SHAPES FIT FOR PURPOSE

Anupriya Balikai[1] , Dr. Paul Rosin[2] , Yi Zhe Song[1] , Dr. Peter M Hall[1]

a.balikai@bath.ac.uk , paul.rosin@cs.cardiff.ac.uk , yzs20@cs.bath.ac.uk , pmh@cs.bath.ac.uk

{1} Media Technology and Research Center, University of Bath
{2} University of Cardiff

Our Motivation

University of Bath has been carrying out research in image classification using structure to be class invariant regardless of their depictive styles [1]. We have been successful in this area and want to extend an image's description to wider classes of objects. For example we want to define an object and its parts not only based on its structure but also its shape.

INTRODUCTION

We would like to describe any image using a collection of known shapes, specifically: ellipses, rectangles and triangles. If none of these fit we name the shape type a robust convex hull. Fitting these shapes to image segments generates a description which offers a high level of abstraction. This can be used in many applications.

BACKGROUND on Fitting SHAPES

Shape fitting has been studied for a very long time. It is usually restricted to a single shape model, like looking for a pattern of circles in an image, or for rectangles which represent buildings, regular polygons etc. However, our task is to choose from amongst several shape families.

Fitting multiple shape models and then selecting the most appropriate is less common due to the problems that can arise while comparing shape measurements.

Many schemes have been proposed to overcome this problem but there is no single agreed way to fit some shape from more than one family. The most common types of fitting are chosen for their mathematical tractability and computational convenience, rather than how well they correspond to perceptual or aesthetic judgments.

Thus if we show that we are able to match photographs to artwork and produce synthetic artwork, we show that we value aesthetic similarity.

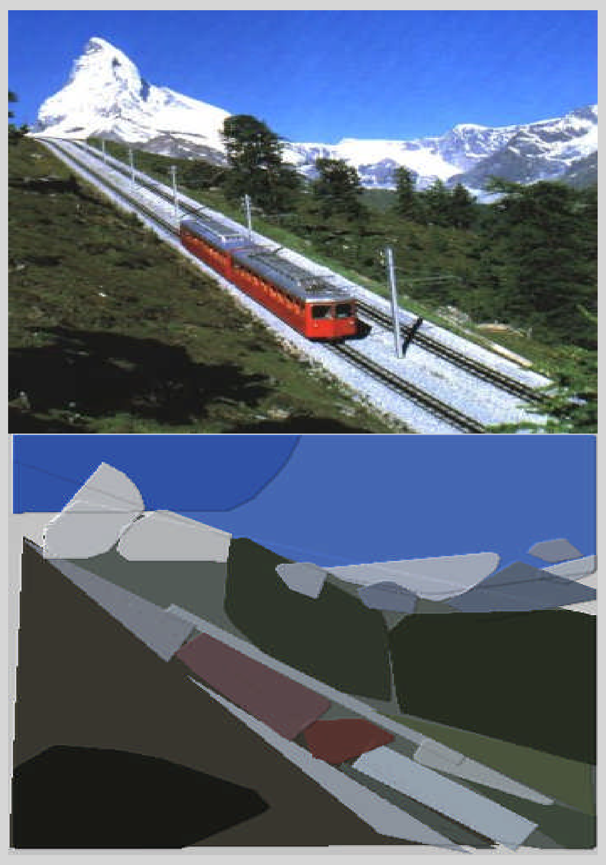

Figure 5: Left, an original photograph, right an abstraction in shapes. This shows that simple shapes are sufficient to capture the essence of a complicated image: the lack of detail can be advantageous and even desirable.

WHAT WE DO:

We fit a shape drawn from a selection of shape families, by a supervised classifier using just 3 images.

Fitting Shapes of a Single Type

We fit four categories of shape. Three of them are 'known' because they can be qualitatively labelled: ellipse, triangle and rectangle. The fourth shape type suggested by Rosin and Mumford [2] is actually a version of a convex hull, which is a convex polygon that maximises the area of overlap with the input polygon. We use it in our algorithm to fit shape types that do not fit as one of the 'known' shape types.

Selecting one Shape from Many

Selecting appropriate shape models is done using a supervised classification paradigm using a C 4.5 decision tree [3].

Performance Data

The classifier was trained using 35 samples of each shape type, obtained from manual annotation of a segmented data set. Segmentation was carried out using normalized cuts with N=3. Figure 1 shows the confusion matrix associated with the classifier for the training set

	Ellipse	Rectangle	Triangle	Convex Hull
Ellipse	32	1	1	1
Rectangle	4	29	1	1
Triangle	3	1	30	1
Convex Hull	1	1	1	32

Figure 1: Confusion Matrix Each row shows the result of classifying a ground truth set of known shape. The number of times a ground truth shape class is classified as some nominated shape class are the values in each column for that row.

