

PAPER • OPEN ACCESS

Augmentations of Forman's Ricci curvature and their applications in community detection

To cite this article: Lukas Fesser *et al* 2024 *J. Phys. Complex.* **5** 035010

View the [article online](#) for updates and enhancements.

You may also like

- [Comment on 'Random sequential packing in square cellular structures'](#)
G C Barker and M J Grimson
- [Physical computing: a category theoretic perspective on physical computation and system compositionality](#)
Nima Dehghani and Gianluca Caterina
- [Modeling chlorine residuals in urban water distribution networks \(Al-Ashrafieh – Homs\)](#)
Ruba D Alsaeed, Bassam Alaji and Lina Khouri



OPEN ACCESS

RECEIVED

12 December 2023

ACCEPTED FOR PUBLICATION

17 July 2024

PUBLISHED

12 August 2024

Original Content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.



PAPER

Augmentations of Forman's Ricci curvature and their applications in community detection

Lukas Fesser^{1,4,*}, Sergio Serrano de Haro Iváñez^{2,4} , Karel Devriendt³, Melanie Weber¹ and Renaud Lambiotte²

¹ Harvard University, Cambridge, MA, United States of America

² University of Oxford, Oxford, United Kingdom

³ Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

⁴ Joint first authors.

* Author to whom any correspondence should be addressed.

E-mail: lukas_fesser@fas.harvard.edu

Keywords: discrete curvature, network analysis, community detection

Abstract

The notion of curvature on graphs has recently gained traction in the networks community, with the Ollivier–Ricci curvature (ORC) in particular being used for several tasks in network analysis, such as community detection. In this work, we choose a different approach and study augmentations of the discretization of the Ricci curvature proposed by Forman (AFRC). We empirically and theoretically investigate its relation to the ORC and the un-augmented Forman–Ricci curvature. In particular, we provide evidence that the AFRC frequently gives sufficient insight into the structure of a network to be used for community detection, and therefore provides a computationally cheaper alternative to previous ORC-based methods. Our novel AFRC-based community detection algorithm is competitive with an ORC-based approach.

1. Introduction

Curvature is a key notion in differential geometry, and Ricci curvature is one of several ways of formalizing it. Intuitively, Ricci curvature measures the deformation of a small ball as we push it along a geodesic (the shortest path between two points), and quantifies how much a shape locally differs from standard, flat Euclidean space. As Ricci curvature encapsulates essential information in differential geometry and related fields (e.g. it is a term of Einstein's field equations in general relativity), there have been multiple attempts to profitably extend its definition to other domains.

The structural insight gained via discrete Ricci curvature has proven valuable for applications in Network Science and Machine Learning. Discrete Ricci curvature relates to many classical network characteristics and network analysis tasks [1] and, as a result, has been applied for the analysis of network data in many domains, including biological networks [2, 3], chemical networks [4] and financial transaction networks [5], among others. In Machine Learning, discrete curvature has been utilized to characterize and mitigate the oversquashing effect in Graph Neural Networks [6] and to identify embedding spaces in Representation Learning [7].

The two curvature versions studied in this paper are respectively due to Ollivier [8, 9] and Forman [10]. Ollivier–Ricci curvature (ORC) is a direct discretization of the original definition, substituting spheres and geodesics by, respectively, node neighborhoods and shortest edge-paths. The construction of Forman–Ricci curvature (FRC) takes a different approach: it follows from Bochner's method [11–13], which in the differential case is used to write the Laplacian as the sum of a Laplacian-like term and a term that depends on the Ricci curvature. Forman's work consists in discretizing both the Laplacian and Bochner's method, which produces a term that, by analogy, can be interpreted as a discrete Ricci curvature. Thus, ORC and FRC are two measures with solid geometric foundations yet presenting significant differences, which makes them two potential ways of gaining a better understanding of the geometry of networks.

Here, we focus on geometric approaches to community detection, which have previously received much attention [14–16], particularly the use of ORC [17–21]. Community structure, i.e. the mesoscale organization of a network into groups of nodes with similar connections, is widely observed in real-world networks and its detection is a classical problem in the study of networks and graphs [22, 23]. Community detection using ORC relies on the observation that positively curved edges tend to be well connected, and hence naturally lie inside communities. Negatively curved edges, on the other hand, can be interpreted as bridges between communities. ORC-based community detection algorithms have been found to perform well on both synthetic and real-world networks. These methods are, however, computationally expensive as the definition of ORC involves the computation of network flows of shortest paths via the Wasserstein distance, which has cubic complexity and is often approximated via the (quadratic) Sinkhorn algorithm [24].

One advantage of FRC over ORC is its significantly smaller computational complexity: in its simplest form, the FRC of an edge only depends on the degrees of its end points. This is often too simple of a measure in practice, and there have been several proposals to enrich the simple FRC by taking into account the cycles that contain an edge [25–27]⁴. In this article, we consider the augmented FRC (AFRC) introduced in [26] and study how different augmentations affect the balance between expressivity and computational complexity in the context of community detection. Based on this evaluation, we identify classes of graphs in which AFRC-based community detection can be recommended compared with ORC-based detection.

The main contributions in this article are

- We make a detailed empirical and theoretical investigation of the augmentations of FRC.
- We discuss the relation of AFRC to ORC and FRC.
- We provide evidence that AFRC distinguishes communities in a manner comparable to ORC for many model and real-world networks.
- We introduce a novel AFRC-based community detection algorithm, with comparable performance to existing ORC-based methods while being computationally cheaper.
- We develop the first codebase for fast and efficient computations with augmentations of the FRC.

The rest of this paper is structured as follows: section 2 formally introduces ORC, FRC, its augmentations, and possible approximations. Section 3 studies how the different curvatures are related in real-world and artificial networks. Section 4 studies the curvature gap—a quantitative indicator for how well a curvature can distinguish between edges within and between communities—for the ORC, FRC, and AFRC. Section 5 introduces our AFRC-based community detection algorithm and demonstrates its performance on several model and real-world networks. Section 6 concludes the article with a summary of the main findings.

2. Curvature definitions

Ricci curvature is a central tool in differential geometry that characterizes the local geometric properties of geodesic spaces by relating the local rate of volume growth to geodesic dispersion. Several discretizations of Ricci curvature have been proposed, including notions by Forman [10], Ollivier [8, 9], Erbar and Maas [29], Bakry *et al* [30], and Devriendt and Lambiotte [31]. In this paper, we focus on edge-level Ollivier and Forman curvature. For the latter, we analyze *augmented* Forman curvature, which explicitly encodes the impact of cycles that contain the respective edge. Variants of augmented Forman curvature have been previously studied in [25, 27, 32], among others.

Notation. We reserve the letters v and u to denote the vertices (equiv., nodes) of a network (equiv., graph), and the symbols e and e' to denote its edges. A node u is a neighbor of v if (v, u) is an edge of the graph. A cycle (of length r) is a sequence of different vertices $\gamma = v_1 v_2 \dots v_r$ such that all cyclically consecutive vertices are connected by an edge in the graph.

2.1. ORC

Ollivier introduced a notion of discrete Ricci curvature for metric spaces in [8]; this setting includes graphs with the shortest-path metric between vertices. Aside from a metric, Ollivier's curvature also requires a choice of probability distributions μ_v —one for each vertex v in the case of a graph—which can be thought of as open balls around points in the space. A common choice for graphs is the uniform distribution over the neighbors $\text{neigh}(v)$ of a vertex [19, 21]

$$\mu_v(u) := \begin{cases} \frac{1}{\deg(v)} & \text{if } u \in \text{neigh}(v) \\ 0 & \text{else} \end{cases},$$

⁴ This operation is reminiscent of measures for the density around an edge, which have been used in Network Science, e.g. to test the strength of weak ties [28].

where $\deg(v) = |\text{neigh}(v)|$ is the number of neighbors of v . With this data, the *Ollivier–Ricci curvature* (ORC) $\mathcal{O}(e)$ of an edge $e = (v_1, v_2)$ is defined as

$$\mathcal{O}(e) := 1 - \frac{W_1(\mu_{v_1}, \mu_{v_2})}{d(v_1, v_2)},$$

where d is the shortest-path distance between vertices (in this case $= 1$) and W_1 denotes the Wasserstein distance (also called earth mover's distance in discrete settings) between measures on a metric space; see [8, 9] for details.

Geometrically, the ORC compares the distance between the end points of an edge to the distance between the balls centered at said end points. If v_1 and v_2 share many neighbors or have many edges between their neighborhoods, the distance between μ_{v_1} and μ_{v_2} will be small and so $\mathcal{O}(e) \approx 1$. On the other hand, if the neighborhoods of v_1 and v_2 are poorly connected, the distance between μ_{v_1} and μ_{v_2} will be large and $\mathcal{O}(e)$ will be small and potentially negative. In the context of network community structure, this explains why we expect larger ORC within than between communities: the neighborhoods of two vertices at the ends of an edge are better connected if the edge is contained in a single community. In particular, the threshold value 0 has been previously used in ORC-based community detection to discriminate between inter- and intra-community edges, see [19] for a more detailed discussion.

2.2. FRC and augmentations

Forman developed a notion of discrete Ricci curvature for weighted CW complexes in [10] in the context of algebraic topology. A CW complex is a type of topological space that includes undirected graphs as a one-dimensional case. A direct translation of Forman's definition to unweighted graphs without multi-edges, as derived in [33], gives the *Forman–Ricci curvature* (FRC) $\mathcal{F}(e)$ of an edge $e = (v_1, v_2)$

$$\mathcal{F}(e) := 4 - \deg(v_1) - \deg(v_2). \quad (1)$$

While this definition is well-supported by Forman's work and is easy to calculate, it is too simple a measure for many practical and theoretical purposes. For instance, $\mathcal{F}(e)$ does not depend on the number of triangles in which e is contained, which is an important local geometric property of graphs [34].

