

Measuring Convexity via Convex Polygons

Paul L. Rosin¹ and Joviša Žunić²

¹ School of Computer Science & Informatics,
Cardiff University,
UK

² Mathematical Institute,
Serbian Academy of Sciences and Arts,
Belgrade,
Serbia

`Paul.Rosin@cs.cf.ac.uk, jovisa_zunic@mi.sanu.ac.rs`

Abstract. This paper describes a general approach to compute a family of convexity measures. Inspired by the use of geometric primitives (such as circles) which are often fitted to shapes to approximate them, we use convex polygons for this task. Convex polygons can be generated in many ways, and several are demonstrated here. These polygons are scaled and translated to ensure that they fit the input shape and produce a meaningful convexity measure. Subsequently, a convexity measure can be computed based on the degree of overlap between the two shapes.

Keywords: shape measure, convexity, classification

1 Introduction

Shape is a characteristic of objects that is useful in both biological perception and computer vision, and many different shape characteristics exist, e.g. linearity, circularity, convexity, ellipticity, elongation, compactness, linearity, sigmoidality, tortuosity, etc. While there also exist many approaches to measure these various characteristics, several approaches are commonly used across a range of shape characteristics. For instance, shapes which can be described by parametric models can be fitted to the object, and the fitting error used to quantify the membership of the object to that class of shape. An example is shown in figure 1b in which a circle has been fitted to minimise the squared Euclidean errors of the data points. Such an approach is attractive, since the fitted model presents the

ideal instance of the model most similar to the data, and therefore the errors fairly represent deviations from the shape characteristic [1]. Another popular approach for simple geometric models is to quantify the error with respect to the circumscribed or inscribed model, see figure 1c&d.

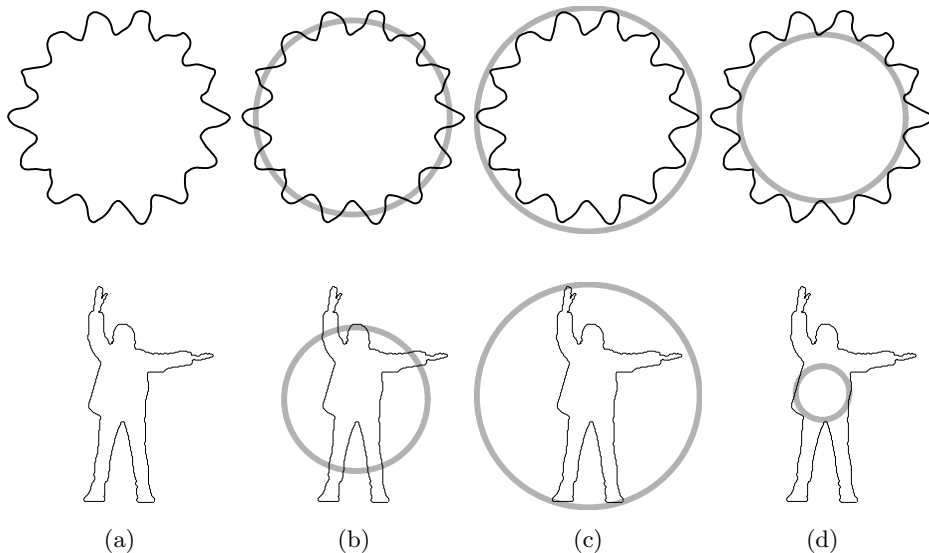


Fig. 1. Various approaches to fitting circles to data. (a) input shape (b) least squares fit of circle to data points (c) circumscribed circle (d) inscribed circle.

