

Ollivier-Ricci curvature convergence in random geometric graphs

Pim van der Hoorn ,^{1,2,3} William J. Cunningham ,^{4,2,3} Gabor Lippner,⁵ Carlo Trugenberger,⁶ and Dmitri Krioukov ^{2,3,5,7}

¹*Department of Mathematics and Computer Science, Eindhoven University of Technology, 5600 MB Eindhoven, Netherlands*

²*Department of Physics, Northeastern University, Boston, Massachusetts 02115, USA*

³*Network Science Institute, Northeastern University, Boston, Massachusetts 02115, USA*

⁴*Perimeter Institute for Theoretical Physics, Waterloo, Ontario N2L2Y5, Canada*

⁵*Department of Mathematics, Northeastern University, Boston, Massachusetts 02115, USA*

⁶*SwissScientific Technologies SA, CH-1204 Geneva, Switzerland*

⁷*Department of Electrical and Computer Engineering, Northeastern University, Boston, Massachusetts 02115, USA*



(Received 6 August 2020; accepted 15 February 2021; published 5 March 2021)

Connections between continuous and discrete worlds tend to be elusive. One example is curvature. Even though there exist numerous nonequivalent definitions of graph curvature, none is known to converge in any limit to any traditional definition of curvature of a Riemannian manifold. Here we show that Ollivier curvature of random geometric graphs in any Riemannian manifold converges in the continuum limit to Ricci curvature of the underlying manifold, but only if the definition of Ollivier graph curvature is properly generalized to apply to mesoscopic graph neighborhoods. This result establishes a rigorous link between a definition of curvature applicable to networks and a traditional definition of curvature of smooth spaces.

DOI: [10.1103/PhysRevResearch.3.013211](https://doi.org/10.1103/PhysRevResearch.3.013211)

I. INTRODUCTION

Curvature is one of the most basic geometric characteristics of space. The original definitions of curvature apply only to smooth Riemannian or Lorentzian manifolds, but there exist numerous extensions of curvature definitions applicable to graphs, simplicial complexes, and other discrete structures. For a variety of reasons, these extensions have recently seen a surge of interest in areas as diverse as network or data science and quantum gravity.

In network science, graph curvature is interesting in general since many real-world networks were found to possess different flavors of geometry [1]. Measuring network curvature is then a way to learn what this geometry really is [2]. On the application side, graph curvature was used to characterize congestion in telecommunication networks [3,4], detect cancer cells [5], predict COVID-19 spreading [6] and stock market fluctuations [7], analyze robustness and other properties of the Internet, financial, and brain networks [8–12], and in a number of classic applications such as community inference [13] and network embedding in machine learning [14]. In a majority of these applications, graph curvature appears as a signature of a (pathological) event or property of interest. How reliable such a signature is for a task at hand depends on whether the used definition of graph curvature does what is intended—that it is indeed curvature, and not something else.

In general relativity, scalar Ricci curvature is a key player since it appears in the Einstein-Hilbert action whose least-action extremization leads to Einstein's equations [15]. Therefore, various extensions of Ricci curvature to discrete structures representing “quantum spacetimes” have been considered. The first step in that direction was based on replacing continuum manifolds with simplicial complexes [16]. To that end, variations of the seminal definition of discrete curvature by Regge [17] were developed and utilized. More recently, a collection of more minimalistic structures was considered, such as causal sets [18–27] and random graphs [28–33]. Yet another class of approaches aimed at explaining spacetime emergence from “quantum bits” appears in space-from-entanglement proposals such as ER=EPR or tensor networks [34–40]. In all these examples, one can hope that a particular discrete structure represents Planck-scale gravitational physics, only if discrete curvature of this structure converges to Ricci curvature of classical spacetime in the continuum limit. If such convergence does not hold in a particular approach, then the approach does not agree with general relativity, so it cannot be considered realistic, regardless of how attractive it is in other respects.

Quite contrary to the continuum world of Riemannian and Lorentzian manifolds where there is no real ambiguity regarding the definition of curvature, there exist not one or two but very many nonequivalent definitions of curvature applicable to graphs. We briefly review the most prominent definitions in the next section. Unfortunately, none of these definitions are known to converge in the continuum limit to any traditional curvature of any Riemannian manifold. Therefore, it is often unclear how one should interpret different measurements of different curvatures in different graphs and networks, raising question concerning reliability and generalizability of

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

predictions based on such measurements. Similarly, the evasiveness of classical limit in some approaches to quantum gravity can be often linked to unclarity concerning whether discrete “quantum” curvature converges to a continuous classical one. These and other problems related to curvature convergence become particularly acute in view of growing evidence that different graph curvature definitions may disagree even about the sign of curvature in some paradigmatic graphs and networks [41,42].

Here we show that Ollivier curvature [43–45] of random geometric graphs [46–48] in any Riemannian manifold converges to Ricci curvature of the manifold in the continuum limit. Unfortunately, this convergence does not and cannot hold for the standard definition of Ollivier curvature of microscopic one-link neighborhoods in sparse unweighted graphs [49]. Fortunately, we find a natural generalization of Ollivier curvature applicable to mesoscopic neighborhoods in unweighted and weighted graphs. This generalization is a not-previously-considered key ingredient that allows us to prove the curvature convergence. Our proofs hold only when the sizes of mesoscopic graph neighborhoods fall within certain windows that depend on graph density and on how graph edges are weighted. We also perform large-scale simulations that agree with our proofs and suggest that the convergence windows are actually much wider than the limits imposed by our proof techniques.

These results link rigorously a definition of graph curvature to the traditional Ricci curvature of a Riemannian manifold. To the best of our knowledge, the closest, in spirit, previous results are Cheeger *et al.*’s proof [50] of Regge’s seminal observation [17] that an angle-defect-based curvature of increasingly finer-grained simplicial triangulations of a manifold converges to its Ricci curvature, and the more recent demonstration [21,22] that the Benincasa-Dowker action of causal sets converges to its classical counterpart in the continuum limit.

To proceed, we first recall the Ollivier curvature definition in perspective of other curvature definitions applicable to graphs and networks, Sec. II. We then describe our generalization of the Ollivier curvature definition to mesoscopic neighborhoods in graphs in general and in random geometric graphs in particular, Sec. III. We then state our main results in Sec. IV, and conclude with a discussion of their implications and caveats in Sec. V.

II. GRAPH CURVATURE DEFINITIONS

There are many nonequivalent definitions of curvature applicable to graphs. It is impossible to list them all. Here we mention some notable ones due to Steiner [51], Regge [17,50], Bakry-Émery [52], Gromov [53,54], Higuchi [55], Eckmann-Moses [56], Forman [57–59], Lin-Yau [60,61], Knill [62,63], Keller [64,65], and hybrids thereof [66]. We first outline the taxonomy of these curvatures and basic ideas behind their definitions to put Ollivier curvature [43–45] in perspective, and then recall the detailed definition of the latter for completeness.

With a few exceptions, the discrete curvature definitions above fall into the following classes.

A. Direct approaches

The first class of definitions comes from relatively direct extensions of continuous curvature to nonsmooth objects such as simplices, simplicial complexes, or more general complexes. Since graphs are 1-dimensional simplicial complexes, any notion of curvature applicable to complexes applies to graphs as well. Casting cliques in graphs as higher-dimensional simplices turns graphs into higher-dimensional complexes, so that curvature definitions applicable to higher-dimensional complexes are also applicable to graphs.

One subclass of such approaches first defines curvature of polyhedra using the Gauss-Bonnet theorem that relates the curvature of a surface to its Euler’s characteristic. This definition is then extended to general simplicial complexes. Curvatures defined by Steiner [51], Regge [17,50], Knill [62,63], and Keller [64,65] all fall into this category.

Forman curvature [57–59] follows similar ideas, but is slightly different in that it uses the Bochner-Weitzenböck formula that relates the Laplacian of a Riemannian manifold to its discrete approximation, which is then used to derive a discrete version of curvature applicable to a very general class of complexes. The definition of Forman curvature for unweighted graphs is one of the simplest. If a graph is considered as a 1-complex, so that only nodes (0-simplices) and links (1-simplices) are considered, then the Forman curvature $F_1(i, j)$ of a link between nodes i and j of degrees k_i and k_j is $F_1(i, j) = 4 - (k_i + k_j)$. If the graph is considered as a 2-complex, so that triangles are also considered, and $\Delta_{i,j}$ is the number of triangles containing the link, then the definition becomes $F_2(i, j) = F_1(i, j) + 3\Delta_{i,j}$.

B. Indirect approaches

Another class of ideas behind discrete curvature is less direct as these are based on more sophisticated curvature-related properties of smooth spaces.

One such idea, pioneered by Bakry and Emery [52], relies on the characterization of Ricci curvature via curvature-dimension inequalities. The curvature-dimension inequality establishes a lower bound for the Ricci curvature of a Riemannian manifold based on the gradient of harmonic functions on it. The inequality is an immediate consequence of Bochner’s identity that relates Ricci curvature to the gradient of harmonic functions. It was an important insight by Bakry and Emery that the curvature-dimension inequality can be lifted from manifolds and considered as a substitute of the lower bound for Ricci curvature in spaces where direct discretizations of Ricci curvature are not possible. In addition to the Bakry-Emery definition [52], the definitions of Lin-Yau [60,61] also fall into this category.

The key point behind this type of approaches to discrete curvature is to base its definitions on some fundamental non-trivial curvature-related properties that transcend well beyond the limits of smooth spaces.

C. Ollivier curvature

Ollivier curvature belongs to the second class of approaches outlined above. The fundamental aspect of Ricci curvature that Ollivier curvature embarks on to sail to the

discrete world is somewhat different from—although related to—the curvature-dimension inequalities. This aspect is of a less analytic and more geometric nature. It deals with how balls shrink or expand under parallel transport. Recall that if Ricci curvature of a space is negative, zero, or positive, then under parallel transport, balls in this space expand, stay the same, or shrink, respectively.

To capture this property, one needs to consider the Wasserstein, also known as transportation, distance $W(\mu_1, \mu_2)$ between two probability distributions μ_1, μ_2 in *any* nice (Polish to be exact) metric space:

$$W(\mu_1, \mu_2) = \inf_{\mu \in \Gamma(\mu_1, \mu_2)} \int d(x, y) \mu(x, y) dx dy, \quad (1)$$

where $d(x, y)$ is the distance between points x and y in the space, and the infimum is taken over all possible transportation plans which are joint probability distributions μ whose marginals are μ_1 and μ_2 . If μ_1, μ_2 are represented by sand piles, then $W(\mu_1, \mu_2)$ is the minimum cost to transport pile μ_1 into pile μ_2 , where the grain of sand at x gets transported to y incurring cost $d(x, y)$.

