

<https://doi.org/10.1038/s44260-024-00013-z>

Feature-aware ultra-low dimensional reduction of real networks

Check for updates

Robert Jankowski^{1,2}, Pegah Hozhabrierdi³, Marián Boguñá^{1,2} & M. Ángeles Serrano^{1,2,4}

In existing models and embedding methods of networked systems, node features describing their qualities are usually overlooked in favor of focusing solely on node connectivity. This study introduces *FiD-Mercator*, a model-based ultra-low dimensional reduction technique that integrates node features with network structure to create D -dimensional maps of complex networks in a hyperbolic space. This embedding method efficiently uses features as an initial condition, guiding the search of nodes' coordinates toward an optimal solution. The research reveals that downstream task performance improves with the correlation between network connectivity and features, emphasizing the importance of such correlation for enhancing the description and predictability of real networks. Simultaneously, hyperbolic embedding's ability to reproduce local network properties remains unaffected by the inclusion of features. The findings highlight the necessity for developing network embedding techniques capable of exploiting such correlations to optimize both network structure and feature association jointly in the future.

Factors relevant to the links' formation process in real-world complex networks are typically numerous, and unknown or difficult to measure. For instance, off-line and on-line social relations of friendship can be sustained by a variety of reasons including individual factors, such as approachability or social skills, and environmental factors, such as geographic proximity and life events and activities. This has not prevented that generative models based on simple connectivity rules are able to capture many of the salient properties of real networks, such as degree heterogeneity, small-worldness, clustering, and community structure.

Such is the case of a family of geometric random graph models^{1,2}, in which pairs of nodes are connected with a probability that depends on their popularity and on their distance in a similarity metric space. For a real network, these variables can be inferred by maximizing the likelihood of the model to reproduce the observed structure, as done by the embedding tool *Mercator*³. The combined use of geometric models and embedding tools prove that hyperbolic space is a natural geometry for real networks⁴, and have been able to explain heterogeneous degree distributions, clustering, small-worldness, percolation, spectral properties, community structure, different forms of self-similarity, preferential attachment in growing networks, the non-trivial coupling between weights and topology, correlations in multilayer networks, and served as the basis for defining a renormalization group for complex networks, see ref. 2 and references therein.

The popularity component of the hyperbolic space where real networks are embedded accounts for nodes' degrees, and distances in the

similarity subspace for affinities not influenced by degrees. Assuming that relational properties of nodes, other than popularity, define this similarity subspace, it is expected that the organization of real networks often correlates with node descriptive metadata. This acts as additional information to the network topology, enriching the description of nodes and helping to discover, identify, and categorize them. Going back to the example of friendship networks, link formation could be strongly affected by age, gender, interests, background, and other social determinants⁵⁻⁷. This influence is often captured in hyperbolic maps of real networks, even though the corresponding information was not deliberately incorporated during the mapping process. For instance, autonomous systems form clusters in the similarity space according to geographic data in hyperbolic maps of the Internet⁸, and functional modules form clusters in the hyperbolic maps of brain connectomes⁹. However, this might not always be the case and descriptive metadata could be irrelevant or only partially correlated with network structure and cannot be treated as ground truth¹⁰.

Under certain conditions, integrating metadata with network structure can enhance the accuracy of community detection methods¹⁰⁻¹⁴ and facilitate a more realistic assessment of the robustness of interconnected systems¹⁵. This poses the question of whether integrating structural information and metadata in network models enables a better understanding of connectivity in complex systems and a better performance in downstream tasks like node classification or prediction of missing links. From a broader perspective, adding features to the network structure leads to a hybridization

¹Departament de Física de la Matèria Condensada, Universitat de Barcelona, Martí i Franquès 1, 08028 Barcelona, Spain. ²Universitat de Barcelona Institute of Complex Systems (UBICS), Universitat de Barcelona, Barcelona, Spain. ³Department of Electrical Engineering and Computer Science, Syracuse University, 223 Link Hall, Syracuse, NY 13244, USA. ⁴ICREA, Pg. Lluís Companys 23, E-08010 Barcelona, Spain. e-mail: marian.serrano@ub.edu

that reveals a distinct type of spatial description through the implementation of a different inductive bias. This bias involves a set of assumptions that consider both features and network structure simultaneously and independently of whether they are correlated.

Embedding techniques in machine learning, particularly graph neural networks (GNNs)^{16–19}, naturally work by aggregating the node-level information extracted from the nodes' features via message passing. They have been shown to achieve good performance on downstream tasks. However, vanilla GNNs work well only if a network is homophilic, i.e., a network where connected nodes have the same class labels and similar features^{20,21}. Many popular GNNs fail to generalize to this setting and are even outperformed by models that ignore the graph structure. It is also worth mentioning that GNNs are data-driven and task-oriented embedding methods. On the other hand, unsupervised machine learning models, such as Node2Vec²² or DeepWalk²³, can be applied to multiple downstream tasks but rely purely on the structure of the graph to generate embeddings.

The integration of network structure and descriptive metadata can be implemented in the embedding process to produce maps of real networks where coordinates of nodes are determined by the two sources of information, and where the relative attention paid to one versus the other can be controlled. The last and more advanced generator of model-based hyperbolic maps is *D*-Mercator²⁴, a generalization of Mercator³ to multi-dimensional hyperbolic embedding. As demonstrated in ref. 25, real-world complex networks exhibit a compelling fit within the framework of hyperbolic spaces characterized by low dimensionality, typically falling within the range of $D \in [2, 10]$. This observation provides a solid rationale for extending the foundation laid by *D*-Mercator as outlined in ref. 24, by incorporating node features into the model.

In this paper, we introduce a model-based dimensional reduction technique named Feature-initialized *D*-dimensional Mercator, *FiD*-Mercator, that embeds networks into an ultra-low dimensional hyperbolic space by combining network topology and the collection of features describing the qualities of nodes. Our overarching goals are to provide an embedding tool that hybridizes network topology and features, and work toward clarifying to what extent node features can help in explaining the structure of real networks, and under which conditions it is advantageous to supplement geometric network models with this information for specific tasks. We used our tool to analyze link prediction and node classification in real networks and we found that the amount of correlation between network connectivity and descriptive metadata determines performance. In general, when this correlation is high, adding features opens a fast track to optimized embeddings for link prediction and improves the node classification results significantly.

Results

The *FiD*-Mercator embedding method

Given a graph-structured dataset consisting of a set of N nodes forming a complex network and a set of N_F features associated with the same set of nodes, the embedding method *FiD*-Mercator finds the hidden variables in the geometric soft configuration \mathbb{S}^D model¹, \mathbb{H}^{D+1} in the isomorphic hyperbolic geometry formulation²⁶, that best reproduces the topology of the network. In the \mathbb{S}^D model, networks exist in an underlying metric space with an effective hyperbolic geometry²⁷.

The coordinates of the N nodes in the latent space represent popularity and similarity dimensions. The popularity dimension of a node is associated with its hidden degree κ , equivalent to a radial coordinate in hyperbolic space and proportional to its observed degree in the network topology. The similarity subspace is represented as a D -sphere, which is a hypersphere in the $(D + 1)$ -dimensional Euclidean space. In this representation, nodes occupy predetermined positions, with angular distances between them representing all the influential factors, aside from degrees, that affect their likelihood to form connections.

Connection probabilities are dictated by effective distances, which are derived from both hidden degrees and angular separations. This follows a gravity law principle, whereby nodes with higher hidden degrees or those

positioned closely in the similarity space are more likely to connect with others. The connection probability is influenced by two parameters: β , which must be greater than D , controls the clustering of the network ensemble and quantifies the level of coupling between the network topology and the metric space; and μ , which dictates the average degree. In this work, we restrict to $D = 2$ for ease of visualization and to constrain the number of variables involved. The similarity space of the \mathbb{S}^2 model is represented as a two-dimensional sphere of radius R , see section “ \mathbb{S}^2 model” for technical details. It is important to note that not all networks are optimally represented in $D = 2$. However, the results could only be enhanced by identifying the most suitable dimension for a given network.

