Affine and Projective Normalization of Planar Curves and Regions. *

Kalle Åström

Dept. of Mathematics, Lund Institute of Technology, Box 118, S-221 00 Lund, Sweden

Abstract. Recent research has showed that invariant indexing can speed up the recognition process in computer vision. Extraction of invariant features can be done by choosing first a canonical reference frame, and then features in this reference frame. These automatically become invariants. In this paper methods are given where a global and stable approach is used both in choosing the canonical reference frame and in the extraction of invariants for affine and projective transformations. These invariants can be used to recognize fairly general feature configurations, and they can be used in a semi-local way to recognize occluded objects. In the affine case a unique distinguished reference frame can be chosen in a continuous way with respect to the Hausdorff metric. In the projective case this is impossible, since arbitrarily close to every pair of closed curves exists a projectively equivalent pair. However, by sacrificing uniqueness it is possible to construct useful projective normalization schemes. In the case where more information is at hand, e.g. when the curve has a concavity, the reference frame can still be chosen in a unique way.

Keywords: Recognition, planar non-algebraic curves, projective invariants, affine invariants.

1 Introduction

The pinhole camera model is a fairly adequate model of a real camera. Points in three dimensions are projected onto a plane. Using this model it is straightforward to predict the image of a collection of objects in specified positions. The inverse problem to identify and to determine the three dimensional positions of possible objects from an image, is however much more difficult. Traditionally recognition has been done by matching each model in a model data base with parts of the image. Recently, model based recognition using viewpoint invariant features of planar curves and point configurations has attracted much attention, cf. [MZ1]. Invariant features are computed directly from the image and used as index functions in a model data base. This gives algorithms which are significantly faster than the traditional methods. These techniques cannot, however, be used to recognise general curves or point features in three dimensions by means of one single image. Additional information, e.g. that the object is planar, is needed. For point configurations the reason is that only trivial invariants exist in the general case, as is shown in [BW1]. In the present paper affine and projective normalization schemes of planar non-algebraic curves are discussed. Using ideas from the 'shape from texture' community, cf. [DH1, Wi1, BY1, Gå1], it is possible to construct invariants for which the choice of a canonical frame is based on global features of a curve segment or region. If

* The work has been supported by the Swedish National Board for Technical and Industrial Development (NUTEK). The work is done within the ESPRIT–BRA project VIVA.

these features are chosen wisely, one can construct affine and projective normalization schemes with good discrimination and high tolerance to digitization errors. However, it has been shown in [Ås3] that for projective invariants of curves, even in the planar case, there are some fundamental difficulties. This has to do with the fact that the euclidean errors in image processing do not interact well with projective equivalence.

2 Background

Invariant based recognition of planar curves under affine or projective transformations has been studied earlier using several different approaches.

1. Differential Invariants, cf. [We1]. Smooth curves can be recognized using differential invariants. These invariant signatures can in theory be used to recognize curves even though almost all of the curve is occluded. Unfortunately such differential invariants under projective transformations require fifth order derivatives. There are substantial numerical problems in estimating these high-order derivatives.

2. Semi-differential Invariants, cf. [GM1]. If distinguished points can be found it is possible to reduce the orders of derivatives needed.

3. Distinguished points. If many distinguished points can be found it is possible to use these directly in the construction of invariants which require no derivatives.

4. Projectively invariant fitting of algebraic curves, cf. [Ca1]. A method has been developed for projectively invariant fitting of ellipses to planar curves. These can be used to construct invariants directly or in the transformation into a canonical reference frame. This construction can also be used to find distinguished points.

5. Canonical frame, cf. [RZ1]. Four distinguished points at a concavity can be used to transform the concavity into a canonical reference frame. Any feature in this reference frame are invariant and can be used to recognize the curve.

6. Normalization. There are no reasons why the reduction into a canonical frame has to be done with point features or derivatives. One can in principle use any kind of features in the reduction into a 'normal' reference frame. A global approach can be justified if either the curve is unoccluded, or if unoccluded chunks can be segmented from the edges detected by an edge detector. The invariant features used are thus global with respect to this chunk, e.g. one concavity, but they are semi-local with respect to the whole curve.

3 Normalization and Invariants as a tool in a recognition system

This section is an attempt to justify and clearify the use of more global normalization schemes and invariants. It is important to understand that normalization and invariants simply are tools that have to be judged in their correct context, i.e. as parts of a complete recognition system. The ideas described hear are due to many of the participants of the ESPRIT-project VIVA, and is heavily influenced by the recognition system that has been developed in Oxford, cf. [RZ1, Ro1].

The goal is to identify non-algebraic planar structures in a scene by inspection of a single gray-scale image. Edges are extracted from the image, e.g. using a Canny-Deriche edge detector.

It is possible that a whole closed contour corresponding to one object can be found in this way, but often occlusion and errors in the edge detection cause problems. Not only might the curve be fragmented so that only a part of the curve is detected, but also might a curve fragment consist partly by a contour caused by one object and partly by something else, e.g. an occluding object or underlying texture.

To cope with these problems the curve is segmented, e.g. using distinguished points like bitangents and inflections. It is still a question of debate, which segments should be used first for recognition. The smallest possible segment, e.g. the curve between the two closest distinguished points, stands a greater chance of coming from the same planar curve, but such a small curve fragment has little discriminatory power. A larger portion of the curve, e.g. one consisting of one or several concavities, has less chance of belonging to the same planar structure but if it does the discriminatory power is much larger. The segmentation of a curve into curve fragments, regions or points are illustrated in Figure 1.

