Image Shape Classification with the Weighted Euler Curve Transform

Dhanush Giriyan
Arizona State University. United States.

Jessi Cisewski-Kehe
Department of Statistics, University of Wisconsin-Madison, United States.

Brittany Terese Fasy
School of Computing & Dept. of Mathematical Sciences, Montana State University, United States.

Abstract

The weighted Euler curve transform (WECT) was recently introduced as a tool to extract meaningful information from shape data, when the shape is equipped with a weight function. In this extended abstract, we provide an experimental investigation on using the WECT for image classification.

2012 ACM Subject Classification Theory of computation → Computational geometry; Mathematics of computing → Geometric topology

Keywords and phrases Topological data analysis, Euler characteristic, classification, image data

Digital Object Identifier 10.4230/LIPIcs.CVIT.2016.23

Funding The authors are supported by a collaborative grant through the National Science Foundation, NSF DMS 1854336 and DMS 2038556.

1 Introduction

One of the primary goals of topological shape analysis is to summarize the shape of data using topological invariants. The *Euler Characteristic* (EC) is a simple invariant that is easily computed from data. For a filtered topological space, this can be extended to an *Euler Characteristic Curve* (ECC). This extended abstract considers a variant of the ECC for a filtered and weighted simplicial complex, called the *Weighted Euler Curve Transform* (WECT) [6].

2 Background

Given a simplicial complex K and a filter function $f\colon K \to \mathbb{R}$, the ECC is the function $\chi_f\colon \mathbb{R} \to \mathbb{Z}$ that maps $x \in \mathbb{R}$ to the EC of the sublevel set $f^{-1}((-\infty, r])$. A common filter function is the *directional filter function*, which, for K embedded in $\mathbb{R}^2$, is the lower-star filter for a direction $s \in \mathbb{S}^1$. Using an ensemble of directional filter functions, one for each direction in $\mathbb{S}^1$, we obtain a parameterized collection of ECCs, called the *Euler curve transform (ECT)*.[1] No two distinct shapes have the same ECT [9]. Since its introduction, the ECT has developed both in theory and in practice [1, 3, 5–7].

The *WECT* is a generalization of the ECT for weighted complexes [6]. Let K be a simplicial complex in $\mathbb{R}^2$ and let $g\colon K \to \mathbb{R}$ be its weight function (g need not be a filter

[1] The ECT can be defined more generally; however, for brevity, we restrict the definition here to our specific problem setting.

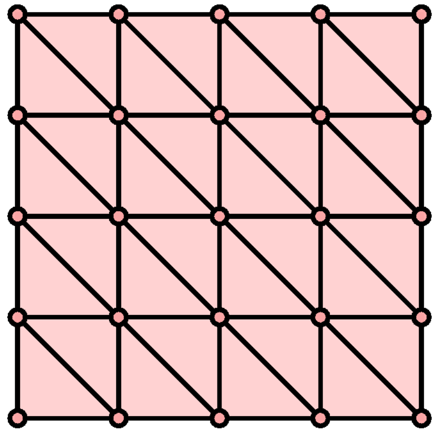

Figure 1 Freudenthal Triangulation of 5×5 image.

function). The WECT for K assigns to each $s \in \mathbb{S}^1$ the function $\omega_{s,g} \colon \mathbb{R} \to \mathbb{Z}$ defined by:

$$\omega_{s,g}(t) = \sum_{d=0}^{2}(-1)^d \sum_{\substack{\sigma \in K_d \\ h_s(\sigma) \leq t}} g(\sigma) = \sum_{\substack{\sigma \in K \\ h_s(\sigma) \leq t}} (-1)^{\dim(\sigma)} g(\sigma),$$

where $h_s \colon K \to \mathbb{R}$ is the lower-star filter in direction s.

3 Experiments

3.1 Data Sets

Synthetic image datasets are generated using the following parameters: shape, pixel intensity, and function extension. The domain of an image is a 64×64 grid, with vertices on an integer lattice, triangulated using the Freudenthal triangulation [4]; see Figure 1. The *shapes* are the domain of the image pixels with positive intensity. Two shapes are considered: the disc (D) and the square (S). The disc D is centered at $(32, 32)$ with a radius of 20 pixels. The square S is also centered at $(32, 32)$ and has an edge length of 35, ensuring that both D and S have approximately the same number of pixels. Let K be the triangulation of either D or S.

Pixel intensities are sampled from Gaussian (G) and uniform (U) density shapes. A bivariate Gaussian density with a peak at the middle of the image and diagonal covariance matrix with standard deviations 100 is used. Points are sampled from this distribution and then assigned to the nearest grid cell to create a 2D histogram. This histogram is rescaled so that the maximum intensity is 255. In the uniform distribution, each pixel intensity is sampled from $[0, 255]$. In what follows, $g \colon K_0 \to \mathbb{R}$ is the function on the vertices of K.