Let x and y be now two points at a small distance $d_M(x, y) = \delta$ in a Riemannian manifold M , and let $\mu_1 = \mu_x$ and $\mu_2 = \mu_y$ be the normalized restrictions of the volume form in M onto the balls $B_M(x, \delta), B_M(y, \delta)$ of radius δ centered at x, y , i.e.,

$$\mu_x(z) = \frac{\text{vol}(z)}{\text{vol}[B_M(x, \delta)]} \quad \text{if } z \in B_M(x, \delta), \quad (2)$$

$$\mu_y(z) = \frac{\text{vol}(z)}{\text{vol}[B_M(y, \delta)]} \quad \text{if } z \in B_M(y, \delta), \quad (3)$$

where $\text{vol}(z) dz$ is the volume element in M . Ollivier curvature between x and y is then defined by

$$\kappa_M(x, y) = 1 - \frac{W_M(\mu_x, \mu_y)}{\delta}. \quad (4)$$

We note that this curvature definition depends on the ball radius δ which can be any positive real number. It was shown in [44] that

$$\lim_{\delta \rightarrow 0} \frac{\kappa_M(x, y)}{\delta^2} = \frac{\text{Ric}(v, v)}{2(D + 2)}, \quad (5)$$

where D is M 's dimension, and $\text{Ric}(v, v)$ is the Ricci curvature at x along v which is the unit tangent vector at x pointing along the geodesic from x to y . Ricci curvature $\text{Ric}(v, v)$ is equal to the average of sectional curvatures at x over all tangent planes containing v .

In words, the important convergence result in (5) says that in Riemannian manifolds, the rescaled Ollivier curvature converges to Ricci curvature in the limit of small ball sizes $\delta \rightarrow 0$.

It is relatively straightforward to extend the definition of Ollivier curvature from manifolds M to simple unweighted graphs G , but there are important differences. The most crucial one is that in smooth spaces M the Ollivier ball radius δ appearing in the Ollivier curvature definition (4) can be any positive real number, which can tend to zero to prove the convergence in (5). In unweighted graphs G , however, the possible nonzero values of the distance—and so of the δ —are positive integers, the hop lengths of shortest paths

between pairs of nodes. The smallest possible nonzero distance is $\delta = 1$, the graph distance $d_G(x, y)$ between directly connected nodes x and y . Presumably because of the “ δ better be small” logic, the Ollivier ball radius δ was routinely set to its smallest possible value $\delta = 1$ in the definition of Ollivier graph curvature appearing in the past literature, even though δ in graphs can certainly be set to any positive integer value.

If $\delta = 1$, then the balls $B_G(x, 1)$ and $B_G(y, 1)$ of radius 1 around nodes x and y are the sets of the neighbors of x and y , while the probability distributions μ_x and μ_y become the uniform distributions over these sets: $\mu_x = 1/k_x$ and $\mu_y = 1/k_y$, where k_x and k_y are the degrees of x and y . These distributions govern the standard random walk in the graph. Thus defined, Ollivier graph curvature was shown to be bounded by -2 and 1 , with the bounds achieved on infinite double stars and complete graphs [49]. The Ollivier graph curvature in this definition was measured in a great variety of synthetic and real-world networks [5,8,9,13,41,42,49,67]. It was also investigated at great depths in connection to quantum gravity [28–33].

However, it is evident that Ollivier curvature of graphs as defined above cannot converge to Ricci curvature of manifolds under any circumstances, simply because δ is no longer real but integer, so that $\lim_{\delta \rightarrow 0}$ makes no sense, and because Ollivier graph curvature is always between -2 and 1 , while Ricci curvature of a Riemannian manifold can be any real number. Therefore, in the next section, we generalize the Ollivier graph curvature definition to make it much more versatile, capturing a much wider spectrum of graph properties that are not necessarily microscopic.

III. MESOSCOPIC GRAPH NEIGHBORHOODS

A. General idea

The main limitation of the Ollivier graph curvature definition in the previous section is that it is “too microscopic.” It limits the curvature-related considerations to graph balls or neighborhoods of the radius of one hop, $\delta = 1$. Such neighborhoods may be too small to “feel” any curvature in a general case. Recall that microscopically, any smooth curved space is locally flat. This observation instructs us to consider much larger mesoscopic graph neighborhoods with $\delta \gg 1$. Unfortunately, this instruction cannot be correct either, because to talk about any curvature convergence, we really need $\delta \ll 1$ since δ must go to zero in the key convergence equation (5), which is the main motivation to set δ to its smallest possible value 1 in graphs.

The combination of the calls for $\delta \ll 1$ and $\delta \gg 1$ looks like a clear sign of unhealthy pathology, suggesting abandoning any hopes for Ollivier-to-Ricci curvature convergence in graphs. And indeed we believe it is obvious that there cannot be any such convergence in unweighted graphs, simply because δ cannot be less than 1 in them.

However, these observations do not preclude the curvature convergence in weighted graphs with positive real weights. In such graphs, the Ollivier ball radius δ is back to being a positive real number, which can tend to zero because link weights can tend to zero with some characteristic rate ϵ . Suppose also that ϵ and δ tend to zero such that $\epsilon \ll \delta$. Then, on

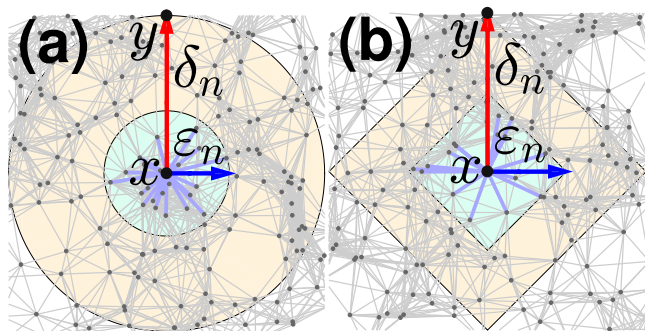


FIG. 1. Random geometric graphs in (a) Riemannian space and (b) Lorentzian spacetime. The figure visualizes the flat two-dimensional Euclidian plane (a) and Minkowski spacetime (b), where distances are the ℓ_2 and ℓ_1 norms, respectively. Nodes are sprinkled randomly via the Poisson point process with rate $r = n$. The blue regions are the balls $B_M(x, \varepsilon_n)$ of radius ε_n around node x located at the origin. In (b) this ball is the intersection of the past light cone centered at time $t = \varepsilon_n$ with the future light cone centered at time $t = -\varepsilon_n$. Node x is connected to all other nodes in the graph that happen to lie within these blue balls. The yellow regions are the Ollivier balls $B_M(x, \delta_n)$ of radius $\delta_n \geq \varepsilon_n$. Node y is at distance δ_n from x . Vector v (not shown) is the tangent vector at x toward y . The probability distribution μ_x is the uniform distribution over all the nodes z that happen to lie within distance δ_n from x in the graph, $z \in B_G(x, \delta_n)$. These are not exactly the nodes lying within the yellow Ollivier balls, but at large n the difference is negligible, Appendix A 2.

the one hand, we have that $\delta \rightarrow 0$ as needed for convergence (5), while on the other hand, the hopwise lengths of the radii of Ollivier’s balls are $\delta/\varepsilon \gg 1$, so that their sizes actually grow in terms of number of nodes. This way we simultaneously satisfy the two requirements above: the Ollivier balls are small ($\delta \ll 1$) and large ($\delta/\varepsilon \gg 1$) at the same time, and hence we call them *mesoscopic*. It is intuitive to think of ε and δ as the microscopic and mesoscopic scales of the system.

The outlined approach is very general and can be implemented in a variety of situations. In particular, the graphs do not really have to be weighted, but δ must necessarily play a role of a properly rescaled graph distance. In the next section, we make these ideas concrete in application to random geometric graphs.

B. Random geometric graphs

To describe our implementation of the general ideas from the previous subsection in application to random geometric graphs (RGGs), we first recall what RGGs are.

Given any Riemannian manifold, which henceforth we will often call just *space*, the RGGs in it are defined constructively via the following two-step procedure: (1) sprinkle points uniformly at random in the space via a Poisson point process of rate $r > 0$ driven by the volume form defined by the metric in the space [68], and (2) connect by edges all pairs of points whose pairwise distances in the space are smaller than threshold $t > 0$, Fig. 1(a). In topology, RGGs are a very fundamental object because they are 1-skeletons of Vietoris-Rips complexes [69] whose topology was proven to converge to the space topology under very mild assumptions [70]. In

that context, our results below show that in a proper sense, the geometry of RGGs also converges to the space geometry.

We are now ready to describe our ε, δ settings with RGGs in the continuum limit $n \rightarrow \infty$ suitable for curvature convergence investigations.

Given any nice D -dimensional Riemannian manifold M , we first fix any point x in it and any unit tangent vector v at x . We are concerned with Ricci curvature at x in the v direction. For a given n , we add the second point y at distance δ_n from x , $d_M(x, y) = \delta_n$, along the geodesic from x in the v direction. The distance δ_n is also the radius of the balls in the Ollivier curvature definition above, so that it goes to zero in the limit, $\delta_n \rightarrow 0$.

We then add the Poisson point process (PPP) of rate $r \sim n$ in M to the two nonrandom points x, y . Henceforth, $f(n) \sim g(n)$ means $f(n)/g(n) \rightarrow c \in (0, \infty)$, and $f(n) \approx g(n)$ means $c = 1$. The connection distance threshold t in RGGs defined on top of this PPP + $\{x, y\}$ is then set to $t = \varepsilon_n \rightarrow 0$. That is, all pairs of points are linked in RGG G if the distance between them in M does not exceed ε_n . In the continuum limit $n \rightarrow \infty$ of this setup, our PPP samples the manifold increasingly densely (as there are on average $r \sim n$ vertices in G per unit volume in M), while the RGGs become increasingly “microscopic” (as their edges span M ’s distances no greater than $\varepsilon_n \rightarrow 0$).

Every created edge is then weighted as described in the next section. For any two vertices z_1 and z_2 in RGG G , we thus have two distances defined: the manifold distance $d_M(z_1, z_2)$ and the graph distance $d_G(z_1, z_2)$, which is the length of the shortest path between z_1 and z_2 in the *weighted graph* G .