On top of leveraging the capacity of the \mathbb{S}^D model to fit realistic network structures, *FiD*-Mercator employs the Uniform Manifold Approximation and Projection (UMAP) algorithm²⁸ for dimensionality reduction (section “The UMAP algorithm”). UMAP allows us to obtain an informed initial guess of the coordinates of the nodes in the similarity subspace using the information encoded into the nodes' features. These coordinates are subsequently adjusted by a maximum likelihood estimation technique to improve the probability that the observed network topology is reproduced by the model. Standard *D*-Mercator is instead initialized with coordinates on top of \mathbb{S}^D that only depend on the connectivity structure of the network. Given the complexity of the optimization process, such initial conditions may lead to reasonable but not optimal embeddings. Thus, the key contribution of *FiD*-Mercator over *D*-Mercator is that we determine a guided initial condition by mapping the high-dimensional vectors of node features on the two-sphere instead of extracting it from the network topology, such that the final embedding incorporates information about both the connectivity structure of the network and features describing nodes' qualities. Notice that other manifold learning techniques could be used instead of UMAP as far as they produce a projection in the two-sphere that is able to discern features. UMAP is particularly convenient for its favorable trade-off between accuracy and computation cost.

***FiD*-Mercator pipeline.** In practical terms, *FiD*-Mercator builds upon the multidimensional hyperbolic embedding tool *D*-Mercator²⁴. This tool employs a machine learning method and a statistical inference technique to pinpoint the hidden degrees and the angular coordinates that maximize the likelihood of accurately reproducing a network's topology through the geometric soft configuration \mathbb{S}^D model. Furthermore, it adjusts the parameters β and μ based on this model. *FiD*-Mercator leverages the maximum likelihood estimation machinery of *D*-Mercator. However, it differs in its approach to obtaining an initial condition from the network topology. Rather than employing the Laplacian Eigenmaps method²⁹ adjusted by the model, *FiD*-Mercator uses the UMAP dimension reduction technique to embed node features onto the two-sphere. UMAP is a well-established tool for representing and visualizing data in ultra-low dimensions. Offering exceptional efficiency with intuitive parameters, it makes a suitable choice for our applications.

The *FiD*-Mercator method requires two primary inputs. The first is the adjacency matrix, denoted as $\{a_{ij}\}$, which encodes the topology of the network, with $a_{ij} = 1$ if the link between nodes i and j exists and $a_{ij} = 0$ otherwise. The second input is the feature matrix, $\{f_{im}\}$. This matrix is formed by the collection of N_F binary features associated with the nodes, where $f_{im} = 1$ if node i exhibits feature m and $f_{im} = 0$ otherwise. The embedding process consists of the following steps: (1) inferring the hidden degrees κ , (2) inferring the inverse temperature β from the network topology, (3) initial embedding of the nodes using node vector features and the UMAP algorithm, (4) adjustment of the angular coordinates by maximizing likelihood (ML) to improve the congruency between the original network and the model, and (5) final adjustment of hidden degrees according to the final angular positions. These steps are explained in more detail in section “The *FiD*-Mercator algorithm”. The schematic representation of the algorithm is shown in Fig. 1.

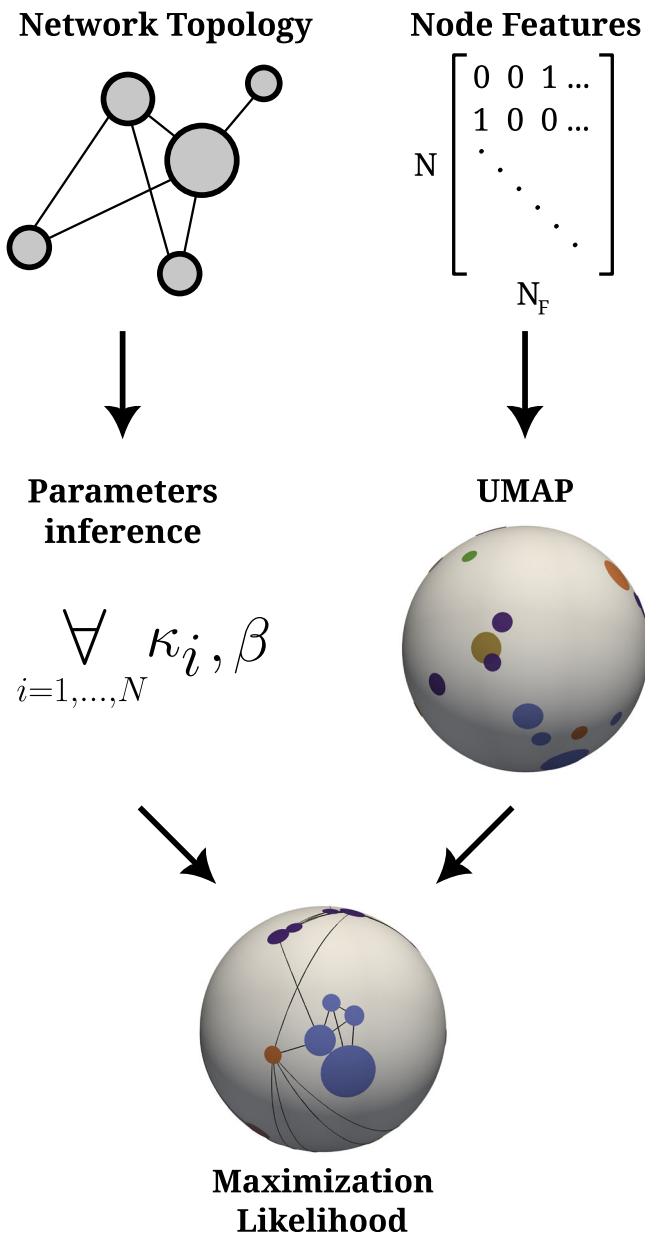


Fig. 1 | Schematic view on the proposed method *FiD-Mercator*. First, from the network we infer the hidden degrees κ and parameter β . Second, we use the UMAP algorithm to map the nodes onto the two-sphere using the feature matrix. Lastly, the initial nodes' coordinates from the last step are used in the maximization likelihood procedure which tries to fit data to the S^2 model. The size of the nodes is proportional to its expected degree and they are colored according to their communities. Black lines on the two-sphere represent connections produced according to the model.

Notice that by using UMAP, we can work with networks containing more than one connected component. This is a situation that cannot be addressed in *D-Mercator*, as the Laplacian Eigenmaps technique used to provide the initial condition there requires having only one connected component. Nevertheless, in this work we keep working with the largest connected component so that results can be compared with the baseline provided by *D-Mercator*.

An implementation of *FiD-Mercator* is publicly available at <https://github.com/networkgeometry/FiD-Mercator>.

Embedding of real networks. We gathered ten datasets, which describe structural connectivity and node features of real networks in various

domains; these range from citation and social networks to goods and webpage networks. It is noteworthy that a significant portion of the networks we analyzed are commonly used as benchmarks in machine learning research^{16,17,30}. Each network is accompanied by descriptive metadata describing the nodes' features as the bag-of-words. A non-zero entry indicates a presence of a given word for the node. For the citation networks, such as Citeseer or Cora, the node represents a publication, and a feature vector is the bag-of-words from the abstract of that publication. On the other hand, for Amazon Photo, the feature vectors are extracted from the product reviews. For more details about the datasets please refer to section "Datasets". In Table 1, we report global statistics for each network, including number of nodes and features, average degree, average number of features per node, and mean clustering coefficient.