To overcome the limitations of the simple formula for FRC, definition (1) can be extended to take into account more of the graph structure around a given edge. The idea is to start from the graph and add solid 2D faces in the spaces enclosed by cycles; see figure 1 for an illustration. This 'augmented graph' is still a CW complex—the new faces are two-dimensional cells with one-dimensional boundaries given by the cycles—and one can apply Forman's original definition to get a discrete edge curvature that also depends on the cycles that contain an edge. The following formula for the *augmented Forman–Ricci curvature* (AFRC) $\mathcal{AF}(e)$ of an edge e was derived in [26] based on Forman's definition [10, theorem 2.2]:

$$\mathcal{AF}(e) = 2 + \Gamma_{ee} - \sum_{e' \sim e} |\Gamma_{ee'} + \Gamma_{e'e} - 1| - \sum_{e' \not\sim e} |\Gamma_{ee'} - \Gamma_{e'e}|, \quad (2)$$

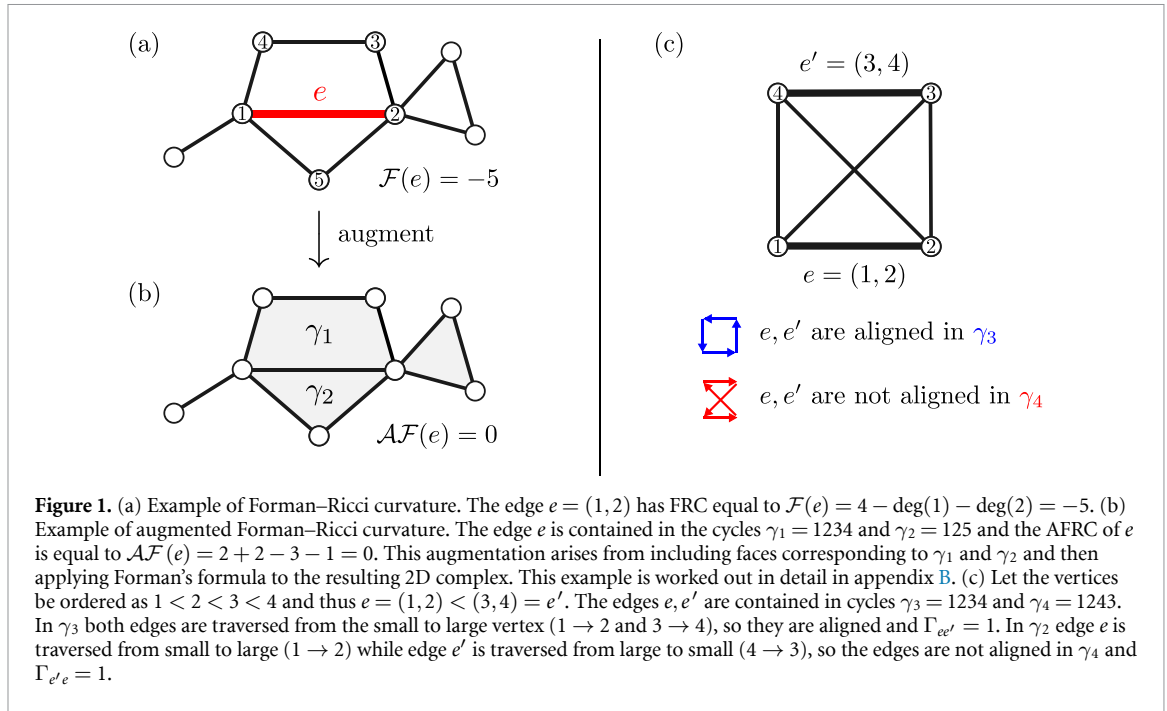
where $e \sim e'$ (resp. $e \not\sim e'$) denotes two distinct edges that share (resp. do not share) a vertex, and where Γ is a matrix (defined in (3) below) of size equal to the number of edges and whose entries count the cycles that contain e and e' . The counting goes as follows: fix an arbitrary order on the vertices of the graph. Two edges e, e' are called *aligned* in a cycle γ if the cycle traverses both edges from smallest to largest vertex (with respect to the order) or both from largest to smallest vertex. The entries of Γ are then defined as

$$\Gamma_{ee'} := \begin{cases} \# \text{ cycles in which } e \text{ and } e' \text{ are aligned} & (\text{if } e < e') \\ \# \text{ cycles in which } e \text{ and } e' \text{ are not aligned} & (\text{if } e > e') \end{cases}$$

$$\Gamma_{ee} := \# \text{ cycles that contain } e. \quad (3)$$

Here, the order $e < e'$ on edges follows from the order on the vertices⁵, and we note that the value of $\mathcal{AF}(e)$ is independent of the choice of vertex order. Figure 1 illustrates (non-)alignment of edges in a cycle and the resulting distinction between $\Gamma_{ee'}$ and $\Gamma_{e'e}$. This distinction is necessary when translating Forman's definition, since it relies on *oriented* 2D faces. We note that when calculating \mathcal{AF} for all edges, the count matrix Γ only needs to be constructed once, by iterating over all cycles in the graph. Appendix B contains a detailed example with the above definitions.

⁵ If we think of an edge as a two-letter word with letters in the (ordered) alphabet of vertices, the induced order on the edges is simply the familiar alphabetical order on the words. See appendix B for an example.



In definition (2) we may choose to only take a subset of all cycles into account; in particular, we will consider the following class of curvatures:

$$\mathcal{AF}_n(e) := \mathcal{AF}(e) \text{ only taking into account cycles of length } \leq n.$$

Restricting AFRC to cycles up to a certain length is supported by the idea that Ricci curvature should be a local property. This is also satisfied by other notions of discrete curvature such as ORC and Bakry–Émery curvature, which only depend on cycles of length $n \leq 5$. In this article we mainly work with \mathcal{AF}_n with $n \in \{3, 4\}$, which are the first augmentations of FRC that give a non-trivial improvement over FRC in a number of applications while still being computationally efficient; in addition to being more expensive, we have found \mathcal{AF}_5 to be less effective in community detection—see appendix figures 8 and 9. We believe the poorer performance is due to the higher number of 5-cycles that contain between-community edges, which blurs the curvature difference between inter- and intra-community edges.

For certain choices of cycles, definition (2) for the AFRC reduces to known extensions of FRC. For instance, the curvature $\mathcal{F}(e) + 3 \cdot \#\{3\text{-cycles} \ni e\} = \mathcal{AF}_3(e)$ was proposed in [25] as a variant of Forman’s curvature in graphs and networks. In [27, 35], it was observed that in graphs where any two cycles share at most one edge—these graphs are called quasi-convex in [10] and independent short-cycle graphs in [35]—, Forman’s curvature simplifies to $\mathcal{F}(e) + \sum_{n \geq 3} (6 - n) \cdot \#\{n\text{-cycles} \ni e\}$. This expression can provide a simple estimate for edges whose neighborhoods are expected to be ‘almost quasi-convex’, but in general it differs from AFRC. In [32], Jost and Münch studied different choices of weighted cycles when augmenting a graph.

Finally, we present a rewriting of equation (2) as an explicit augmentation of FRC:

$$\mathcal{AF}(e) = \mathcal{F}(e) - \Gamma_{ee} + 2 \cdot \sum_{e' \sim e} \min\{\Gamma_{ee'} + \Gamma_{e'e}, 1\} - \sum_{e' \not\sim e} |\Gamma_{ee'} - \Gamma_{e'e}|. \quad (4)$$

Here the second term again counts the number of cycles that contain e , the third term counts the number of edges $e' \sim e$ that share a cycle with e and the fourth term again counts the aligned and unaligned cycles through e and $e' \not\sim e$.

Equation (4) reveals how the addition of cycles affects the value of AFRC compared to FRC. Including cycles of length 3, we find that $\mathcal{AF}_3(e) = \mathcal{F}(e) + 3 \cdot \#\{3\text{-cycles} \ni e\}$ and thus that AFRC is higher for edges whose end points share many neighbors. In networks with community structure, we thus expect \mathcal{AF}_3 to be large within communities and small (possibly negative) between communities, similar to ORC. Including cycles of length 4 and above, however, we expect a different behavior. Now the positive term in (4) due to the neighbors of e is limited by the degrees of the end points of e , while the negative terms (Γ_{ee} and the sum over non-neighbors) can become increasingly negative as more cycles are included. As a result, we can expect edges within communities to have lower $\mathcal{AF}_{n \geq 4}$ values compared to edges between, opposite to

Table 1. Comparison of Ollivier–Ricci curvature (OR) with Forman–Ricci curvature (FR) and augmentations of the Forman–Ricci curvature (\mathcal{AF}_3 and \mathcal{AF}_4) of edges in artificial networks. In this table, we list the Pearson correlation between the edge curvatures.

Network	OR, FR	OR, \mathcal{AF}_3	OR, \mathcal{AF}_4	FR, \mathcal{AF}_3	FR, \mathcal{AF}_4
ER(1000, 0.003)	0.84	0.87	0.87	0.99	0.98
ER(1000, 0.007)	0.42	0.48	0.59	0.98	0.94
ER(1000, 0.01)	−0.01	0.12	0.30	0.98	0.89
SBM(10, 5, 0.7, 0.05)	0.27	0.79	0.76	0.64	0.59
SBM(10, 10, 0.7, 0.05)	0.07	0.83	−0.66	0.49	0.31
SBM(10, 15, 0.7, 0.05)	0.08	0.90	−0.84	0.28	0.19
SBM(10, 20, 0.7, 0.05)	−0.03	0.90	−0.91	0.33	0.19
BG(50, 0.03)	0.74	0.74	0.84	1.0	0.91
BG(50, 0.07)	0.01	0.01	0.58	1.0	0.69
BG(50, 0.1)	0.29	0.29	0.27	1.0	0.79
HBG(50, 0.5, 0.05)	−0.01	−0.01	−0.72	1.0	0.55
HBG(50, 0.5, 0.1)	−0.11	−0.11	−0.25	1.0	0.61
HBG(50, 0.5, 0.15)	−0.01	−0.01	0.00	1.0	0.74

what is observed for ORC and \mathcal{AF}_3 ; in later experiments on networks with community structure, we confirm this analysis. Finally, we remark that un-augmented FRC (1) does not distinguish edges within and between communities when different communities have the same typical degrees.

We note that, to the best of our knowledge, there is no natural threshold to discriminate edges according to their AFRC (unlike the curvature value 0 for ORC). In fact, our experiments show that the potential curvature threshold varies drastically between networks—see e.g. figure 3; we discuss how to choose an adequate case-by-case value in section 5.

3. Correlations between curvatures

Since ORC and AFRC are discretizations of the same smooth concept, their comparison is a natural and important question.

On the theoretical side, [35] found evidence that ORC and \mathcal{AF}_3 are equivalent in quasi-convex graphs and [32] found that the maximal Forman curvature over all weighted augmentations of a graph is equal to ORC. In general, the relation between different notions of discrete curvature remains an important open question.