The Ollivier balls centered at x, y are then the sets of G ’s vertices z lying within *weighted graph* distances $d_G \leq \delta_n$ from vertices x, y : $B_G(x, \delta_n) = \{z \in G : d_G(x, z) \leq \delta_n\}$, $B_G(y, \delta_n) = \{z \in G : d_G(y, z) \leq \delta_n\}$. The probability distributions μ_x, μ_y are the uniform distributions on the vertex sets $B_G(x, \delta_n), B_G(y, \delta_n)$:

$$\mu_x(z) = \frac{1}{|B_G(x, \delta_n)|} \quad \text{if } z \in B_G(x, \delta_n), \quad (6)$$

$$\mu_y(z) = \frac{1}{|B_G(y, \delta_n)|} \quad \text{if } z \in B_G(y, \delta_n). \quad (7)$$

Finally, Ollivier *graph* curvature $\kappa_G(x, y)$ is defined by the same equation (4), except that probability distributions μ_x, μ_y are given by Eqs. (6) and (7), Fig. 1(a):

$$\kappa_G(x, y) = 1 - \frac{W_G(\mu_x, \mu_y)}{\delta_n}. \quad (8)$$

We reiterate that as opposed to previous literature on Ollivier graph curvature, we do not require the RGG connection radius ε_n to be equal to the Ollivier ball radius δ_n , Fig. 1(a). However, we do not exclude this possibility either. In fact, we will see in the next section that the Ollivier-to-Ricci curvature convergence does hold even with $\varepsilon_n = \delta_n$ if links are weighted by the manifold distances between the linked nodes, and if graphs are sufficiently dense. That is, our only requirements to ε_n and δ_n are that they both must go zero, and that $\varepsilon_n \leq \delta_n$.

IV. RESULTS

We obtain curvature convergence results for two different schemes of edge weighting in our random geometric graphs. We consider them separately in the next two sections.

A. Space distance weights

We first state our results for the settings in the previous section with the weights w_{ij} of all edges i, j in RGGs G set equal to the distances between i and j in manifold M :

$$w_{ij} = d_M(i, j). \tag{9}$$

With these settings, we can show that in the continuum limit, rescaled Ollivier curvature in random graph G converges to Ricci curvature in manifold M ,

$$\lim_{n \rightarrow \infty} \left\langle \left| \frac{\kappa_G(x, y)}{\delta_n^2} - \frac{\text{Ric}(v, v)}{2(D+2)} \right| \right\rangle = 0, \tag{10}$$

if the RGG connectivity radius ε_n and Ollivier ball radius δ_n shrink to zero with n as

$$\varepsilon_n \sim n^{-\alpha}, \quad \delta_n \sim n^{-\beta}, \tag{11}$$

with exponents α, β satisfying

$$0 < \beta \leq \alpha, \quad \alpha + 2\beta < \frac{1}{D}. \tag{12}$$

The connectivity and Ollivier balls can be equal and can shrink at the same rate,

$$\varepsilon_n \sim \delta_n \sim n^{-\alpha}, \tag{13}$$

in which case (12) becomes

$$\alpha < \frac{1}{3D}. \tag{14}$$

The expectation $\langle \cdot \rangle$ in (10) is with respect to the RGG ensemble, in which $\kappa_G(x, y)$ is random. The result in (10) is strong in the sense that it implies not only the convergence of the expected value of Ollivier graph curvature to Ricci curvature generalizing (5) to graphs,

$$\lim_{n \rightarrow \infty} \frac{\langle \kappa_G(x, y) \rangle}{\delta_n^2} = \frac{\text{Ric}(v, v)}{2(D+2)}, \tag{15}$$

but also the concentration of random $\kappa_G(x, y)$ around its expected value,

$$\lim_{n \rightarrow \infty} \text{Prob} \left[\left| \frac{\kappa_G(x, y)}{\delta_n^2} - \frac{\text{Ric}(v, v)}{2(D+2)} \right| > \varepsilon \right] = 0 \tag{16}$$

for any $\varepsilon > 0$.

We note that the conditions (12) allow graphs to be arbitrarily sparse, but not exactly ultrasparse. We call graphs *dense*, *sparse*, and *ultrasparse* (also *truly sparse*) if their expected average degree

$$\bar{k}_n \sim n \text{vol}[B_M(\cdot, \varepsilon_n)] \sim n \varepsilon_n^D \sim n^{1-\alpha D} \tag{17}$$

is $O(n)$, $o(n)$, and $\sim \text{constant}$, respectively. Observe that the exponent β in (12) cannot be exactly zero because the Ollivier balls must shrink for the sake of convergence, but it can be arbitrarily close to zero. The closer the β to zero, the closer the α can be to $1/D$ according to (12), while $\alpha = 1/D$ corresponds

to the ultrasparse case with $\bar{k}_n \sim \text{constant}$. We note though that the closer the β to zero, the slower the convergence, simply because the Ollivier balls shrink too slowly in this case.

Another important observation is that the case with $\varepsilon_n = \delta_n$ corresponds to the traditional definition of Ollivier graph curvature with microscopic one-hop neighborhoods. In this case, Eq. (14) says that we cannot prove the convergence anywhere close to the ultrasparse limit with $\alpha = 1/D$, and our simulations suggest that there is no convergence there indeed.

In general, the sparser the graphs, the more problematic the convergence becomes. This is because the sparser the graphs, the smaller the connection radius ε_n . As discussed in Sec. III, if ε_n is too small, and if $\delta_n = \varepsilon_n$, then all balls are of the radius of one hop, and such microscopic neighborhoods may not “feel” any curvature; all they can “feel” is locally Euclidean flatness.

1. Proofs

Our proof of (10), outlined in Appendix A, is not trivial. It consists of three major steps.

First, we extend the weighted shortest path distance in RGGs G to space M . We do so by defining the following auxiliary distance in M . For any pair of points $x, y \in M$ we add x and y as nodes to the graph G , and connect them to all other nodes in G within some distance $\lambda_n \ll \varepsilon_n$ from x, y in M . The auxiliary manifold distance between x, y is then the distance between them in the extended graph. We must select λ_n very carefully, to make sure the auxiliary distance is a sufficiently good approximation of the true distance in M .

Once we have this auxiliary distance in M , we need to approximate the Ollivier balls in RGGs G with the corresponding balls in M . That is, on the one hand we have sets of nodes within graph distance δ_n from node x , while on the other hand we have sets of nodes within manifold distance δ_n from x . We work out this approximation by deriving an upper bound on the Wasserstein distance between the uniform probability distributions on these two balls. This allows us to work with balls defined by manifold distances, instead of graph distances.

Finally, the most nontrivial step, we must relate the discrete uniform probability distribution defined on finite sets of nodes in balls in G to the corresponding nonuniform continuous distributions defined on the continuum of points in balls in M . To this end, using a result on matching two Poisson point processes, we first bound the Wasserstein distance between these two distributions in Euclidean space. We finally extend this result to non-Euclidean M , using the local flatness of the space together with the fact that the neighborhoods that we consider shrink as $\delta_n \rightarrow 0$.

Overall, at each step of this proof, we introduce an approximation error to the relation between the Wasserstein distances in graphs G and manifold M . The crux of each step is that it must be done with care to ensure that the introduced error is upper bounded by $\delta_n^3 \rightarrow 0$. This is because if the error is that small, then (10) follows from (5). The further details behind how we achieve this at each step of the proof can be found in Appendix A.

2. Simulations

The fact that we can prove the Ollivier \rightarrow Ricci convergence only under the conditions (12) does not mean that there is no convergence outside of this parameter region, and this is indeed what we observe in simulations.

For simulations we select the sphere, torus, and the Bolza surface as the three 2D manifolds of constant Ricci curvature +1, 0, and -1 . The Bolza surface [71,72] is the simplest hyperbolic 2D manifold of genus 2 with no boundaries. We consider the Bolza surface, versus something simpler of negative curvature, such as a hyperbolic disk, because we want our manifolds to have no boundaries. We want this, because computing Ollivier curvature in simulations is a major challenge at large n , so that we do not want any boundary effects at any n .

We then fix x and y as described above, sprinkle n points on the three manifolds uniformly at random according to the manifold volume form, and link all pairs of points at distances $\leq \varepsilon_n$ on the manifold. We then find the Ollivier balls $B_G(x, \delta_n), B_G(y, \delta_n)$, set up the uniform probability distributions $\mu_x = 1/|B_G(x, \delta_n)|, \mu_y = 1/|B_G(y, \delta_n)|$ on them, and compute all the pairwise distances $d_G(x_i, y_j)$ in graph G between nodes $i \in B_G(x, \delta_n)$ and $j \in B_G(y, \delta_n)$ whose coordinates are x_i and y_j . All these data allow us to compute the Wasserstein distance $W_G(\mu_x, \mu_y)$ by solving the linear program

$$W_G(\mu_x, \mu_y) = \min_{\rho_{ij}} \sum_{i,j} d_G(x_i, y_j) \rho_{ij} \mu_x, \text{ such that}$$

$$\sum_i \rho_{ij} \mu_x = \mu_y, \quad 0 \leq \rho_{ij} \leq 1, \quad \sum_j \rho_{ij} = 1, \quad (18)$$

where the minimization is over transportation plans ρ_{ij} whose entries describe pairwise movements of probability masses. Having computed this $W_G(\mu_x, \mu_y)$, Ollivier curvature is finally given by (8). Further details on the simulations are in Appendix C.

The results are shown in Fig. 2. Remarkably, we observe that rescaled Ollivier curvature in our RGGs converges to Ricci curvature of the underlying manifold not only in the parameter regime accessible to our proofs, Fig. 2(a), but also well outside of this regime, for much sparser graphs, Fig. 2(b). In particular, we observe that the convergence is as good for truly sparse graphs with constant average degree as for denser graphs.