Convexity is a commonly used shape characteristic, and many convexity measures have been developed (e.g. [2–10]). Since it is not straightforward to determine what general convex geometric model is appropriate to fit in the least squared sense to the data, the dominant approach is to use instead the convex hull, as shown in figure 2a&c. This could be considered the analogue of the circumscribed circle presented in figure 1b. The convex skull (i.e. the largest convex polygon contained within the object) can also be used to measure convexity [8] – see figure 2b&d, but is rarely used due to the high cost of computing it (i.e. $O(n^7)$, see [11]), which is in contrast to the convex hull, for which efficient algorithms exist. However, a disadvantage is that for shapes not close to an instance of the perfect model, the circumscribed/inscribed model of the convex hull does




Fig. 2. Two approaches to fitting convex polygons to data. (a)&(c) convex hull (b)&(d) convex skull.

not provide a representative instance of the model, and therefore errors with respect to this model will not accurately reflect the shape characteristic.

As an alternative to the convex hull, Rosin and Mumford [8] described two symmetric convexity measures that were based on a “robustified” version of the convex hull, which was defined as the convex polygon that maximised its overlap with the input polygon. It can be seen in figure 3a that this polygon neither circumscribes or inscribes the input shape, and for the circle example it is rather similar to the fitted circle in figure 1b. In this paper we present several alternative approaches which also produce convex polygons that partially overlap the input polygon (see figure 3b–f). Again, for the circle example they are similar to the fitted circle in figure 1b.

2 New Convexity Measures

The basic approach taken in this paper to compute a convexity measure is:

1. generate a convex polygon based on the input shape; several options are presented in this paper,
2. transform the convex polygon so that it corresponds as best as possible to the input shape; in this paper we have applied a scaling such that the areas of the input and convex polygons are the same, and have also translated the polygons to maximise their area of overlap
3. compute a convexity measure; in this paper we have used the ratio of the area of overlap to the area of the input polygon.

We now describe the alternative approaches we have taken to step 1:

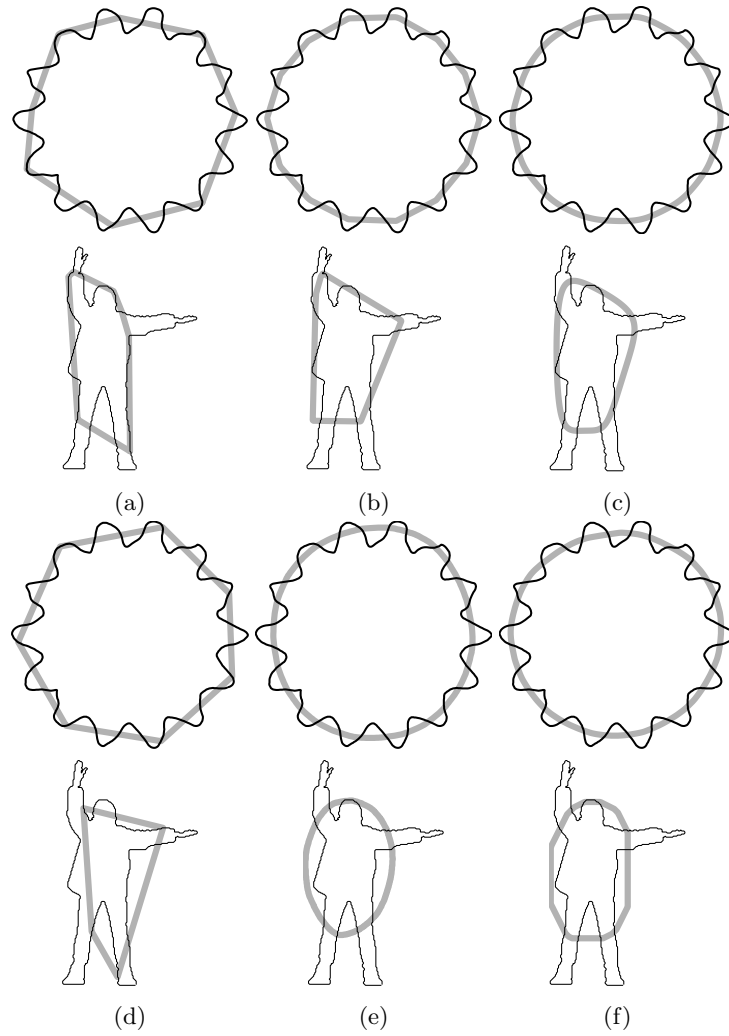


Fig. 3. Some alternative approaches to fitting convex polygons to data. (a) Rosin and Mumford (b) rescaled convex hull (c) smoothed input polygon (d) polygonal simplification of the input polygon (e) convexification (f) convexification using extended flipturns.

