is an affine projection of a n dimensional curve ϕ_k . Reconstructing ϕ from its projections ψ_i is identical to the curve reconstruction problem in computer vision.

Another desirable property of shape models is that they should be invariant to the position and orientation of the curve in the image.

In this Section it is assumed that m curves $\{c_i\}_{i=0}^{m-1}$ are given. Our goal here is to formulate a shape description, which is defined for curves, is independent of parameterisation and ordering of the curves and also invariant to affine transformations.

We want our curves to be linear combinations of some basis $\phi_i (i = 1, \dots, n)$. This is the model ϕ we are looking for.

We start with a derivation of affine shape for points. This is then generalised to curves. A new shape description model is defined using this theory and an algorithm for reconstructing the model ϕ from example curves $\{\psi_i\}_{i=0}^{m-1}$ are given.

2.1 Affine Shape of Point Configurations

In this Section affine shape for point configurations is described. For more details see [4].

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

$$\mathcal{X} = (p_1, p_2, \dots, p_m) \in \mathbb{R}^{m \cdot n}$$

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, $a : \mathbb{R}^n \rightarrow \mathbb{R}^n$, is meant a map of the form,

$$a(x) = Mx + t, \quad (1)$$

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

$$a \circ \mathcal{X} = a \circ (p_1, p_2, \dots, p_m) = (a(p_1), a(p_2), \dots, a(p_m)). \quad (2)$$

Now let A be the group of nonsingular affine transformations $\mathcal{C}_m^n \rightarrow \mathcal{C}_m^n$. 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

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

be the natural projection. Then, $s(\mathcal{X}) \in \mathcal{C}_m^n/A$ is called the affine **shape** of $\mathcal{X} \in \mathcal{C}_m^n$.

It can be shown, [18], that

$$s(\mathcal{X}) = \left\{ \xi \mid \sum_i \xi_i p_i = 0, \sum_i \xi_i = 0, \mathcal{X} = (p_1, p_2, \dots, p_m) \right\}, \quad (3)$$

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 $\mathcal{Y} = a(\mathcal{X})$. We will now extend shape to more general sets than finite point configurations, as for example curves and surfaces in \mathbb{R}^3 .

2.2 Extension of shape

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

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

be defined by $\phi_{\mathcal{X}}(i) = p_i, i \in I$. The definition (3) can then be rewritten as,

$$s(\mathcal{X}) = \left\{ f : I \rightarrow \mathbb{R} \mid \sum_{x \in I} f(x) \phi_{\mathcal{X}}(x) = 0, \sum_{x \in I} f(x) = 0 \right\}.$$

An extension to more general sets is done by replacing the sums with integrals.

When dealing with finite point configurations, as in Section 2.1, \mathcal{X} denotes an ordered set of points. The ordering can be considered as a parameterisation of the set. To deal with more general point configurations, it is convenient to let \mathcal{X} denote point sets, without ordering or parameterisation. This is done below.

Definition 2.1. Let $\mathcal{X} \subseteq \mathbb{R}^n$. By a **parameterisation** of \mathcal{X} is meant a surjective map

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

where $I = [0, 1]$ if we consider open curves (and $I = \mathbb{S}^1$ for closed curves). By the affine shape is meant the following linear space of functions

$$s(\phi_{\mathcal{X}}) = \left\{ f \mid \begin{array}{l} \langle f | \phi_{\mathcal{X}} \rangle = 0, \langle f | 1 \rangle = 0, \\ f \in L^2 \end{array} \right\}.$$

Here we consider square integrable functions ($f, g \in L^2$) with the scalar product

$$\langle f | g \rangle = \int_I f(x)g(x)dx.$$

Often the subscript \mathcal{X} of $\phi_{\mathcal{X}}$ is dropped, when it is clear from the context which configuration is meant.

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

Using terminology borrowed from the finite dimensional case, cf. [19], we define the depth space.