We produced three different maps for each network: a *FiD-Mercator* embedding that incorporates information about the network topology and the node features; a *D-Mercator* embedding that only relies on the network connectivity structure; and the UMAP projection based solely on the node vector features. As an example, Fig. 2 shows the three embeddings of the LastFM network. Visualizations for the rest of the datasets are given in Supplementary Figs. 12–20.

The noticeable difference of nodes' layout between the *FiD-Mercator* and *D-Mercator* embeddings, as shown in Fig. 2, can be quantified by tracking how the global log-likelihood of the embeddings changes during the maximum likelihood optimization steps. Figure 3a, b (see also Supplementary Figs. 1 and 2) illustrates the time evolution of the global log-likelihood at each step of the maximum likelihood algorithm. In general, *FiD-Mercator* achieves a higher final value of global log-likelihood compared to *D-Mercator*. Thus, initialization with UMAP can lead to a local optimum that could improve the results of initializing with LE.

Moreover, we carried out experiments to observe how the nodes' coordinates change over time. We computed the average of the angular distances separating the positions of the nodes at $t = 0$ and every ML step. Figure 3c, d shows the results for Amazon Photo and LastFM datasets (see also Supplementary Figs. 3 and 4), where the average is computed for nodes with $k > 2$. The figures show that the average angular distance is increasing over time. However, at some moment, it saturates, meaning that the coordinates of nodes with $k > 2$ remain stable. This analysis can be used to assess how far the embedding is from the initial condition and specify when to stop the maximum likelihood process.

The embeddings can be combined with the probability of connection in the S^2 model to generate synthetic surrogates to be compared with the original networks, see Fig. 4 and Supplementary Figs. 21–29. One can observe that both *FiD-Mercator* and *D-Mercator* are able to reproduce the degree distribution, the clustering spectrum, and the average nearest neighbors degree with high fidelity. However, these embeddings are not identical and nodes in the *FiD-Mercator* maps show a greater tendency to expand and fill in the similarity subspace. In contrast, the UMAP embedding performs poorly in reproducing network properties as expected, since it contains no information about the network structure. We conclude that adding node features to network structure did not disrupt the quality of the model-based hyperbolic embedding in terms of its capacity at explaining local network topological properties. Next, we are going to explore how adding features can affect downstream tasks that also depend on the mesoscopic and global organization of the networks.

Downstream tasks

In this section, we show that enriching network topology with node features can have a mixed effect depending on how the two are correlated and on the extent to which each category is relevant for the downstream task to be performed. Next, we report results for link prediction and node classification tasks.

Link prediction. Given a partially observed network, link prediction (LP) aims to infer the most likely missing links based on the available ones³¹. Typically, absent links are ranked in decreasing order of some likelihood

Table 1 | Properties of real networks

Network	N	$\langle k \rangle$	\bar{c}	N_F	$\langle k_F \rangle$	N_L^{Class}	$N_L^{Topology}$	$N_L^{Features}$	β	μ	$corr(\mathcal{G}, F)$
Citeseer	2110	3.48	0.23	3703	32.06	6	37	4	2.98	0.0375	0.763
Cora	2485	4.08	0.28	1433	18.30	7	27	5	3.23	0.0373	0.650
DBLP	2728	31.44	0.67	2000	42.14	4	18	4	7.21	0.0089	0.506
IMDB	3228	19.46	0.55	2000	76.95	3	27	5	5.12	0.0125	0.172
Amazon Photo	7487	31.8	0.42	745	260.47	8	14	7	3.41	0.0052	0.270
Cornell	184	3.03	0.29	1703	94.20	5	14	4	2.40	0.0200	0.169
Wisconsin	251	3.59	0.28	1703	95.84	5	12	5	2.49	0.0204	0.201
Texas	183	3.05	0.32	1703	83.42	5	12	4	2.24	0.0123	0.116
LastFM	7624	7.29	0.29	7842	395.38	18	29	5	2.96	0.0175	0.465
Twitch PTBR	1912	32.74	0.34	3169	19.88	2	8	5	2.39	0.0018	-0.053

The N_F represents the number of features, $\langle k_F \rangle$ the average number of features per node and N_L the number of labels. We used the giant connected component for all networks.

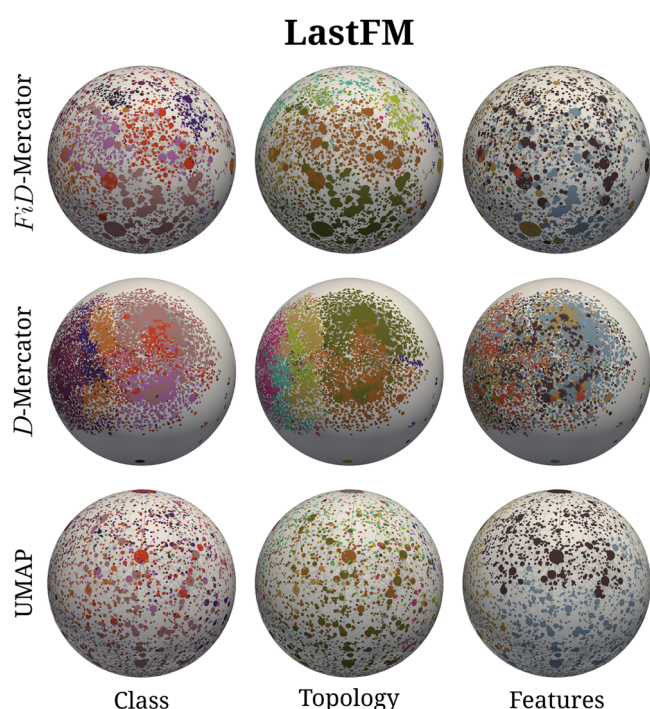


Fig. 2 | Two dimensional maps of the LastFM network. Each row corresponds to the different embedding methods whereas each column is for a different assignment of labels. The size of a node is proportional to its expected degree, and its color indicates the community it belongs to. For the sake of clarity, only the connections with probability (Eq. (1)) $p_{ij} > 0.999$ are shown.

value and those at the top are selected as candidates for missing links. The likelihood can be estimated based on some heuristic metric given the patterns of connectivity observed in the network structure or using an underlying network model. When it comes to graph models that employ latent space, the task of link prediction essentially involves ranking pairs of nodes in decreasing order of their connection probability computed using their latent distances. Results from ref. 32 indicate that as the complexity of a specific link prediction task increases, the utilization of hyperbolic geometry should be more seriously considered.

We conducted tests on the link prediction task in two scenarios. In the first scheme, we randomly removed $L = q \cdot E$ links from the original graph with q defined as the fraction of missing links. Then, we used the coordinates and parameters obtained from the three different embeddings of the complete network to compute the connection probability as given by S^2 (Eq. (1)).

In the second scheme, we selected the giant connected component (GCC) of the network after removing L links, counted the number of these missing links L' within the GCC, and embedded this subgraph to compute the coordinates of its nodes. We then calculated the likelihood of each non-existing link in the GCC.

In both cases, we sorted the links in decreasing order of probability, and selected the top L and L' , respectively, as predicted links. Even though the Area Under the ROC Curve (AUC) is a popular measure in the literature to assess the goodness of a link prediction method, the precision is a more stringent measure since it concentrates on the top part of the ranking where, under limited resources, more detection efforts would be directed³³⁻³⁵. The precision was calculated as the fraction of correctly identified links as presented in ref. 35.