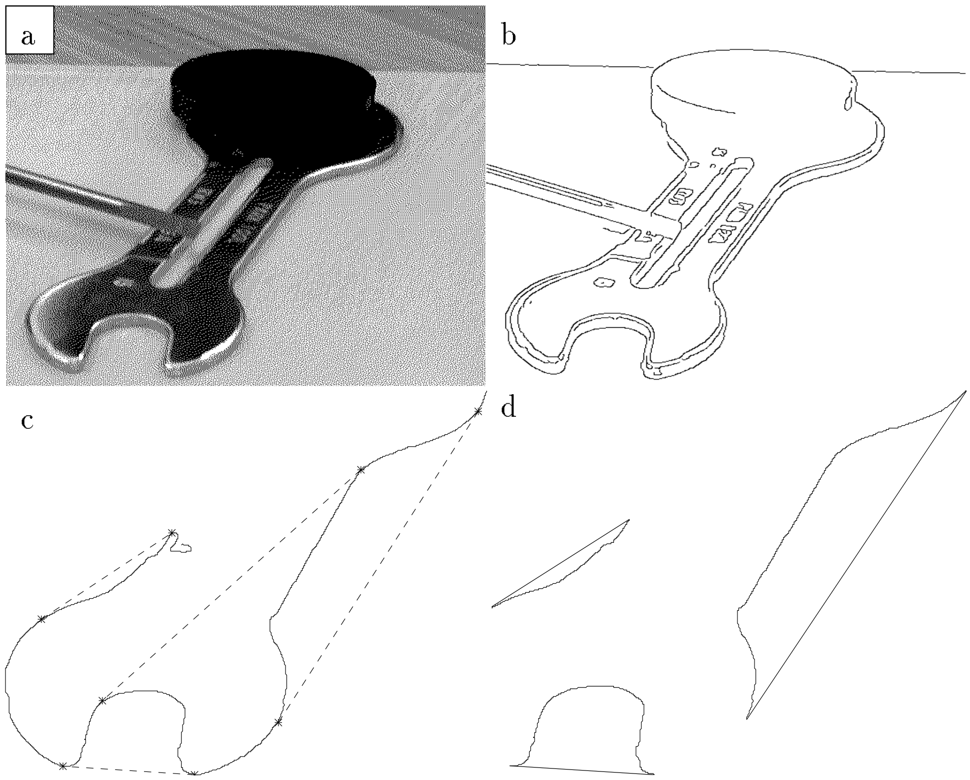

Fig. 1. 1.a. A gray scale image of a scene with a roughly planar object. 1.b. Edges are extracted using a Canny-Deriche edge detector. 1.c. Distinguished points on one segment can be used to segment a curve into pieces in a projectively invariant way. Both points and curve segments can be used in recognition. 1.d. Distinguished points and lines can also be used to extract small regions in a projectively invariant way. Three such regions are shown in the figure. These regions can be used to recognize the curve.

The key point here is that once the curve is segmented into small chunks, these can be recognized using invariant indexing. The chunks are then assumed to belong to the same object and global invariants can be used to recognize them. It might therefore be useful to have tools to recognize fairly general chunks, like points, curve fragments, regions or a combination of these. Thus although the techniques described in the following sections are global, they can be used as tools in a semi-local way.

The general idea is to choose a canonical reference frame based on global features. The segment can easily be recognized in this canonical frame by linear search through a database of similarly normalized models, but this is often inefficient. It is often more efficient to extract invariant indices. Once in a canonical reference frame any feature extracted defines an invariant. The invariants used in recognizing curve fragments are designed to give highest discrimination rate, e.g. using principal components decomposition. Recognition of a curve segment now consists of a simple table lookup using these invariant indices. This gives a direct reduction of possible hypotheses into a smaller amount which then can be searched through linearly and be accepted or rejected at verifications at different levels.

A recognition system thus consists of many parts:

- Edge extraction
- Segmentation
- Normalization of Segments.
- Invariant Index Extraction and Table-lookup
- Verification at different levels.

In the remaining sections of the paper, only the problem of normalization and invariant index extraction will be addressed. The task is thus to extract invariant indices from a segment, where a segment may constist one or several points, curve fragments or regions.

4 Normalization, Invariants, Recognition and Pose

In this section some notations are introduced. Normalization is presented as a way of constructing invariants. The process can be seen as the introduction of an abstract coordinate system, on a configuration set. Invariants of objects consisting of a finite combination of points and lines are relatively easy to construct. A survey of different methods can be found in [G1]. It is also relatively easy to model the configuration space, usually R^n or P^n_R, with full topological and differentiable structures.

One difficulty, when working with curves and regions is that their structure is much more complicated than that for point configurations. We will first consider our objects, i.e. curves and regions, as elements in an abstract set Ω. Using practically no structure at all it is still meaningful to talk about invariants and normalization under group action. Adding topology on Ω and differentiable structure on the group G it becomes meaningful also to talk about continuity.

4.1 Preliminaries

The notion of a group acting on a set is briefly presented. This group action defines equivalence classes of objects which are said to have the same shape. Invariants are seen as mappings that are constant on the equivalence classes.

A group G is said to act on a set Ω if there exists a mapping

$$(G, \Omega) \ni (g, \omega) \longrightarrow g(\omega) \in \Omega$$

with the following properties

$$1(\omega) = \omega, \qquad \forall \omega \in \Omega$$

and

$$g_1(g_2(\omega)) = (g_1 g_2)(\omega), \qquad \forall \omega \in \Omega, \qquad \forall g_1, g_2 \in G$$

The notation for group action is either $g\omega$ or $g(\omega)$.

Two elements ω_1 and ω_2 are said to have the same shape if $\omega_1 = g\omega_2$ for some transformation $g \in G$. This is an equivalence relation, because of the group structure of G. We write

$$\omega_1 \sim \omega_2 \iff \exists g \in G, \quad \omega_1 = g\omega_2 \tag{1}$$

The equivalence relation divides Ω into disjoint equivalence classes. Denote the equivalence class containing ω by $G\omega = \{g\omega | g \in G\}$. This is also called the orbit of ω under the group action G.

Let $T : \Omega \longrightarrow W$ be a function defined on Ω with values in some feature set W. This function is called an **invariant** if

$$\omega_1 \sim \omega_2 \Longrightarrow T(\omega_1) = T(\omega_2) \tag{2}$$

An invariant is called **complete** if

$$\omega_1 \sim \omega_2 \Longleftrightarrow T(\omega_1) = T(\omega_2) \tag{3}$$

In computer vision invariants are used to determine the shape, under some transformation group, of a configuration in the image. Completeness is an important property. For complete invariants no information is lost when going from ω to $T(\omega)$. Any function of $T(\omega)$ is also an invariant. Furthermore all invariant features of ω can be calculated from $T(\omega)$. Useful invariants should also be easy to compute and stable under 'small' distortions in ω, where the notion of small will be commented upon below.

Example 1. The trivial mapping $T \equiv 0$ is an invariant, but it cannot be used to discriminate between any elements of different shape.

Example 2. The mapping $T : \omega \longrightarrow G\omega$ is a complete invariant. It is however difficult to work with if the set $G\omega$ is large.