Definition 2.2. Let $\phi_{\mathcal{X}} = (\phi_1, \dots, \phi_n, \phi_{n+1})$, with $\phi_{n+1} \equiv 1$, be a parameterisation 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^2 ,

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

An important property of these spaces are given by the following theorem

Theorem 2.1 (Affine shape theorem). *Let $\phi_{\mathcal{X}} : I \rightarrow \mathcal{X} \subseteq \mathbb{R}^n$ and $\phi'_{\mathcal{X}'} : I \rightarrow \mathcal{X}' \subseteq \mathbb{R}^n$ be parameterisations as in Definition 2.1. If $a : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is an affine transformation, then*

$$\begin{aligned} a : \phi_{\mathcal{X}} \rightarrow \phi'_{\mathcal{X}'} &\iff s(\phi_{\mathcal{X}}) \subseteq s(\phi'_{\mathcal{X}'}) \\ &\iff \\ d(\phi_{\mathcal{X}}) &\supseteq d(\phi'_{\mathcal{X}'}). \end{aligned}$$

Corollary 2.1. *Under the assumptions of Theorem 2.1,*

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

with a nonsingular affine transformation

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

Proof. Apply Theorem 2.1 (i) to a and a^{-1} . ■

For the finite dimensional versions of these, see [19].

3 Formulating Affine Invariant Active Shape

3.1 Extensions to shape variation

For the moment assume that the curves are in fact affine projections of a higher dimensional model curve \mathcal{X} and that point correspondences are known. Denote the parameterisations by $\phi : I \rightarrow \mathcal{X}$ and $\psi_i : I \rightarrow \mathcal{Y}_i$, so that for some affine projection a_i

$$\psi_i(x) = a_i \circ \phi(x), \quad i = 0, 1, \dots, m-1$$

holds for all $x \in I$. Then, by Theorem 2.1

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

Together this implies

$$s(\phi) \subseteq \bigcap_0^{m-1} s(\psi_i).$$

Given enough different projections there is in general an equality in the above equation.

Let \mathbf{P}_{ϕ} , \mathbf{P}_i , \mathbf{Q}_{ϕ} and \mathbf{Q}_i be the orthogonal projections from $L^2(I)$ onto $s(\phi)$, $s(\psi_i)$, $d(\phi)$ and $d(\psi_i)$, respectively.

These projection operators can be explicitly written using orthonormal bases. For example, using extended coordinates, $\phi = (\phi_1, \dots, \phi_n)$, where $\phi_n \equiv 1$, let $\{\tilde{\phi}_1, \dots, \tilde{\phi}_n\}$ be an orthonormal basis for $\text{linhull}(\phi_1, \dots, \phi_n)$. Then

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

and

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

The operator \mathbf{Q}_i is the orthogonal projection on the depth space $d(\psi_i)$, which is contained in the depth space $d(\phi)$. The operator \mathbf{P}_{ϕ} annihilates every function in $d(\phi)$ and thus also every function in $d(\psi_i)$. Thus the operator $\mathbf{P}_{\phi}\mathbf{Q}_i$ is the zero operator for every i .

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

Any criterion that is based on the projection operator 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, is 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

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

Here HS stands for the Hilbert-Schmidt norm, see [11], 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 $d(\psi_i)$ (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 $d(\psi_i)$, $i = 0, \dots, m-1$, it follows that

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

Other proximity measures are derived in [4].

4 Algorithm for non-closed curves

In the following, let $I = [0, 1]$ and $\{\mathcal{Y}_i\}_0^{m-1}$ be a sequence of curves. We wish to find an n -dimensional curve \mathcal{X} that represents the shape and shape variations of these curves. Furthermore, let all the parameterisations $\phi : I \rightarrow \mathcal{X}$ and $\psi_i : I \rightarrow \mathcal{Y}_i$, $i = 0, \dots, m-1$, be expressed in extended coordinates.

