1 From Networks to Maps

Throughout history, maps have been at the center of political, economic, and geostrategic decisions and have become a critical piece of our everyday lives, serving as a precise and relevant information source. Maps provide an accurate way of visualizing, storing, and communicating information, help to recognize locational distributions, spatial patterns, and relationships, and allow us to track processes that operate through space at different length scales. Our work in the last decade led us to prove that many real complex networks are natural geometric objects and can be mapped into hidden low-dimensional metric spaces with hyperbolic geometry, where distances determine the likelihood of the interactions and encode the different intrinsic attributes determining how similar the elements of the system are (Allard et al. [2017]; Boguñá, Papadopoulos, and Krioukov [2010]; García-Pérez et al. [2016]; García-Pérez, Boguñá, and Serrano [2018]; Kleineberg et al. [2016]; Krioukov et al. [2012]; Krioukov et al. [2010]; Papadopoulos et al. [2012]; Serrano, Boguñá, and Sagues [2012]; Serrano, Krioukov, and Boguñá [2008]). We took advantage of the large amount of empirical data available and the current explosion in computing power to create meaningful geometric maps of large real networks by embedding them in an underlying space that ought not to be geographical or spatially obvious. In this Element, we review our most relevant research on this topic, with a special focus on models and applications to real networks. These results triggered the field of network geometry to become one of the fundamental areas within network science devoted to the discovery and modeling of nontrivial geometric properties of complex networks (Boguñá et al. [2020]).

Complex networks typically have been studied as topological objects (Dorogovtsev and Mendes [2003]; M. E. J. Newman [2010]), graphs where elements are represented as nodes and their interactions as links. Graphs of real networks are not regular lattices nor are they completely disordered or random, and their structure is imprinted with universal features. One of the most paradigmatic examples is the small-world phenomenon, connecting every pair of nodes in a network, on average, by a small number of intermediate links (Amaral et al. [2000]; Watts and Strogatz [1998]). Other ubiquitous properties are scale-free, or heavy tailed, distributions of the number of connections per node (degree) (Barabási and Albert [1999]), with a few nodes linked to an enormous number of neighbors (hubs with very high degrees, while most other nodes are poorly connected), modularity, and hierarchical structure (Amaral [2008]). These and other prevalent features are not a mere curiosity but arise as the outcome of evolutionary pressures or functional needs and affect the dynamics that characterize or take place within and between networks (Barrat, Barthélemy, and Vespignani [2008]).

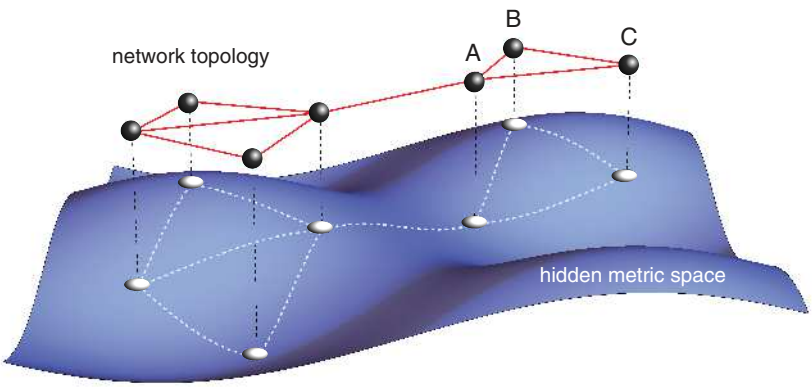

Figure 1 Hidden metric spaces help to understand the structure and function of complex networks. The smaller the distance between two nodes in the hidden metric space – the more similar they are – the more likely they are connected in the observable network topology. If node A is close to node B, and B is close to C, then A and C are necessarily close because of the triangle inequality in the metric space. Therefore, triangle ABC exists in the network topology with high probability, which explains the strong clustering observed in real complex networks.