The results are reported in Fig. 5 and Supplementary Figs. 9 and 10. Again, UMAP alone performs poorly whereas *FiD-Mercator* and *D-Mercator* give improved results. In the first scheme, in seven out of the 10 analyzed real networks *FiD-Mercator* outperforms *D-Mercator* by a significant amount. In the remaining three networks, the two schemes reach similar levels of precision. In the second scheme, the overall precision is lower due to the increased difficulty of the task. We excluded three networks (Cornell, Wisconsin, and Texas) from this analysis due to their small size after the removal of links. In three out of the remaining seven networks, *FiD-Mercator* surpasses *D-Mercator*, and in two networks, both methods yield similar results.

The improvement of *FiD-Mercator* over *D-Mercator* in the LP task can only be explained if node features are correlated with the network topology. To quantify such correlation, we computed the relative difference between the average cosine similarities of features of connected nodes in the original network and features of nodes that were not directly connected (see section “Correlation between nodes’ features and network structure” and Table 1). As one can notice, higher correlation values for Amazon Photo and DBLP lead to improvement in the link prediction task. On the other hand, when the correlation is weak, both *FiD-Mercator* and *D-Mercator* perform similarly, as in the case of the Twitch PTBR dataset.

Node classification. Node classification (NC) is a task where the goal is to categorize nodes into different sets described by specific labels. The categories can be based on different sources of information. For each real network, we considered three different ways of assigning labels to nodes.

- **Class:** using descriptive metadata, different from features, which characterizes some aspect that divides nodes in a limited number of classes, for instance product types, professional roles, research fields, or film genres. The metadata defining the classes is typically provided along the network graph. In this partition, we label nodes with the class names.

Fig. 3 | Evolution of *FiD*-Mercator and *D*-Mercator embeddings during the maximum likelihood optimization process. a, b Evolution of the global log-likelihood during the maximum likelihood (ML) optimization steps. Each ML step takes a subset of nodes for which new coordinates are proposed. The nodes are ordered through the onion decomposition⁴⁰. Notice that we plot the negative log-likelihood here, hence the lower the value the better. **c, d** The average angular distance between each node of the initial embedding (for *D*-Mercator: Laplacian Eigenmaps; for *FiD*-Mercator: UMAP) and the embedding after each ML step. We computed the average angular distance only for nodes with $k > 2$.

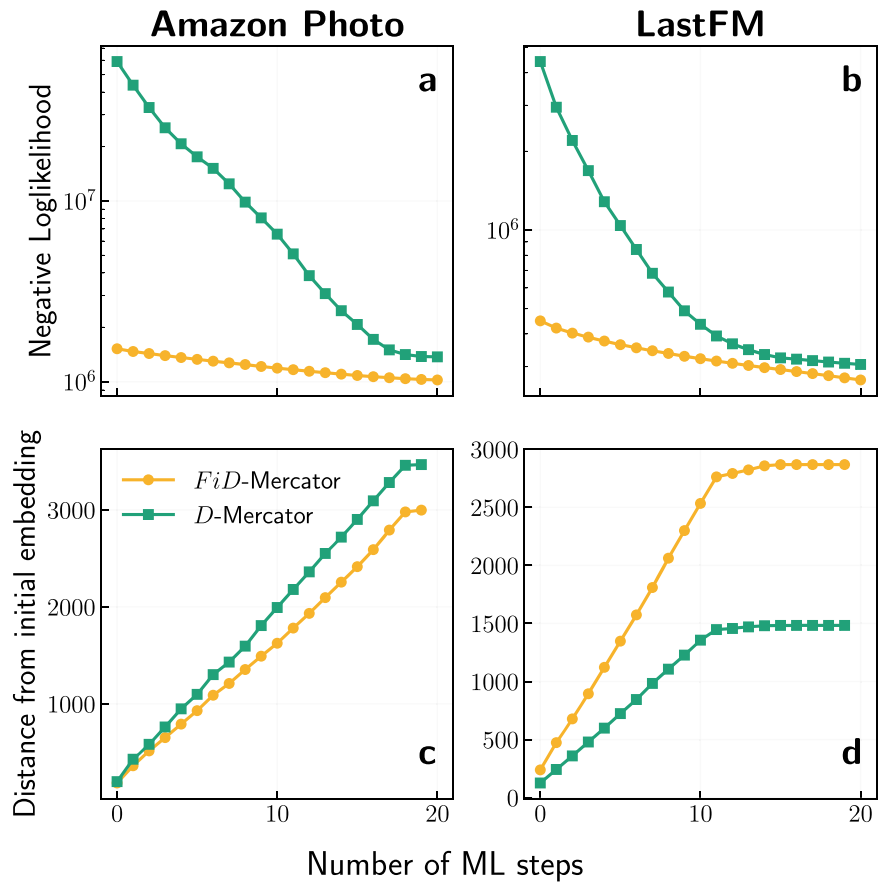


Fig. 4 | Validation of the embeddings of the Amazon Photo network. a shows the complementary cumulative degree distribution and **b** the clustering spectrum $\bar{c}(k)$. Symbols correspond to the values in the original network. The lines indicate an estimate of the expected values in the ensemble of random networks in the different embeddings. The S^2 model was used to generate 100 synthetic networks with the parameters and positions inferred by *FiD*-Mercator, *D*-Mercator, or UMAP. The error bars show the 2σ confidence interval around the expected value. **c** shows the average nearest neighbors degree $\bar{k}_{nn}(k)$. **d** displays the comparison of the expected connection probability based on the estimated β (expected) and the actual connection probability computed with the inferred hidden variables.