We propose the following algorithm, which is based on repeatedly finding $s(\phi)$ and the parameterisations 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 , $i = 0, \dots, m-1$. To obtain an n -dimensional reconstruction, that is finding an n -dimensional shape model, the problem is to find the parameterisations ψ_i

- I Initialisation:** Choose one parameterisation ψ_i in each image curve, for example by using the image based arc-length. Set $d(\psi_i) = \text{linhull}(\psi_i)$, $i = 0, \dots, m-1$.
- II Update $d(\phi)$:** Keeping $d(\psi_i)$ fixed for all i , find \mathbf{P}_ϕ that minimises μ .
- III Update parameterisation:** Keeping $d(\phi)$ and $d(\psi_i)$ fixed, find a continuous bijection $\gamma_i : I \rightarrow I$, such that $d(\psi_i \circ \gamma_i)$ minimises μ . Set $d(\psi_i) := d(\psi_i \circ \gamma_i)$, $i = 0, \dots, m-1$ and go to II

It is difficult to minimise μ with respect to all parameters simultaneously. It is, however, reasonably fast to solve steps II and III 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 . Notice however that each step lowers the same error criteria μ .

4.0.1 Step I. Initialisation

Let each image curve

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

be parameterised using scaled image arc-length $t \in I$, so that $\psi_i(0)$ and $\psi_i(1)$ are the endpoints, and such that $(\psi'_{i1})^2 + (\psi'_{i2})^2$, $i = 0, \dots, m-1$, is constant. Initially, let $d(\psi_i) = \text{linhull}(\psi_{i,1}, \psi_{i,2}, \psi_{i,3})$.

4.0.2 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 3-dimensional linear space $d(\psi_i)$. The n -dimensional linear space $d(\phi)$, corresponding to the n -dimensional-curve to be reconstructed, is then the linear span of all basis functions $\psi_{i,k}$, $i = 0, \dots, m-1$, $k = 1, 2, 3$, i.e.

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

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

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

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

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

Compute a singular value 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 n . In the case of measured data, the matrix which is closest in Frobenius norm to a matrix of rank n , is $\hat{M} = US_n V^T$, where S_n is obtained by setting all but the n largest diagonal elements in S to zero. An orthonormal basis for $d(\phi)$ can be shown to be

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

This $d(\phi)$ solves the optimisation problem (4).

4.0.3 Step III. Reparameterisation of the image curves

Let $\{\psi_{i1}, \psi_{i2}, \psi_{i3}\}$ be an orthonormal basis for $d(\psi_i)$, and let $\{\phi_1, \dots, \phi_n\}$ be an orthonormal basis for the n -dimensional linear space $d(\phi)$, corresponding to the curve to be reconstructed. We want to find a reparameterisation 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 minimised over some set of reparameterisations. Again, we drop the index i for convenience.

Parameterise γ by using a finite basis g_j , according to

$$I \times \mathbb{R}^n \ni (t, x) \rightarrow \gamma(t, x) = t + \sum_j x_j g_j(t), \quad (5)$$

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

$$|x|^2 < \min_{t \in I} \frac{1}{\sum_j |g'_j(t)|^2}. \quad (6)$$

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 minimisation problem

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

is obtained from the normal equations

$$-Ax = b, \quad (7)$$

where

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

and

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

If the solution of (7) gives an x not fulfilling (6), 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 (7) for x , (5) defines a reparameterisation 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 $\psi'_{ij} \in L^2(I)$.

5 Experimental Validation

In the experimental validation we validate our algorithm on three data sets. One data set of 17 curves of the letter g, one data set of 23 contours of a hand and one data set of 32 contours of femurs. The curves of the letter g are data sampled by a device for handwriting recognition. The hand outlines were segmented out semi-automatically from a video stream of one hand. The hand were filmed on a dark background; making it easy to segment out using Dijkstras algorithm. The femurs are taken from X-rays in the supine projection. The validation will be to show the effectiveness of the proposed algorithm in locating a dense correspondence among a set of contours. We use the algorithm in automatic model building and compare it to using the MDL-approach [9]. We have been using Thodberg's efficient implementation of MDL [21]. MATLAB source code and test data are available from hht@imm.dtu.dk.