One of the main consequences of the small-world effect is the apparent lack of a metric structure defined on the system. Certainly, in a small-world network, the distribution of shortest path lengths among pairs of nodes is sharply peaked around its average and, therefore, any pair of nodes is roughly separated by the same minimal number of intermediate links. This is the reason why complex networks are often considered as pure topological objects and are difficult to map. Yet, many networks are embedded in metric spaces. Some are explicit (Barthélemy [2011]) – like in airport networks (Barrat et al. [2004]; Guimerà et al. [2005]), power grids, or urban networks – whereas some are hidden yet shaping the network topological structure (Boguñá et al. [2010]; Krioukov et al. [2010]; Krioukov et al. [2009]; Serrano et al. [2008]); see Figure 1. This idea led to hidden metric space models of complex networks with an underlying effective hyperbolic geometry. These models are able to explain universal features observed in real-world systems, including the small-world property, scale-free degree distributions, clustering, and also fundamental mechanisms like preferential attachment in growing networks (Papadopoulos et al. [2012]), the emergence of communities (Zuev et al. [2015]), and multiscale self-similarity (García-Pérez, Boguñá, and Serrano [2018]). The discovery of the hidden geometry of real complex networks led to the emergence of the field of network geometry (Boguñá et al. [2020]), a major research area within network science.

The hidden metric space network models of complex networks couple their topology to an underlying geometry through a probabilistic connectivity law depending on distances in the space, which combine popularity and similarity dimensions in such a way that more popular and similar nodes have more chance to interact (Krioukov et al. [2010]; Papadopoulos et al. [2012]; Serrano et al. [2008]). The basic assumptions in our model are that there exists some similarity between nodes which, along degrees, plays an important role in how connections are established and that, since similarity is transitive, geometry is an appropriate mathematical formalism to encode it. The clue for the connection between topology and geometry is then clustering – transitive relationships, or triangles – which arises as a reflection in the topology of the network of the triangle inequality in the underlying hidden metric space. These models can be combined with statistical inference techniques to find the coordinates of the nodes in the underlying metric space that maximize the likelihood that the topology of the network is reproduced by the model (Blasius et al. [2018]; Boguñá et al. [2010]; García-Pérez et al. [2019]; Papadopoulos, Aldecoa, and Krioukov [2015]). One of the key properties of these maps is that the shortest paths in the topology of the networks follow closely geodesic lines in the underlying space. This ensures that networks highly congruent with the hidden metric space model are navigable, where navigability is understood as efficient routing of information based on the metric embedding (Allard and Serrano [2020]; Boguñá and Krioukov [2009]; Boguñá Krioukov, and Claffy [2009]; Boguñá et al. [2010]; Gulyás et al. [2015]; Krioukov et al. [2010]; Papadopoulos et al. [2010]).

One example of the power of this geometric approach is the discovery of the hyperbolic plane as the effective geometry of many real networks (see Fig. 2), including such disparate systems as the Internet at the Autonomous Systems level (Boguñá, Papadopoulos, and Krioukov [2010]), genome-scale reconstructions of metabolic networks in the cell (Serrano et al. [2012]), the World Trade Web from 1870 to 2013 (García-Pérez et al. [2016]), and brains of different species (Allard and Serrano [2020]). In the case of the Internet, the metric space provides a solution to the scalability limitations of current inter-domain routing protocols. For metabolic networks, it allows us to redefine the concept of biological pathways and to quantify their crosstalk. In international trade, the maps provide information about the long-term evolution of the system, unraveling the role of globalization, hierarchization, and localization forces. Finally, the effective geometry of human and nonhuman brain structures is also better described as hyperbolic than Euclidean, thus implying that hyperbolic embeddings are universal and meaningful maps of brain structure that allow for an efficient routing of information.

Figure 2 Hidden metric space maps of real complex networks in the hyperbolic plane. (a) Embedding of the human connectome, (b) Internet graph at the Autonomous Systems level, and (c) embedding of the World Trade Web as of year 2013. In all cases, the size of a node is proportional to the logarithm of its degree, and the radial coordinate decreases with increasing degree so that higher degree nodes are placed closer to the center of the disk. In (a) and (b) brain regions' and countries' names are located at the average angular position of all nodes belonging to the same region/country.