On the experimental side, Samal *et al* [25] extensively studied the correlation between ORC and both FRC and \mathcal{AF}_3 in model and real-world networks. In their results, ORC tended to greatly correlate with both curvatures, and it generally correlated higher with \mathcal{AF}_3 . We extend this study here by including correlations of these curvatures with \mathcal{AF}_4 . Tables 1 and 2 show the Pearson correlation coefficient between the different curvature measures in, respectively, model and real-world networks.

Model networks: We study four different network models: the Erdős–Rényi (ER) graph, which is a completely homogeneous model and should not exhibit community structure; a bipartite version (BG), which does not contain any 3-cycles and thus $\mathcal{F}(e) = \mathcal{AF}_3(e)$ for all edges; and the stochastic block model (SBM) and hierarchical bipartite graph (HBG), which do have ground-truth communities. Details on these network models can be found in appendix A.

Table 1 shows that the FRC vs \mathcal{AF}_3 and \mathcal{AF}_4 correlations exhibit the expected behavior for model networks: the fewer edges—and thus cycles—present in a network, the more similar the curvatures are (as they are equal in a network with no cycles). In particular, note the correlation of 1 between FRC and \mathcal{AF}_3 for the bipartite networks—as expected from networks with no 3-cycles. The only outlier of this behavior is the HBG, as the correlation FRC vs \mathcal{AF}_4 seems to increase with the number of edges.

Regarding ORC vs the different versions of AFRC, all exhibit a similar behavior on the ER and BG graphs, being quite correlated in sparser networks and losing this correlation as the connection probability p increases—although in general ORC appears best correlated to \mathcal{AF}_4 . This same behavior can be observed in the SBM for the ORC vs FRC correlation, whereas the ORC vs \mathcal{AF}_3 and \mathcal{AF}_4 are notably distinct. We observe how the correlation ORC vs \mathcal{AF}_3 increases positively with the community size, whereas ORC vs \mathcal{AF}_4 grows negatively with it. This behavior agrees with our discussion at the end of section 2, and is illustrated later in figure 3, and appendix figures 8–9. Finally, in the HBG the ORC vs \mathcal{AF}_4 correlation initially shows a similar behavior to that of the SBM, but we see that they become uncorrelated as the inter-community connection probability q grows (while the intra-community probability p is held fixed). We will further explore this last effect in section 4, where we show that ORC does not appear to detect communities in HBGs with noisy

Table 2. Comparison of Ollivier–Ricci curvature (OR) with Forman–Ricci curvature (FR) and augmentations of the Forman–Ricci curvature (\mathcal{AF}_3 and \mathcal{AF}_4) of edges in real-world networks. The two sections of the table correspond respectively to networks without and with community metadata. In this table, we list the Pearson correlation between the edge curvatures.

Network	OR, FR	OR, \mathcal{AF}_3	OR, \mathcal{AF}_4	FR, \mathcal{AF}_3	FR, \mathcal{AF}_4
Dolphin	0.08	0.75	−0.15	0.25	0.77
Power Grid	0.48	0.72	0.61	0.82	0.89
Word Adjacency	−0.04	0.18	−0.24	0.89	0.89
Southern Women	0.47	0.47	−0.01	1.0	0.71
Corporate Interlocks	0.48	0.48	−0.05	1.0	0.64
Karate Club	0.56	0.81	0.58	0.62	0.84
College Football	−0.09	0.94	−0.76	−0.07	0.37
School 1	−0.01	0.84	−0.26	0.38	0.94
School 2	0.02	0.81	−0.27	0.36	0.92

structure—and so it is natural for ORC and \mathcal{AF}_4 not to be correlated when only one depends on the communities of the network.

Real-world networks: We study a variety of networks, including two bipartite (Southern Women and Corporate Interlocks) and four with community structure metadata (second half of table 2). All networks have between 10 and 250 nodes, except for the US Power Grid, which has around 5000. Detailed descriptions of the networks can be found in appendix A.

As seen in table 2, all networks without a baseline community structure (except for the Power Grid) exhibit similar curvature behaviors. The main observation is that ORC and \mathcal{AF}_4 are slightly negatively correlated, an effect we only observe in the models *with* an underlying community structure. On the other hand, the Power Grid network behaves akin to the ER, with all curvatures moderately correlated. We believe this is related to the network being substantially sparser (average degree ~ 1.3) and with few 3- and 4-cycles (651 and 979, respectively), which causes all curvatures to behave similarly.

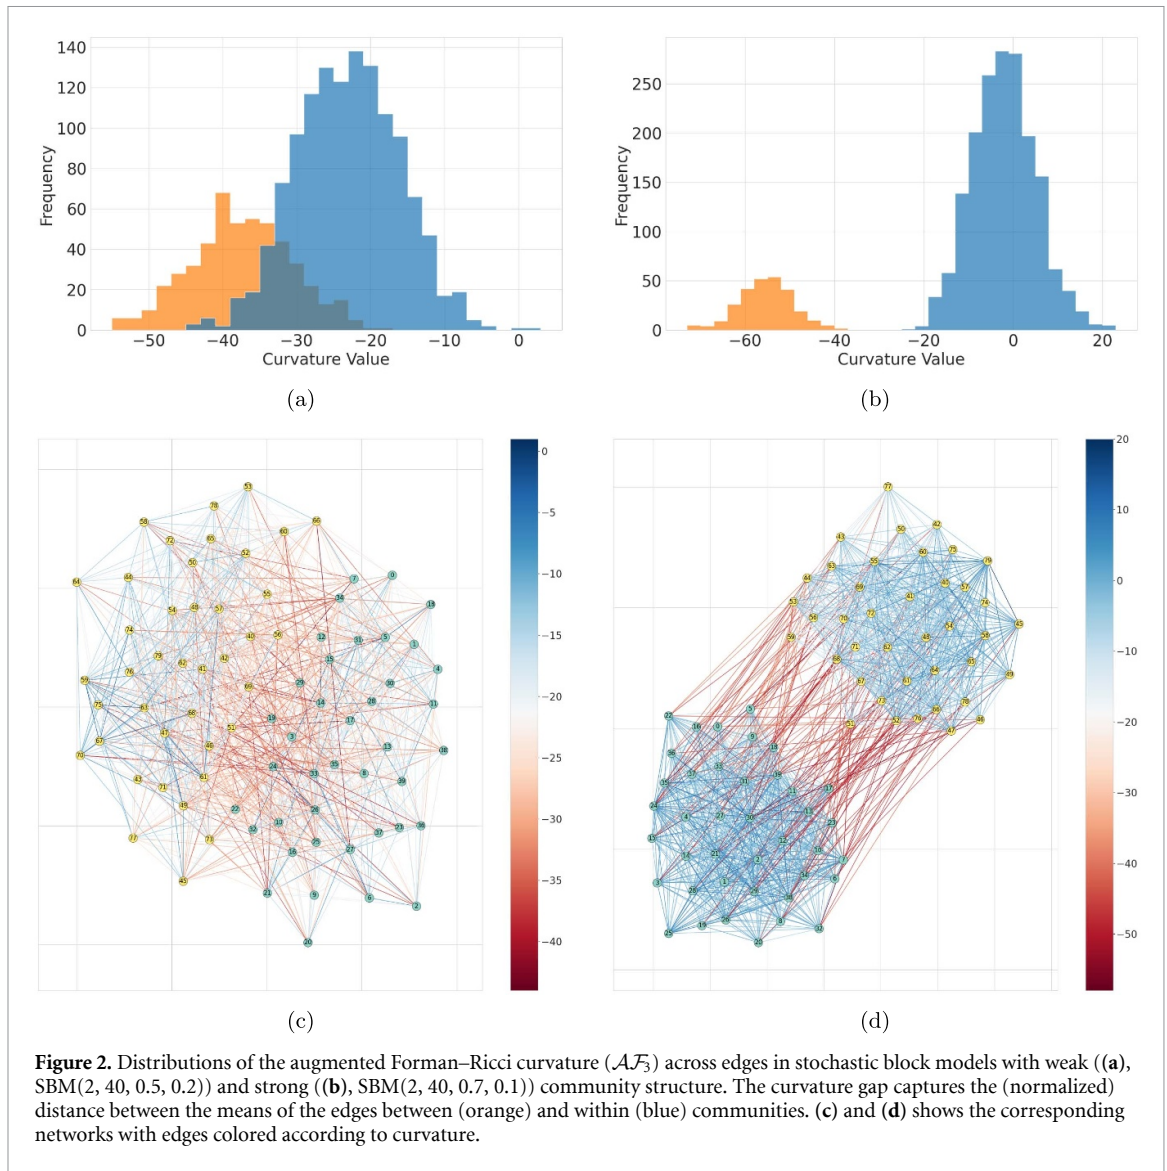
Regarding the networks with ground-truth community structure, all but the Karate Club show a positive (resp. negative) correlation between ORC and \mathcal{AF}_3 (resp. \mathcal{AF}_4), as expected. We attribute the different behavior of the Karate Club to the fact that its ground-truth community structure is slightly blurred in the context of network analysis, which is reflected by the relatively low ratio $\hat{p}/\hat{q} \approx 6.57$, where \hat{p} and \hat{q} are estimates of the probabilities of edges forming within and between communities, respectively. By comparison, this ratio is about 19.51 in the College Football network. We obtain these estimates by dividing the number of edges within (between) communities by the number of possible edges. For the Karate Club and School networks, we note that \mathcal{AF}_4 is much more highly correlated to FRC than ORC, while \mathcal{AF}_3 exhibits the converse behavior. This could be an indicator that in these cases the simpler augmentation \mathcal{AF}_3 is better for community detection, as ORC is expectedly effective in detecting community structure whereas FRC is not.

To give a summary of our correlation results, we have experimentally verified the discussion of section 2. In particular, we have shown how the amount of edges and cycles of a network have different effects on the various versions of FRC, and we have also provided evidence for the expected behaviors of the different curvatures in networks with an underlying community structure. The high correlation between ORC and different versions of AFRC are a first indication that the latter have a potential use in community detection; our results also suggest that the effectiveness of the different augmentations may depend on the characteristics of the network.

4. Curvature gaps

In this section we compare the performance of the studied curvatures in community detection—both for model and real-world networks with an underlying community structure.