Example 3. A Non-trivial example. If Ω is the set of ordered linear 4-point configurations, then the cross ratio is a complete invariant under the projective transformation group.

4.2 Normalization

Normalization schemes can be constructed using properties that change 'a lot' under transformations in G. These latter properties will enable us to select a small number of representatives from each equivalence class, defining the invariants.

A typical property could be obtained in the following way. Take any mapping $P : \Omega \longrightarrow R^n$. This will in general not be an invariant. The element ω is said to have the property P if $P(\omega) = 0$. Assume that P can be chosen so that at least one element from each equivalence class has the property P and let

$$\Omega_P = \{\omega | P(\omega) = 0\}.$$

The projection of ω onto Ω_P along the equivalence class $G\omega$,

$$T(\omega) = G\omega \cap \Omega_P \tag{4}$$

is a complete invariant. Here $T(\omega)$ is a set-valued function, containing the representatives for the equivalence class $G(\omega)$. In the sequel, the term **normal** reference frame will be used for elements in Ω_P.

Assume for simplicity that $T(\omega)$ contains only one element. An element ω_1 can then be represented as

$$\omega_1 = g_1 \omega_1^{inv} \tag{5}$$

with $g_1 \in G$ and $\omega_1^{inv} = T(\omega_1)$. The element ω_1 is thus represented as a product of two factors, one transformation and one element of Ω. One way to view this is to say that we have introduced an abstract coordinate system on $\Omega = G \times \Omega_P$. An element ω_1 has shape coordinate ω_1^{inv} and group coordinate g_1. Notice that ω_1^{inv} carries all information about the shape of ω_1 and that g_1 carries all information about where on the equivalence class ω_1 lies. A second element ω_2 is also represented as

$$\omega_2 = g_2 \omega_2^{inv}$$

with $g_2 \in G$ and $\omega_2^{inv} = T(\omega_2)$. These two elements have same shape if and only if $\omega_1^{inv} = \omega_2^{inv}$. If this is the case, then $\omega_1 = g_1 g_2^{-1} \omega_2$, as is seen in the following diagram.

$$\begin{array}{ccc} \omega_1 & \longleftarrow & \omega_2 \\ \uparrow g_1 & & \uparrow g_2 \\ \omega_1^{inv} & = & \omega_2^{inv} \end{array}$$

In this way, we have developed a tool to determine the transformation between two elements of the same shape.

Example 4. Affine normalization of four points or affine coordinates. Let Ω be the class of ordered 4-point configurations,

$$\omega = \{(x_1, y_1), (x_2, y_2), (x_3, y_3), (x_4, y_4)\}$$

in R^2, let G be the group of planar affine transformations acting on ω pointwise, and let the normal reference frames be defined by

$$\Omega_P = \{\omega | (x_1, y_1) = (0,0), (x_2, y_2) = (1,0) \text{ and } (x_3, y_3) = (0,1)\}.$$

If the first three points of ω are not colinear there is exactly one representative in each equivalence class. Notice that other normal reference frames are obtained by choosing other points as basis points.

Example 5. Affine normalization of curve fragments. Shape from texture. Let Ω be the class of curves with finite arclength, let G be the group of planar affine transformations acting on each curve pointwise. Then the curves can be normalized into an essentially unique normal reference frame using either

- isotropy, cf. [Wi1, Gå1].

$$\Omega_P = \{\omega | \omega \text{ is weakly isotropic}\}$$

- maximization of compactness, cf. [BY1].

$$\Omega_P = \{\omega | l^2(\omega)/A(\omega) \text{ is minimal in the equivalence class}\},$$

where A is area enclosed by the curve and l is arclength of the curve.

4.3 Adding structure: Continuity and Probabilistic Behaviour

So forth we have only considered the set Ω without any structure. Suppose that we have some notion of how close two elements are to each other, e.g. a metric or topology. Then it becomes meaningful to discuss continuity of the normalization scheme.

Useful invariants should be complete, continuous and easy to compute. These considerations give claims on the choice of Ω_P, or alternatively on the features used to define Ω_P.

- $\Omega_P \cap G\omega$ should be a small set, but not empty
- $\Omega_P \cap G\omega$ should be insensitive to distortion in ω
- $\Omega_P \cap G\omega$ should be easy to compute

In the sequel we will say that a normalization scheme has the uniqueness property if there is only one normal reference frame. A normalization scheme is called continuous if the normal reference frames depend continuously on ω.

One direct result of the implicit function theorem is that g is defined implicitly by $F_\omega(g) = P(g^{-1}\omega) = 0$ and will depend continuously on ω, in a neighborhood of (g, ω) if the following conditions hold.

- $F_\omega(g) = 0$
- F_ω is continuous in both g and ω
- The differential dF_ω/dg is non-singular and continuous in g and ω

The differential also gives some information about the robustness of the normal reference frame.

5 A Continuous Affine Normalization Scheme. Experiments.

One simple example where these ideas work is the following affine normalization scheme on the class Ω of compact subsets ω of R^n.

Definition 1. Introduce the Hausdorff metric on Ω, i.e.

$$d(\omega_1, \omega_2) = \max_{z_1 \in \omega_1} \min_{z_2 \in \omega_2} ||z_1 - z_2|| + \max_{z_2 \in \omega_2} \min_{z_1 \in \omega_1} ||z_1 - z_2||. \quad (6)$$

Let the moments be defined as

$$\begin{aligned} m_0(\omega) &= \textstyle\int_{x \in \omega} dx_1 \dots dx_n \\ m_1(\omega)_i &= \textstyle\int_{x \in \omega} x_i \, dx_1 \dots dx_n \\ m_2(\omega)_{ij} &= \textstyle\int_{x \in \omega} x_i x_j \, dx_1 \dots dx_n \end{aligned}$$

Notice that m_0 is a scalar, m_1 a vector and m_2 a matrix. Notice also that these features depend continuously on ω in the Hausdorff metric. With I denoting the identity matrix, let the class of 'normal' elements be

$$\Omega_P = \{\omega | \quad m_0(\omega) = 1, m_1(\omega) = 0, m_2(\omega) = aI, a \in R\}. \quad (7)$$

This gives a normalization scheme with the following properties.

- Uniqueness. In (5), g and ω^{inv} are unique up to rotation.
- Continuity. Both g and ω^{inv} depend continuously on ω.