Source: Panel (a) modified from Zheng et al. (2020). Panel (b) reprinted from Boguñá et al. (2010). Panel (c) reprinted from García-Pérez et al. (2016).

These results suggest that the geometric paradigm improves our knowledge of the basic principles underlying the organization, function, and evolution of complex systems. But, in the long run, it also will transform research on how to model, predict, and control them, with potential implications for a large number of current challenges. These include efficient recommendation systems and search engines, prediction of epidemic spreading, and drug design in cancer and brain research.

2 Geometric Models for Static Topologies

Our first remarkable observation was to identify clustering – a measure of the number of triangles in a graph – as the key connection between complex networks and an underlying hidden geometry. Indeed, the triangle inequality in a metric space induces clustering in the structure of the graph, as illustrated in Figure 1. In Serrano et al. (2008), we analyzed the clustering coefficient of several real complex networks and found that their topological structure was compatible with an underlying hidden metric space. This finding led us to introduce the $\mathbb{S}^1$ class of network models (Serrano et al. [2008]). In these models, nodes are embedded in a metric space and connections exist with a gravity-law-like connection probability balancing the distance between nodes and their degrees; see Figure 3a. The connection probability encodes, in a simple and general way, the two major forces at play, namely, the effect of a similarity distance and the effect of the importance of the nodes involved in the connections. In this way, the model is able to generate scale-free, small-world, and clustered graphs very similar to those found in real complex networks, where the heterogeneity in the distribution of the number of contacts per node can be controlled independently of the level of clustering that measures the coupling with the metric space.

The $\mathbb{S}^1$ model is a mixed model in the sense that it combines a metric component and a topological component. Nodes are given coordinates in a metric similarity space but are also given degrees, determining their number of neighbors. At first glance, it seems difficult to combine, in a purely geometric framework, the small-world and scale-free properties that we observe in real networks. The major complication arises as a consequence of the small-world effect. This effect implies an exponential expansion of space, that is, the number of nodes within a disk of a certain radius grows exponentially with the radius (up to the finite size of the system). This behavior is in stark contrast to what happens in Euclidean spaces, where space grows as a power of the radius, but it is similar to what happens in hyperbolic geometry. In Krioukov et al. (2009, 2010), we developed the theory of random geometric graphs in hyperbolic geometry; see Figure 3b. Interestingly, scale-free graphs are the

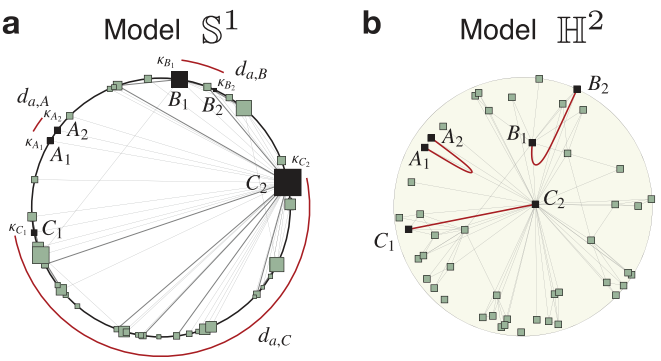

Figure 3 (a) Model $\mathbb{S}^1$. The similarity distances d_a between pairs of nodes A1-A2, B1-B2, and C1-C2 have been highlighted. The size of a node is proportional to its hidden degree κ. (b) Model $\mathbb{H}^2$ in the hyperbolic plane. Nodes in the different pairs are separated by the same hyperbolic distance. Nodes are equally sized, but nodes with higher hidden degree are positioned closer to the center. The similarity distance is the same in the two representations.

Source: Modified from the Supplementary Information in García-Pérez et al. [2016].

natural outcome of the formalism, indicating that this geometry is the most appropriate to model complex networks. Indeed, it produces in a natural way scale-free, small-world, and clustered graphs. However, the most surprising result is that this class of models, which we call $\mathbb{H}^2$, is isomorphic to our previous $\mathbb{S}^1$ version (Krioukov et al. [2010]; Serrano et al. [2008]).