The essential observation for ORC-based community detection, as discussed in section 2 and used for instance in [19], is that the ORC of edges within a community tends to be higher than the ORC of edges between communities. As a consequence, the histogram of curvature values in a network with community structure tends to be bimodal with two separate peaks or clusters. Figure 2 shows this effect for \mathcal{AF}_3 in two graphs with a clear community structure. Figure 3 furthermore shows that the clusters in the histogram become more pronounced as the community structure becomes stronger. One way to quantify the separation between the two clusters is the *curvature gap*, introduced in [17]. Given a partition of a network into communities, let κ_{within} denote the average curvature of all edges contained in a community, κ_{between} the average curvature of edges between communities, and σ_{within} , σ_{between} the standard deviation of the respective curvatures. The curvature gap of this partition of the network is defined as



$$\Delta\kappa := \frac{1}{\sigma} |\kappa_{\text{within}} - \kappa_{\text{between}}|,$$

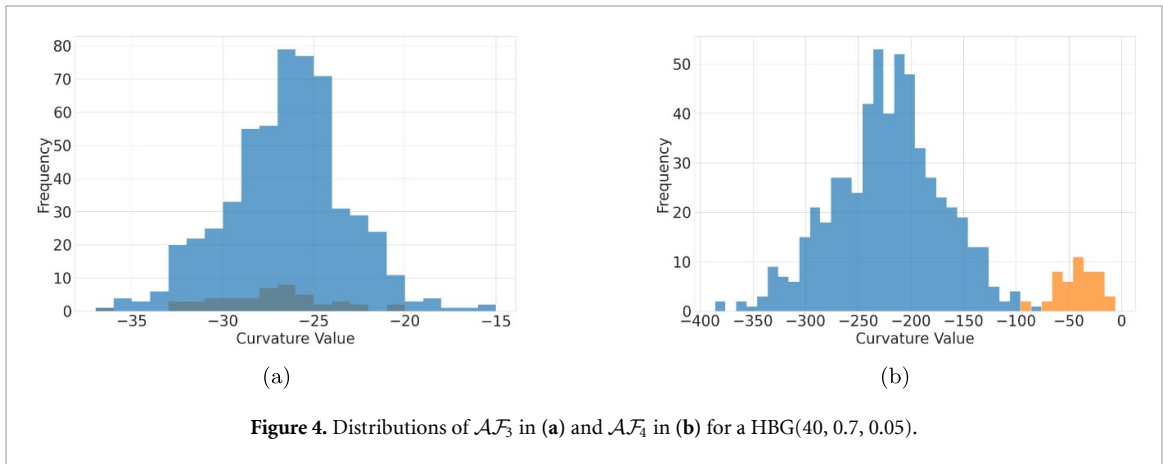
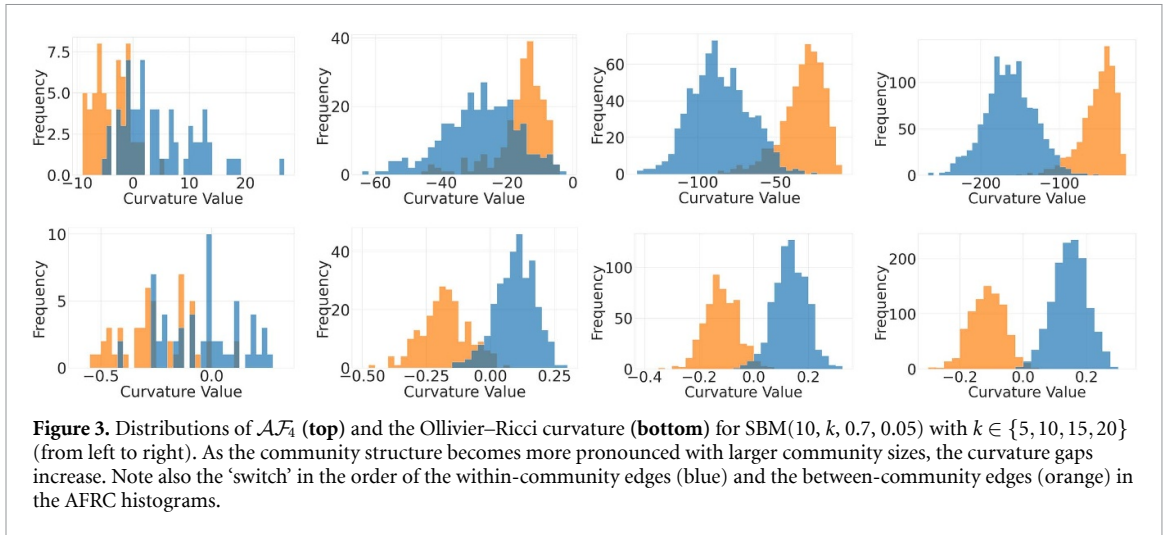
where $\sigma = \sqrt{\frac{1}{2}(\sigma_{\text{within}}^2 + \sigma_{\text{between}}^2)}$. Intuitively, the resulting number can be interpreted as ‘the number of standard deviations between the mean curvatures’. Since the curvature gap reflects how well the curvature distinguishes edges within and between communities, we will use it as a proxy for how well any given curvature detects the community structure. Tables 3 and 4 show the curvature gap obtained from each curvature measure in model and real-world networks, respectively. Importantly, these curvature gaps also depend on how clear the community structure is in a graph (i.e. the ratio

p/q mentioned before, or the size vs the number of the communities); this dependency is illustrated in figure 3 and appendix figures 8–11.

Model networks: We study the previously introduced SBM and HBG models. The latter is particularly interesting because its lack of triangles means that $\mathcal{A}\mathcal{F}_3$ no longer augments the FRC. In this case, $\mathcal{A}\mathcal{F}_4$ is the first non-trivial augmentation.

As discussed in section 2, FRC should not be able to detect communities, which is confirmed by the very poor curvature gap observed in table 3 for both models. Similarly as expected, $\mathcal{A}\mathcal{F}_3$ performs poorly detecting communities in bipartite networks. This is also illustrated in figure 4.

On the other hand, ORC attains good curvature gaps. As noted, we observe that the gap depends on how clear the community structure is: it decreases when we reduce the number of inside-community edges or increase the number of between-community edges; moreover, the gap increases with the size of the communities of the SBM. Finally, $\mathcal{A}\mathcal{F}_4$ (and $\mathcal{A}\mathcal{F}_3$ for the unipartite model) performs similarly to ORC: it shows a relevant gap in all the studied models, which also depends on the clearness of the community



structure. We observe how ORC and \mathcal{AF}_3 tend to perform similarly and outperform \mathcal{AF}_4 in the studied SBMs (although the latter still attains curvature gaps of a similar magnitude), whereas in the studied HBGs \mathcal{AF}_4 attains relevant curvature gaps even in the cases where ORC performs poorly—see appendix figures 10 and 11 for some plots of this phenomenon.

Real-world networks: Next, we study the curvature gap in a range of real-world network with a known community structure. We remark that all networks studied in this section are unipartite.

As with the model networks, FRC performs poorly in all of the studied networks, and both ORC and the augmentations of FRC offer more relevant curvature gaps; \mathcal{AF}_3 offers the best overall performance in our network sample, and \mathcal{AF}_4 performs worst (see table 4). Similarly to section 3, the higher the \hat{p}/\hat{q} ratio of the ground-truth communities, the bigger the curvature gap: the Football network ($\hat{p}/\hat{q} \approx 19.51$) presents a substantial curvature gap for both ORC and the FRC augmentations, whereas for the Karate ($\hat{p}/\hat{q} \approx 6.57$) and School networks ($\hat{p}/\hat{q} \approx 5.89$ and ≈ 6.76 , resp.) the curvature gaps are smaller—notably, \mathcal{AF}_4 offers a curvature gap comparable (and in some cases even smaller) to that of FRC. We remark this as an example of how further augmentations of AFRC do not always improve the effectiveness of the community detection. We believe this could be the same effect we observed for \mathcal{AF}_5 : in networks with low \hat{p}/\hat{q} , the inter- and intra-community edges are contained in a number of 4-cycles similar enough that the two curvature distributions overlap significantly.

To summarize the experimental results on the curvature gap, we have seen that plain FRC is not able to detect community structure—which was theoretically expected—whereas both its augmentations and ORC do. The performance of all curvatures depends on how strongly pronounced the community structure is, but the characteristics of each network have different effects on the various curvature versions. In particular, we have found that ORC and \mathcal{AF}_3 tend to outperform \mathcal{AF}_4 in unipartite networks, which suggests \mathcal{AF}_3 as the better (in terms of performance/cost) alternative for curvature based-community detection for this type of network. On the other hand, \mathcal{AF}_4 seems the most effective in the study of bipartite networks. In the

Table 3. Comparison of curvature gaps attained by Ollivier–Ricci curvature (OR), Forman–Ricci curvature (FR), and augmentations of the Forman–Ricci curvature (\mathcal{AF}_3 and \mathcal{AF}_4) of edges in artificial networks with community structure.

Network	OR	FR	\mathcal{AF}_3	\mathcal{AF}_4
SBM (10, 5, 0.7, 0.05)	2.14	0.58	1.93	1.34
SBM (10, 10, 0.7, 0.05)	3.30	0.35	3.54	1.88
SBM (10, 15, 0.7, 0.05)	4.10	0.33	4.02	3.31
SBM (10, 20, 0.7, 0.05)	4.92	0.23	4.61	4.33
SBM (10, 20, 0.5, 0.05)	2.66	0.16	2.16	2.01
SBM (10, 20, 0.3, 0.05)	0.97	0.07	0.79	0.44
SBM (10, 20, 0.1, 0.05)	0.06	0.18	0.17	0.15
HBG (50, 0.5, 0.05)	3.83	0.19	0.19	4.77
HBG (50, 0.5, 0.1)	0.40	0.12	0.12	3.05
HBG (50, 0.5, 0.15)	0.03	0.16	0.16	1.62
HBG (50, 0.4, 0.1)	0.62	0.06	0.06	1.69
HBG (50, 0.3, 0.1)	0.83	0.14	0.14	0.76
HBG (50, 0.2, 0.1)	0.27	0.07	0.07	0.09

Table 4. Comparison of curvature gaps attained by Ollivier–Ricci curvature (OR), Forman–Ricci curvature (FR), and augmentations of the Forman–Ricci curvature (\mathcal{AF}_3 and \mathcal{AF}_4) of edges in real-world networks.