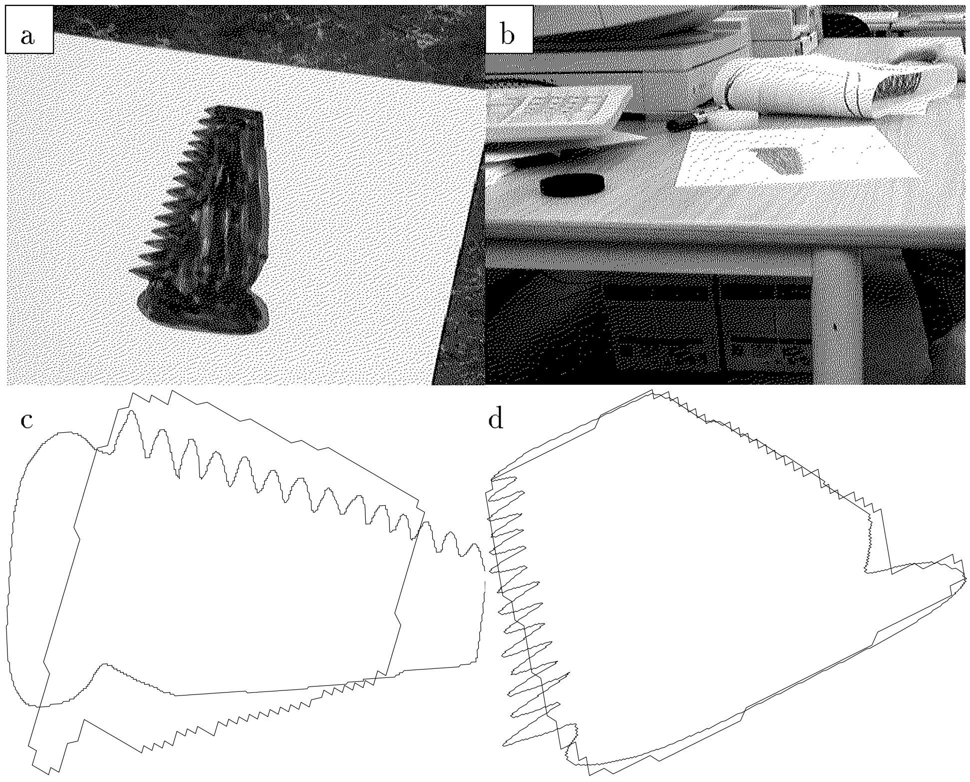

Fig. 2. 2.a. and 2.b. Images of a sawlike planar curve. In 2.b., interesting features disappear because of low resolution. 2.c. The two curves can be affinely normalized using maximum compactness or weak isotropy, but the normal reference frame depends crucially on how the curve is approximated. In this figure maximum compactness have been used. 2.d. Moment based normalization on the other hand is very robust to these kind of digitization errors. Notice the stability of this normal reference frame.

- Easy to compute. The transformation g can be directly computed from the moments of order 0, 1 and 2 of the region ω.
- Robust to digitization errors.
- No distinguished points are needed.

This scheme can be very useful in the recognition of planar curve fragments. The affine approximation can be very strong for these small chunks. Example of its use are given in Figures 2, 3 and 4. It can be seen that the method has good robustness properties, also in comparison with maximum compactness and weak isotropy. Figure 3 shows different concavities in their affine normal reference frame, and Figure 4 is a plot of invariant R^3-features extracted from the concavities in their normal reference frame. The result from the edge detector is jagged because of digitization effects. These edges could be smoothed before normalization but this has not been done in the Figures, to illustrate the size of digitization errors in different images. Completely analogous normalization schemes can

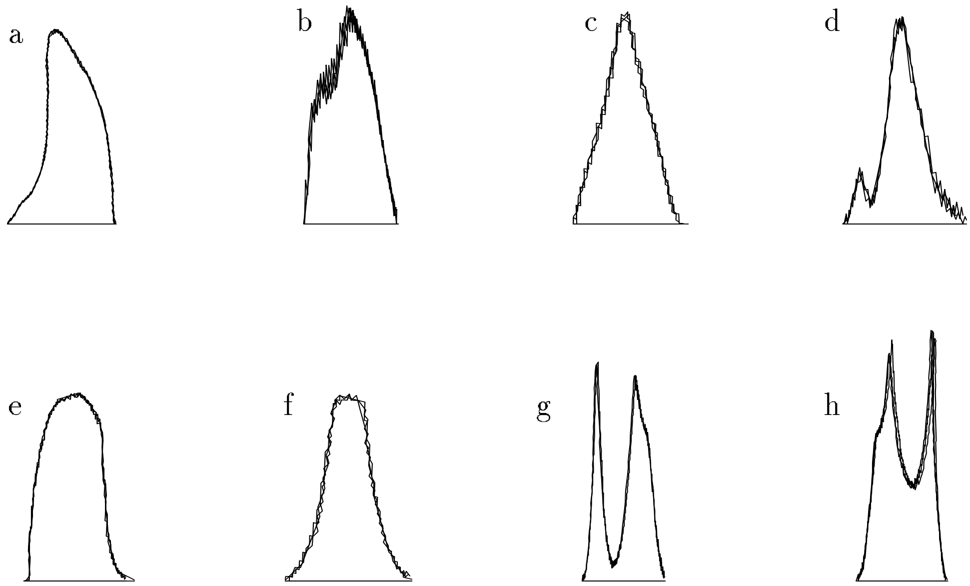

Fig. 3. Extracted edges of eight concavities from three different grayscale images, in their affine normal reference frame. In each of the eight plots there are three superimposed concavities. The edges from the edge detector was not smoothed. These jagged edges does not affect the moment based normalisation scheme. Notice particularly the digitization errors in Figure 3.b., 3.c., 3.d. and 3.f.

be used for point configurations, curve segments and any combination of points curve segments and regions.

6 Why is projective normalization inherently difficult?

It would be nice if it was possible to extend the affine scheme of section 5 to projective transformations. A continuous normalization scheme for compact regions of P_R^2 under projective transformations with unique normal reference frame is perhaps desirable. There are however fundamental reasons why this is not possible, as will be described below.

6.1 All smooth curves can be mapped arbitrarily close to a circle by a projective transformation

It has been proposed to use maximum compactness as a criterion for choosing normal reference frames for curves under projective transformations, i.e.

$$\Omega_P = \{\omega| \quad l^2(\omega)/A(\omega) \text{ is minimal in the equivalence class}\}, \tag{8}$$

where l denotes arclength and A is area enclosed by the curve, cf. [BS1]. One problem with this approach is that the optimization is done over a non-compact parameter space.