Convex Hull – C_A^S : A natural and efficient way to compute a convex polygon is to use the convex hull, see figure 2a&c.

Smoothing – C_B : It is known that if a curve is smoothed using the geometric heat flow equation it becomes more and more circular, eventually shrinking to a circular point in finite time [12]. We iteratively apply smoothing to the input shape until the result becomes convex, see figure 3c. Since the data has been heavily smoothed it is sufficient to test for convexity in a simple manner by checking, for a counterclockwise ordered curve, that at each vertex the sign of its subtended angle is positive. Alternatively more sophisticated methods (e.g. estimating curvature by fitting splines [13]) could be employed.

Polygonal Simplification – C_Y : There are many algorithms available. We use an optimal dynamic programming approach [14] applied to a subsampled version of the contour, which is reasonably efficient and robust. This provides a polyline containing n line segments, and a binary search is applied the convex polygon with the maximum value of n , see figure 3d.

Convexification – C_F : The polygonal convexification process was described by Paul Erdős [15, 16], and involves reflecting a polygon’s concavities about their corresponding edges in the convex hull. Iterating this process results in a convex polygon, which has previously been used to measure convexity [17]. Two such convexified polygons are shown in figure 3e.

Convexification with flipturn – C_{FT} : A modification of the above scheme, in which the concavities are reflected about the complete edge of the convex hull on which the concavity is located, and also have the order of their vertices reversed (an “extended flipturn”) was also considered, see figure 3f.

For those methods that generate a convex polygon by enlarging the input shape (such as the convex hull or convexification) the rescaling applied in step 2 provides a balance between protrusion and intrusion irregularities. This is not otherwise present in the most common convexity measures $\text{area}(S)/\text{area}(\text{CH}(S))$ and $\text{perimeter}(\text{CH}(S))/\text{perimeter}(S)$, which are therefore much more sensitive to protrusions than to intrusions.

For highly concave shapes there may be a large difference in areas between the original shape and its convex hull. Consequently, when the convex hull is shrunk to match the original shape’s area there is the possibility that it will not intersect the original shape, leading to a measured convexity value of zero. Since it is undesirable for the convexity measure to produce zero (as shapes

are assumed to have non-zero area) the shrunk convex hull is translated to ensure an overlap with the original shape. Although it would be possible to use more general transformations (e.g. rotation, affine with area preservation, etc.), if the convex hull example is considered then translation is the minimal transformation that produces a sensible convexity measure. That is, no non-zero area shape should produce a measured convexity equal to 0. It is generally most appropriate to use the simplest solution that is effective.

The translation parameters are determined using Powell’s method [18]. To ensure that the estimation process converges even in situations where there is no intersection between the shrunk convex hull and the original shape, the convex hull is shrunk gradually (10 steps have been found to be sufficient for all our data). Optimisation of the translation parameters is then interleaved with iterations of shrinking. The effectiveness of the optimisation is demonstrated in figure 4.

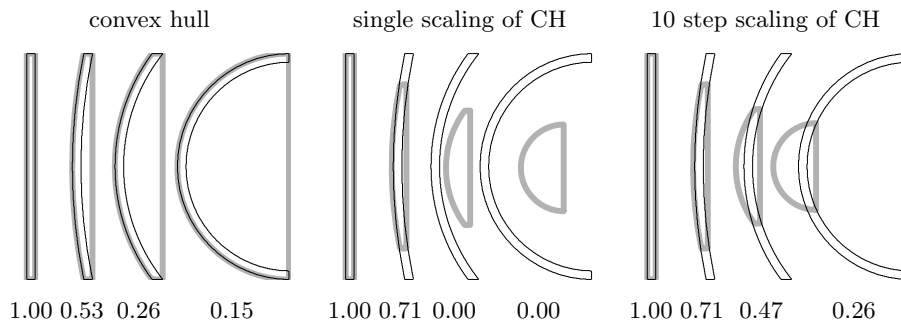


Fig. 4. Convex hull based convexity measure. Four shapes with increased bending are shown, as well as their convex hulls or rescaled convex hulls. Underneath each shape is shown its measured convexity value.