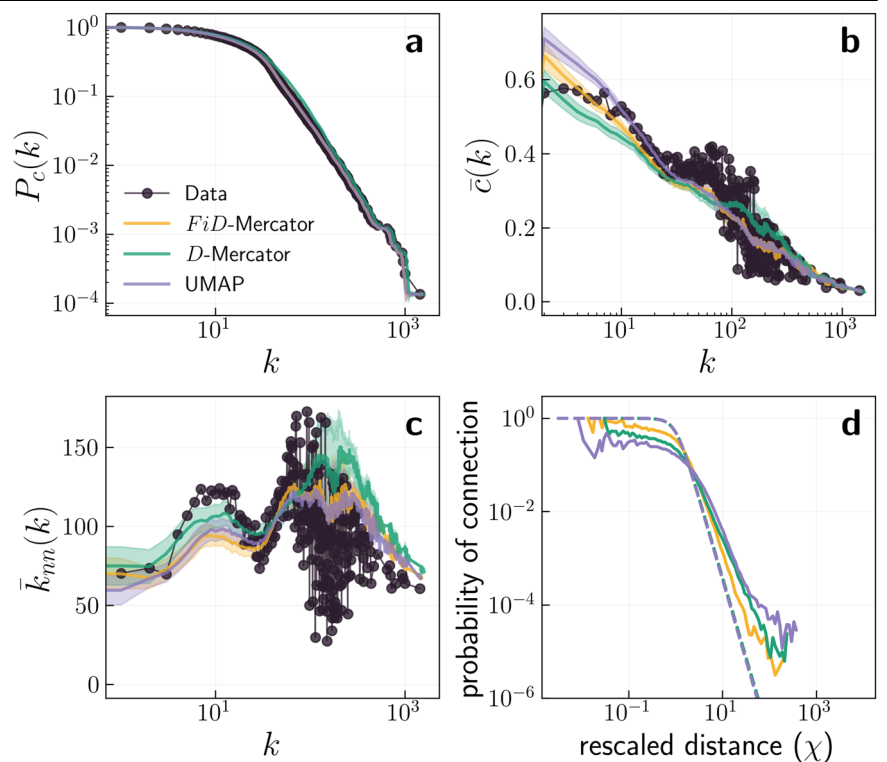
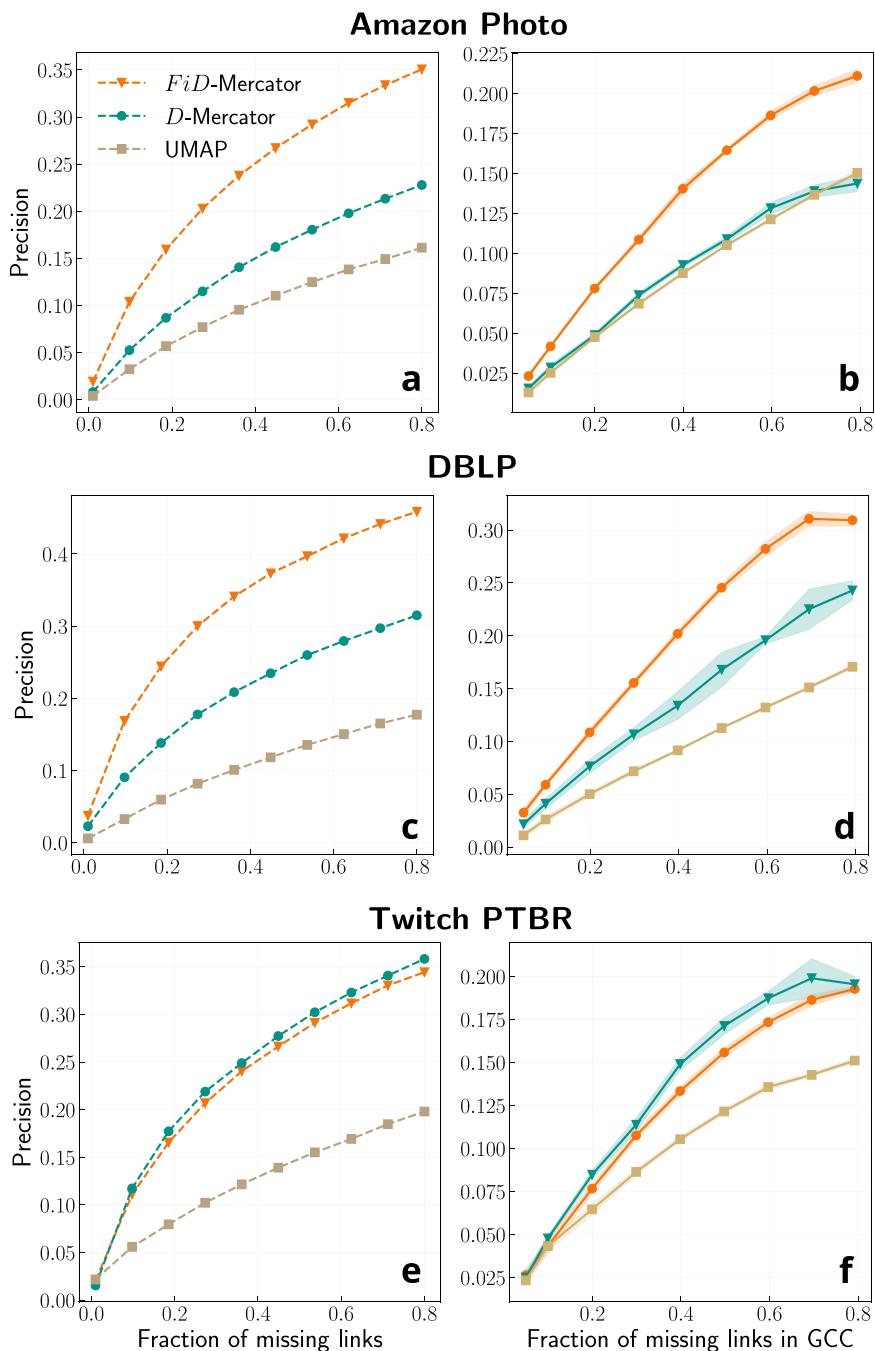


Fig. 5 | Precision of a link prediction task as a function of the fraction of missing links in real networks. **a, b** Amazon Photo, **(c, d)** DBLP, and **(e, f)** Twitch PTBR. **a, c, e** show the results for the scheme in which the embeddings are of the complete network, whereas **b, d, f** are for the second scheme embedding the GCC of the incomplete network. Results are averaged over five different realizations.



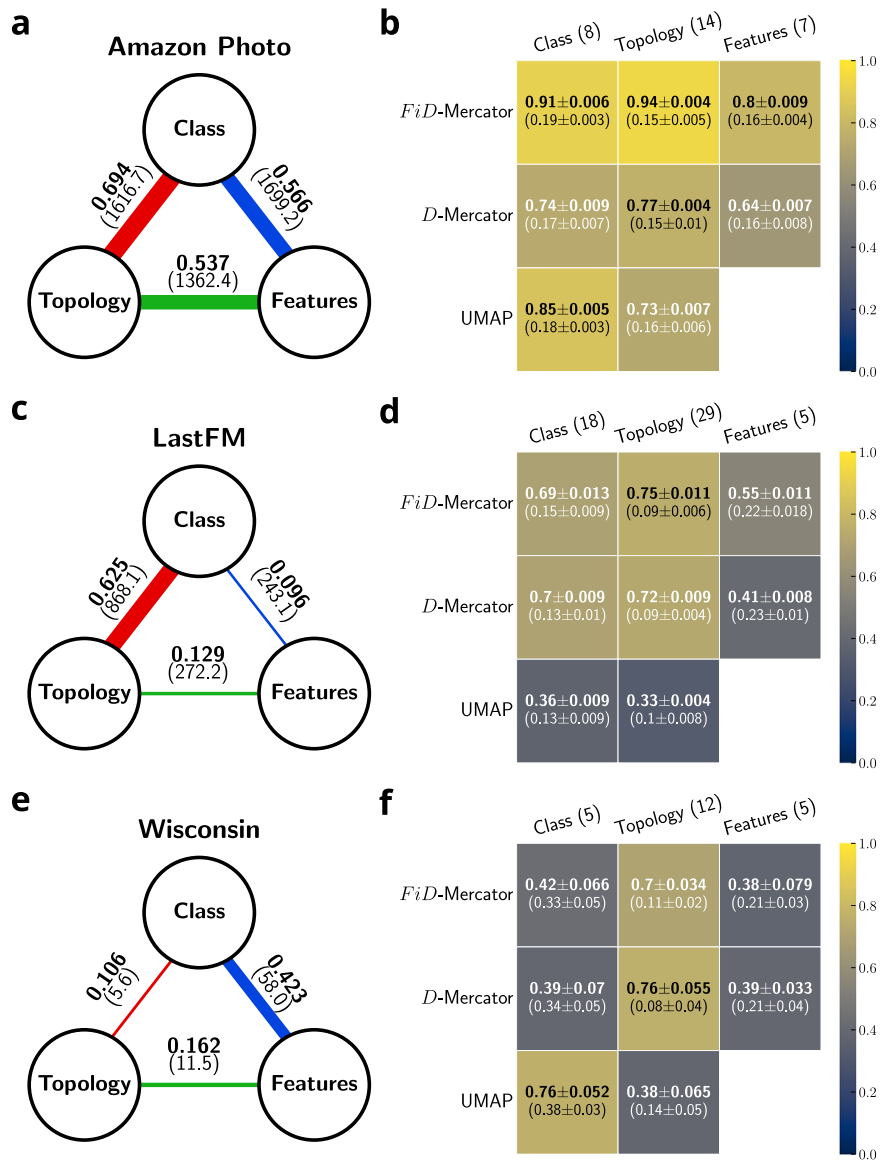
- **Topology:** by leveraging the network topology to construct the communities of densely connected neighbors. Labels based on topological communities were produced using the Louvain method³⁶, which is a greedy modularity optimization technique. The method is fast and detects the number of communities without predefining it.
- **Features:** based on similarities between node features. Feature-based labels were defined from the UMAP embedding used in *FiD-Mercator* to fix the initial conditions. We applied the agglomerative clustering algorithm from the scikit-learn Python library³⁷ to the UMAP projection. This algorithm is a bottom up hierarchical clustering method in which each node starts in its own cluster, and clusters are successively merged together following a linkage criterium that takes the mean distance between elements of each cluster, also called average linkage clustering used in UPGMA³⁸. The nodes projected by UMAP in

close positions of the similarity space, i.e., separated by a small angular distance, have a tendency to merge into one cluster. To fix the number of labels in the hierarchical clustering algorithm, we employed the geometric community concentration measure used in ref. 24. This measure computes the level of clusterization of nodes in the underlying two-sphere embedding. The number of clusters is tuned to the one that produces a maximum value of the geometric concentration measure, see section “Community concentration” for more details.