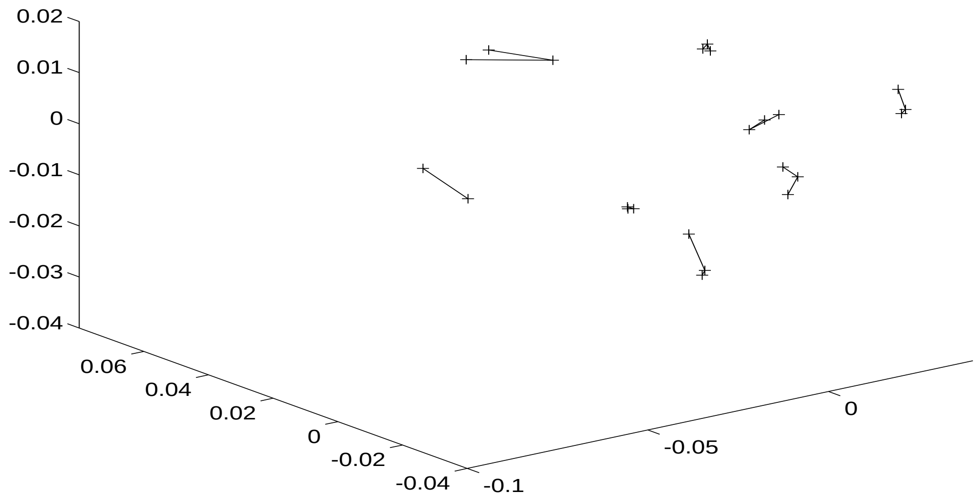

Fig. 4. Three principal components of the invariant features extracted from the 24 concavities in Figure 3. These are used to index a recognition table. Corresponding invariant of the same shape are joined with a line. Notice the nice separation between the eight different shapes.

It is therefore possible that the infimum is not attained at any point. In [Ås3] it is shown that this is indeed the case if the convex hull has a part that is smooth and curved. To explore this we use the following metric.

Definition 2.

$$\tilde{d}(\omega_1, \omega_2) = \max_{z_1 \in \omega_1} \min_{z_2 \in \omega_2} ||z_1 - z_2|| + \max_{z_2 \in \omega_2} \min_{z_1 \in \omega_1} ||z_1 - z_2|| + |l(\omega_1) - l(\omega_2)|, \qquad (9)$$

where $||z||$ is the euclidean norm and l arc length.

This metric is a modification of the Hausdorff metric on compact subsets of the plane, i.e. the max-min parts, the modification being that also the arc lengths should be compared. Being close to a given curve in this metric is a strong restriction. Two curves that are close really look alike, i.e. they would produce almost the same image on a CCD screen or on the retina of an eye. All points on one curve are close to some point of the other and their arc lengths are almost equal. The following theorem is true both for the original Hausdorff metric (6) and for the modified metric (9).

Theorem 3. *Let Γ_0 be the unit circle and let ω be a curve with finite arc length and having a smooth curved part on the convex hull. Then*

$$\inf_{\omega' \in G\omega} \tilde{d}(\omega', \Gamma_0) = 0$$

The theorem is a simple consequence of the fact that by projective transformations any part of the convex hull, can be magnified to any degree. An immediate corollary is that for such curves ω we have

$$\inf_{\omega' \in G\omega} \frac{l(\omega')^2}{A(\omega')} = 4\pi$$

Here we need the modified metric (9) to control the arclength of the curve. The projective transformations involved when approaching the limit are, however, quite extreme. A detailed proof is given in [Ås3].

Since most curves can be approximated by smooth curves, Theorem 3 has some serious consequenses. If a curve is modeled by a cubic spline the infimum is not attained. If a curve is modeled as a piecewise linear curve then the minimum is attained but the normalized reference frame will depend crucially on how the boundary is approximated. This is illustrated in Figure 5. In this figure the same shape is approximated with polygons to different accuracy. The reference frame in which they have maximal compactness is indeed different.

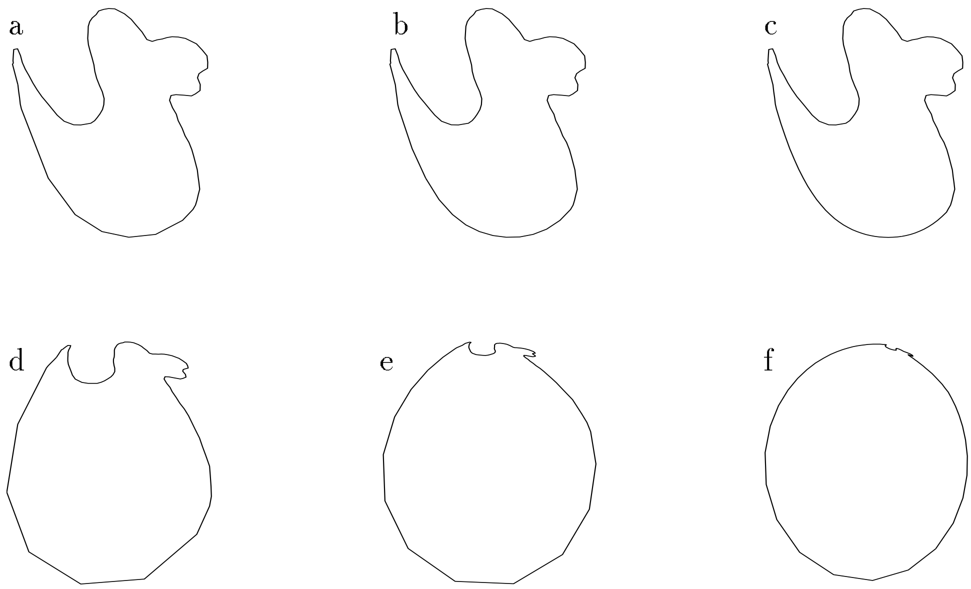

Fig. 5. Three different approximations of a planar curve (a, b and c) and their normal projective reference frames with respect to maximal compactness (d, e and f). Notice how much the normal reference frame depends on the representation of the boundary.