B. ε_n weights

On the one hand, we may be not too happy that the weighted RGGs considered in the previous section contain too much information about the manifold in their edge weights $w_{ij} = d_M(i, j)$, helping the convergence “too much.” On the other hand, we cannot be happy at all if edges are unweighted, $w_{ij} = 1$, because in this case there cannot be any curvature convergence whatsoever for the reasons discussed in Sec. III. The maximum-happiness point appears to be the middle-ground possibility discussed in Sec. III—all edges are weighted by the same weight decaying to zero and playing the role of a system scale, à la the Planck scale. This “Planck scale” in our case is the RGG connectivity radius ε_n . That is,

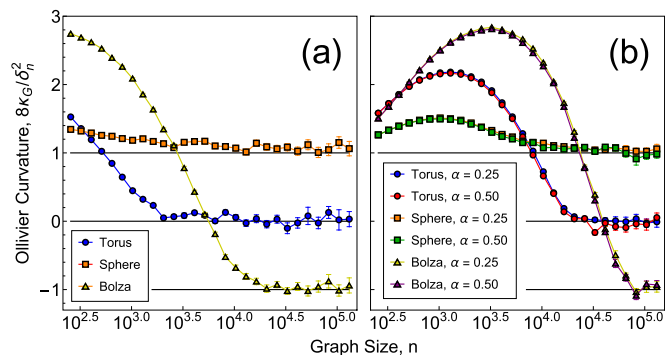


FIG. 2. Ollivier-Ricci curvature convergence in constant-curvature random geometric graphs. Panel (a) shows the simulation data for rescaled Ollivier curvature (8) in (9)-weighted random geometric graphs on the ($D = 2$)-dimensional torus, sphere, and Bolza surface. The scaling exponents of the connectivity and the Ollivier ball radii (11) in the graphs are set to $\alpha = \beta = 0.16 < 1/6$ lying within the proof-accessible regime (14). Panel (b) shows apparent convergence for $\alpha = \{1/4, 1/2\}$ and $\beta = 1/4$. Both of these two settings lie well outside of the proof-accessible regime (12), while $\alpha = 1/2$ corresponds to truly sparse graphs with constant average degree. The error bars represent the standard error $\sigma/\sqrt{n_s}$, where σ is the standard deviation and $n_s = 10 \times 2^{17}/n$ is the number of sampled random graphs of size n . The smallest graph size n is chosen to ensure the graphs are connected, while the largest size $n = 2^{17}$ is bounded by memory constraints of the MOSEK package.

the weights w_{ij} of all edges i, j in our RGGs in this section are

$$w_{ij} = \varepsilon_n. \quad (19)$$

This weighting scheme is equivalent to saying that we actually work with *unweighted* graphs, except that we multiply by ε_n all distances in them, which are hop lengths of shortest paths. If $d_G(z_1, z_2)$ is the shortest path hop distance between vertices z_1 and z_2 in the unweighted RGG G , then the rescaled graphs distance $d'_G(z_1, z_2) = \varepsilon_n d_G(z_1, z_2)$ is actually a good approximation to the manifold distance $d_M(z_1, z_2)$. Unfortunately, how good this approximation is (characterized by *stretch* d'_G/d_M) has been rigorously quantified only for the RGGs in the Euclidean plane [73]. Since the distances we approximate are bounded by the Ollivier ball radius δ_n that tends to zero, and since any Riemannian manifold is locally Euclidean, our results below relying on [73] hold for any 2D manifold. In fact, it turns out that the results in [73] are strong enough for our proof strategy in the weighted case to work essentially without a change in this unweighted case as well, Appendix B. While our result below holds only for $D = 2$, analogous results for higher dimensions can be worked out as soon as stretch results for higher-dimensional RGGs are obtained.

The result in Appendix B is that with weights in (19), the Ollivier \rightarrow Ricci convergence holds in the same strong sense (10) if

$$0 < \beta < \frac{1}{9}, \quad 3\beta < \alpha < \frac{1 - 3\beta}{2}. \quad (20)$$

Similarly to the weighted case, these conditions allow for arbitrarily sparse but not ultrasparse graphs ($\beta \rightarrow 0, \alpha \rightarrow$

1/2). Unlike the weighted case, these conditions imply that, as expected, the convergence is slower: $\beta < 1/9$ ($2D$ unweighted) versus $\beta < 1/6$ ($2D$ weighted). Another difference with the weighted case is that these conditions do not allow for the connectivity radius ε_n to be the same as the Ollivier ball radius δ_n : $\alpha \geq 3\beta \Rightarrow \delta_n \gg \varepsilon_n$. We believe this is not a deficiency in our proof techniques, but a reflection of reality. It seems obvious that one must consider mesoscopic neighborhoods of size $\delta_n \gg \varepsilon_n$ in unweighted graphs since one-hop microscopic ε_n neighborhoods in these graphs are too small to feel any curvature.

V. DISCUSSION

We have established a rigorous connection between curvatures of discrete and continuous objects—random graphs and Riemannian manifolds. Anecdotally, this connection is reminiscent of what Riemann had in mind working on the foundations of Riemannian geometry [74]. The closest result to ours, which deals with convergence of curvature of simplicial triangulations of manifolds in Regge calculus [17,50] and its many derivatives [75], is very different. Informally, while any simplex in any simplicial triangulation of any manifold is a chunk of space, graphs are much more primordial objects in that they do not have any space attached to them whatsoever.

We saw why the convergence of discrete curvature to a continuous one is such a delicate matter. If we use the standard simple definition of Ollivier graph curvature, then the immediately obvious observation is that there cannot be any convergence. We really have to consider mesoscopic regions in graphs with properly rescaled distances for the convergence to take place, an observation that has been entirely overlooked in the past. Even then the graph density and mesoscopic region sizes cannot be arbitrary, but should enter in a delicate interplay. These nontrivial delicate points are one reason why there are so few results on discrete \rightarrow continuous curvature convergence.

Our convergence proofs can likely be improved since our simulations suggest that the convergence holds well outside of the parameter regimes accessible by our proof techniques. Our simulations can be improved as well. We could not simulate graphs large enough to see any convergence in ε_n -weighted graphs, where the convergence is slow. In any case, the investigation of the exact boundaries between regimes where the convergence does and does not take place is a wide open problem, both for Ollivier curvature and for all other discrete curvature definitions mentioned in Sec. II, for most of which it is not even known under what circumstances, if any, they do converge.

If some discrete curvature does not converge to a continuous one for some discrete structure, then it may very well happen that this discrete curvature tells us something very different about the structure from what is intended. In such a case, why should we call this “curvature” *curvature*? This problem is not semantic but practical. If we are sure that graph curvature converges to manifold curvature, then we can be also sure that our graph curvature can be used to learn reliably manifold geometry—an exciting bit of news for manifold learning and many other network or data science applications mentioned in the introduction. As an example of the other

extreme, if we are sure that curvature of some tensor network does not converge to Ricci curvature of classical spacetime, then we can be also sure that this tensor network is not a good model of quantum spacetime.

As far as possible applications to quantum gravity are concerned, our results are an important but very high-level step that does not touch on many critical aspects. We comment on some of them.

First of all, we proved the convergence of curvature only for graph links and paths. Their curvatures first need to be averaged to result in scalar Ollivier-Ricci curvature at a vertex, and then integrated over vertex sets to yield a discrete version of the Einstein-Hilbert action. Is there any convergence in this case, and if so, then under what circumstances? Such convergence results would indicate unambiguously that the discrete Einstein-Hilbert action can be constructed from Ollivier curvature.

A much more difficult problem is the following. What we have solved is the direct problem of graph curvature convergence to space curvature when space is there. The problem in quantum gravity is actually the inverse problem of geometrogenesis [28,29,76]—there is no space, only a discrete quantum structure, which must “look like” space in the classical limit. That is, an illusion of continuous space must somehow emerge in the limit.

One strategy to address this problem is conceptually similar to [77]. One can consider a canonical maximum-entropy Gibbs ensemble of random graphs with fixed values of Ollivier curvature on links, paths, or vertices. The questions then are as follows: is this ensemble equivalent (in the continuum limit) to the ensemble of random geometric graphs on a Riemannian manifold with the corresponding values of Ricci curvature, and do we have any (second-order) phase transitions in this ensemble? If the answer to the first question is *yes*, then the graphs in this ensemble do look like spatial graphs, even though there is no space in their definition. If the connection between Ollivier curvature and the Einstein-Hilbert action is there, then the temperature parameter in this ensemble must be analogous to the gravitational coupling constant. In that context, the settings considered in this paper are particularly interesting. Indeed, since Ollivier curvature (8) is well defined for any shortest path between any pairs of nodes x, y located at any distance in the graph, one can probe graph curvature at *any* scale for renormalization and other purposes.

Finally, yet another critical issue is the issue of time, causality, and Lorentzian metric signature. There are two schools of thought concerning the issue. The first school believes that time and Lorentzian signature must be also emergent. This route is followed, for example, in the combinatorial quantum gravity program [28,29], where time and causality are expected to appear only in the continuum limit above a certain scale, below which spacetime is Riemannian, turning into a random graph at a second-order phase transition defining quantum gravity nonperturbatively.

The second school builds in time and causality *ab initio*, as in causal sets or causal dynamical triangulation. For our results to be applicable in these settings, one needs first to extend them from Riemannian spaces to Lorentzian spacetimes. To the best of our knowledge, no notion of random geometric

graphs or Rips complexes has been defined for Lorentzian manifolds. Since intersections of past and future light cones, Fig. 1(b), known as *causal diamonds* or *Alexandroff sets*, play the role of balls in Lorentzian geometry—specifically, the topology defined by Alexandroff sets agrees with the base topology in any nice Lorentzian spacetime [78]—we propose to define Lorentzian random geometric graphs as shown in Fig. 1(b). This definition is very different from causal sets, where any two time-like-separated nodes are linked, and there are no other links. In the Fig. 1(b) definition of Lorentzian random geometric graphs, there are both time-like and space-like links but only between nodes within causal diamonds of a finite size. Can Ollivier curvature be defined for these graphs and shown to converge to Ricci curvature of the underlying Lorentzian manifold?

ACKNOWLEDGMENTS

We thank R. Loll, B. Dittrich, S. Surya, D. Rosset, E. Andersen, Y. Zhou, J. Jost, and A. Aranovich for useful discussions, suggestions, and comments. This work was supported by ARO Grants No. W911NF-16-1-0391 and No. W911NF-17-1-0491, and by NSF Grants No. IIS-1741355 and No. DMS-1800738. Research at Perimeter Institute is supported by the Government of Canada through the Department of Innovation, Science, and Economic Development and by the Province of Ontario through the Ministry of Colleges and Universities. Experiments were conducted using the Discovery cluster, supported by Northeastern University’s Research Computing team, and the B eluga, Graham, and Cedar clusters, supported by Compute Canada and its regional partners Calcul Qu ebec, Compute Ontario, and Westgrid.

