The Topological "Shape" of Brexit

Bernadette J. Stolz, Heather A. Harrington, and Mason A. Porter[‡]

October 30, 2016

Abstract

Persistent homology is a method from computational algebraic topology that can be used to study the "shape" of data. We illustrate two filtrations — the weight rank clique filtration and the Vietoris–Rips (VR) filtration — that are commonly used in persistent homology, and we apply these filtrations to a pair of data sets that are both related to the 2016 European Union "Brexit" referendum in the United Kingdom. These examples consider a topical situation and give useful illustrations of the strengths and weaknesses of these methods.

1 Introduction

Persistent homology (PH) [5, 8–11, 16] is a method from computational algebraic topology that can be used to study the "shape" of data. The studied data typically consists of networks or high-dimensional point clouds. PH is concerned with topological invariants, such as connectedness and holes in high-dimensional objects, over a range of different scales. Topological features that persist over many scales are often considered to be significant for data shape, and features that persist only over a small range of scales are construed as noise. PH has been applied to study an increasingly diverse set of phenomena in a large variety of subjects. A small set of the myriad examples include granular materials [14], protein binding sites [13], brain-artery trees [4], and functional brain networks [7, 12, 18].

^{*}Mathematical Institute, University of Oxford

[†]Mathematical Institute, University of Oxford

[‡]Department of Mathematics, UCLA; Mathematical Institute, University of Oxford; and CAB-DyN Complexity Centre, University of Oxford

In the present article, we illustrate two filtrations that are commonly used in PH: the weight rank clique filtration (WRCF) [17] and the Vietoris–Rips (VR) filtration [10]. We apply these filtrations to a pair of data sets that are both related to the recent European Union (EU) referendum in the United Kingdom (UK) [3]. The first data set consists of a weighted network based on the year that each European country joined the EU and the second data set consists of two point clouds that were constructed from voting results of the UK referendum. We give a detailed description of both data sets in Subsections 2.1 and 3.1 respectively.

2 The EU Network

2.1 Constructing a Network

We define a network of EU countries based on the geographical location of the countries and the year in which they joined the European Union. Each current EU country (as of June 2016) is a node in the network. We connect two countries by an edge if they are neighbors via a common border (either in Europe or via overseas territories), a bridge (Denmark–Sweden), or a tunnel (UK–France). We also connect Malta to Italy and connect Cyprus to Greece because of their geographical proximity. We show the resulting network in Fig. 1. We define the edge weight $w_{i,j}$ between countries *i* and *j* to be the later of the two years in which those two countries joined the EU. That is,

 $w_{i,j} = \max \{ \text{year in which country } i \text{ joined the EU}, \text{year in which country } j \text{ joined the EU} \}.$

2.2 Weight Rank Clique Filtration

We use a WRCF [15, 17] on the above weighted network to obtain a sequence of embedded graphs, on which we examine PH. One constructs a WRCF as follows:

- 1. Define filtration step 0 as the set of all nodes.
- 2. Rank all edge weights $\{\nu_1, \ldots, \nu_{\text{end}}\}$, with $\nu_1 = \max_{1 \le i,j \le 28} w_{i,j}$ and $\nu_{\text{end}} = \min_{1 \le i,j \le 28} w_{i,j}$.
- 3. In filtration step t, threshold the graph at weight ν_t to create a binary (i.e., unweighted) graph. After thresholding, edges with weight at least ν_t are present, and edges with smaller weights are absent.



Figure 1: Network of EU countries. We connect two countries by an edge if they are considered neighbors via a border (either in Europe or abroad), a bridge, or a tunnel. We define the edge weight between two countries to be the later of the two years in which they joined the EU. [We obtained the employed map from Wikimedia [1].]

4. Find all maximal c-cliques for $c \in \mathbb{N}$, and define them to be c-simplices.

We perform the WRCF on the EU network before the EU referendum and on a network based on the future EU after the referendum results are implemented. In this future EU, we have removed the two edges between the UK and the other EU countries as if they never existed in the first place, but we keep all nodes as before. We run the filtration from years 1953–2016 for the current EU network and 2016 onwards for the hypothetical future EU network. To visualize the results of the PH computations, we show barcode diagrams [10] in Fig. 2.