First an experiment on the hand contours. We sample 500 landmarks on each hand contour and build an Active Shape Model of the (by the proposed algorithm) reparameterised curves. The results are compared to a hand-built model and a model built from 500 equally spaced landmarks. On the hand built model 11 landmarks were placed by hand and in between further points were equally spaced - making a total of 500 landmarks.

To measure the quality of the model, the compactness and the ability of the model to generalise is measured. For a model to be compact the total variation V_T and the variance of each mode of variation should be low, see Table 1.

In Figure 2 we plot the results of running leave-one-out reconstructions on each model of the hand outlines. The model is built with all but one example and then fitted to the unseen example. The plot shows the mean squared approximation error against number of modes used. This measures

Modes	Model		
	AIAS	Hand Built	Arc-length
1	0.40	0.67	2.03
2	0.17	0.35	0.65
3	0.09	0.21	0.31
4	0.04	0.07	0.22
5	0.04	0.06	0.10
6	0.02	0.02	0.08
7	0.01	0.01	0.05
V_T	0.79	1.43	3.59

Table 1: Table of the total variation and variation of each mode of the different models.

the ability of the model to represent unseen shape instances of the object. From Table 1 we see that the AS-model built from the AIAS-model is the most compact and out of Figure 2 we get that it also generalises the best.

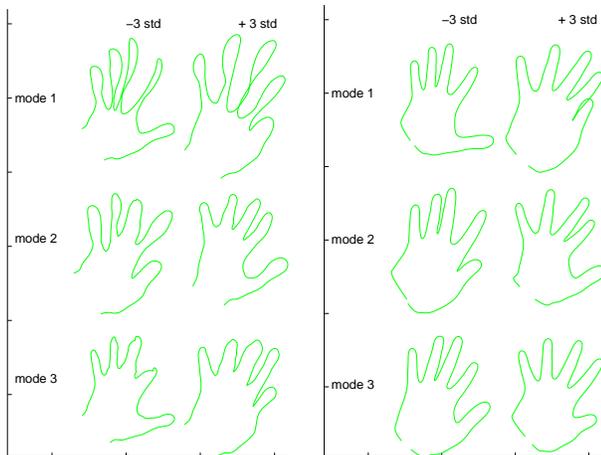


Figure 1: To the left the ± 3 std of the three first modes of variation of the arc-length parameterised Active Shape model and to the right the three first modes of variation of the AIAS based Active Shape model.

Our next experiment is to compare the results of building shape models automatically using the proposed algorithm versus the MDL-approach. For the hand contours we now use 128 landmarks for both methods. For the MDL-method nine nodes, for reparameterising the curves, are set on each contour. In Figure 5 it can be seen that the proposed algorithm decreases the description length (DL) very fast. In MATLAB on a 1.4 GHz computer we run 72 iterations on both methods. This takes 34 seconds for our method and 100 seconds for the MDL-method. Our method is less computational expensive and relatively good solutions are obtained after just a few iterations. In the first iteration the DL goes down from 80 to 64 for the hand and from 58.4 to

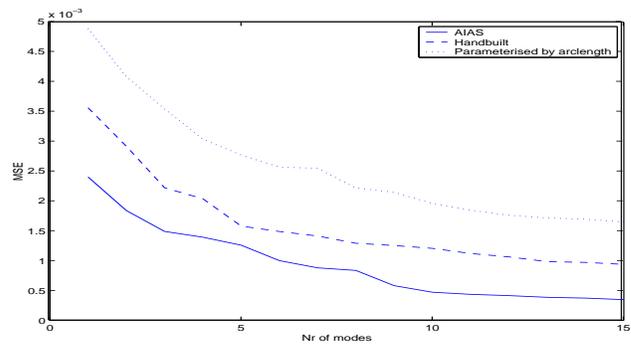


Figure 2: A plot of how the mean square error depends on number of modes.