Each node in the real networks was, thus, assigned three labels—a class label, a topological community label, and a feature-based label—one from each of the three partitionings. This labeling is visualized through color-coding in Fig. 2, which depicts three different maps of the LastFM network, and in Supplementary Figs. 12–20 for the remaining networks. It is worth noting that labels denoting classes and labels based on network topology and

Fig. 6 | Correlation between node labeling and accuracy in node classification task.

a, c, e Normalized mutual information (NMI) between each of two different sets of labels. The value in bracket represents the z-score which quantifies the significance of the correlation between two label sets. A higher z-score indicates a stronger and more statistically significant correlation between the two sets of labels (see section “z-score”). The width of the line is proportional to NMI. **b, d, f** Accuracy heatmaps of node classification task. Each row represents a different embedding method, whereas each column a different set of labels. In brackets, the number of labels is shown. The performance of node classification on shuffled labels are displayed in brackets as the second row in each entry. Train-test size split: 80/20. Results are averaged over 5 different train-test splits.



on node features do not necessarily correlate with one another. For instance, the absence of correlation is the reason why the positions of nodes in different feature-labeled sets overlap in the embeddings obtained by *FiD*-Mercator and *D*-Mercator.

In order to see how the addition of features helps hyperbolic embedding we need to characterize the level of correlation between the three sets of labels. For this, we used the Normalized Mutual Information (NMI) to evaluate the agreement between the labels associated to each node in each class. Results are reported in Fig. 6a, c, e and in Supplementary Figs. 7 and 8. There is always a significant amount of correlation between the labels. However, particular associations are more significant in specific real networks. For instance, all correlations are more important in Amazon Photo as compared to Wisconsin; in LastFM the association between classes and topology-based communities is higher than between classes and feature-based groups. Contrarily, Wisconsin exhibits the reverse trend, with classes being highly correlated with labels derived from features.

To evaluate the significance of the observed correlations, we calculate the z-score between two sets of labels. This measure compares the NMI of the original label sets with the NMI obtained when one set of labels is shuffled. A higher z-score indicates a more significant correlation between the two sets of labels (see section “z-score”). In all the analyzed networks, the observed correlations are significant (z-score > 5), except for the Texas and

Cornell datasets, where the correlation between the classes and the topology-based communities is not statistically relevant.

These correlations explain the performance of a node classification task defined on real network embeddings for the different labeling protocols (we omitted the evaluation of classifying feature-based labels in UMAP embeddings since those labels were crafted from that embedding). First, we split the nodes into training and test sets. The popular choice is 80% of nodes into the training set and the rest 20% into the test set³⁹. Using the information of node labels in the training set, we predicted the labels of nodes in the test set using the *K*NeighborsClassifier from the scikit-learn library³⁷. This neighbors-based classification method assigns the data category which has the most representatives within the *k* nearest neighbors of the node. In our experiments, we set *k* = 5 (we also checked *k* = 10, which gave similar results) and used the angular distance between the nodes in the similarity subspace as the metric to define nearest neighbors. The results were averaged over 5 train-test splits.

Heatmaps showing the achieved accuracy in predicting the nodes’ labels of Amazon Photo, LastFM and Wisconsin are shown in Fig. 6b, d, f. The results for the remaining networks are shown in Supplementary Fig. 11. When the classes are correlated with the topology-based communities, *FiD*-Mercator tends to improve the accuracy in assigning classes to nodes as compared to *D*-Mercator and UMAP (Fig. 6b, d, f, first column). Moreover,

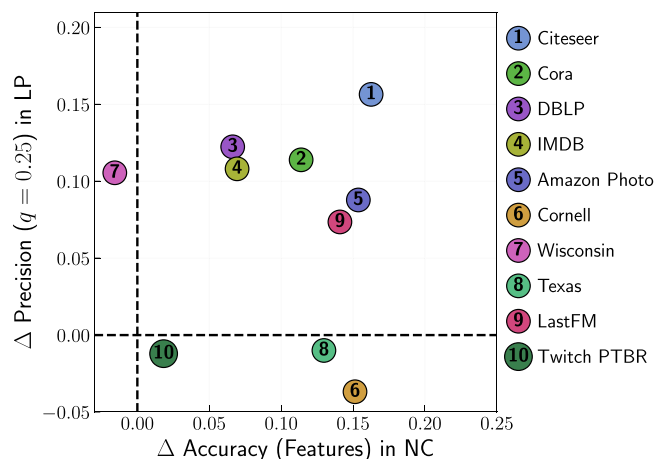


Fig. 7 | The difference of precision and accuracy between the *FiD-Mercator* and *D-Mercator* for real networks. For link prediction the precision was computed at 0.25 fraction of missing links. For node classification the accuracy was calculated for feature-based labels. Train-test size split: 80/20. Results are averaged over five different train-test splits.

FiD-Mercator outperforms in correctly classifying topology-based labels (Fig. 6b, d, second column). On the contrary, when there is a lack of correlation between the two categories, *D-Mercator* performs better (Fig. 6f, second column). The *FiD-Mercator* embedding achieves higher accuracy when predicting labels constructed from features, as shown in Fig. 6b, d (third column). Furthermore, when classes correlate with feature-based labels, incorporating node features into the embedding method enhances accuracy, which aligns with expectations. This is illustrated in Fig. 6b, d (first column). In all cases, UMAP achieves low accuracy predicting nodes' labels when the correlation between classes and feature-based groups is weak, and performs very poorly in classifying on topology-based communities. It is worth mentioning that for all embedding methods, the accuracy is higher than in randomized surrogates, in which the labels were shuffled, i.e., random assignment of the labels, while keeping the original number of communities.

All together, the experiments demonstrate a nuanced effect of incorporating feature data into network structure for node classification tasks based on geometric network embeddings. It can be detrimental when the correlation between features and topology is very low. However, when there is a high correlation between network connectivity and features, adding them significantly enhances the results.

Versatility. To clarify whether incorporating features into the embedding method provides an advantage in a variety of tasks, we show in Fig. 7 the improvement in LP offered by *FiD-Mercator* over *D-Mercator* as a function of the improvement in NC using feature-based labels, for all datasets. In 60% of the real networks the benefit is obvious in both tasks, while in three out of 10 networks the results for LP, but not for NC, are slightly worse and in one network the result for NC, but not for LP, is slightly negative. Taken together, these results suggest that feature-aware embedding methods can improve a variety of downstream tasks beyond the ones explored here.

Discussion

Real-world complex networks are typically represented as graphs where connections are often enriched with attributes, such as the intensities of the interactions or their directionality. These attributes are incorporated in the network analysis process to obtain a better approximation to the actual behavior of the system. However, when node descriptive metadata different from connectivity is available, in the form of node features describing their qualities (e.g. in social systems, age, gender, interests, background, and other social determinants), the common practice in network science is to not

incorporating it into the analysis but to use the information as an external source of verification. For instance, in validating aggregation methods like community detection, determining node categories based on network structure. By contrast, advanced machine learning techniques in computer science, for instance many deep neural networks, incorporate information about both network structure and node features. In this case, the difficulty comes from the fact that the approaches are data-driven and the methods are black boxes in which the blending process becomes opaque and it is impossible to discern the precise role of every contributing factor.