APPENDIX A: OLLIVIER → RICCI CONVERGENCE IN DISTANCE-WEIGHTED RGGs

Here we outline the main steps and ingredients in our curvature convergence proof. All further lower-level details can be found in [79].

Let μ_x^M and μ_y^M denote the normalized restrictions of the volume form in manifold M onto the balls $B_M(x, \delta_n), B_M(y, \delta_n)$ in M , while μ_x^G and μ_y^G denote the uniform probability distribution over nodes lying within the balls $B_G(x, \delta_n), B_G(y, \delta_n)$ in weighted RGG G . Observe that if

$$\left| |W_G(\mu_x^G, \mu_y^G) - W_M(\mu_x^M, \mu_y^M)| \right| \ll \delta_n^3, \tag{A1}$$

then (10) follows from (5). Indeed, if (A1) holds, then

$$\begin{aligned} \left| |\kappa_G(x, y) - \kappa_M(x, y)| \right| &= \frac{\left| |W_G(\mu_x^G, \mu_y^G) - W_M(\mu_x^M, \mu_y^M)| \right|}{\delta_n} \\ &\ll \delta_n^2, \end{aligned}$$

so that

$$\left\langle \left| \frac{\kappa_G(x, y)}{\delta_n^2} - \frac{\kappa_M(x, y)}{\delta_n^2} \right| \right\rangle \ll 1.$$

The main idea of the proof is thus to show (A1). We proceed in three steps.




FIG. 3. Illustration of the construction of the extended graph distance \tilde{d}_M . The blue points are nodes of the RGG and the two black points are the selected points on the manifold. The blue circles indicate the connection radius ϵ_n , while the black circles have radius λ_n . In this example x is connected to node v_1 and y to node v_6 .

1. Approximating graph distances by manifold distances

The first difficulty we face is that the Wasserstein distances W_G and W_M are in different spaces with different distances: on the one hand we have RGG G with weighted shortest path distance d_G , and on the other hand we have the manifold M with distances d_M . How can we compare the two then? We simply extend d_G to a new distance \tilde{d}_M on M , such that the difference between the new and original Wasserstein distances \tilde{W}_M and W_M is $\ll \delta_n^3$.

This extension is accomplished by the following procedure. Let

$$\lambda_n = \ln(n)^{2/D} n^{-1/D}, \tag{A2}$$

so that

$$n\lambda_n^D \rightarrow \infty. \tag{A3}$$

If the conditions (12) hold, then we have

$$\lambda_n \ll \epsilon_n, \tag{A4}$$

$$\lambda_n \ll \delta_n^3. \tag{A5}$$

Given any two points x, y in M we extend the RGG G by adding x and y to the set of G ’s nodes, and connecting both x and y to all other nodes in G whose distance from x and y , respectively, in the manifold is less than λ_n , Fig. 3. We then define the new manifold distance $\tilde{d}_M(x, y)$ as the weighted shortest path distance in the extended graph. This new manifold distance \tilde{d}_M is well defined, with probability converging to 1, because the expected number of points of a PPP of rate n in a ball of radius λ_n is $n \text{vol}[B_M(x, \lambda_n)] \sim n\lambda_n^D$, so that the probability to find a graph vertex within distance λ_n from any point x in M is given by $1 - e^{-n \text{vol}[B_M(x, \lambda_n)]}$ which converges to 1 thanks to (A3). Observe that if x and y happen to be already nodes in G , then the new manifold distance between them is the graph distance, $\tilde{d}_M(x, y) = d_G(x, y)$.

We now have to show that for any two points x and y in M , the distance $\tilde{d}_M(x, y) = d_M(x, y) + o(\delta_n^3)$. As evident from Fig. 3, $\tilde{d}_M(x, y) \geq d_M(x, y)$. For the upper bound, let us partition the geodesic between x and y into $m = \lceil 3d_M(x, y)/\epsilon_n \rceil$ sections of equal length, and let $c_1 = x, c_2, \dots, c_m, c_{m+1} = y$




FIG. 4. Approximating the manifold distance by the extended weighted shortest path graph distance.

denote the $m + 1$ end points of this partition, Fig. 4. Consider the $m + 1$ balls $B_M(c_i, \lambda_n/4)$. The probability that every ball contains a graph vertex converges to 1. Denote those $m + 1$ vertices by v_1, \dots, v_{m+1} . We then observe the following:

(1) the node v_1 is at distance at most $\lambda_n/2$ from x , and so is v_m from y ;

(2) the distance between each consecutive pair of nodes v is bounded by

$$d_M(v_i, v_{i+1}) \leq d_M(x, y)/m + \lambda_n \leq \varepsilon_n/3 + \lambda_n;$$

$$(3) \quad d_G(v_i, v_{i+1}) \leq d_M(c_i, c_{i+1}) + \lambda_n.$$

From these observations and (A4), we conclude that each consecutive pair of nodes v is connected in the RGG and

$$\begin{aligned} \tilde{d}_M(x, y) &\leq \lambda_n + \sum_{i=1}^m d_M(v_i, v_{i+1}) \\ &\leq \lambda_n + d_M(x, y)(1 + 3\lambda_n/\varepsilon_n). \end{aligned}$$

If $d_M(x, y) \leq C\delta_n$, for some $C > 0$, then

$$|d_M(x, y) - \tilde{d}_M(x, y)| \leq d_M(x, y) \frac{3\lambda_n}{\varepsilon_n} + \lambda_n \quad (\text{A6})$$

$$\leq 3C\lambda_n\delta_n\varepsilon_n^{-1} + \lambda_n \quad (\text{A7})$$

$$\sim \ln(n)^{\frac{2}{\beta}} n^{-\frac{1}{\beta} - \beta + \alpha} + \lambda_n. \quad (\text{A8})$$

The second term in the last equation is $\ll \delta_n^3$ by (A5), and so is the first term if (12) holds. We therefore conclude that

$$\tilde{d}_M(x, y) = d_M(x, y) + o(\delta_n^3).$$

Finally, the Wasserstein distances W_M and \tilde{W}_M are defined by distances d_M and \tilde{d}_M , respectively. Therefore we immediately conclude that

$$|W_M(\mu_1, \mu_2) - \tilde{W}_M(\mu_1, \mu_2)| \ll \delta_n^3$$

for any two probability distributions μ_1 and μ_2 on M . If now x, y are also any two nodes of our RGG G , then $d_G(x, y) = d_M(x, y)$, so that it follows that if μ_1 and μ_2 are defined on G 's nodes, then

$$|W_M(\mu_1, \mu_2) - W_G(\mu_1, \mu_2)| \ll \delta_n^3. \quad (\text{A9})$$

The calculations above thus show that from now on we can always work only with the Wasserstein distance W_M in the space M , and this indeed is what we do in the next steps.

2. Approximating probability distributions in graphs

The next hurdle we need to overcome is comparing the uniform probability distribution μ_x^G on $B_G(x, \delta_n)$ with μ_x^M , the normalized restriction of the volume form to $B_M(x, \delta_n)$. The main difficulty here lies in that $d_M(x, z) \leq \delta_n$ does not necessarily imply that $d_G(x, z) \leq \delta_n$. In general,

the shape of the intersection $B_G(x, \delta_n) \cap B_M(x, \delta_n)$ can be highly nontrivial. We tame this shape by introducing a new probability distribution $\hat{\mu}_x^G$ which is the uniform probability distribution on the set $B_G := B_M(x, \delta_n) \cap G$ which is the set of G 's nodes that happen to lie within distance δ_n from x in the space, i.e., all nodes in the rose ball in Fig. 1(a). Observe that $B_G(x, \delta_n) \subseteq B_G$ since $d_G(x, y) \geq d_M(x, y)$. The main goal of this step is to show that $\langle W_M(\mu_x^G, \hat{\mu}_x^G) \rangle \ll \delta_n^3$.

First, consider (A6). If the conditions (12) hold, then $\lambda_n \ll \delta_n^3$ and $\lambda_n/\varepsilon_n \ll \delta_n^2$. It follows then that there exists a $\xi_n \ll \delta_n$ such that

$$|d_M(x, y) - d_G(x, y)| \leq d_M(x, y)\xi_n^2 + \xi_n^3. \quad (\text{A10})$$

For instance, we can take the ξ_n to be

$$\xi_n = \max \{ \sqrt{3\lambda_n/\varepsilon_n}, \lambda_n^{1/3} \}.$$

Define now the new radii $\delta_n^\pm = (\delta_n \pm \xi_n^3)/(1 \mp \xi_n^2)$, and note that $\delta_n^- < \delta_n < \delta_n^+$, so that

$$B_M(x, \delta_n^-) \subset B_M(x, \delta_n) \subset B_M(x, \delta_n^+).$$

Let $B_G^\pm := B_M(x, \delta_n^\pm) \cap G$ and μ_x^\pm denote the uniform probability distribution on B_G^\pm . We then obtain an upper bound for $W_M(\hat{\mu}_x^G, \mu_x^+)$ by considering the following joint probability distribution, also known as transport plan, μ for $(u, v) \in B_G \times B_G^+$:

$$\mu(u, v) = \begin{cases} \frac{1}{|B_G^+|}, & \text{if } u = v \in B_G, \\ \frac{1}{|B_G||B_G^+|}, & \text{if } u \in B_G, v \in B_G^+ \setminus B_G, \\ 0, & \text{otherwise.} \end{cases}$$

Since the Wasserstein distance $W_M(\hat{\mu}_x^G, \mu_x^+)$ is the infimum over all joint distributions (1), and since $\hat{\mu}_x^G$ and μ_x^+ are discrete, it follows that

$$\begin{aligned} W_M(\hat{\mu}_x^G, \mu_x^+) &\leq \sum_{u \in B_G, v \in B_G^+} d_M(u, v) \mu(u, v) \\ &= \sum_{u \in B_G, v \in B_G^+ \setminus B_G} \frac{d_M(u, v)}{|B_G||B_G^+|} \\ &\leq 2\delta_n^+ \frac{|B_G^+ \setminus B_G|}{|B_G^+|}. \end{aligned}$$

Since we are dealing with a PPP, both $|B_G^+ \setminus B_G|$ and $|B_G^+|$ are Poisson distributed with diverging means. Although they are not completely independent, it can be shown that

$$\left\langle \frac{|B_G^+ \setminus B_G|}{|B_G^+|} \right\rangle \sim \frac{\langle |B_G^+ \setminus B_G| \rangle}{\langle |B_G^+ \setminus B_G| \rangle + \langle |B_G| \rangle} \sim \frac{(\delta_n^+)^D - (\delta_n^-)^D}{\delta_n^D}.$$