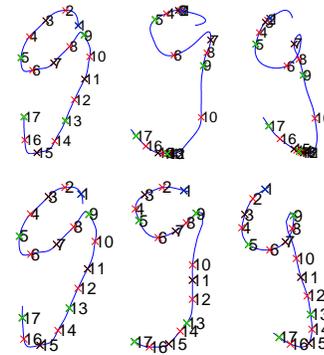


Figure 3: The parameterisations of the letter g. Running the MDL alone the reparameterisations tend to collapse into a few points (above). The DL goes down to 45.7 compared to 51.0 for the combined approach. The combined approach (below) does, however, not collapse.

47.8 for the femurs. It takes 5 and 8 iterations respectively before the MDL-approach has passed this value.

One must also keep in mind that our method minimises the proximity measure and not the DL. The description length of the model is just evaluated after each iteration of our algorithm for comparison.

In Figure 6 we compare the results of running leave-one-out reconstructions on both the MDL method and the proposed method. Approximately equally good models; our method in one third of the time.

Both the MDL and our approach can get stuck at local minima. Our algorithm is less computational expensive and relatively good solutions are obtained after a few iterations. The MDL is, however, better suited at fine-tuning the parameters given good initial estimates to the problem. The MDL method tends sometimes to collapse the shapes into one point (since the DL then will be zero). Given good initial parameterisations this can be prevented, see Figure 3.

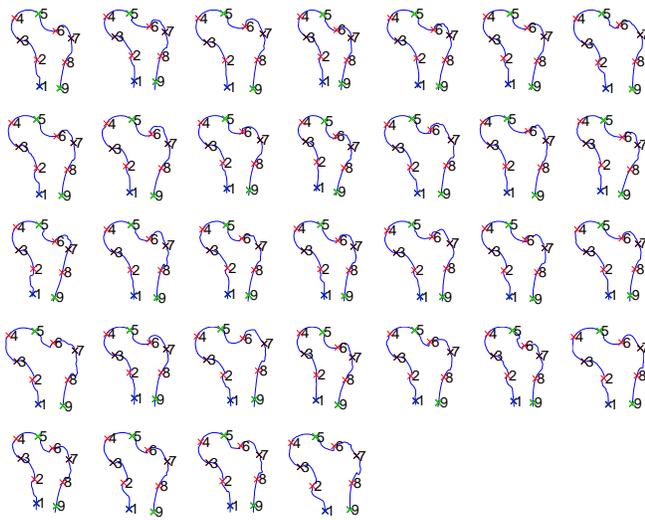


Figure 4: All 32 examples of the femurs after the combined search.

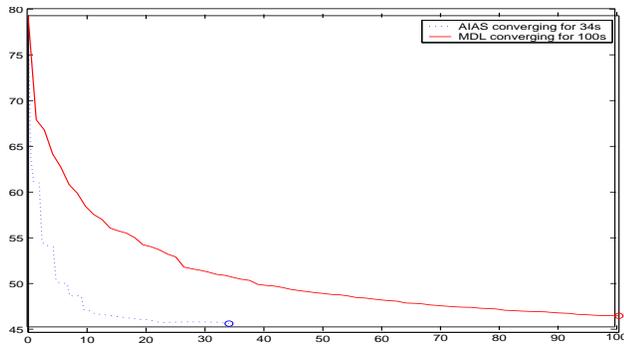


Figure 5: Comparison of how the DL decreases running the proposed algorithm for 34 seconds or the MDL for 100 seconds.

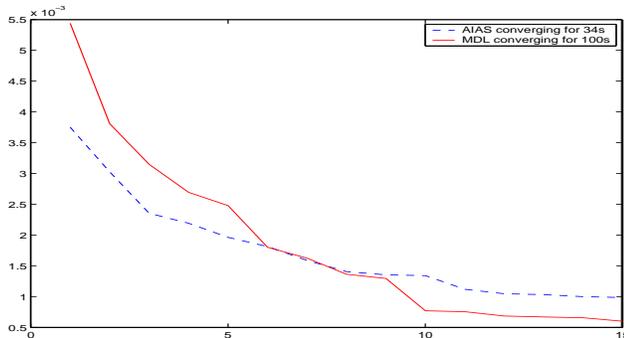


Figure 6: A comparison of how the mean square error depends on the number of modes after the proposed algorithm has been run for 34 seconds or the MDL for 100 seconds on the hand contours.