In this work, we surmount these drawbacks and merge network structure with node features in a model-based environment to obtain a more advantageous ultra-low dimensional reduction technique for real networks. This hybridization technique, that we named *FiD-Mercator*, produces *D*-dimensional geometric representations of real networks in hyperbolic space. The features determine a convenient initial condition to guide more efficiently the embedding of the network topology toward a better local optimum. The relative attention paid to the network structure versus the features can be controlled by early stoppage guided by the evolution of the log-likelihood function. As a result, the hybridization reveals a different type of spatial description by implementing a different inductive bias, that is, by considering a set of assumptions that involve both features and network structure simultaneously, and independently of whether those are or not correlated. Moreover, we found that adding node features to network structure does not disrupt the quality of the hyperbolic embedding in terms of its capacity at explaining local node properties, determining the degree distribution or the clustering spectrum.

However, its impact can vary depending on the correlation between the two and the relevance of each piece of information to a specific downstream task. In link prediction, the major improvement of *FiD-Mercator* is achieved in cases when the correlation is high but not perfect. Perfect correlation means that the information contained in the nodes' features would be redundant. Instead, high correlation below one implies that there is new information encoded in the features that can be used to infer structural properties of the network. When the correlation is very low, features and network connectivity are hardly related and the prediction of missing links cannot benefit from adding features. This suggests that there is an optimal correlation value for which better results are obtained. For node classification tasks, adding features can be detrimental when the correlation between features and topology is very low. In contrast, adding features significantly enhances the results when the correlation is high.

Nevertheless, there may be tasks where adding features is not optimal, independently of the level of correlation, which is difficult to know a priori. Other open questions for future research are how to smartly filter the set of features to adjust the correlation between features and network connectivity for optimal task performance, and the effects of incorporating non-binary features into the analysis, which could be easily mapped with UMAP. In any case, our results indicate that, overall, features are in general correlated with network connectivity and that the level of correlation is a good indicator of their relevance to improve many downstream tasks based on geometric network embeddings. This emphasizes the need for the development of new network embedding techniques that simultaneously optimize the joint probability of network structure and feature association in the near future.

Methods

\mathbb{S}^2 model

In the \mathbb{S}^2 model, a node *i* is endowed with a hidden degree κ_i and a position in the 2-similarity space, $\mathbf{v}_i = \{x_i, y_i, z_i\}$ with $\|\mathbf{v}_i\| = R$ where *R* is the radius of the two-sphere. The connection probability between a node *i* and a node *j* has form:

$$p_{ij} = \frac{1}{1 + \chi_{ij}^\beta}, \text{ with } \chi_{ij} = \frac{R\Delta\theta_{ij}}{\sqrt{\mu\kappa_i\kappa_j}}. \tag{1}$$

We fix the density of nodes in the two-sphere to one so that:

$$R = \sqrt{\frac{N}{4\pi}} \tag{2}$$

where N is the number of nodes. The angular distance is defined as $\Delta\theta_{ij} = \arccos(\frac{\mathbf{v}_i \cdot \mathbf{v}_j}{R^2})$. The parameter $\beta > 2$, named inverse temperature, controls the clustering in the network. Finally, the parameter μ controls the average degree of the network and is defined as:

$$\mu = \frac{\beta \sin \frac{2\pi}{\beta}}{2\pi^2 \langle k \rangle}. \tag{3}$$

The hidden degrees can be generated randomly from an arbitrary distribution, or taken as a set of prescribed values. The model has the property that the expected value of the degree of a node with hidden variable κ is $\bar{k}(\kappa) = \kappa$.

The FiD-Mercator algorithm

For a complete description of the inference of the hidden degrees and parameter β , the likelihood maximization and technique, and the final adjustment of hidden degrees please refer to ref. 24.

Inferring the hidden degrees and parameter β . The inference of the hidden degree and the inverse temperature β is implemented as an iterative process. We start with the approximate guess for $\beta \in (2, 3)$ and initialize the hidden degrees as the observed degrees $\{k_i, i, \dots, N\}$ in the original network. The estimation aims to adjust the hidden degrees such that the expected degree of each node in the model matches the observed degree in the original network. After the hidden degrees are computed, the theoretical mean of the local clustering coefficient of networks in \mathbb{S}^2 ensemble can be evaluated. If its value varies from the original network, \bar{c} , the value of β is adjusted. Then the process is rerun using the current estimation of hidden degrees until a predetermined precision is reached.

Initial embedding of the nodes from UMAP. The UMAP takes as the input the nodes' features matrix $N \times N_f$ with the distance metric. In our case, we use the Haversine metric, which is designed to calculate distances on a spherical surface. This process positions the nodes within the embedding space. For more details about the UMAP algorithm, see section "The UMAP algorithm".

Likelihood maximization. The nodes' coordinates in the similarity space inferred using UMAP are fine-tuned by Maximum Likelihood Estimation (MLE) to optimize the probability that the observed network is generated by the \mathbb{S}^2 model. Nodes are visited sequentially by the network's onion decomposition⁴⁰. For each node, we propose the candidate's positions near its neighbors. The most favorable proposed position, i.e., maximizing the local log-likelihood, is selected, and the process is repeated until the local log-likelihood function reaches a plateau.

Final adjustment of hidden degrees. Lastly, we correct the hidden degree to compensate deviations from $\bar{k}(\kappa_i) = \kappa_i$, which might have been introduced in estimating the coordinates of nodes in the similarity spaces.

The UMAP algorithm

UMAP (Uniform Manifold Approximation and Projection)²⁸ is a dimension reduction technique. By default, it embeds data in Euclidean space. This is a convenient and widely used option for visualization and clustering^{41,42}. However, there are no significant constraints that prevent the algorithm from working with non-Euclidean spaces.

UMAP allows you to choose the dimension of the embeddings. We have verified that with $D = 1$ performs poorly because we first need to map the nodes to a 2D plane and then place them on a circle. As this procedure mixes communities, we chose $D = 2$ for UMAP dimension.

UMAP maps the node feature vectors on top of the 2-sphere, which can be treated as a similarity space of the \mathbb{S}^2 model. It is worth noting that the UMAP algorithm is unsupervised, meaning it does not require any node labels.

The UMAP algorithm, being stochastic, can produce different results in various runs. While the inferred positions of the nodes might slightly vary between these runs, our tests, involving running the UMAP algorithm twice, indicate that the obtained embeddings lead to consistent results in downstream tasks.

Correlation between nodes' features and network structure

Let us define a complementary graph of the network topology, $\mathcal{G}' = (V, E')$ with $E' \subseteq \{(u, v) \mid u, v \in V \text{ and } (u, v) \notin E\}$ where V is the set of nodes in the original network and E' is a set of edges non-existing in the original network. Computing cosine similarity between all non-existing links in \mathcal{G} is computationally expensive due to its sparsity. As an alternative, we randomly selected $|E|$ edges from the complementary graph and averaged the results over five realizations of the selection process for further analysis.

For each realization, we computed the average cosine similarity between connected nodes in graph \mathcal{G} and between node pairs in the selected samples of \mathcal{G}' . The relative difference between them quantifies the strength of the correlation between the nodes' features and the networks structure:

$$\text{corr}(\mathcal{G}, F) = \frac{\overline{\text{CS}}(\mathcal{G}) - \overline{\text{CS}}(\mathcal{G}'_E)}{\overline{\text{CS}}(\mathcal{G})}. \tag{4}$$

The higher the obtained value the more correlated are features with the network topology.