In dimension 0, PH detects the connected components of the network in every filtration step. In our case, every bar in the 0-dimensional barcode represents one EU country. The countries merge as the EU grows over time until there is one large EU



Figure 2: Barcodes for dimensions 0 (connected components) and 1 (loops) obtained from a weight rank clique filtration performed on current and hypothetical future EU networks. See Table 1 for a listing of EU enlargements.

Founding countries	Belgium, Germany, France, Italy, Luxembourg, and The
	Netherlands
First enlargement	Denmark, Ireland, and the United Kingdom
Fourth enlargement	Austria, Finland, and Sweden
2004 enlargement	Czech Republic, Estonia, Cyprus, Latvia, Lithuania, Hun-
	gary, Malta, Poland, Slovakia, and Slovenia

Table 1: EU enlargements listed in the barcodes.

component in 2013. Because of how we construct the edge weights in the network, we note that although Greece joined the EU in 1981, it only joins the large EU component in 2007 when its neighbor Bulgaria joined the EU. Similarly, in the barcode after the 2016 EU referendum, both the UK and Ireland are in separate components because the latter (in our hypothetical "post-Brexit" world) has lost its connection to the main EU component via the UK even though it is still a member of the EU.

In dimension 1, PH detects loops in the network. In the case of the WRCF these loops consist of at least four edges. In other words, edges between three pairwise neighboring countries are not registered as loops, but the loop created by edges between Switzerland's neighboring countries of leads to a bar in the 1-dimensional barcode starting in 1995 when Austria joins the EU. We remark that this is also the only occurrence of a loop in this network.

3 The Voting Point Cloud

3.1 Data Acquisition

We obtained UK voting data via the website "Number Cruncher Politics" [2]. We extracted coordinates for one city per voting district by combining python scripts and the Google Maps API. We thereby construct two point clouds: one, which we call the *leave point cloud*, contains the coordinates of cities in voting districts that voted to leave the EU; the other, which we call the *remain point cloud*, contains the coordinates of cities in voting districts that point clouds in Voting districts that voted to stay in the EU. We show both point clouds in Fig. 3.



Figure 3: Point clouds based on the 2016 EU referendum voting result in the UK. We show the coordinates of "leave" districts in blue and the coordinates of "remain" districts in red.

3.2 Vietoris–Rips Filtration

We apply a VR filtration [10] to our point-cloud data. One constructs a VR filtration as follows:

- 1. Choose a sequence ϵ of increasing distances: $\epsilon = \{\epsilon_1, \ldots, \epsilon_n\}$.
- 2. In the *i*th filtration step, define k-simplices using unordered (k+1)-tuples whose pairwise distance is at most ϵ_i .



Figure 4: (Top two rows) Barcodes for dimensions 0 and 1 from a Vietoris–Rips filtration on (left) the leave point cloud and (right) the remain point cloud. (Bottom row) UK referendum voting map including the loops in dimension 1 of the Vietoris–Rips filtration.

We show the resulting barcodes in Fig. 4.

Our calculations reveal multiple loops in dimension 1 of the leave district barcodes that are more persistent than other features of the barcode. In particular, London and the region around Oxfordshire, Reading, and The Cotswolds stand out as two holes in the leave district point cloud, as they voted in the opposite way as their surrounding voting districts. However, we also note that the VR filtration detects a persistent hole around Manchester, where there ought to be a non-persistent hole. This is due to the fact that the points representing the surrounding voting districts are located far away from Manchester, and this leads to an artificially large hole. In the remain point cloud, we find a persistent hole around Birmingham in dimension 1.

4 Conclusions

Our two examples illustrate the strengths and limitations of the applied PH methods. For the weighted rank clique filtration on the country network, we find a single large (and relevant) network loop around Switzerland in dimension 1. The WRCF in dimension 0 gives a good idea of the years when countries join the EU, but the results also depend heavily on the neighbors in the network, which leads to the method suggesting erroneously that Greece joined the EU later than it actually did and Ireland being illustrated as a separate component in the EU network after the UK leaves the EU.