3 Experimental Results

In this section we provide several experiments in order to illustrate the behaviour and effectiveness of the new convexity measures proposed in this paper.

3.1 Lesions

The first experiment is to classify 40 lesions as either benign or malignant melanomas. The lesions have been rated by 14 dermatologists from a four point scale according to their probability of being a melanoma. Lee *et al.* [19] presented this data, and introduced an “overall irregularity index” which they showed provided a Spearman rank correlation of 0.88 against the mean expert rating. Previously we found that convexity measures also performed well on this task. The standard area based convex hull convexity measure C_A had a similar correlation value of 0.888, while both Rosin and Mumford’s [8] convex hull based measures C_P and C_Q achieved the high correlation value of 0.958.




Fig. 5. Lesions rated by experts as having a low (top row) and high (bottom row) probability of being a melanoma. The convex polygons corresponding to C_Y , C_B , C_A^S , C_F and C_{FT} are overlaid.

C_Y	C_B	C_A^S	C_F	C_{FT}
0.938	0.933	0.933	0.446	0.299

Table 1. Spearman rank correlation of the proposed convexity measures against the mean expert rating of the probability of a lesion being a melanoma.

Figure 5 shows two lesions overlaid with the convex polygons. As table 1 shows, most of the new convexity measures proposed in this paper also outperform the overall irregularity index, and come close to the high correlation value achieved by Rosin and Mumford’s C_P and C_Q measures.

3.2 Greebles

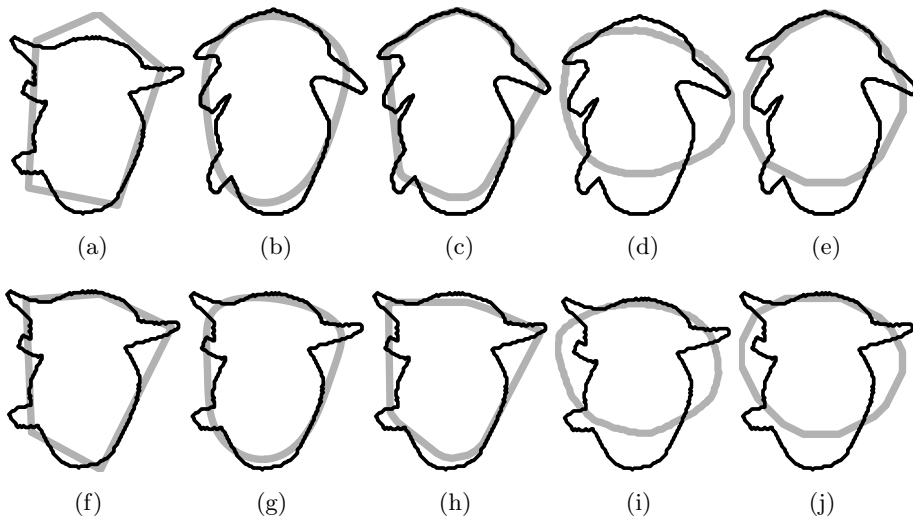


Fig. 6. Two examples of greebles; top row: female, bottom row: male. The convex polygons corresponding to C_Y , C_B , C_A^S , C_F and C_{FT} are overlaid.

C_Q	C_P	C_B	C_Y	C_S	C_A	C_A^S	C_J	C_F	C_{FT}	C_L
68.9	69.0	69.6	73.3	78.7	89.3	92.20	96.5	101.28	103.70	113.6

Table 2. Convexity measures applied to 1137 polygons and ranked left to right by the standard deviation of the greeble ranks.