Together, the shape and distribution define four image classes: square-Gaussian (SG), square-uniform (SU), disc-Gaussian (DG), and disc-uniform(DU), with 250 sampled images in each; see Figure 2. Given an image (a function $g \colon K_0 \to \mathbb{R}$), we consider two extensions (functions $F \colon K \to \mathbb{R}$). First, the maximum extension is defined by mapping a simplex to the maximum value on its vertices. Second, the average extension takes average instead of the maximum. For each image class, we test both function extensions, creating eight data sets. We refer to an image, along with its extension to all simplices an *extended image*.

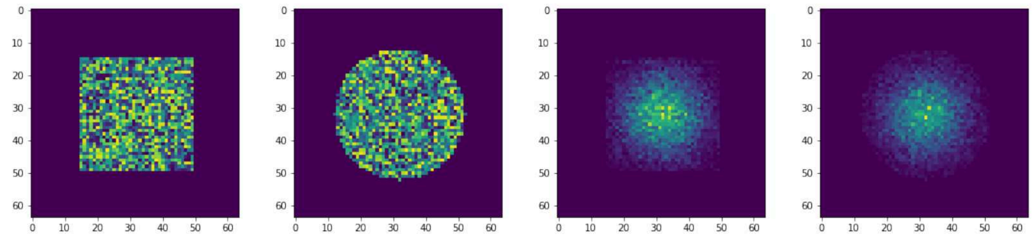

Figure 2 Four image classes.

Table 1 Pairwise classification results with maximum weight function extension.

	SG	SU	DG	DU
SG	0.500 ± 0.000			
SU	1.000 ± 0.000	0.500 ± 0.000		
DG	0.902 ± 0.023	0.998 ± 0.006	0.500 ± 0.000	
DU	1.000 ± 0.000	1.000 ± 0.000	0.998 ± 0.006	0.500 ± 0.000

3.2 Methods

We evaluate the WECT on the task of binary classification. Given two data sets, we first compute the WECT for each extended image. To approximate the WECT (a function over $\mathbb{S}^1$), we sample eight equally-spaced directions from $\mathbb{S}^1$, which was sufficient given the simplicity of the image features. To discretize each WECC, 50 threshold values are sampled, thus each image results in a vectorized WECT of length 400. Next, we train a Support Vector Classifier (SVC) [2] separately for each binary classification model. The set of WECTs is partitioned into training and test sets, with an 80/20 train-test split. All SVCs in this work use a regularization parameter of value 20. This was implemented in Python, using data structures from the Dionysus 2 library [8]. Experiments were conducted on an Intel Xeon E5-2680 (128 GB RAM).

3.3 Results and Conclusion

For each pair of image classes, we compute the WECTs for both the maximum and the average function extensions. Then, we fit SVC models on 10 unique collections of 250 images of each class and report the average accuracies and the standard errors. Table 1 and Table 2 provide the percentage of correctly classified test images.

In the simple data sets that we considered, we found the WECT to be suitable for the task of binary classification. The weakest classification result was comparing the SG and DG classes under the average weight function extension. Because Gaussian density decreases toward zero in the extremes, the corners of SG have weights that are close to zero making

Table 2 Pairwise classification results with average weight function extension.

	SG	SU	DG	DU
SG	0.500 ± 0.000			
SU	1.000 ± 0.000	0.500 ± 0.000		
DG	0.830 ± 0.048	1.000 ± 0.000	0.500 ± 0.000	
DU	1.000 ± 0.000	0.948 ± 0.034	1.000 ± 0.000	0.500 ± 0.000

SG resemble the DG. Why the maximum weight function extension performed better than the average is not immediately clear. We plan to extend the set of shapes, distributions, and function extensions in order to explore this observation.

References

1. Eric Berry, Yen-Chi Chen, Jessi Cisewski-Kehe, and Brittany Terese Fasy. Functional summaries of persistence diagrams. *Journal of Applied and Computational Topology*, 4(2):211–262, 2020.
2. Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 20:273–297, 1995.
3. Lorin Crawford, Anthea Monod, Andrew X Chen, Sayan Mukherjee, and Raúl Rabadán. Predicting clinical outcomes in glioblastoma: an application of topological and functional data analysis. *Journal of the American Statistical Association*, 115(531):1139–1150, 2020.
4. Hans Freudenthal. Simplizialzerlegungen von beschrankter flachheit. *Annals of Mathematics*, pages 580–582, 1942.
5. Robert Ghrist, Rachel Levanger, and Huy Mai. Persistent homology and Euler integral transforms. *Journal of Applied and Computational Topology*, 2(1-2):55–60, 2018.
6. Qitong Jiang, Sebastian Kurtek, and Tom Needham. The weighted Euler curve transform for shape and image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 844–845, 2020.
7. Lewis Marsh, Felix Y Zhou, Xiao Quin, Xin Lu, Helen M Byrne, and Heather A Harrington. Detecting temporal shape changes with the euler characteristic transform, 2022. arXiv preprint arXiv:2212.10883.
8. Dmitriy Morozov. Dionysus 2. `https://mrzv.org/software/dionysus2/`.
9. Katharine Turner, Sayan Mukherjee, and Doug M. Boyer. Persistent homology transform for modeling shapes and surfaces. *Information and Inference: A Journal of the IMA*, 3(4):310–344, 2014.