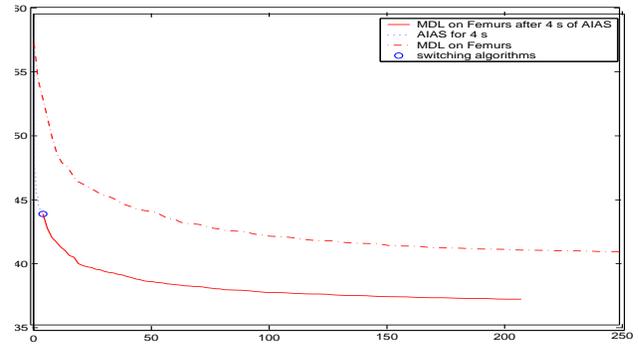


Figure 7: Comparison of how the DL decreases running the combined algorithm for 208 seconds or the MDL for 250 seconds. Using only MDL the DL goes down to 41.2 after 250 seconds. The combined passes this already after 15 seconds.

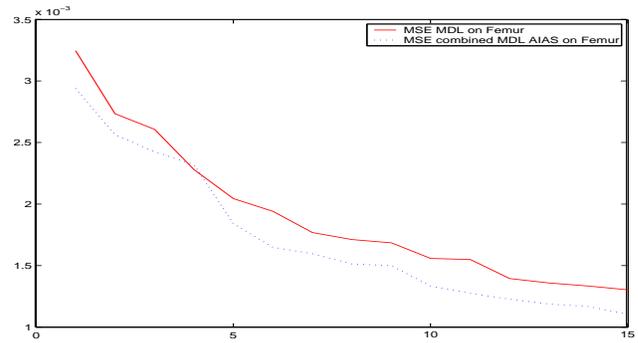


Figure 8: A comparison of how the mean square error depends on the number of modes after the combined have been run 208 seconds and the MDL for 250 seconds.

Combining the two algorithms give a more robust algorithm. First the proposed algorithm is run on the set of femurs for 4 seconds (11 iterations) and then the MDL is run for 204 seconds. In Figure 7 the DL is plotted for the proposed optimisation scheme and compared to only running MDL for 250 seconds on the femurs. In Figure 8 the MSE is plotted of the two femur models.

6. Summary and Conclusions

In this paper a novel method to model shape variation is proposed. Affine Shape was formerly a theory developed to construct 3D models of curves and surfaces. Here we extend the theory of Affine Shape. Combining the main ideas of Affine Shape and Active Shape a new approach, called Affine Invariant Active Shape is obtained. It is invariant of affine transformation, invariant to curve parameterisation and ordering of the curves, and also locates a dense corre-

spondence between the curves.

It also turns out that there is a relation between constructing 3D models and models of shape variation.

The objective function to be optimised is differentiable and can be written explicitly. This means that each step the objective function is decreased. In just a few iterations it converges and the algorithm is quite fast.

The idea of using proximity measures that are invariant under choice of coordinate systems is very appealing.

The main contributions of this paper are: 1 novel theory to explain shape variation, 2 show a connection between 3D-reconstruction and shape modelling, 3 algorithm that in just a few iterations locates a parameterisation, 4 compared to MDL our approach can faster reach the same results as MDL.

In our next study we will extend our work to closed curves, we will test our theory on surfaces and we will try to improve our models by using algorithms that have been successful in 3D-reconstruction.

Acknowledgments

First of all the authors would like to thank Hans Henrik Thodberg for his implementation of MDL [21], which is free on the web. Pronosco is acknowledged for providing the contours of femurs. We also like to thank Johan Karlsson for the curves of his hand. This work has been financed by the Swedish Research Council (VR), project TFR 2000-221-606. The authors thank Gunnar Sparr and Rikard Berthilsson for valuable discussions.