This duality allows us to use either model indistinctly, depending on the particular application. The $\mathbb{S}^1$ version is especially convenient for theory development, analytical calculations, and the implementation of embedding techniques – which estimate the coordinates that maximize the likelihood of the observed structure being produced by the model. Instead, the $\mathbb{H}^2$ version is well suited for visualization purposes, to analyze navigation properties (Allard and Serrano [2020]; Boguñá and Krioukov [2009]; Boguñá et al. [2009]; Boguñá et al. [2010]; Gulyás et al. [2015]; Krioukov et al. [2010]; Papadopoulos et al. [2010]), or to define hierarchies within the network (García-Pérez et al. [2016]).

An interesting aspect is that our geometric class of models corresponds to an entropy-maximizing probabilistic mixture of grand canonical network ensembles, where network links can be thought of as noninteracting fermions whose energies depend on distances on the underlying geometry, with the particular choice of the functional form of this dependency defining network properties. At present, these models provide the simplest class of models capturing sparsity, the small-world property, power-law degree distributions,

and nonvanishing clustering in a geometric framework with explicit symmetry structure (Boguñá et al. [2020]).

2.1 The $\mathbb{S}^1$ Model

In the $\mathbb{S}^1$ model (Serrano et al. [2008]), a node i is assigned two hidden variables: a hidden degree κ_i quantifying its popularity; and an angular position θ_i in a one-dimensional sphere (or circle), the similarity space, where distances with the other nodes serve as a proxy for their similarity. The radius of the circle is adjusted to $R = N/2\pi$, where N is the number of nodes, so that the density is set to 1 without loss of generality. The probability of connection between any pair of nodes takes the form of a gravity law, whose magnitude increases with the product of the hidden degrees (i.e., their combined popularities) and decreases with the angular distance between the two nodes. In other words, more-similar nodes are angularly closer and are, therefore, more likely to be connected, whereas not-so-similar pairs of nodes have a high probability of being connected only if they are popular. Specifically, nodes i and j are connected with probability

$$p_{ij} = \frac{1}{1+\chi_{ij}^{\beta}} = \frac{1}{1+\left(\frac{d_{ij}}{\mu\kappa_i\kappa_j}\right)^{\beta}}, \tag{2.1}$$

where μ controls the average degree of the network, β controls its level of clustering, and $d_{ij} = R\Delta\theta_{ij}$, and $\Delta\theta_{ij} = \pi - |\pi - |\theta_i - \theta_j||$ is the angular distance between nodes i and j. Notice that there are no constraints on the distribution of hidden variables κ and θ. The angular distribution could be nonhomogeneous, and both hidden variables could even be correlated. This is an important observation because such angular inhomogeneities or correlations can explain the emergence of communities and other nontrivial topological patterns observed in real networks (Allard and Serrano [2020]).

A priori, the functional form of the connection probability could be any integrable function of argument $f\left(\frac{d_{ij}}{\mu\kappa_i\kappa_j}\right)$. However, the Fermi–Dirac form of the connection probability in Eq. (2.1) is the only possible choice that defines maximally random ensembles of graphs that are simultaneously sparse,[1] heterogeneous, clustered, small-worlds, and maximally degree–degree uncorrelated (Boguñá et al. [2020]).[2] Besides, with this choice, parameter β has full control

[1] By sparse networks we mean ensembles of networks with size-independent average degree.
[2] By maximally degree–degree uncorrelated we mean that the probability of a node with hidden variable κ having a neighbor with hidden variable κ' is independent of κ. Yet, for heterogeneous scale-free networks, some level of degree–degree correlation is unavoidable, as shown in Boguñá, Pastor-Satorras, and Vespignani (2004).

of the level of clustering without affecting the degree distribution. It can be shown that the model undergoes a structural phase transition at $\beta = 1$ so that, for $\beta < 1$, networks are unclustered, whereas for $\beta > 1$, the ensemble generates networks with finite clustering in the thermodynamic limit (Serrano et al. [2008]).