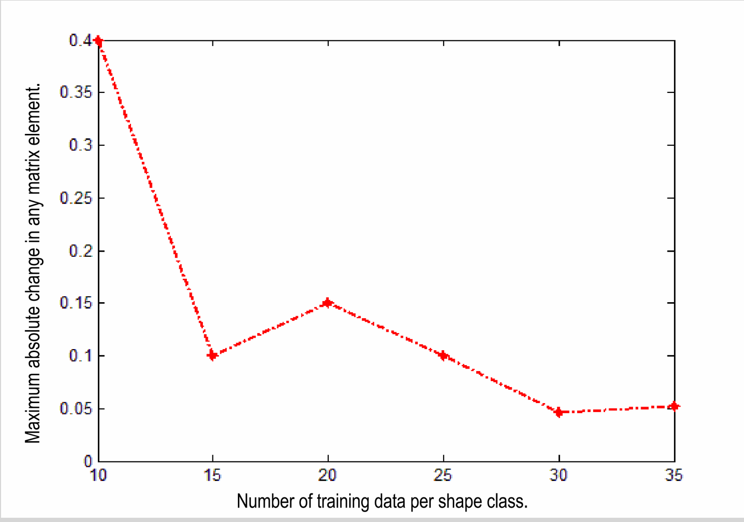

Figure 2: The maximum absolute change in any normalised confusion matrix element as a function of the number of training data in each shape class.

We use results from the classifier to match photographs and artwork of particular objects using a few qualitative shapes.

We have used our shape fitter in two applications: painting and matching.

Our matcher accepts two images as input and returns a list of matched regions. The regions are produced by the N-cut segmentor [4], which cuts an image into N segments. Given two images we manually select a particular level of N from just one of them, to in order to pick our foreground object. This foreground object then acts as a query image. The matcher is to locate this object in the tree representing the second image. We now find the matches between two sub trees. Each node in the tree has a shape fitted to it, using the shape classifier, so it only stores a label which is an element of {ELL, RECT, TRI, RCH}.

Figure 3 shows how photographs are matched to artwork. This is of interest in further applications, such as content based image retrieval.

This classifier is later on used to process photographs into abstract synthetic artwork.

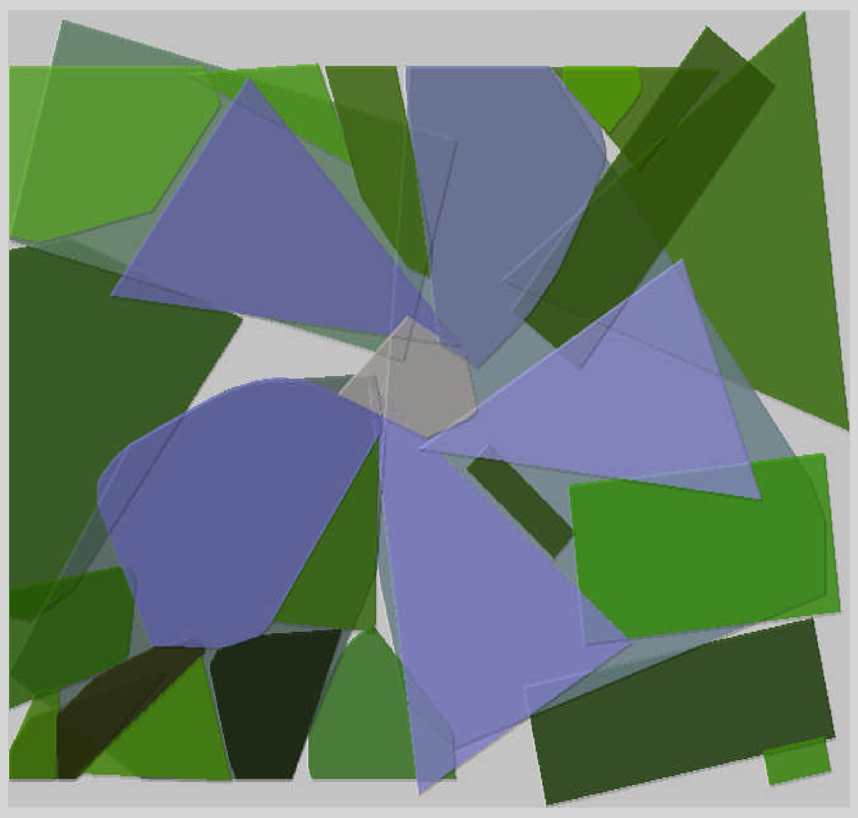

Figure 4: Left, a photograph of a flower is rendered as transparent shapes. Right. the photograph of a car has been rendered as paper cut-outs, which shows the shape fitted to each region.

In our second application, we use the shapes fitted by the same classifier as used by the matcher to create synthetic artworks of an abstract nature. Many of the abstract artworks we produce are largely motivated by artists such as Kandinsky and late Matisse, who used pure geometric shapes as primitives to create art.

We segment the image into different granularities using the same segmentor used in the matcher to show the performance of the classifier. Finer segments are generated by diving the image into many segments and courser segments are generated by dividing the image into only some segments. Finer segments are then rendered on top of the coarser ones, after filtering out some of the detailed/finer shapes.