6.2 Transforming a rabbit to a duck

In the proof of Theorem 3 and as suggested by Figure 5, one notices that the main part of the curve is squeezed into a neighbourhood of a point. This can be exploited, yielding the following theorem.

Theorem 4. *Given $\Gamma_1, \ldots \Gamma_m$, closed continuous curves with finite arclength. To every $\epsilon > 0$, there exists a curve C and projective transformations $q_1, \ldots, q_m$ so that*

$$d(q_i(C), \Gamma_i) < \epsilon, \qquad i = 1, \ldots, m$$

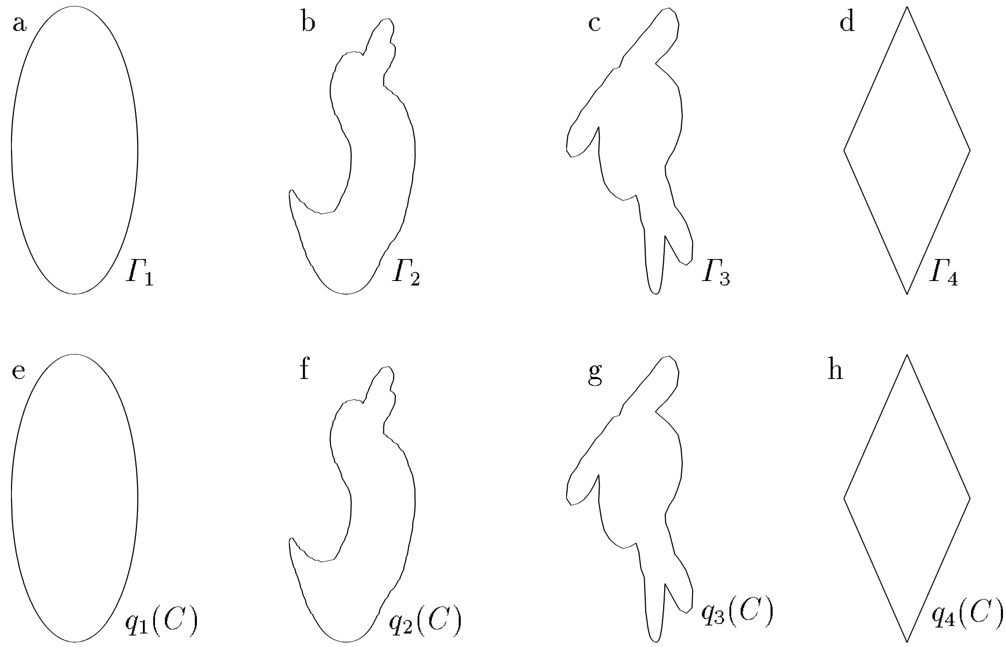

Fig. 6. 6.a-d. Four planar curves with different projective shapes. 6.e-h. Four planar curves that are projectively equivalent, i.e. have the same shape. Errors in the plotting device, printing, copying and viewing conditions make them look pairwise equal (a and e, b and f, c and g, d and h).

Note that the curves Γ_i do not have to be smooth. The theorem is illustrated by Figure 6. A detailed proof is given in [Ås3]. Notice that in Figure 6 the lower four curves are projectively equivalent. The only errors are in the plotting, printing, copying and viewing of the curves. The theorem is in itself somewhat surprising and unintuitive at first, but it is a simple trick of hiding a shape along the convex hull of another shape. The reason it works is the use of extreme, but non-singular, projective transformations. The consequences are perhaps more important. Consider a fixed curve ω and all curves in an arbitrary small neighborhood O_ω of ω. The theorem says that the orbit of this open set is dense in ω. A corollary is thus.

Corollary 5. *Let T be a projectively invariant mapping from the set of closed continuous curves with finite arclength to R^n, then T maps all curves at which it is continuous to the same value.*

The problems that cause these kind of anomalies can be experienced in almost normal viewing conditions. It is the simple fact that euclidean scale is changed under projective transformations so that what seems like a small scale perturbation from one viewpoint is large scale shape from another. This is illustrated in Figure 7.

Suppose that we have a continuous projective normalization scheme that gives a unique representative from each equivalence class. It would then be possible to construct

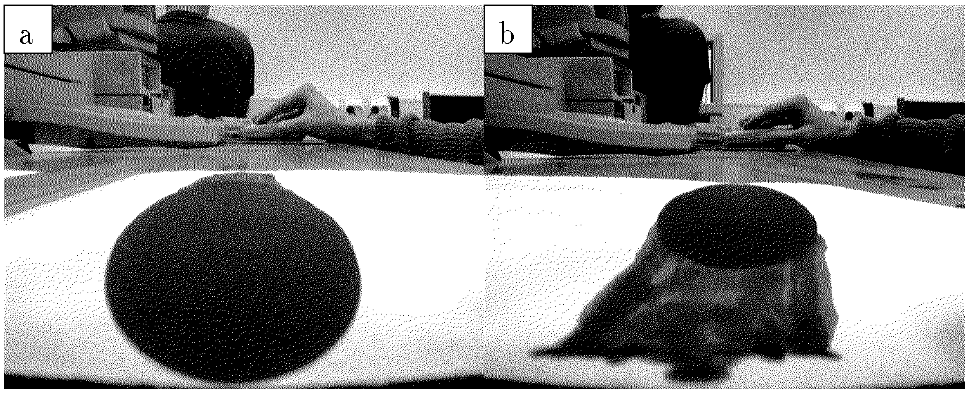

Fig. 7. It is difficult to talk about scale when dealing with projective equivalence of planar curves. What seams like a small scale pertubation from one viewpoint might look like large scale shape from another.

a continuous and non-constant projective invariant mapping from the set of curves with the hausdorff metric to the real line. This is impossible according to Corollary 5. The conclusion is that projective normalization schemes on the set of planar curves cannot both be continuous and give a unique representative from each equivalence class. Either continuity or uniqueness has to be sacrificed. Corollary 5 seams to rule out all attempts to use projective invariants in recognition of planar non-algebraic curves. However, it will be shown in the next section that it is possible to construct something almost as good.