Next, we replicate an experiment from Rosin and Mumford [8] in which they computed convexity for a set of 1137 polygons containing a variety of shapes as well as 53 “greebles”. The latter are synthetic test objects used as stimuli in psychological tests, and are shapes that have a qualitatively similar appearance

with some minor variations in their four protrusions; two examples are shown in figure 6. The convexity measures are compared by ordering the full set of 1137 polygons according to each of the convexity measures. The standard deviation of the greeble ranks is computed, and low values indicate the effectiveness of a measure since it implies that the measure is stable and consistent across minor variations in shape. Results are shown in table 2, and it can be seen that several of the new convexity measures (in particular C_B) are essentially as effective as the convexity measures based on the robust convex hull and the convex skull (i.e. C_P and C_Q) [8].

3.3 Diatoms




Fig. 7. Outer boundary and interior contours of examples of diatoms from different taxa.

	C_A	C_L	$C_J(B)$	$C_J(I)$	$C_J(B) \& C_J(I)$
convexity	16.83	11.26	16.83	21.41	45.17
convexity and other features	87.75	86.88	87.75	90.59	90.47
	C_Y	C_B	C_A^S	C_F	C_{FT}
convexity	20.79	17.95	21.04	17.33	13.00
convexity and $C_J(I)$	47.52	47.40	45.54	45.67	40.22
convexity and other features	88.61	87.75	87.75	85.77	87.25

Table 3. Classification percentage accuracies for 808 diatoms with 38 classes using convexity plus additional general shape measures.

This experiment uses the data from the ADIAC project. 808 contours were taken from images of diatoms (unicellular algae) which come from 38 taxa (classes) that were manually determined by an expert. Not only the boundary contours, but also each diatom’s ornamentation, which consists of zero or more (mainly open) curve sections in the interior, are available – see figure 7. Previously we have classified this data using several convexity measures both alone and in combination with the following set of descriptors [10]: circularity, ellipticity, rectangularity, triangularity [20] aspect ratio, compactness, convexity, eccentricity, the first four rotation, translation, and scale moment invariants, four rotation, translation, and scale moment invariants [21], the first three affine moment invariants [22]. This experiment has been rerun using a nearest neighbour classifier with Mahalanobis distances and leave-one-out cross validation rather than oblique decision trees (as in [10]), and the new (improved) classification accuracies are shown in the top half of table 3. The lower half of table 3 shows the results of classification using the convexity measures proposed in this paper. It can be seen that the performance of the two sets of convexity measures is similar: while any individual convexity measure has a low classification rate, it is significantly boosted when combined with the (single) interior convexity measure, and is even more substantially increased when the 14 general shape descriptors are added. A further improvement is achieved by combining the general shape descriptors with the interior based convexity measure $C_J(I)$ and the new boundary based convexity measure C_{FT} to obtain 91.58%.

It is interesting to note that although the various convexity measures are nominally measuring the same shape characteristic, in fact they are capturing different aspects of convexity. Therefore when used together the different convexity measures provide independent information that enables their combination to outperform any individual convexity measure. The convexity measures previously tested in Žunić and Rosin [10], namely C_A , C_L , $C_J(B)$ and $C_J(I)$, when combined provide 62.13% accuracy, a three-fold increase on the individual convexity measures. If both the old and new convexity measures are combined,³ namely C_L , $C_J(B)$, $C_J(I)$, C_Y , C_B , C_A^S , C_F and C_{FT} then the classification rate increases to 77.48%.

³ Including C_A with all the other convexity measures caused a small drop in accuracy to 76.61%.

4 Conclusions

In this paper we have described a new approach to compute a family of convexity measures. The procedure involves 1/ generating a convex polygon, 2/ fitting the convex polygon to the input shape by applying a scaling and translation, and 3/ computing the convexity measure based on the input shape and the convex polygon.