References

- [1] H. Aanaes and F. Kahl. Estimation of deformable structure and motion. In *Workshop on Vision and Modelling of Dynamic Scenes, ECCV'02*, 2002.
- [2] A. Baumberg and Hogg D. Learning flexible models from image sequences. In *Proc. European Conf. on Computer Vision, ECCV'94*, pages 299–308, 1994.
- [3] A. Benayoun, Ayache N., and Cohen I. Adaptive meshes and nonrigid motion computation. In *Proc. International Conference on Pattern Recognition, ICPR'94*, pages 730–732, 1994.
- [4] K. Berthilsson, R. Astrom. Extensions of affine shape. *Journal of Mathematical Imaging and Vision*, 2(11):119–136, 1999.
- [5] F.L. Bookstein. Landmark methods for forms without landmarks: Morphometrics of group differences in outline shape. *Medical Image Analysis*, 3:225–243, 1999.
- [6] T.F. Cootes, A. Hill, C.J Taylor, and Haslam J. The use of active shape models for locating structure in medical images. *IEEE Trans. medical imaging*, 12(6):355–365, 1994.
- [7] T.F Cootes and C.J. Taylor. *Statistical Models of Appearance for Computer Vision*. University of Manchester, 2001.
- [8] Rhodri H. Davies, Tim F. Cootes, and Chris J. Taylor. A minimum description length approach to statistical shape modeling. In *Information Processing in Medical Imaging*, 2001.
- [9] Rhodri H. Davies, Carole J. Twining, Tim F. Cootes, John C. Waterton, and Chris J. Taylor. A minimum description length approach to statistical shape modeling. *IEEE Trans. medical imaging*, 21(5):525–537, 2002.
- [10] Ian L. Dryden and Kanti V. Mardia. *Statistical Shape Analysis*. John Wiley & Sons, 1999.
- [11] N Dunford and J. T. Schwarz. *Linear operators. Spectral theory*, 2. Interscience, 1963.
- [12] A. Hill and C. J. Taylor. Automatic landmark generation for point distribution models. In *Proc. British Machine Vision Conference*, pages 429–438, 1994.
- [13] C. Kambhamettu and D.B. Goldgof. Points correspondences recovery in non-rigid motion. In *Proc. Conf. Computer Vision and Pattern Recognition, CVPR'92*, pages 222–237, 1992.
- [14] A. Kelemen, G. Szekely, and Gerig G. Elastic model-based segmentation of 3d neuroradiological data sets. *IEEE Trans. medical imaging*, 18(10):828–839, 1999.
- [15] A.C.W. Kotcheff and C. J. Taylor. Automatic construction of eigenshape models by direct optimization. *Medical Image Analysis*, 2:303–314, 1998.
- [16] T. Sebastian, P. Klein, and B. Kimia. Constructing 2d curve atlases. In *IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, pages 70–77, 2000.
- [17] G. Sparr. Projective invariants for affine shapes of point configurations. In *Proceedings ESPRIT-DARPA Invariants Workshop*, 1991.
- [18] G. Sparr. A common framework for kinetic depth, reconstruction and motion for deformable objects. In *Proc. European Conf. on Computer Vision, ECCV'94*, volume 801, pages 471–482. Springer-Verlag, 1994.
- [19] G. Sparr. Structure and motion from kinetic depth. Technical Report ISRN LUTFD2/TFMA–95/7016–SE, Dept of Mathematics, Lund Institute of Technology, 1995. in Proc of the Sophus Lie International Workshop on Computer Vision and Applied Geometry.
- [20] H.D. Tagare. Shape-based nonrigid correspondence with application to heart motion analysis. *IEEE Trans. medical imaging*, 18:570–579, 1999.
- [21] H. H. Thodberg. Minimum description length shape and appearance models. Technical Report IMM TECHNICAL REPORT 2003-01, Technical University of Denmark, 2003.
- [22] Y. Wang, B.S. Peterson, and L.H Staib. Shape-based 3d surface correspondence using geodesics and local geometry. In *Proc. Conf. Computer Vision and Pattern Recognition, CVPR'00*, pages 644–651, 2000.