7 Projective Normalization. Experiments.

A consequence of the previous section is that no matter what method is used in normalizing a curve with respect to projective transformations in a continuous way, there are curves with more than one normal reference frame. In this section an example of such a normalization scheme is presented. Only closed contours or rather the regions enclosed by the contours will be considered. Let therefore Ω be the class of compact regions in R^2 with positive measure, bounded by a piecewise smooth curve. Let the moments of a region ω be defined according to

$$\begin{aligned} m_0(\omega) &= \int_{x \in \omega} dx_1 dx_2 \\ m_1(\omega)_i &= \int_{x \in \omega} x_i \, dx_1 dx_2 \\ m_2(\omega)_{ij} &= \int_{x \in \omega} x_i x_j \, dx_1 dx_2 \\ m_3(\omega)_{ijk} &= \int_{x \in \omega} x_i x_j x_k \, dx_1 dx_2 \end{aligned}$$

Consider the group of planar projective transformations G acting on ω. In the sequel only physically realizable transformations will be considered, i.e. those which do not send any of the points in ω to infinity. Now base a normalization scheme on the following 'normal' reference frames

$$\Omega_P = \{\omega | \quad m_0(\omega) = 1, m_2(\omega) = aI, m_3(\omega) = 0, a \in R\} \tag{10}$$

This gives a normalization scheme with several representatives from each equivalence class. The number of solutions may vary, but each solution can be continuous in the Hausdorff metric. One way of locking the rotation is to demand that the maximum distance of a point in ω to the origin occurs at the x_1-axis. This method can be used also with convex curves.

The normalization scheme has been implemented and an experimental session will be presented. In this experiment grayscale images of roughly planar objects are taken with a digital camera. Polygon approximations of contours in the image are obtained using a Canny-Deriche edge detector. These curves are then normalized according to the proposed method, see Figure 8. Notice the good performance in Figures 8.b and

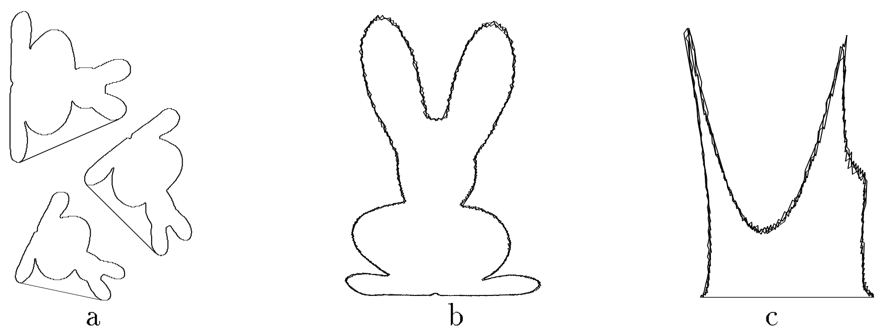

Fig. 8. Three images of rabbits, cf. Fig. 8.a, are normalized into a reference frame in which their third moment tensor is zero and the second moment matrix is the identity, cf. Fig. 8.b. The same normalization is applied to the three regions enclosed by a bitangent and part of a contour, cf. Fig. 8.c.

8.c. The three normalized curves lie practically on top of each other in spite of possible nonlinearities in the camera, errors in segmentation and in edge detection. The rabbit covered roughly 150×200 pixels in the image.

The normalization scheme is quite general and it is possible to normalize convex curves as well. This is illustrated in Figure 9. Sixteen images are taken of a convex shape. These are transformed into the unique normal affine reference frame, given by (7) and into one normal projective reference frame, given by (10). The affine approximation holds reasonably well for several of these images, but for the more extreme views, like Figure 9.a., projective normalization is more adequate. Projective invariant R^2-features extracted from one of several possible normal projective reference frames is shown in Figure 10.

8 Conclusions

In this paper projective normalization schemes of planar non-algebraic curves are discussed. Such schemes should be continuous, and preferably give a unique representative from each equivalence class. In the affine case it has been shown that this can be achieved.

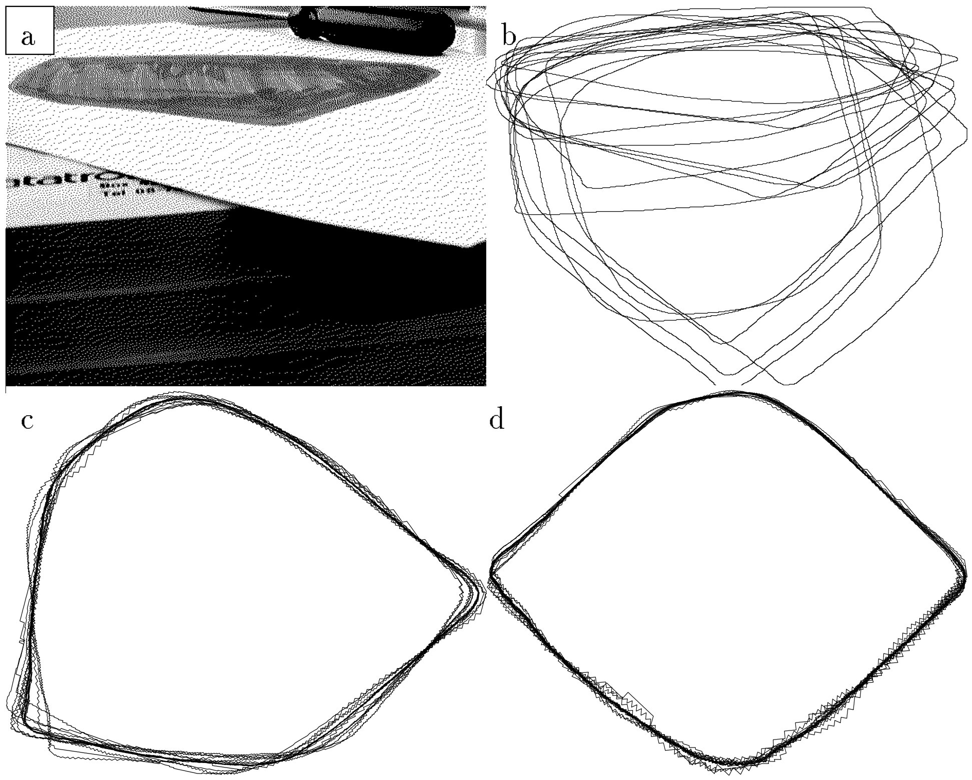

Fig. 9. Figure 9.a shows one of 16 images taken of a planar convex curve. All 16 extracted edges are shown in figure 9.b. In figure 9.c. the edges have been transformed into their unique affine normal reference frame, and in figure 9.d one of several possible projective normal reference frames have been used.