Network	OR	FR	\mathcal{AF}_3	\mathcal{AF}_4
Karate	1.01	0.31	0.94	0.43
College Football	2.95	0.34	2.87	2.91
School 1	1.16	0.74	1.78	0.57
School 2	1.04	0.85	1.72	0.65

following section we further discuss the performance in community detection of these curvature measures, both in terms of accuracy and computational cost.

5. Community-detection via sequential edge-deletion

In this section, we further analyze the performance in community detection of our curvature measures, both in terms of accuracy and computational cost. Our algorithm proceeds by analogy with the Girvan–Newman method and with the ORC-based community detection method proposed in [19], in that we incrementally delete edges to disconnect the network. We provide pseudo-code of the algorithm in figure 5.

The process can be divided into the following steps, with only minimal differences between using \mathcal{AF}_3 and \mathcal{AF}_4 : given a simple, unweighted, and undirected graph $G(V, E)$, (i) calculate \mathcal{AF}_3 (\mathcal{AF}_4) for all edges in the network, (ii) based on the \mathcal{AF}_3 (\mathcal{AF}_4) distribution, choose a curvature threshold Δ , (iii) remove the most negative \mathcal{AF}_3 (positive \mathcal{AF}_4) edge, (iv) re-compute the edge \mathcal{AF}_3 (\mathcal{AF}_4) only for those edges affected by the removal, (v) check if all \mathcal{AF}_3 (\mathcal{AF}_4) are above (below) Δ , otherwise repeat steps (iii) and (iv) until the condition is satisfied. It is clear that our algorithm must terminate in at most $|E|$ steps, as there will be no more edges to delete—for a sensible threshold such as the low-point between the modes in the \mathcal{AF}_3 (\mathcal{AF}_4) histogram, it will in fact terminate much sooner. To choose the threshold Δ , we fit a Gaussian mixture model with two modes to the curvature distribution. For normal distributions $\mathcal{N}_1(\mu_1, \sigma_1)$ and $\mathcal{N}_2(\mu_2, \sigma_2)$ determined in this way, we then compute $\Delta = \frac{\sigma_2}{\sigma_1 + \sigma_2} \mu_1 + \frac{\sigma_1}{\sigma_1 + \sigma_2} \mu_2$, based on the minimum-error decision boundary between these Gaussians [36, chapter 2.6].

Different choices for Δ are possible and may improve the community detection accuracy. Empirically, we find that when there is less community structure, \mathcal{N}_1 and \mathcal{N}_2 can overlap significantly, which makes finding a good value for Δ harder and tends to degrade accuracy. Choosing Δ more aggressively, i.e. closer to the upper (lower) mode when using \mathcal{AF}_3 (\mathcal{AF}_4) seems to work well in these cases.

In general, it is possible to use *a priori* knowledge of a network to choose an appropriate augmentation of the Forman curvature to detect communities in a computationally efficient way, but with accuracy similar to that achieved using the ORC. For example, as proposed in section 4, for unipartite, SBM-like networks with relatively clear community structure (i.e. large ratio \hat{p}/\hat{q}), we can use \mathcal{AF}_3 . As figure 6 shows, the accuracy of this method is comparable to that of the ORC-based method (panel (a)), while being orders of magnitude faster to run (panel (b)).

For networks which we suspect to be (approximately) bipartite, we can include the 4-cycles, which makes the algorithm more costly to run. However, as can be seen in figure 6, panels (c) and (d), this is still significantly faster than using the ORC, while suffering only minor drops in accuracy. We also note that the

Require: A graph $G(V, E)$ and a threshold Δ

Ensure: A list of tuples (v, l) , with vertices v and labels l

For all $e \in E$, compute \mathcal{AF}_3

while there exists an edge with $\mathcal{AF}_3 < \Delta$ in G **do**

if there is a unique edge e_{\min} with minimal \mathcal{AF}_3 **then**

 remove e_{\min} from G

else

 choose an edge e_{\min} from the edges with minimal \mathcal{AF}_3 uniformly at random and remove it from G

end if

 re-calculate \mathcal{AF}_3 for all affected edges in G

end while

assign the same label l to each vertex v in a connected component of G

return a list of tuples (v, l)

Figure 5. AFRC community detection algorithm, here with \mathcal{AF}_3 .

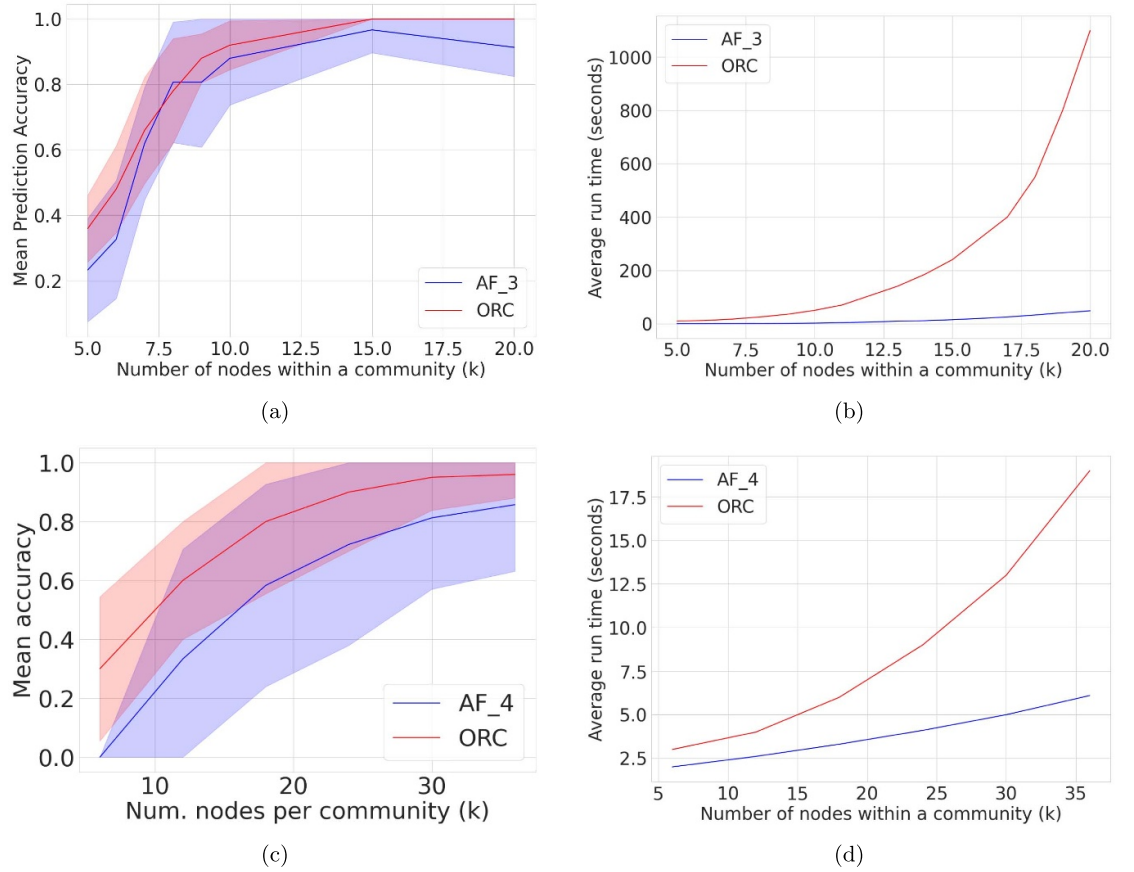


Figure 6. Mean accuracy (percentage of correctly identified communities) and standard deviation (shaded areas) of ORC-based (red) and \mathcal{AF}_3 -based community-detection (blue) in (a), and the time required for the algorithm to terminate (in seconds) in (b) for an SBM with parameters $(10, k, 0.7, 0.05)$. (c) and (d) depict the same using the \mathcal{AF}_4 for a hierarchical bipartite graph with parameters $(k + k, 0.5, 0.05)$.

variance in our HBG results is quite large, which we believe is because we only have two communities. Nonetheless, for a community size of ~ 30 , $\mathcal{AF}_4(e)$ and ORC are on average competitive. The computational efficiency of our algorithm follows from the fact that for an edge $e = (u, v)$ with $d_{\max} = \max\{\deg(u), \deg(v)\}$, computing $\mathcal{AF}_3(e)$ scales as $O(d_{\max})$, $\mathcal{AF}_4(e)$ as $O(d_{\max}^2)$, and $\mathcal{O}(e)$ as $O(d_{\max}^3)$. The number of edges deleted by the algorithm is approximately equal for all three curvatures, so this difference in computational complexities explains the increased speed observed when using $\mathcal{AF}_3(e)$ or $\mathcal{AF}_4(e)$.

In appendix figures 13 and 14, we present additional results on real-world networks and also compare our algorithm against two well-established community detection algorithms: the Girvan–Newman

Algorithm and the Louvain Method. On the American College Football network, the ORC-based algorithm correctly identifies 6 out of 12 communities, while the \mathcal{AF}_3 - and \mathcal{AF}_4 -based algorithms each identify 5 communities correctly. Girvan–Newman performs best and correctly detects 7 communities, while Louvain only identifies 3 ground-truth communities. On the Karate Club, all five algorithms fail to correctly identify the two underlying communities, but the figure 5 algorithm with \mathcal{AF}_3 and Girvan–Newman arguably perform best as both only mislabel two nodes.

We also conduct some additional experiments to compare the performance of our algorithm with the three curvatures around the community detection threshold, and find that for networks with few or no cycles within communities, all three curvatures yield small curvature gaps, which makes community detection difficult or even impossible. We quantify this using the notion of tree-SBMs, i.e. SBMs where the communities are not ER random graphs, but randomly chosen trees with a small probability p of having additional edges within communities. For p close to zero, only the ORC picks up on the emerging community structure. As we increase p , and our tree-SBM becomes more similar to a standard SBM, the augmentations of the FRC begin to also yield increasingly large curvature gaps—see appendix table 6.

Similarly, we find that for standard SBMs where the probability p of an edge within a community is only marginally larger than the probability q of an edge between communities, the ORC yields larger curvature gaps and more accurate community detection results than our augmentations of the FRC. As we increase p , all three curvatures begin to yield similarly good accuracies—see a comparison of ORC vs \mathcal{AF}_3 in appendix figure 12. This underlines the idea that for networks with a particularly difficult community structure, especially in the very sparse regime, the ORC is generally the best, and sometimes only, curvature for community detection. In most cases however, the augmentations of the FRC give similarly good results, which together with their computational efficiency makes them preferable to the ORC.