A Simple Algorithm to Generate Networks from the $\mathbb{S}^1$ Ensemble

The algorithm below generates networks from the $\mathbb{S}^1$ ensemble in the limit $N \gg 1$, in the simple scenario of uncorrelated hidden variables κ and θ, and with the similarity coordinate homogeneously distributed.

1. Fix the number of nodes N, parameter $\beta > 1$, and the target average degree $\langle k \rangle$
2. Set μ to
$$\mu = \frac{\beta}{2\pi \langle k \rangle} \sin\left(\frac{\pi}{\beta}\right)$$
3. Assign a hidden degree κ to every node from $\rho(\kappa)$ so that $\langle \kappa \rangle = \langle k \rangle$. Assign also an angular position θ to each node sampled uniformly at random within the interval $[0, 2\pi]$.
4. Connect every pair of nodes with probability given by Eq. (2.1).

With this parametrization – and in the thermodynamic limit – the expected degree of a node with hidden degree κ is simply $\bar{k}(\kappa) = \kappa$, which justifies the name of hidden degree. Indeed, the expected degree of any node i with hidden variables (κ_i, θ_i) can be evaluated as $\bar{k}(\kappa_i, \theta_i) = \sum_j p_{ij}$, where the connection probability is given in Eq. (2.1). If the network is homogeneous with respect to the similarity space, $\bar{k}(\kappa_i, \theta_i)$ is independent of θ_i. Thus, the expected degree of any node with hidden degree κ, located without loss of generality at $\theta = 0$, can be expressed as

$$\bar{k}(\kappa) = 2\mu\kappa \int \kappa' \rho(\kappa') \left[\int_0^{\frac{N}{2\mu\kappa\kappa'}} \frac{dt}{1+t^\beta} \right] d\kappa' = \tag{2.2}$$

$$= N \int \rho(\kappa') {}_2F_1\left(1, \frac{1}{\beta}, 1+\frac{1}{\beta}, -\left[\frac{N}{2\mu\kappa\kappa'}\right]^\beta\right) d\kappa',$$

where ${}_2F_1(1, \frac{1}{\beta}, 1+\frac{1}{\beta}, -x^\beta)$ is the hypergeometric function, whose asymptotic behavior when $x \to \infty$ is ${}_2F_1(1, \frac{1}{\beta}, 1+\frac{1}{\beta}, -x^\beta) \sim \pi \csc(\pi/\beta)/(\beta x)$.

Using this result, we recover the proportionality between expected and hidden degrees.

The degree distribution of the model can be evaluated as

$$P(k) = \frac{1}{k!}\int \kappa^k e^{-\kappa}\rho(\kappa)d\kappa, \tag{2.3}$$

that is, a mixture of Poisson distributions weighted by $\rho(\kappa)$. Eq. (2.3) shows that the model generates nodes with degree zero with probability $P(0) = \langle e^{-\kappa}\rangle$, so that the expected number on nonzero degree nodes is $N_{obs} = N[1 - P(0)]$, whereas the observable average degree (counting only nodes with nonzero degree) is $\langle k\rangle_{obs} = \langle k\rangle/[1 - P(0)]$.

In the case of scale-free networks, we consider $\rho(\kappa)$ to be a power-law distribution of the form

$$\rho(\kappa) = (\gamma-1)\kappa_0^{\gamma-1}\kappa^{-\gamma} \;\; ; \;\; \kappa > \kappa_0 = \frac{\gamma-2}{\gamma-1}\langle k\rangle \;\; ; \;\; \gamma > 2. \tag{2.4}$$