The Vietoris–Rips filtration points towards large groups of regions that have voted differently than their surrounding regions, but we also find that the holes in dimension 1 depend on the choice of points to represent the voting districts.

The two methods that we have illustrated are of course not the only topological methods that one could apply to the data set. For example, one could build a single filtration for the EU network over all years that allows edges to disappear from one filtration step to the next by using zigzag persistent homology [6, 19], or one could focus on counting the number of persistent loops in the point-cloud data and compare this to point clouds of other UK votes to see whether there are recurring vote shapes (which could represent recurring voting patterns).

5 Acknowledgements

We thank Matt Singh from Number Cruncher Politics for providing us with the referendum voting data. We also thank Joshua Bull for his help with Python webcrawling and Jared Tanner for helpful discussions. BJS also gratefully acknowledges the EP-SRC and MRC (grant number EP/G037280/1), and she thanks F. Hoffmann–La Roche AG for funding her doctoral studies. HAH acknowledges funding from EPSRC Fellowship EP/K041096/1.

References

 [1] European union. Available at https://commons.wikimedia.org/wiki/ European_Union, accessed 13 August 2016.

- [2] Number cruncher politics. Available at http://www.ncpolitics.uk, accessed 13 August 2016.
- [3] United kingdom european union membership referendum, 2016. Available at https://en.wikipedia.org/wiki/United_Kingdom_European_Union_ membership_referendum, 2016, accessed 13 August 2016.
- [4] Paul Bendich, J. S. Marron, Ezra Miller, Alex Pieloch, and Sean Skwerer. Persistent homology analysis of brain artery trees. arXiv:1411.6652, 2014.
- [5] Gunnar Carlsson. Topology and data. Bulletin of the American Mathematical Society, 46:255–308, 2009.
- [6] Gunnar Carlsson and Vin de Silva. Zigzag persistence. Foundations of Computational Mathematics, 10:367–405, 2010.
- [7] Carina Curto. What can topology tell us about the neural code? arXiv:1605.01905, 2016.
- [8] Herbert Edelsbrunner and John L. Harer. Persistent homology A survey. Contemporary mathematics, 453:257–282, 2008.
- [9] Herbert Edelsbrunner and John L. Harer. Computational Topology. American Mathematical Society, Providence R. I., 2010.
- [10] Robert Ghrist. Barcodes: The persistent topology of data. Bulletin of the American Mathematical Society, 45:61–75, 2008.
- [11] Robert Ghrist. Elementary Applied Topology. Createspace, 1.0 edition, 2014. Available at https://www.math.upenn.edu/ ghrist/notes.html.
- [12] Chad Giusti, Robert Ghrist, and Danielle S. Bassett. Two's company and three (or more) is a simplex. Journal of Computational Neuroscience, 41(1):1–14, 2016.
- [13] Violeta Kovacev-Nikolic, Peter Bubenik, Dragan Nikolic, and Giseon Heo. Using persistent homology and dynamical distances to analyze protein binding. *Statistical Applications in Genetics and Molecular Biology*, 15(1):1–27, 2016.
- [14] M. Kramár, A. Goullet, L. Kondic, and K. Mischaikow. Persistence of force networks in compressed granular media. *Physical Review E*, 87:042207, 2013.

- [15] Hyekyoung Lee, Hyejin Kang, Moo K. Chung, Bung-Nyun Kim, and Dong Soo Lee. Weighted functional brain network modeling via network filtration. In NIPS Workshop on Algebraic Topology and Machine Learning, 2012.
- [16] Nina Otter, Mason A. Porter, Ulrike Tullmann, Peter Grindrod, and Heather A. Harrington. A roadmap for the computation of persistent homology. arXiv:1506.08903, 2015.
- [17] Giovanni Petri, Martina Scolamiero, Irene Donato, and Francesco Vaccarino. Topological strata of weighted complex networks. *PLoS ONE*, 8(6):e66505, 2013.
- [18] Bernadette J Stolz, Heather A Harrington, and Mason A Porter. Persistent homology of time-dependent functional networks constructed from coupled time series. arXiv:1605.00562, 2016.
- [19] Andrew Tausz and Gunnar Carlsson. Applications of zigzag persistence to topological data analysis. arXiv:1108.3545, 2011.