Several specific instances of this approach were demonstrated using the convex hull, smoothing, convexification, etc. to generate the convex polygon, while for the convexity measure we used the ratio of the area of overlap of the two shapes to the area of the input polygon. Moreover, many additional possibilities exist for each of these two steps, which further expands the number of convexity measures that can be generated.

The resulting convexity measures have the desirable properties that only convex shapes have measured convexities equal to 1, and that no non-zero area shape should produce a measured convexity equal to 0.

Experiments showed that the proposed convexity measures are effective for classification, and that multiple convexity measures can be combined with each other and with non-convexity measures to improve classification rates.

Acknowledgement

This work is partially supported by the Ministry of Science of the Republic of Serbia, projects OI174008/III044006.

References

1. Neal, F., Russ, J.: *Measuring Shape*. Taylor & Francis (2012)
2. Boxer, L.: Computing deviations from convexity in polygons. *Pattern Recognition Letters* **14** (1993) 163–167
3. Corcoran, P., Mooney, P., Winstanley, A.C.: A convexity measure for open and closed contours. In: *Proc. British Machine Vision Conf.* (2011) 1–11
4. Kakarala, R.: Testing for convexity with Fourier descriptors. *Electronics Letters* **34** (1998) 1392–1393
5. Martin, R., Rosin, P.: Turning shape decision problems into measures. *Int. Journal of Shape Modelling* **10** (2004) 83–113

6. Pao, H., Geiger, D., Rubin, N.: Measuring convexity for figure/ground separation. In: *Int. Conf. Computer Vision*. (1999) 948–955
7. Rahtu, E., Salo, M., Heikkilä, J.: A new convexity measure based on a probabilistic interpretation of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **28** (2006) 1501–1512
8. Rosin, P., Mumford, C.: A symmetric convexity measure. *Computer Vision and Image Understanding* **103** (2006) 101–111
9. Stern, H.: Polygonal entropy: a convexity measure. *Pattern Recognition Letters* **10** (1989) 229–235
10. Žunić, J., Rosin, P.: A new convexity measurement for polygons. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **26** (2004) 923–934
11. Chang, J., Yap, C.: A polynomial solution for the potato-peeling problem. *Discrete Comput. Geom.* **1** (1986) 155–182
12. Grayson, M.: The heat equation shrinks embedded plane curves to round points. *Journal of Differential Geometry* **26** (1987) 285–314
13. Kozera, R., Noakes, L.: Smooth interpolation with cumulative chord cubics. In: *Computer Vision and Graphics*. Volume 32. (2006) 87–94
14. Perez, J., Vidal, E.: Optimum polygonal approximation of digitized curves. *Pattern Recognition Letters* **15** (1994) 743–750
15. Aichholzer, O., Cortés, C., Demaine, E., Dujmović, V., Erickson, J., Meijer, H., Overmars, M., Palop, B., Ramaswami, S., Toussaint, G.: Flipturning polygons. *Discrete & Computational Geometry* **28** (2002) 231–253
16. Erdős, P.: Problem number 3763. *Amer. Math. Monthly* **42** (1935) 627
17. Rosin, P.: Classification of pathological shapes using convexity measures. *Pattern Recognition Letters* **30** (2009) 570–578
18. Press, W., Flannery, B., Teukolsky, S., Vetterling, W.: *Numerical Recipes in C*. Cambridge University Press (1990)
19. Lee, T., D.McLean, Atkins, M.: Irregularity index: A new border irregularity measure for cutaneous lesions. *Medical Image Analysis* **7** (2003) 47–64
20. Rosin, P.: Measuring shape: Ellipticity, rectangularity, and triangularity. *Machine Vision and Applications* **14** (2003) 172–184
21. Sonka, M., Hlavac, V., Boyle, R.: *Image Processing, Analysis, and Machine Vision*. Thomson-Engineering (2007)
22. Flusser, J., Suk, T.: Pattern recognition by affine moment invariants. *Pattern Recognition* **26** (1993) 167–174