Plugging this expression into Eq. (2.3), the degree distribution reads

$$P(k) = (\gamma-1)\kappa_0^{\gamma-1}\frac{\Gamma(k+1-\gamma,\kappa_0)}{k!}, \tag{2.5}$$

where $\Gamma(k+1-\gamma,\kappa_0)$ is the incomplete gamma function, so that the asymptotic behavior of the degree distribution is $P(k) \sim k^{-\gamma}$. To simulate sparse scale-free networks with $\gamma < 2$ (as found, for instance, in airport networks) we need to introduce a cutoff in the distribution of hidden degrees κ_c. In particular, we choose a hard cutoff of the form

$$\rho(\kappa) = \frac{(\gamma-1)\kappa_0^{\gamma-1}}{1-\left(\frac{\kappa_c}{\kappa_0}\right)^{1-\gamma}}\kappa^{-\gamma} \;\; \text{with} \;\; \kappa_0 < \kappa < \kappa_c, \tag{2.6}$$

where the lower cutoff is the solution of the equation

$$\langle k\rangle = \frac{\gamma-1}{\gamma-2}\kappa_0\frac{1-\left(\frac{\kappa_c}{\kappa_0}\right)^{2-\gamma}}{1-\left(\frac{\kappa_c}{\kappa_0}\right)^{1-\gamma}}. \tag{2.7}$$

Equations (2.6) and (2.7) can also be used to compensate for finite size effects in scale-free networks with $\gamma \gtrsim 2$. Indeed, to prevent extreme fluctuations arising when γ is very close to 2, instead of generating values of κ from the unbounded distribution Eq. (2.4), we introduce a hard cutoff whose value is the same as the natural cutoff of the unbounded distribution, which can be approximated by $\kappa_c = \kappa_0 N^{1/(\gamma-1)}$ (Boguñá, Pastor-Satorras, and Vespignani [2004]). Then, we generate values of κ from Eq. (2.6) with parameter κ_0 equal to

$$\kappa_0 = \frac{1-N^{-1}}{1-N^{\frac{2-\gamma}{\gamma-1}}}\frac{\gamma-2}{\gamma-1}\langle k\rangle. \tag{2.8}$$

Notice that when γ is very close to 2, finite size effects can be very important even for large networks. However, notice that this is not the only source of finite size effects. To fully account for finite size effects, we must also consider the effects coming from the upper limit in the integral in Eq. (2.2), as done in García-Pérez et al. (2019). However, in many practical applications, the correction in Eq. (2.8) is enough.

The $\mathbb{S}^1$ model can be used to produce synthetic ensembles with controllable structural features or for high-fidelity replication of a specific real network. To that end, observed degrees in the real network can be taken as good proxies of hidden degrees, and parameters μ and β can be tuned to reproduce the average degree and clustering of the real network. This procedure is not very accurate for heterogeneous networks due to strong fluctuations. Actual hidden degrees could be estimated from real data to avoid the mismatch between hidden and observed degrees, but this operation can be demanding and, besides, there is no guarantee that all nodes end up with the exact same degree they had in the real network. An alternative is the implementation of the geometric randomization model introduced in Starnini, Ortiz, and Serrano (2019), which preserves exactly the degree sequence of the input network while producing a version of the network maximally congruent with the $\mathbb{S}^1$ model.

The geometric randomization model assumes the same form of the connection probability as in the $\mathbb{S}^1$ model. Given a real network, nodes are given angular coordinates in the similarity space uniformly at random. Then, the network is rewired in order to maximize the likelihood that the new topology is generated by the $\mathbb{S}^1$ model while preserving the observed degrees and, thus, the total number of edges. After selecting a value of β, for instance, the one that replicates the level of clustering of the original network, the rewiring procedure is conducted by executing a Metropolis–Hastings algorithm as follows.

Geometric Randomization Model

1. Assign each node an angular coordinate uniformly at random.
2. Choose two links at random, say between nodes i and j and between nodes l and m.
3. Compute the probability of rewiring (connecting i and l and j and m) as

$$p_r = \min\left[1, \frac{\mathcal{L}_{new}}{\mathcal{L}_{old}}\right] = \min\left[1, \left(\frac{\Delta\theta_{ij}\Delta\theta_{lm}}{\Delta\theta_{il}\Delta\theta_{jm}}\right)^{\beta}\right], \tag{2.9}$$

where $\mathcal{L}_{new}$ corresponds to the value of the likelihood function after the swap and $\mathcal{L}_{old}$ before the swap, both evaluated using Eq. (3.1) (see next section) and the probability of connection in Eq. (2.1). Notice that p_r only requires information about the angular coordinates of nodes.