A canonical frame, that does not depend on ordering of points or choice of affine basis points, can be chosen for general feature configurations like compact regions, curve fragments, and point configurations. In the projective case things are more difficult. It has been shown that it is impossible to achieve both uniqueness and continuity for projective normalization of general non-algebraic curves. A normalization scheme is presented, where uniqueness has been sacrificed for robustness. The normalization scheme is applied to closed curves, if possible, otherwise to parts that have been segmented in a projectively invariant manner. The resulting invariants have good discriminatory properties and are robust to digitization errors. There exist however curves which are inherently difficult to normalize.

The work has focused on simple algebraic and topological properties and can be extended in several directions. First of all there are several questions of how these invariants should be used in a recognition system. There seems to be curves that are generically difficult to normalize with respect to projective transformations. It might be possible to examine these critical sets. The continuity of the proposed invariants should

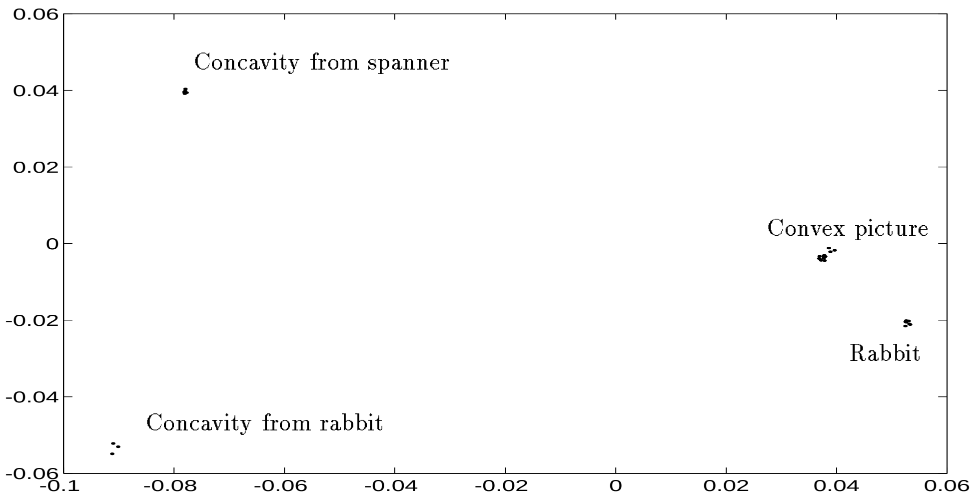

Fig. 10. Invariant features extracted after projective normalization. Sixteen images of the convex shape in figure 9, seven images of the whole rabbit in figure 8, three images of one of the concavities from the rabbit in figure 8 and nine images of one concavity from a spanner. Notice the high discriminatory power.

be investigated further. It would also be interesting to incorporate probabilistic models for image distortions. This could give valuable insight into the effectiveness of the normalization schemes and their use in recognition.

Acknowledgements

I would like to thank my supervisor Gunnar Sparr for inspiration and guidance. I would also like to thank my fellow students Anders Heyden and Carl-Gustav Werner for their help.

References

[BM1] Blake, A., Marinos, C.: Shape from Texture: Estimation, Isotropy and Moments. Artificial Intelligence **45** (1990) 332-380

[BS1] Blake, A., Sinclair, D.: On the projective normalisation of planar shape. Technical Report OUEL Oxford Great Britain (1992)

[BY1] Brady, M., Yuille, A.: An Extremum Principle for Shape from Contour. PAMI-6 **3** (1984) 288-301

[BW1] Burns J. B., Weiss R. S., Riseman E. M.: The Non-existence of General-case View-Invariants. Geometrical Invariance in Computer Vision, Mundy, J. L. and Zisserman, A. editors, MIT Press (1992)

[Ca1] Carlsson S.: Projectively Invariant Decomposition and Recognition of Planar Shapes, Proc. ICCV4, May, Berlin, Germany (1993) 471-475

[DH1] Duda, R. O. and Hart, P. E.: Pattern Classification and Scene Analysis. Wiley-Interscience (1973)

[GM1] Van Gool, L., Moons, T., Pauwels, E. and Oosterlinck, A.: Semi-differential Invariants. Proc. First DARPA-ESPRIT Workshop on Invariance, Reykjavik, Iceland (1991) 359-386
[Gå1] Gårding, J.: Shape from Surface Markings. Ph. D. thesis, Dept. of Numerical Analysis and Computer Science, Royal Institute of Technology, Stockholm, Sweden (1991)
[G1] Gros, P., and Quan L.: Projective Invariants for Vision. Technical Report RT 90 IMAG - 15 LIFIA, LIFIA-IRIMAG, Grenoble, France (1992)
[LS1] Lamdan, Y., Schwartz, J. T., and Wolfson, H. J.: Affine Invariant Model-based Object Recognition. IEEE Journal of Robotics and Automation **6** (1990) 578-589
[MZ1] Mundy, J. L., and Zisserman A. (editors): Geometric invariance in Computer Vision. MIT Press, Cambridge Ma, USA (1990)
[RZ1] Rothwell, C. A., Zisserman, A., Forsyth, D. A. and Mundy J. L.: Canonical Frames for Planar Object Recognition. Proc. ECCV2 Genova Italy (1992) 757-772
[Ro1] Rothwell, C. A.: Hierarchical Object Description Using Invariants. Proc. Second ARPA/NSF-ESPRIT Workshop on Invariance, Ponta Delgada, Azores (1993) 287-302
[We1] Weiss, I.: Noise-resistant Invariants of Curves. Proc. First DARPA-ESPRIT Workshop on Invariance, Reykjavik, Iceland (1991) 359-386
[Wi1] Witkin, A. P.: Recovering Surface Shape and Orientation from Texture. J. of Artificial Intelligence **17** (1981) 17-45
[Ås1] Åström, K.: A Correspondence Problem in Laser-Guided Navigation. Proc. Swedish Society for Automated Image Analysis, Uppsala, Sweden (1992) 141-144
[Ås2] Åström, K.: Affine Invariants of Planar Sets. Proc. SCIA8, Tromsö, Norway (1993) 769-776
[Ås3] Åström, K.: Fundamental Difficulties with Projective Normalization of Planar Curves. Proc. Second ARPA/NSF-ESPRIT Workshop on Invariance, Ponta Delgada, Azores (1993) 377-389

This article was processed using the LaTeX macro package with ECCV92 style