6. Conclusion

Discrete notions of curvature have recently become the subject of interest in network science, both for theoretical investigations and in applications such as community detection. Much of this interest has been focused on the ORC: to the best of our knowledge, all previous curvature-based community-detection algorithms have used the ORC. While this yields state-of-the-art results for many networks, it is also computationally expensive and quickly becomes inapplicable for larger graphs.

In this work, we chose a different approach and provided a detailed study of augmentations of the Ricci-curvature discretization proposed by Forman (AFRC). We discussed the relation of these augmentations to the FRC and ORC both empirically and theoretically—in particular, we provided evidence that in many artificial and real-world networks, AFRC can be used to attain similar or even better community detection results than the traditionally used ORC. Established ORC-based algorithms for community detection are therefore not without well-performing and computationally cheaper alternatives.

The advantages of AFRC over ORC in community detection suggest that the former could represent important improvements in other areas where different discrete curvature measures have already found an application—for instance network embedding or graph learning. Moreover, the fact that AFRC and ORC behave differently in the studied networks indicates that the former is not just a computational improvement over the latter, but it also reflects different geometric properties of networks. The augmentations proposed in this paper may thus provide further insight into curvature-related problems, and create new possibilities in contexts where other curvature measures are ineffective. For instance, we found the augmentation based on 4-cycles to significantly outperform ORC in community detection in bipartite networks, which opens opportunities for the design of curvature-based methods for this kind of networks.

Data availability statement

The data that support the findings of this study will be openly available following an embargo at the following URL/DOI: https://github.com/Weber-GeoML/AFRC_Community_Detection.

Acknowledgments

The authors declare that they have no conflict of interest.

M W acknowledges partial support from NSF Award 2112085.

R L acknowledges support from the EPSRC Grants EP/V013068/1 and EP/V03474X/1.

Appendix A. Studied networks

Model networks

- **Erdős–Rényi graph:** $ER(m, p)$. Network of m nodes where each edge has an independent probability p of being present. The bipartite version $BG(m, p)$ consists on two groups of m nodes, with the inter-group edges having a probability p of appearing.
- **Stochastic Block Model:** $SBM(l, k, p, q)$. Network of $l \cdot k$ nodes divided into l communities of size k . Each intra-group (resp. inter-group) edge has an independent probability p (resp. q) of being present.
- **Hierarchical Bipartite Graph:** $HBG(m, p, q)$. Bipartite network of $4m$ nodes divided into 2 communities of size $2m$, where each community contains m nodes of each partition. Each intra-group (resp. inter-group) edge has an independent probability p (resp. q) of being present.


Real-world networks

We present an overview of the real-world networks studied in this article in table 5.

Table 5. Summary of the studied real-world networks. The three sections correspond respectively to networks without additional structure, bipartite networks, and networks with community metadata.

Network	Nodes	Edges	Triangles	Squares	\hat{p}/\hat{q}
Dolphin [37]	62	159	95	278	—
Power Grid [38]	4941	6594	651	979	—
Word Adjacency [39]	112	425	284	2579	—
Southern Women [40]	32	89	0	341	—
Corporate Interlocks [41]	244	358	0	245	—
Karate Club [42]	34	78	45	154	6.57
College Football [43]	115	613	810	3915	19.51
School 1 [44]	236	5899	48 401	1678 496	5.89
School 2 [44]	238	5539	46 687	1560 642	6.76

Appendix B. Example: computing $\mathcal{AF}(e)$ in a small graph



We revisit the example from section 2 and calculate the FRC and AFRC of edge $e = e_{12} = (1, 2)$ in figure 7.

For the FRC (1), we only have to count the degrees of the end points of edge e_{12} . These are $\deg(1) = 4$ and $\deg(2) = 5$ and thus we obtain

$$\mathcal{F}(e) = 4 - \deg(1) - \deg(2) = -5.$$

For the AFRC, we need to compute the count matrix Γ . Since we focus on edge e_{12} it suffices to calculate the e_{12} th row and column of Γ . We fix the vertex order $1 < 2 < \dots < 7 < 8$ which induces the order $e_{12} < e_{14} < e_{15} < e_{23} < e_{25} < e_{34}$ on the edges. Since e_{12} is contained in two cycles, $\gamma_1 = 1234$ and $\gamma_2 = 125$, we have

$$\Gamma_{e_{12}e_{12}} = 2.$$

Cycle $\gamma_1 = 1234$ traverses e_{12}, e_{23}, e_{34} from small to large, and e_{14} from large to small. Applying definition (3), we find

$$\begin{cases} \Gamma_{e_{12}e_{23}} = 1 & (\text{since } e_{12}, e_{23} \text{ are aligned in } \gamma_1, \text{ and } e_{12} < e_{23}) \\ \Gamma_{e_{12}e_{34}} = 1 & (\text{since } e_{12}, e_{34} \text{ are aligned in } \gamma_1, \text{ and } e_{12} < e_{34}) \\ \Gamma_{e_{14}e_{12}} = 1 & (\text{since } e_{12}, e_{14} \text{ are not aligned in } \gamma_1, \text{ and } e_{14} > e_{12}). \end{cases}$$

Cycle $\gamma_2 = 125$ traverses e_{12}, e_{25} from small to large and e_{15} from large to small. Applying definition (3), we find

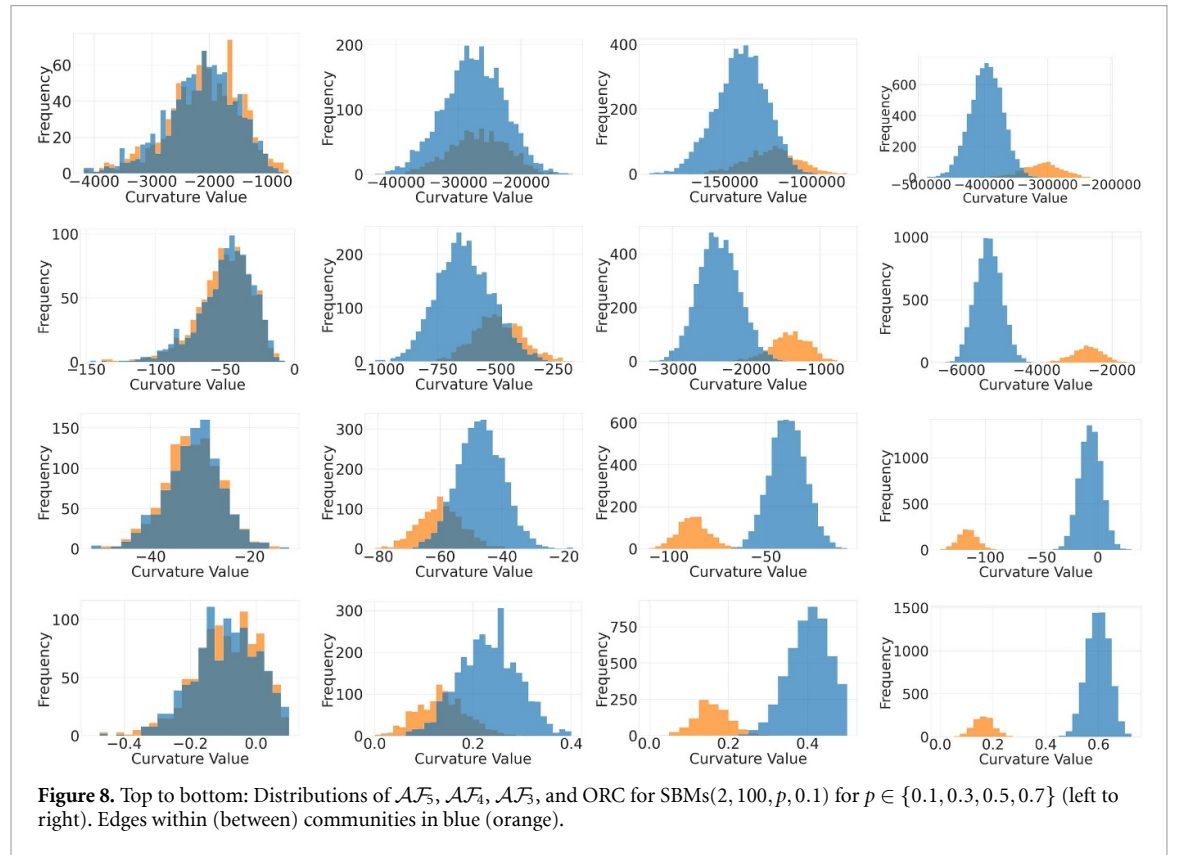
$$\begin{cases} \Gamma_{e_{12}e_{25}} = 1 & (\text{since } e_{12}, e_{25} \text{ are aligned in } \gamma_2, \text{ and } e_{12} < e_{25}) \\ \Gamma_{e_{15}e_{12}} = 1 & (\text{since } e_{12}, e_{15} \text{ are not aligned in } \gamma_2 \text{ and } e_{15} > e_{12}). \end{cases}$$

All other entries in the e_{12} th row and column of Γ are zero since they correspond to edges that do not share a cycle with e_{12} (e.g. $\Gamma_{e_{12}e_{16}} = 0$). Finally, the neighbors of e_{12} are $\{e_{14}, e_{15}, e_{16}, e_{23}, e_{25}, e_{28}, e_{27}\} \sim e_{12}$ and the non-neighbors are $\{e_{34}, e_{78}\} \not\sim e_{12}$. The AFRC is calculated from expression (2) as

$$\begin{aligned} \mathcal{AF}(e_{12}) &= 2 + \Gamma_{e_{12}e_{12}} - \sum_{e' \sim e_{12}} |\Gamma_{e_{12}e'} + \Gamma_{e'e_{12}} - 1| - \sum_{e' \not\sim e_{12}} |\Gamma_{e_{12}e'} - \Gamma_{e'e_{12}}| \\ &= 2 + 2 - (4 \cdot |1 - 1| + 3 \cdot |0 - 1|) - (|1| + |0|) \\ &= 0. \end{aligned}$$

The code to calculate the AFRC in a graph will be made available along with the publication of the paper. This includes a calculation of $\mathcal{AF}(e_{12})$ in the example above using the codebase.