This fraction is of the order $\xi_n^2 \ll \delta_n^2$, so that we conclude

$$\langle W_M(\hat{\mu}_x^G, \mu_x^+) \rangle \ll \delta_n^+ \delta_n^2 \sim \delta_n^3.$$

Finally, let node z lie in the manifold ball of radius δ_n^- around node x , $d_M(x, z) \leq \delta_n^-$. Then (A10) implies that the graph distance between x and z is upper bounded by δ_n :

$$d_G(x, z) \leq d_M(x, z)(1 + \xi_n^2) + \xi_n^3 \leq \frac{\delta_n - \xi_n^3}{1 + \xi_n^2} (1 + \xi_n^2) + \xi_n^3 = \delta_n.$$

This means that $z \in B_G(x, \delta_n)$, so that we have the following sandwich:

$$B_G^- \subseteq B_G(x, \delta_n) \subseteq B_G^+.$$

Using calculations similar to the ones above for this sandwich, we obtain

$$\langle W_M(\mu_x^G, \mu_x^+) \rangle \ll \delta_n^3.$$

Therefore, by the triangle inequality of the W_M distance,

$$\langle W_M(\mu_x^G, \hat{\mu}_x^G) \rangle \leq \langle W_M(\mu_x^G, \mu_x^+) \rangle + \langle W_M(\mu_x^+, \hat{\mu}_x^G) \rangle \ll \delta_n^3. \tag{A11}$$

The calculations above thus show that from now on we can work with the probability distribution $\hat{\mu}_x^G$ instead of μ_x^G . The former is more convenient to work with than the latter because it is over G 's nodes that happen to lie within manifold ball $B_M(x, \delta_n)$, the rose ball in Fig. 1(a).

3. Going from discrete to continuous probability distributions

The final step is to go from the discrete probability distribution $\hat{\mu}_x^G$ to the continuous distribution μ_x^M . Similarly to the goal of the previous section, the task is to show that

$$\langle W_M(\hat{\mu}_x^G, \mu_x^M) \rangle \ll \delta_n^3. \tag{A12}$$

This is done by applying results on the matching distance between a PPP and points of a grid on the same space [80–83].

Note that both $\hat{\mu}_x^G$ and μ_x^M are now defined on $B_M(x, \delta_n)$. We now need to devise a transport plan from $\hat{\mu}_x^G$ to μ_x^M that assigns to any (measurable) set $A \subseteq B_M(x, \delta_n)$ how much mass from each PPP point in $B_G = G \cap B_M(x, \delta_n)$ is used to make up $\mu_x^M(A)$. We do so as follows.

First, suppose that space M is flat, the D -dimensional Euclidean space. We construct our transport plan in three steps:

- (1) we first place a grid on the space;
- (2) we next find a minimal matching between this grid and the PPP points in B_G ; and
- (3) finally, for each (measurable) set A in $B_M(x, \delta_n)$, we find all the PPP points in B_G that are matched to the grid points that lie in A .

Observe that this transport plan is such that every PPP point, i.e., graph vertex in B_G contributes an equal fraction of its mass to make up $\mu_x^M(A)$. This plan is also such that the largest distance any amount of mass has to move is given by the largest matching distance. Relying on the results from [80–83], it can then be shown that

$$\langle W_M(\hat{\mu}_x^G, \mu_x^M) \rangle \sim \ln(n)n^{-1/D} \ll \delta_n^3, \tag{A13}$$

where the last inequality holds if the conditions (12) are satisfied.

Next, we relax the condition that the space M is flat. Indeed, it can be any nice Riemannian manifolds because $\delta_n \rightarrow 0$, so that we can map the ball $B_M(x, \delta_n)$ to the flat D -dimensional tangent space at x using the exponential map. We have to be careful here though, because the exponential map does not preserve distances. Still, by fixing a sufficiently small neighborhood U around the origin of the tangent space, we can ensure that for large enough n , $\exp^{-1} B_M(x, \delta_n) \subset U$,

and the distances are distorted by a fixed small amount:

$$B_D\left(0, \frac{\delta_n}{1+\xi}\right) \subseteq \exp^{-1} B_M(x, \delta_n) \subseteq B_D\left(0, \frac{\delta_n}{1-\xi}\right) \subset U,$$

where $B_D(0, \delta)$ is the D -dimensional Euclidean ball of small radius δ around the origin in the tangent space.

Finally, relying on this mapping and on the PPP mapping theorem [68] that applies to U , we see that the mapped PPP is still a PPP with intensity $\sim n$, now relative to the Euclidean volume form. Moreover, since the mapped ball is sandwiched between two balls whose radii scale as δ_n , the results for the D -dimensional Euclidean space above yield matching lower and upper bounds for the Wasserstein distance, from which (A12) follows.

The combination of (A12), (A11), and (A9) yields (A1), completing the proof.

APPENDIX B: OLLIVIER → RICCI CONVERGENCE IN ε_n -WEIGHTED RGGs

We first observe that the last step in the proof outlined in Appendix A does not rely on any graph distances at all. Further, the key element in the second step is (A10), which is also the fundamental ingredient in the first step in showing that distance \tilde{d}_M is a good approximation to distance d_M , leading to (A9). Therefore, the convergence proof for ε_n -weighted graphs is complete as soon as (A10) is established for the new graph distance $d'_G = \varepsilon_n d_G$, where d_G is now the vanilla shortest path (hop count) distance in graph G . We also note that we have to do this only for the 2-dimensional Euclidean space, since the neighborhoods shrink and we can apply techniques similar to the ones in Appendix A3 to extend the result to any curved manifold.

We first observe that the conditions (20) imply that $\varepsilon_n \gg \sqrt{\ln(n)/n}$. In this case, the main stretch result from [73] reads

$$|d'_G(x, y) - d_M(x, y)| \leq d_M(x, y)\gamma_n + \varepsilon_n, \tag{B1}$$

where $\gamma_n = \max\{\gamma_n^{(1)}, \gamma_n^{(2)}, \gamma_n^{(3)}\}$, and where the three γ_n 's have somewhat complicated explicit expressions with different constants and different scalings with n which can be shown to be

$$\gamma_n^{(1)} \sim \left(\frac{\ln n}{n\varepsilon_n^2}\right)^{\frac{2}{3}}, \quad \gamma_n^{(2)} \sim \left(\frac{\ln n}{n\varepsilon_n^2}\right)^2, \quad \text{and} \quad \gamma_n^{(3)} \sim \left(\frac{1}{n\varepsilon_n^2}\right)^{\frac{2}{3}},$$

if the conditions (20) hold. Since these conditions also imply that $n\varepsilon_n^2 \rightarrow \infty$, it follows that

$$\gamma_n \sim \left(\frac{\ln n}{n\varepsilon_n^2}\right)^{\frac{2}{3}} \tag{B2}$$

in our case.

The ε_n in (B1) is $\ll \delta_n^3$ if the lower bound in the second condition in (20) holds. Therefore, to translate (B1) to (A10), all we need to do is to show that $\gamma_n \ll \delta_n^2$. If this holds, then we can simply take the ξ_n in (A10) to be $\xi_n = \max\{\sqrt{\gamma_n}, \varepsilon_n^{1/3}\}$. The requirement $\gamma_n \ll \delta_n^2$ is satisfied if

$$\frac{\ln n}{n\varepsilon_n^2\delta_n^3} \sim \ln(n)n^{2\alpha+3\beta-1} \ll 1, \tag{B3}$$

which is indeed true if the upper bound in the second condition in (20) holds.

APPENDIX C: NUMERICAL METHODS

The computation of Ollivier graph curvature in simulations consists of three tasks: graph construction, distance matrix computation, and Wasserstein distance computation. We rely on a hybrid multicore solution which parallelizes each of these tasks, with the first two utilizing custom CUDA methods split among multiple GPUs, and the last utilizing multiple CPU cores via the MOSEK software package.

1. Graph construction

We construct random geometric graphs of size n by first sampling n points uniformly at random, according to the manifold volume form, on the surface of a unit torus, sphere, and Bolza surface. This is trivial for the first two manifolds, and less so for the Bolza surface.

The simplest representation of the Bolza surface [71,72] is the hyperbolic octagon with vertices at complex coordinates

$$o_k = R \exp \left[\frac{i\pi}{4} \left(\frac{1}{2} + k \right) \right], \quad k = 0, 1, \dots, 7, \quad (C1)$$

where $R = 2^{-1/4}$ in the Poincaré disk model of the hyperbolic plane. When the opposite sides of this octagon are glued, the Bolza surface of constant curvature -1 is formed. To sprinkle points uniformly at random onto this octagon, we first sprinkle them uniformly at random, according to the volume form in the Poincaré model, which in the polar coordinates is

$$dV = \frac{4r dr d\theta}{(1 - r^2)^2}, \quad (C2)$$

onto the Poincaré disk of radius R , and then remove those points that lie in this disk but do not lie in the octagon. We do this removal in the Klein disk model of the hyperbolic plane, because geodesics—octagon sides, in particular—are straight lines there. The map to go from the Poincaré (r, θ) to Klein (r_K, θ) polar coordinates is $(r_K, \theta) = (2r/(1 + r^2), \theta)$, and in the latter coordinates, the coordinates of points that lie in the octagon meet the condition $r_K < r_c[\phi(\theta)]$, where

$$r_c(\phi) = R_K \frac{\cos(\pi/8)}{\cos(\pi/8 - \phi)}, \quad \phi \in [0, \pi/4], \quad (C3)$$

$$\phi(\theta) = \left(\theta - \frac{\pi}{8} [1 + 2k(\theta)] \right) \text{ mod } 2\pi, \quad (C4)$$

$$k(\theta) = \left\lfloor \frac{4}{\pi} \left(\theta - \frac{\pi}{8} \right) \right\rfloor \text{ mod } 8, \quad (C5)$$

where $R_K = 2R/(1 + R^2) = 2^{5/4}/(1 + \sqrt{2})$ is the octagon radius in the Klein model. The nodes that do not pass this test are thrown out. The remaining nodes lie in the octagon, and we work back with their Poincaré coordinates, $(r, \theta) = (1/r_K - \sqrt{1/r_K^2 - 1}, \theta)$.