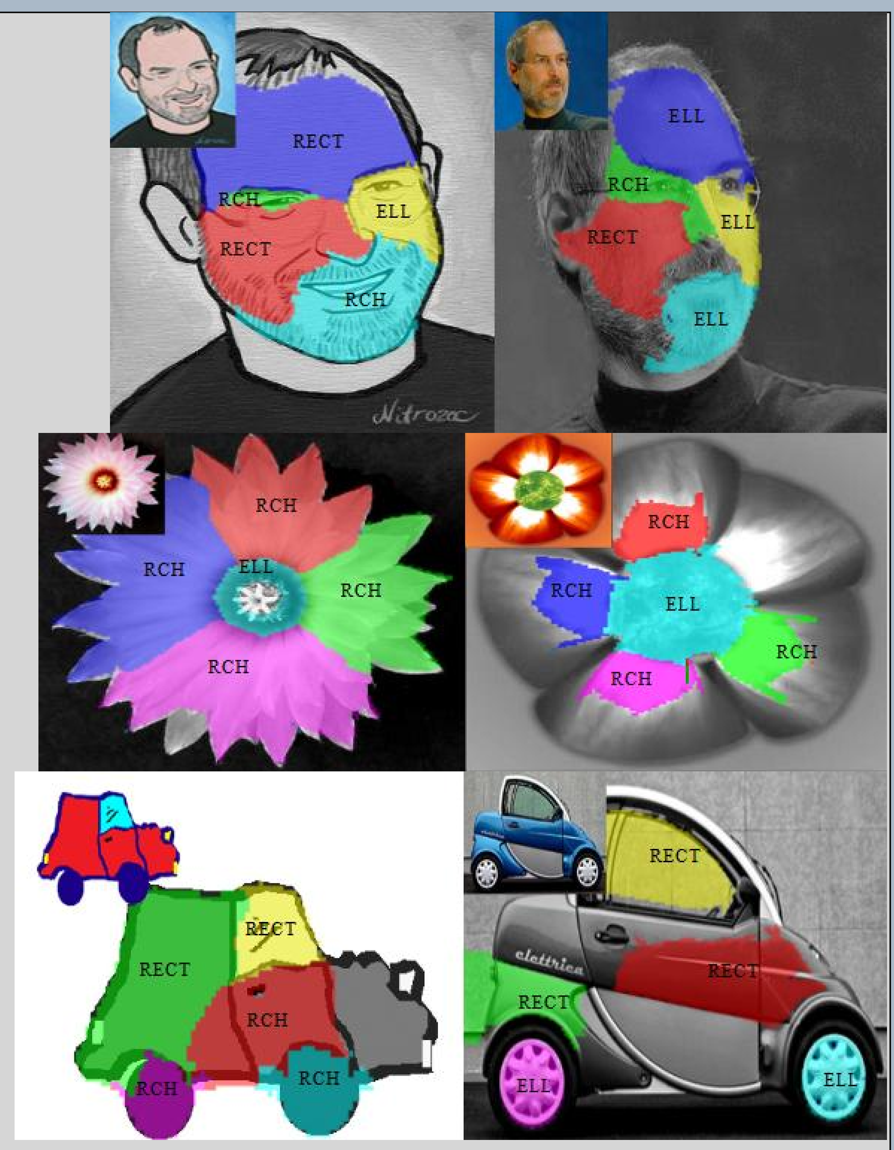

Figure 3: Results after matching photographs and artwork. The parts that are matched with each other are shown by colour coded regions. The shape labels that the individual parts were identified as are shown by the shape labels RCH – standing for Robust Convex Hull, RECT – Rectangle, ELL – Ellipse, TRI – Triangle.

CONCLUSION

We provide a novel method to fit not just a single shape to a region, but a way to classify a region as some shape from amongst several shape families. The classifier is extendible to shapes other than those we have chosen here – for example, super ellipses can be classified too. We have restricted ourselves to simple shapes based on the precedent of early 20th century art..

The description in images that come from our classifier have been put to use in both discriminative tasks (matching) and generative tasks (synthesis). Both applications offer novelty and both could find use elsewhere, so are utilitarian too. A qualitative shape matcher, as ours, might be used to initialise a more complex, quantitative shape matcher that will include measurement data.
We conclude that a shape based image description offers a level of abstraction that is of significant value to computer vision.

REFERENCES

[1] X. Bai, Y. Z. Song and P. M. Hall. Learning object classes from Structure. In BMVC 172-181, 2007.
[2] P.L. Rosin and C.L. Mumform. A symmetric convexity measure. Computer Vision Image Understanding 103(2), 101-111, 2006.
[3] J.R. Quinlan, C4.5: Programs for Machine Learning, Morgan Kaufmann, Los Altos, CA (1993).
[4] T. Cour, F. Benezit and J. Shi. Spectral Segmentation with Multiscale Graph Decomposition. In CVPR 1124-1131, 2005.