Appendix C. Additional numerical experiments



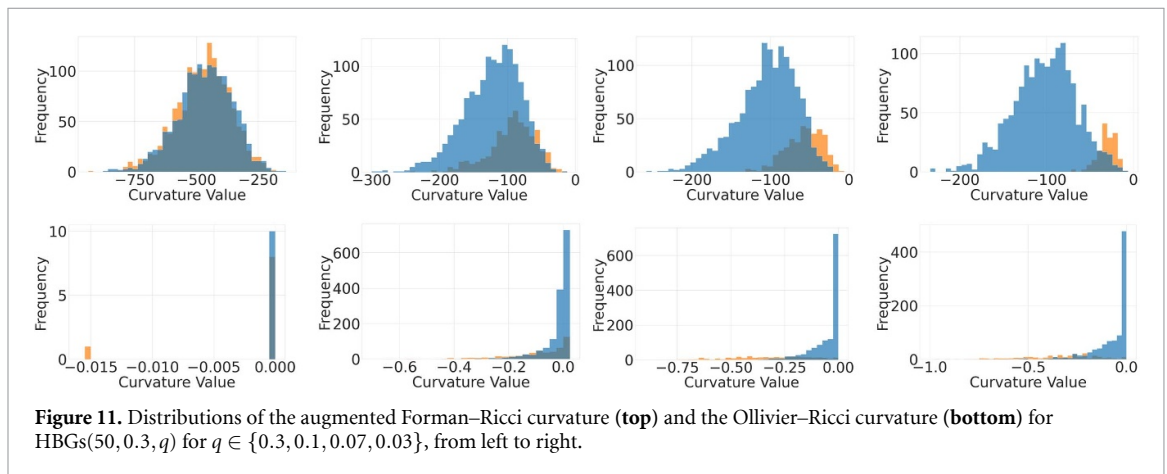
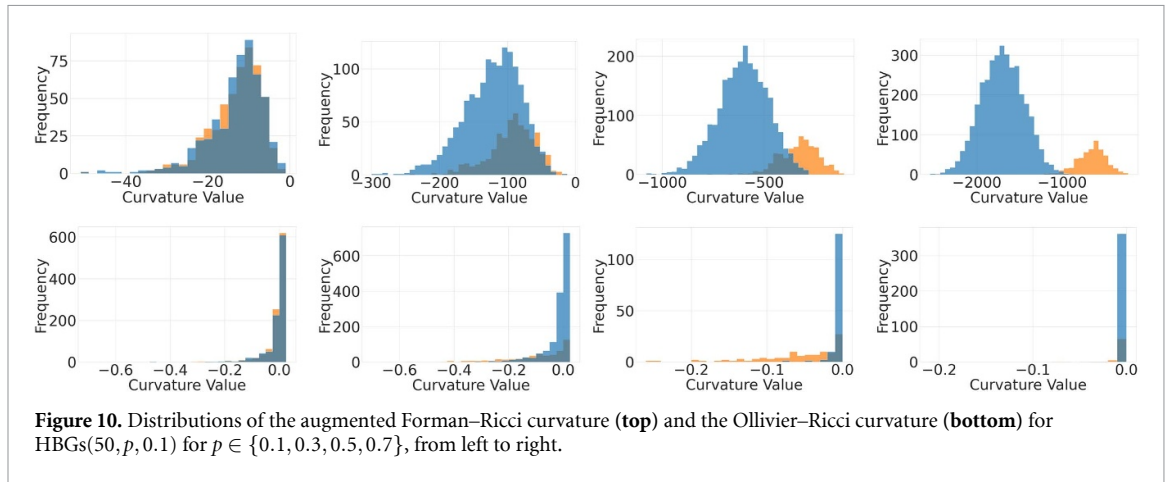
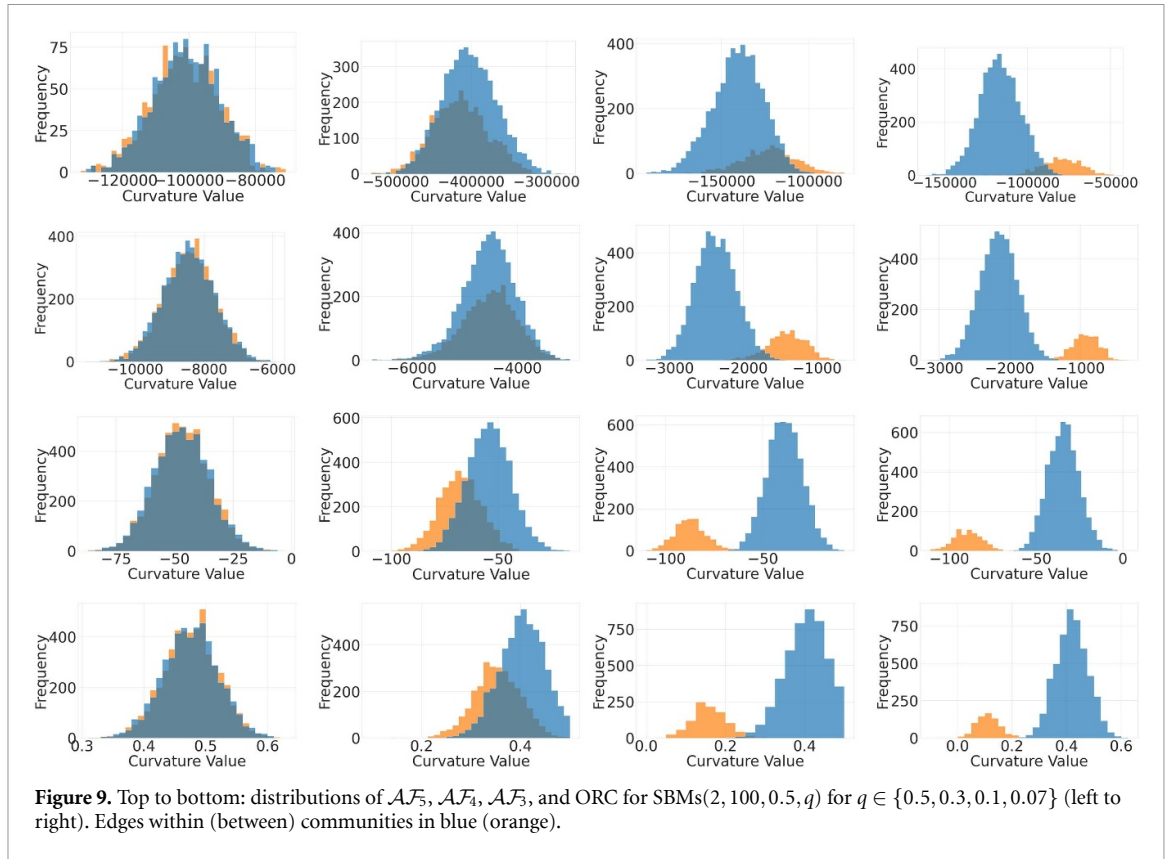


Table 6. Comparison of curvature gaps attained by the Ollivier–Ricci curvature and augmentations of the Forman–Ricci curvature for tree-SBMs with increasing probability of additional edges within communities.

Network	ORC	\mathcal{AF}_3	\mathcal{AF}_4
t-SBM(2, 50, 0, 0.03)	0.52	0.46	0.40
t-SBM(2, 50, 0.1, 0.03)	2.14	0.93	0.53
t-SBM(2, 50, 0.2, 0.03)	3.98	1.72	1.41
t-SBM(2, 50, 0.3, 0.03)	6.18	3.04	3.29




Figure 12. Mean accuracy (percentage of correctly identified communities) and standard deviation (shaded areas) of 15 ORC-based (red) and \mathcal{AF}_3 community-detection (blue) for an SBM with parameters (5, 20, p , 0.05) for $p \in \{0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6\}$.




Figure 13. Community detection in the American College Football network: ground truth communities (a), and communities detected by sequential edge-deletion using the ORC (b), \mathcal{A}_{F_3} (c), and \mathcal{A}_{F_4} (d). We also display the communities detected by the Girvan–Newman (e) and Louvain (f) algorithms.