After all the n nodes are sprinkled, we add two additional nodes x, y separated by distance δ_n , which are the centers of the two balls $B_G(x, \delta_n), B_G(y, \delta_n)$. Given the coordinates of all the $n + 2$ nodes, all node pairs are linked whenever

their pairwise distance on the surface is below the connection threshold ε_n . The distances on the torus and the sphere are

$$d_T(x_1, y_1; x_2, y_2) = \left[\left(\frac{1}{2} - \left| \frac{1}{2} - |x_1 - x_2| \right| \right)^2 + \left(\frac{1}{2} - \left| \frac{1}{2} - |y_1 - y_2| \right| \right)^2 \right]^{1/2}, \quad (C6)$$

$$d_S(\theta_1, \phi_1; \theta_2, \phi_2) = \arccos [\cos \theta_1 \cos \theta_2 + \sin \theta_1 \sin \theta_2 \cos(\phi_1 - \phi_2)]. \quad (C7)$$

On the Bolza surface, the distance formula is more complicated. We describe it next.

The Bolza surface can also be considered as the factor space \mathbb{H}^2/\mathcal{F} [72], where \mathbb{H}^2 is the hyperbolic plane, and the Fuchsian group \mathcal{F} is defined by its eight generators

$$g_k = \begin{pmatrix} a & be^{ik\pi/4} \\ be^{-ik\pi/4} & a \end{pmatrix}, \quad k = 0, 1, \dots, 7, \quad (C8)$$

where $a = 1 + \sqrt{2}$ and $b = \sqrt{a^2 - 1}$. This group acts on the hyperbolic plane by linear fractional transformations. Let \mathcal{N} be the set of the 48 elements of \mathcal{F} given by

$$\hat{g}_{k,\ell} = \prod_{t=0}^{\ell} g_{k+3t}, \quad (C9)$$

where $k = 0, 1, \dots, 7$ and $\ell = 0, 1, \dots, 5$. One can check that these elements map the original octagon to its 48 either side- or vertex-adjacent octagons in the tessellation of the hyperbolic plane induced by these identical octagons. Define $\mathcal{N}'_0 := \mathcal{N} \cup I$ to be these 48 elements plus the identity element I . One can show that the distance on the Bolza surface between two points with complex coordinates z_1, z_2 ($z = re^{i\theta}$) is given by

$$d_B(z_1, z_2) = \min_{\hat{g} \in \mathcal{N}'_0} d_H(z_1, \hat{g}z_2), \quad \text{where} \quad (C10)$$

$$d_H(z_1, z_2) = 2 \operatorname{arctanh} \left| \frac{z_1 - z_2}{1 - z_1^* z_2} \right| \quad (C11)$$

is the hyperbolic distance between the points in the Poincaré model.

We compute what nodes are linked in parallel using a multi-GPU solution, where each thread works on one pair of nodes, except in the case of the Bolza surface where each thread computes one or two of 49 distances and each thread block works on two node pairs. The adjacency matrix is first tiled such that (1) the data required to generate each tile fit on a single GPU and (2) the total number of tiles is a multiple of the number of GPUs available. This decomposition provides a scalable solution independent of the number or type of GPUs present in the system. The algorithm is improved by using the shared L1 memory cache to store the Bolza generators and partial results, warp shuffling to accumulate results, function templating to eliminate kernel branches, and asynchronous CUDA calls to pipeline data transfers across different GPUs.

2. Distance matrix computation

After constructing the graph, we identify the balls $B_G(x, \delta_n)$ and $B_G(y, \delta_n)$. When $\varepsilon_n = \delta_n$, these are simply the nearest neighbors of x and y , respectively; otherwise, they

are calculated using Dijkstra’s algorithm [84]. For the largest graphs we consider, the size of these balls can be up to tens of thousands of nodes, resulting in over billions of pairwise distances between the nodes in the two balls. The overall simulation bottleneck is the computation of the matrix of the weighted shortest-path distances between these large sets of nodes.

We employ a custom multi-GPU A^* search algorithm using the methods described in [85]. The standard A^* algorithm works by constructing a priority queue of visited nodes z using a binary heap. In computing the shortest path distance between a source node x_i and destination node y_j , priorities are assigned to node z in the queue according to the heuristic function

$$f(z) = d_G(x_i, z) + h(z, y_j), \quad (\text{C12})$$

where $d_G(x_i, z)$ is the weighted graph distance between z and source $x_i \in B_G(x, \delta_n)$, while $h(z, y_j)$ is the lower-bound estimate of the distance between z and destination $y_j \in B_G(y, \delta_n)$. As soon as new node $z' \in B_G(y, \delta_n)$ is added to the priority queue, its weighted graph distance $d_G(x_i, z')$ is added to the distance matrix.

The A^* algorithm is implemented on a single GPU by utilizing multiple priority queues, one per CUDA thread, so

that it is efficient for graphs with large average degrees of the order of hundreds. Each priority queue extracts multiple states, after which we detect duplicates using a technique called parallel hashing with replacement, which is a modification of the cuckoo hashing scheme that avoids hash conflicts by allowing some duplicates to remain [85,86]. During this step, the heuristic for extracted nodes is updated, and they are re-added to the priority queues using a parallel heap insertion. This procedure continues until the destination node y_j has been extracted by at least one of the priority queues.

3. Wasserstein distance computation

Having in place the distance matrix between the sets of nodes in the two balls $B_G(x, \delta_n)$ and $B_G(y, \delta_n)$, we compute the Wasserstein distance between the uniform probability distributions on these balls by solving the linear program (18). For the simulations presented here, we found it sufficient to use the MOSEK package [87] as long as the number of variables, given by $|B_G(x, \delta)||B_G(y, \delta)|$, is roughly less than 1.3×10^9 , past which we run out of memory. In most cases with smaller balls, however, we solve the linear program quite quickly using the standard primal-dual interior point method [88].

-
- [1] M. Boguñá, I. Bonamassa, M. De Domenico, S. Havlin, D. Krioukov, and M. Á. Serrano, Network geometry, *Nat. Rev. Phys.* **3**, 114 (2021).
- [2] S. Lubold, A. G. Chandrasekhar, and T. H. McCormick, Identifying the latent space geometry of network models through analysis of curvature, [arXiv:2012.10559](https://arxiv.org/abs/2012.10559).
- [3] O. Narayan and I. Sanjeev, Large-scale curvature of networks, *Phys. Rev. E* **84**, 066108 (2011).
- [4] E. Jonckheere, M. Lou, F. Bonahon, and Y. Baryshnikov, Euclidean versus hyperbolic congestion in idealized versus experimental networks, *Internet Math.* **7**, 1 (2011).
- [5] R. Sandhu, T. Georgiou, E. Reznik, L. Zhu, I. Kolesov, Y. Senbabaoglu, and A. Tannenbaum, Graph curvature for differentiating cancer networks, *Sci. Rep.* **5**, 12323 (2015).
- [6] D. B. de Souza, F. Santos, E. Figueiroa, J. Correia, H. da Silva, J. L. d. L. Filho, and J. Albuquerque, Using curvature to infer COVID-19 fractal epidemic network fragility and systemic risk, [medRxiv \(2020\)](https://medrxiv.org/abs/2020.04.01.20047225), doi:10.1101/2020.04.01.20047225.
- [7] Ö. Akgüller and M. A. Balci, Geodesic convex boundary curvatures of the communities in stock market networks, *Physica A (Amsterdam)* **505**, 569 (2018).
- [8] C.-C. Ni, Y.-Y. Lin, J. Gao, X. David Gu, and E. Saucan, in *2015 IEEE Conference on Computer Communications (INFOCOM)* (IEEE, Hong Kong, 2015).
- [9] R. S. Sandhu, T. T. Georgiou, and A. R. Tannenbaum, Ricci curvature: An economic indicator for market fragility and systemic risk, *Sci. Adv.* **2**, e1501495 (2016).
- [10] B. Tadić, M. Andjelković, and M. Šuvakov, Origin of hyperbolicity in brain-to-brain coordination networks, *Front. Phys.* **6**, 7 (2018).
- [11] B. Tadić, M. Andjelković, and R. Melnik, Functional geometry of human connectomes, *Sci. Rep.* **9**, 12060 (2019).
- [12] H. Farooq, Y. Chen, T. T. Georgiou, A. Tannenbaum, and C. Lenglet, Network curvature as a hallmark of brain structural connectivity, *Nat. Commun.* **10**, 4937 (2019).
- [13] J. Sia, E. Jonckheere, and P. Bogdan, Ollivier-Ricci curvature-based method to community detection in complex networks, *Sci. Rep.* **9**, 9800 (2019).
- [14] A. Gu, F. Sala, B. Gunel, and C. Ré, in *7th International Conference on Learning Representations (ICLR)*, 2019.
- [15] S. M. Carroll, *Spacetime and Geometry* (Cambridge University Press, Cambridge, 2019).
- [16] J. Ambjørn, A. Görlich, J. Jurkiewicz, and R. Loll, Nonperturbative quantum gravity, *Phys. Rep.* **519**, 127 (2012).
- [17] T. Regge, General relativity without coordinates, *Nuovo Cimento* **19**, 558 (1961).
- [18] L. Bombelli, J. Lee, D. Meyer, and R. D. Sorkin, Space-Time as a Causal Set, *Phys. Rev. Lett.* **59**, 521 (1987).
- [19] J. Henson, The causal set approach to quantum gravity, in *Approaches to Quantum Gravity*, edited by D. Oriti (Cambridge University Press, Cambridge, 2009), pp. 393–413.
- [20] S. Surya, The causal set approach to quantum gravity, *Living Rev. Relativ.* **22**, 5 (2019).
- [21] D. M. T. Benincasa and F. Dowker, Scalar Curvature of a Causal Set, *Phys. Rev. Lett.* **104**, 181301 (2010).
- [22] A. Belenchia, D. M. T. Benincasa, and F. Dowker, The continuum limit of a 4-dimensional causal set scalar d’Alembertian, *Class. Quantum Gravity* **33**, 245018 (2016).
- [23] L. Glaser and S. Surya, The Hartle-Hawking wave function in 2D causal set quantum gravity, *Class. Quantum Gravity* **33**, 065003 (2016).