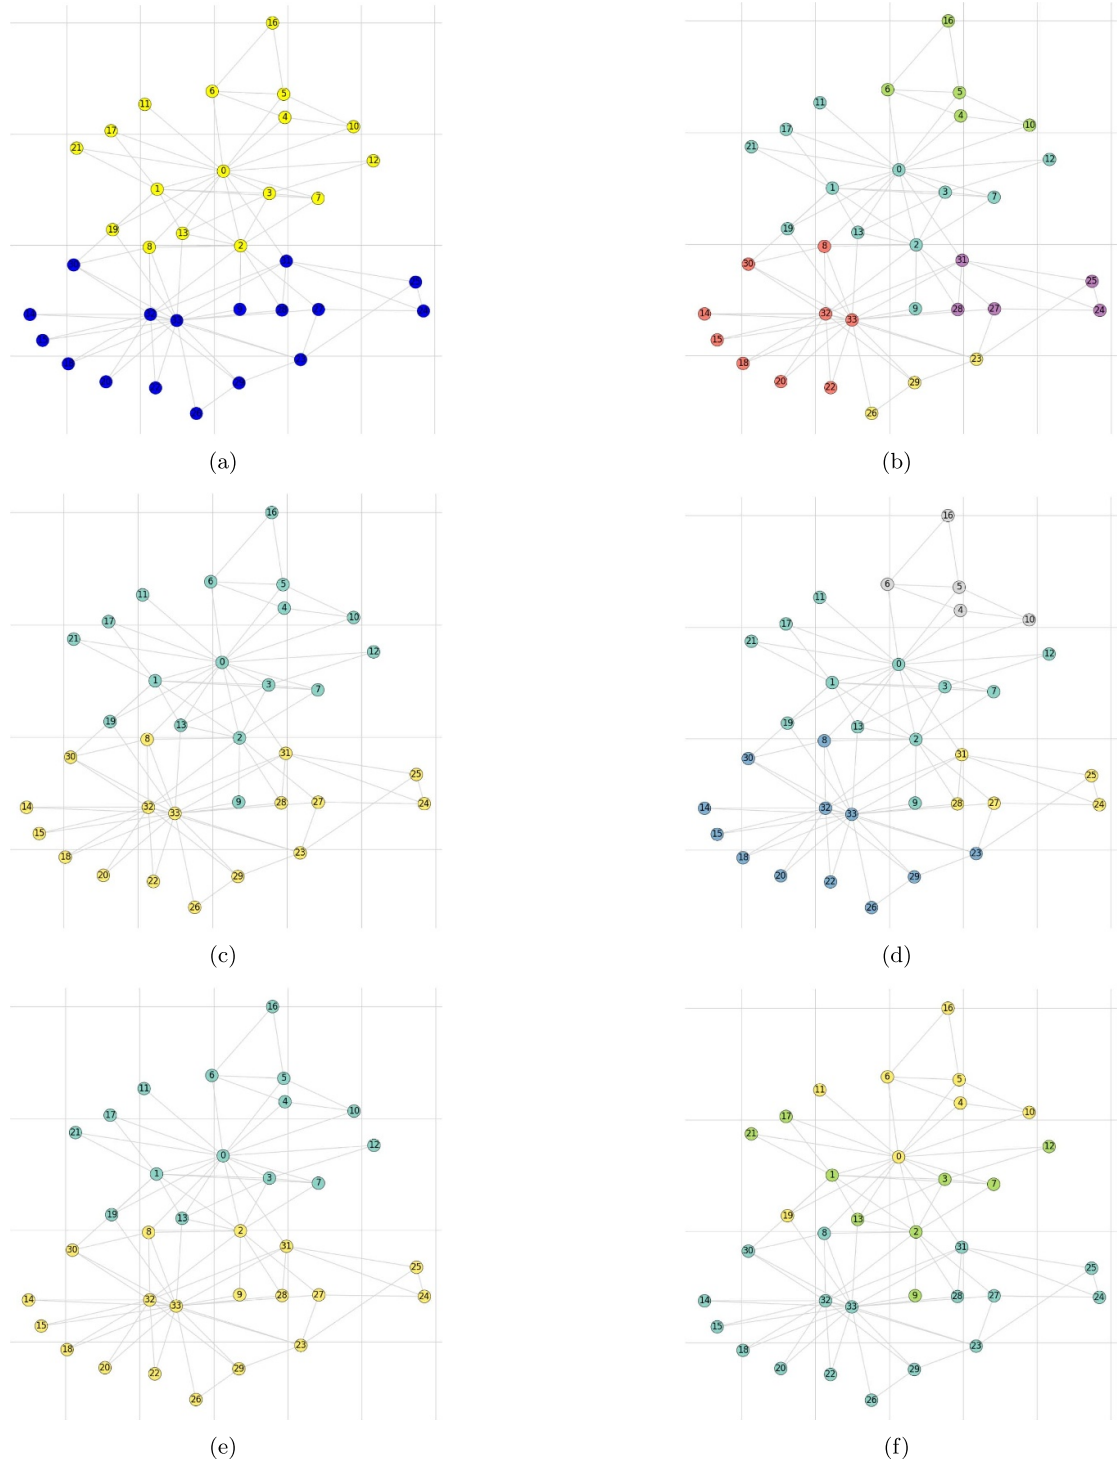


Figure 14. Community detection in the Karate Club network: ground truth communities (a), and communities detected by sequential edge-deletion using the ORC (b), \mathcal{A}_3 (c) and \mathcal{A}_4 (d). We also display the communities detected by the Girvan–Newman (e) and the Louvain (f) algorithms.

ORCID iDs

Sergio Serrano de Haro Iváñez  <https://orcid.org/0009-0000-3711-0121>

Renaud Lambiotte  <https://orcid.org/0000-0002-0583-4595>

References

- [1] Weber M, Saucan E and Jost J 2017 Characterizing complex networks with Forman-Ricci curvature and associated geometric flows *J. Complex Netw.* **5** 527–50
- [2] Tannenbaum A, Sander C, Zhu L, Sandhu R, Kolesov I, Reznik E, Senbabaoglu Y and Georgiou T 2015 Ricci curvature and robustness of cancer networks (arXiv:1502.04512)
- [3] Weber M, Stelzer J, Saucan E, Naitsat A, Lohmann G and Jost J 2017 Curvature-based methods for brain network analysis (arXiv:1707.00180)
- [4] Leal W, Restrepo G, Stadler P F and Jost J 2021 Forman-Ricci curvature for hypergraphs *Adv. Complex Syst.* **24** 2150003
- [5] Sandhu R S, Georgiou T T and Tannenbaum A R 2016 Ricci curvature: an economic indicator for market fragility and systemic risk *Sci. Adv.* **2** e1501495
- [6] Topping J, Di Giovanni F, Chamberlain B P, Dong X and Bronstein M M 2021 Understanding over-squashing and bottlenecks on graphs via curvature (arXiv:2111.14522)
- [7] Weber M 2020 Neighborhood growth determines geometric priors for relational representation learning *Int. Conf. on Artificial Intelligence and Statistics* vol 108 pp 266–76
- [8] Ollivier Y 2007 Ricci curvature of metric spaces *C. R. Acad. Sci., Paris* **345** 643–6
- [9] Ollivier Y 2009 Ricci curvature of Markov chains on metric spaces *J. Funct. Anal.* **256** 810–64
- [10] Forman R 2003 Bochner’s method for cell complexes and combinatorial Ricci curvature *Discrete Comput. Geom.* **29** 323–74
- [11] Bochner S 1946 Vector fields and Ricci curvature *Bull. Am. Math. Soc.* **52** 776–97
- [12] Bochner S 1948 Curvature and Betti numbers *Ann. Math.* **49** 379–90
- [13] Bochner S 1949 Curvature and Betti numbers. II *Ann. Math.* **50** 77–93
- [14] Kovács B and Palla G 2021 The inherent community structure of hyperbolic networks *Sci. Rep.* **11** 16050
- [15] Xue Y and Bogdan P 2017 Reliable multi-fractal characterization of weighted complex networks: algorithms and implications *Sci. Rep.* **7** 7487
- [16] Yang R, Sala F and Bogdan P 2021 Hidden network generating rules from partially observed complex networks *Commun. Phys.* **4** 199
- [17] Gosztolai A and Arnaudon A 2021 Unfolding the multiscale structure of networks with dynamical Ollivier-Ricci curvature *Nat. Commun.* **12** 4561
- [18] Ni C C, Lin Y Y, Luo F and Gao J 2019 Community detection on networks with Ricci flow *Sci. Rep.* **9** 1–12
- [19] Sia J, Jonckheere E and Bogdan P 2019 Ollivier-Ricci curvature-based method to community detection in complex networks *Sci. Rep.* **9** 1–12
- [20] Sia J, Zhang W, Jonckheere E, Cook D and Bogdan P 2022 Inferring functional communities from partially observed biological networks exploiting geometric topology and side information *Sci. Rep.* **12** 10883
- [21] Tian Y, Lubberts Z and Weber M 2023 Mixed-membership community detection via line graph curvature *Symmetry and Geometry in Neural Representations* (PMLR) pp 219–33
- [22] Abbe E 2018 Community detection and stochastic block models: recent developments *J. Mach. Learn. Res.* **18** 1–86
- [23] Porter M A, Onnela J-P and Mucha P J 2009 Communities in networks *Not. Am. Math. Soc.* **56** 1082–97
- [24] Cuturi M 2013 Sinkhorn distances: lightspeed computation of optimal transport *Advances in Neural Information Processing Systems* vol 26
- [25] Samal A, Sreejith R P, Gu J, Liu S, Saucan E and Jost J 2018 Comparative analysis of two discretizations of Ricci curvature for complex networks *Sci. Rep.* **8** 8650
- [26] Serrano de Haro Iváñez S 2022 Comparative analysis of Forman-Ricci curvature versions applied to the persistent homology of networks (arXiv:2212.01357)
- [27] Weber M, Saucan E, Jost J and Estrada E 2018 Coarse geometry of evolving networks *J. Complex Netw.* **6** 706–32
- [28] Onnela J-P, Saramäki J, Hyvönen J, Szabó G, Lazer D, Kaski K, Kertész J and Barabási A-L 2007 Structure and tie strengths in mobile communication networks *Proc. Natl Acad. Sci.* **104** 7332–6
- [29] Erbar M and Maas J 2012 Ricci curvature of finite Markov chains via convexity of the entropy *Arch. Ration. Mech. Anal.* **206** 997–1038
- [30] Bakry D, Gentil I and Ledoux M 2014 *Analysis and geometry of Markov diffusion operators Grundlehren der Mathematischen Wissenschaften* vol 348 (Springer)
- [31] Devriendt K and Lambiotte R 2022 Discrete curvature on graphs from the effective resistance *J. Phys. Complex.* **3** 025008
- [32] Jost J and Münch F 2021 Characterizations of Forman curvature (arXiv:2110.04554)
- [33] Sreejith R P, Mohanraj K, Jost J, Saucan E and Samal A 2016 Forman curvature for complex networks *J. Stat. Mech.* **063206**
- [34] Jost J and Liu S 2013 Ollivier’s Ricci curvature, local clustering and curvature-dimension inequalities on graphs *Discrete Comput. Geom.* **51** 300–22
- [35] Tee P and Trugenberger. C A 2021 Enhanced Forman curvature and its relation to Ollivier curvature *Europhys. Lett.* **133** 60006
- [36] Duda R O, Hart P E and Stork D G 2012 *Pattern Classification* (Wiley)
- [37] Lusseau D, Schneider K, Boisseau O J, Haase P, Slooten E and Dawson S M 2003 The bottlenose dolphin community of doubtful sound features a large proportion of long-lasting associations: can geographic isolation explain this unique trait? *Behav. Ecol. Sociobiol.* **54** 396–405
- [38] Watts D J and Strogatz S H 1998 Collective dynamics of ‘small-world’ networks *Nature* **393** 440–2
- [39] Newman M 2006 Finding community structure in networks using the eigenvectors of matrices *Phys. Rev. E* **74** 036104

- [40] Davis A, Gardner B B and Gardner M R 1941 *Deep South: a social anthropological study of caste and class* (Phoenix Books) (University of Chicago Press) p 204
- [41] Scott J, Hughes M and Mackenzie J 1980 *The Anatomy of Scottish Capital: Scottish Companies and Scottish Capital, 1900-1979* (McGill-Queen's University Press)
- [42] Zachary. W W 1977 An information flow model for conflict and fission in small groups *J. Anthropol. Res.* **33** 452–73
- [43] Girvan M and Newman M 2002 Community structure in social and biological networks *Proc. Natl Acad. Sci. USA* **99** 7821–6
- [44] Stehlé J *et al* 2011 High-resolution measurements of face-to-face contact patterns in a primary school *PLoS One* **6** e23176