- [24] L. Glaser, D. O'Connor, and S. Surya, Finite size scaling in 2d causal set quantum gravity, *Class. Quantum Gravity* **35**, 045006 (2018).
- [25] A. Eichhorn, S. Surya, and F. Versteegen, Spectral dimension on spatial hypersurfaces in causal set quantum gravity, *Class. Quantum Gravity* **36**, 235013 (2019).
- [26] A. Eichhorn, Steps towards Lorentzian quantum gravity with causal sets, *J. Phys.: Conf. Ser.* **1275**, 012010 (2019).
- [27] W. J. Cunningham and S. Surya, Dimensionally restricted causal set quantum gravity: Examples in two and three dimensions, *Class. Quantum Gravity* **37**, 054002 (2020).
- [28] C. A. Trugenberger, Combinatorial quantum gravity: Geometry from random bits, *J. High Energy Phys.* **09** (2017) 045.
- [29] C. Kelly, C. A. Trugenberger, and F. Biancalana, Self-assembly of geometric space from random graphs, *Class. Quantum Gravity* **36**, 125012 (2019).
- [30] C. Kelly, C. Trugenberger, and F. Biancalana, Emergence of the circle in a statistical model of random cubic graphs, *Class. Quantum Grav.* **38**, 075008 (2021).
- [31] N. Klitgaard and R. Loll, Introducing quantum Ricci curvature, *Phys. Rev. D* **97**, 046008 (2018).
- [32] N. Klitgaard and R. Loll, Implementing quantum Ricci curvature, *Phys. Rev. D* **97**, 106017 (2018).
- [33] N. Klitgaard and R. Loll, How round is the quantum de Sitter universe? *Eur. Phys. J. C* **80**, 990 (2020).
- [34] J. Maldacena and L. Susskind, Cool horizons for entangled black holes, *Fortschr. Phys.* **61**, 781 (2013).
- [35] M. van Raamsdonk, Building up spacetime with quantum entanglement, *Gen. Relativ. Gravit.* **42**, 2323 (2010).
- [36] C. J. Cao, S. M. Carroll, and S. Michalakis, Space from Hilbert space: Recovering geometry from bulk entanglement, *Phys. Rev. D* **95**, 024031 (2017).
- [37] N. Bao, C. J. Cao, S. M. Carroll, and A. Chatwin-Davies, de Sitter space as a tensor network: Cosmic no-hair, complementarity, and complexity, *Phys. Rev. D* **96**, 123536 (2017).
- [38] B. Swingle, Spacetime from entanglement, *Annu. Rev. Condens. Matter Phys.* **9**, 345 (2018).
- [39] R. Orús, Tensor networks for complex quantum systems, *Nat. Rev. Phys.* **1**, 538 (2019).
- [40] J. S. Cotler, G. R. Penington, and D. H. Ranard, Locality from the spectrum, *Commun. Math. Phys.* **368**, 1267 (2019).
- [41] A. Samal, R. P. Sreejith, J. Gu, S. Liu, E. Saucan, and J. Jost, Comparative analysis of two discretizations of Ricci curvature for complex networks, *Sci. Rep.* **8**, 8650 (2018).
- [42] L. Prokhorenkova, E. Samosvat, and P. van der Hoorn, Global graph curvature, in *Algorithms and Models for the Web Graph*, Vol. 2 (Springer International Publishing, Cham, 2020), pp. 16–35.
- [43] Y. Ollivier, Ricci curvature of metric spaces, *C. R. Math* **345**, 643 (2007).
- [44] Y. Ollivier, Ricci curvature of Markov chains on metric spaces, *J. Funct. Anal.* **256**, 810 (2009).
- [45] Y. Ollivier, in *Probabilistic Approach to Geometry* (Mathematical Society of Japan, Tokyo, 2010), pp. 343–381.
- [46] E. N. Gilbert, Random plane networks, *J. Soc. Ind. Appl. Math.* **9**, 533 (1961).
- [47] J. Dall and M. Christensen, Random geometric graphs, *Phys. Rev. E* **66**, 016121 (2002).
- [48] M. Penrose, *Random Geometric Graphs* (Oxford University Press, Oxford, 2003).
- [49] J. Jost and S. Liu, Ollivier's Ricci curvature, local clustering and curvature-dimension inequalities on graphs, *Discrete Comput. Geom.* **51**, 300 (2014).
- [50] J. Cheeger, W. Müller, and R. Schrader, On the curvature of piecewise flat spaces, *Commun. Math. Phys.* **92**, 405 (1984).
- [51] J. Steiner, Über parallele Flächen, in *Jacob Steiner's Gesammelte Werke*, edited by K. Weierstrass (Cambridge University Press, Cambridge, 1882), pp. 171–176.
- [52] D. Bakry and M. Émery, Diffusions hypercontractives, in *Séminaire de Probabilités XIX 1983/84* (Springer, Berlin, 1985), pp. 177–206.
- [53] M. Gromov, Hyperbolic groups, in *Essays in Group Theory*, Mathematical Sciences Research Institute Publications, Vol. 8 (Springer, New York, 1987), pp. 75–263.
- [54] E. Jonckheere, P. Lohsoonthorn, and F. Bonahon, Scaled Gromov hyperbolic graphs, *J. Graph Theory* **57**, 157 (2008).
- [55] Y. Higuchi, Combinatorial curvature for planar graphs, *J. Graph Theory* **38**, 220 (2001).
- [56] J.-P. Eckmann and E. Moses, Curvature of co-links uncovers hidden thematic layers in the World Wide Web, *Proc. Natl. Acad. Sci. USA* **99**, 5825 (2002).
- [57] R. Forman, Bochner's method for cell complexes and combinatorial Ricci curvature, *Discrete Comput. Geom.* **29**, 323 (2003).
- [58] R. P. Sreejith, K. Mohanraj, J. Jost, E. Saucan, and A. Samal, Forman curvature for complex networks, *J. Stat. Mech.: Theory Exp.* (2016) 063206.
- [59] M. Weber, E. Saucan, and J. Jost, Coarse geometry of evolving networks, *J. Complex Networks* **6**, 706 (2018).
- [60] Y. Lin and S.-T. Yau, Ricci curvature and eigenvalue estimate on locally finite graphs, *Math. Res. Lett.* **17**, 343 (2010).
- [61] Y. Lin, L. Lu, and S.-T. Yau, Ricci curvature of graphs, *Tohoku Math. J.* **63**, 605 (2011).
- [62] O. Knill, A discrete Gauss-Bonnet type theorem, *arXiv:1009.2292*.
- [63] O. Knill, On index expectation and curvature for networks, *arXiv:1202.4514*.
- [64] M. Keller, Curvature, geometry and spectral properties of planar graphs, *Discrete Comput. Geom.* **46**, 500 (2011).
- [65] M. Keller and N. Peyerimhoff, Cheeger constants, growth and spectrum of locally tessellating planar graphs, *Math. Z.* **268**, 871 (2011).
- [66] M. Kempton, G. Lippner, and F. Münch, Large scale Ricci curvature on graphs, *Calculus of Variations and Partial Differential Equations* **59**, 166 (2020).
- [67] B. B. Bhattacharya and S. Mukherjee, Exact and asymptotic results on coarse Ricci curvature of graphs, *Discrete Math.* **338**, 23 (2015).
- [68] G. Last and M. Penrose, *Lectures on the Poisson Process* (Cambridge University Press, Cambridge, 2017).
- [69] M. Kahle, Random geometric complexes, *Discrete Comput. Geom.* **45**, 553 (2011).
- [70] J. Latschev, Vietoris-Rips complexes of metric spaces near a closed Riemannian manifold, *Arch. Math.* **77**, 522 (2001).
- [71] R. Aurich and F. Steiner, On the periodic orbits of a strongly chaotic system, *Physica D (Amsterdam)* **32**, 451 (1988).
- [72] J. G. Ratcliffe, *Foundations of Hyperbolic Manifolds*, Graduate Texts in Mathematics, Vol. 149 (Springer International Publishing, Cham, 2019).

- [73] J. Díaz, D. Mitsche, G. Perarnau, and X. Pérez-Giménez, On the relation between graph distance and Euclidean distance in random geometric graphs, *Adv. Appl. Probab.* **48**, 848 (2016).
- [74] B. Riemann, On the hypotheses which lie at the bases of geometry, *Nature (London)* **8**, 14 (1873).
- [75] R. Loll, Discrete approaches to quantum gravity in four dimensions, *Living Rev. Relativ.* **1**, 13 (1998).
- [76] Z. Wu, G. Menichetti, C. Rahmede, and G. Bianconi, Emergent complex network geometry, *Sci. Rep.* **5**, 10073 (2015).
- [77] D. Krioukov, Clustering Implies Geometry in Networks, *Phys. Rev. Lett.* **116**, 208302 (2016).
- [78] E. H. Kronheimer and R. Penrose, On the structure of causal spaces, *Math. Proc. Cambridge Philos. Soc.* **63**, 481 (1967).
- [79] P. van der Hoorn, G. Lippner, C. Trugenberger, and D. Krioukov, Ollivier curvature of random geometric graphs converges to Ricci curvature of their Riemannian manifolds, [arXiv:2009.04306](https://arxiv.org/abs/2009.04306).
- [80] T. Leighton and P. Shor, Tight bounds for minimax grid matching with applications to the average case analysis of algorithms, *Combinatorica* **9**, 161 (1989).
- [81] P. W. Shor, J. E. Yukich *et al.*, Minimax grid matching and empirical measures, *Ann. Probab.* **19**, 1338 (1991).
- [82] M. Talagrand, Matching random samples in many dimensions, *Ann. Appl. Probab.* **2**, 846 (1992).
- [83] M. Talagrand, Matching theorems and empirical discrepancy computations using majorizing measures, *J. Am. Math. Soc.* **7**, 455 (1994).
- [84] E. W. Dijkstra, A note on two problems in connexion with graphs, *Numer. Math.* **1**, 269 (1959).
- [85] Y. Zhou and J. Zeng, in *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI '15* (AAAI Press, Austin, 2015), pp. 1248–1254.
- [86] R. Pagh and F. F. Rodler, in *Algorithms–ESA 2001*, edited by F. M. auf der Heide (Springer, Berlin, 2001), pp. 121–133.
- [87] E. D. Andersen and K. D. Andersen, The MOSEK interior point optimizer for linear programming: An implementation of the homogeneous algorithm, in *High Performance Optimization*, edited by H. Frenk, K. Roos, T. Terlaky, and S. Zhang (Springer US, Boston, MA, 2000), pp. 197–232.
- [88] S. Mehrotra, On the implementation of a primal-dual interior point method, *SIAM J. Optim.* **2**, 575 (1992).