Community concentration

Introduced in ref. 24, the geometric concentration of a community l around a node i has a form:

$$\rho_{i,l} = \frac{n_{i,l}}{n_{i,g}} \frac{N}{N_L}, \tag{5}$$

where $n_{i,l}$ is the number of nodes in community l out of $n_{i,g}$ considered nodes, and N_L is the total number of nodes in community l . We define $n_{i,g}$ as top geometrically closest neighbors. The single scalar ρ for a given network embedding is averaged over all nodes in the same community. Finally, we calculated the community concentration as $c_C = \rho(n_{i,l} = N/10)$, i.e., the geometric concentration at 10% of top geometrically closest nodes.

z-score

To measure the significance of the correlation between two sets of labels, we used z-score. Let us define the quantity:

$$z\text{-score} = \frac{\text{NMI}(X, Y) - \overline{\text{NMI}(X, Y)}_{\text{rand}}}{\sigma(\text{NMI}(X, Y)_{\text{rand}})}, \tag{6}$$

where X and Y are sets of labels. We compute the random case of by comparing set X with a shuffled version of set Y , and this process is repeated 100 times. The higher z-score the more significant is the correlation between two sets of labels.

Datasets

In this work we analyzed ten networks from different domains. Their properties are summarized in Table 1. For each dataset, the nodes' feature vector is a high-dimensional vector containing 0s or 1s, i.e., the presence or absence of a given feature:

- Citeseer⁴⁵: the citation network of Machine Learning papers where each publication is described by a 0 or 1 valued word vector indicating the absence or the presence of the corresponding word from the dictionary.

The dictionary consists of 1433 unique words. The publications are classified into six classes.

- Cora⁴⁴: similar to Citeseer, however the publications are split into seven classes: Case Based, Genetic Algorithms, Neural Networks, Probabilistic Methods, Reinforcement Learning, Rule Learning, Theory.
- DBLP⁴⁵: the citation network of Machine Learning papers. Nodes are classified into four classes: Data Mining, Artificial Intelligence, Computer Vision and Natural Language Processing.
- IMDB⁴⁵: the Movie-Actor-Movie relation dataset. Movies are categorized into three classes (Action, Comedy, Drama).
- Amazon Photo⁴⁶: nodes represent goods and edges represent that two goods are frequently bought together. The node features are bag-of-words encoded product reviews, and class labels are given by the product category.
- Cornell, Wisconsin, Texas⁴⁷: web graphs crawled from three Computer Science departments in 1998, with each page manually classified into one of seven categories: course, department, faculty, project, staff, student, or other.
- LastFM⁴⁸: a social network collected from the public API in March 2020. Nodes are LastFM users from Asian countries and edges are mutual follower relationships between them. The vertex features are extracted based on the artists liked by the users. This node label was derived from the country field for each user.
- Twitch PTBR⁴⁹: user-user networks of gamers who stream in Portuguese language. Nodes are the users themselves and the links are mutual friendships between them. Vertex features are extracted based on the games played and liked, location and streaming habits. The node labels represent whether a streamer uses explicit language.

Data availability

The network datasets used in this study are available from the sources referenced in the manuscript and the Supplementary Materials.

Code availability

The open-source code of FiD-Mercator is available on GitHub at <https://github.com/networkgeometry/FiD-Mercator>.

Received: 31 January 2024; Accepted: 21 June 2024;

Published online: 02 September 2024

References

- Serrano, M. Á., Krioukov, D. & Boguñá, M. Self-similarity of complex networks and hidden metric spaces. *Phys. Rev. Lett.* **100**, 078701 (2008).
- Boguñá, M. et al. Network geometry. *Nat. Rev. Phys.* **3**, 114–135 (2021).
- García-Pérez, G., Allard, A., Serrano, M. Á. & Boguñá, M. Mercator: uncovering faithful hyperbolic embeddings of complex networks. *N. J. Phys.* **21**, 123033 (2019).
- Krioukov, D., Papadopoulos, F., Kitsak, M., Vahdat, A. & Boguñá, M. Hyperbolic geometry of complex networks. *Phys. Rev. E* **82**, 036106 (2010).
- Hays, R. B. A longitudinal study of friendship development. *J. Personal. Soc. Psychol.* **48**, 909 (1985).
- Haynie, D. L., Doogan, N. J. & Soller, B. Gender, friendship networks, and delinquency: a dynamic network approach. *Criminology* **52**, 688–722 (2014).
- Bigelow, B. J. Children's friendship expectations: a cognitive-developmental study. *Child Dev.* **48**, 246–253 (1977).
- Boguñá, M., Papadopoulos, F. & Krioukov, D. Sustaining the Internet with hyperbolic mapping. *Nat. Commun.* **1**, 62 (2010).
- Allard, A. & Serrano, M. Á. Navigable maps of structural brain networks across species. *PLoS Comput. Biol.* **16**, e1007584 (2020).
- Peel, L., Laremore, D. B. & Clauset, A. The ground truth about metadata and community detection in networks. *Sci. Adv.* **3**, e1602548 (2017).
- Newman, M. E. J. & Clauset, A. Structure and inference in annotated networks. *Nat. Commun.* **7**, 1–16 (2016).
- Bassolas, A., Holmgren, A., Marot, A., Rosvall, M. & Nicosia, V. Mapping nonlocal relationships between metadata and network structure with metadata-dependent encoding of random walks. *Sci. Adv.* **8**, eabn7558 (2022).
- Emmons, S. & Mucha, P. J. Map equation with metadata: varying the role of attributes in community detection. *Phys. Rev. E* **100**, 022301 (2019).
- Smith, L. M., Zhu, L., Lerman, K. & Percus, A. G. Partitioning networks with node attributes by compressing information flow. *ACM Trans. Knowl. Discov. Data* **11**, 1–26 (2016).
- Artime, O. & De Domenico, M. Percolation on feature-enriched interconnected systems. *Nat. Commun.* **12**, 2478 (2021).
- Kipf, T. N. & Welling, M. Semi-Supervised Classification with Graph Convolutional Networks. *Int. Conf. Learn. Represent.* (2017).
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P. & Bengio, Y. Graph attention networks. *Int. Conf. Learn. Represent.* (2018).
- Hamilton, W. L., Ying, R. & Leskovec, J. Inductive representation learning on large graphs. *Adv. Neural Inf. Process. Syst.* **30** (2017).
- Xu, K., Hu, W., Leskovec, J. & Jegelka, S. How powerful are graph neural networks? *Int. Conf. Learn. Represent.* (2019).
- Ma, Y., Liu, X., Shah, N. & Tang, J. Is Homophily a Necessity for Graph Neural Networks? *Int. Conf. Learn. Represent.* (2022).
- Zhu, J. et al. Beyond homophily in graph neural networks: current limitations and effective designs. *Adv. Neural Inf. Process. Syst.* **33**, 7793–7804 (2020).
- Grover, A. & Leskovec, J. node2vec. *ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.* (2016).
- Perozzi, B., Al-Rfou, R. & Skiena, S. Deepwalk: online learning of social representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 701–710 (2014).
- Jankowski, R., Allard, A., Boguñá, M. & Serrano, M. Á. The D-Mercator method for the multidimensional hyperbolic embedding of real networks. *Nat. Commun.* **14**, 7585 (2023).
- Almagro, P., Boguñá, M. & Serrano, M. Á. Detecting the ultra low dimensionality of real networks. *Nat. Commun.* **13**, 6096 (2022).
- Budel, G., Kitsak, M., Aldecoa, R., Zuev, K. & Krioukov, D. Random hyperbolic graphs in d.1 dimensions. *Phys. Rev. E* **109**, 054131 (2024).
- Krioukov, D., Papadopoulos, F., Vahdat, A. & Boguñá, M. Curvature and temperature of complex networks. *Phys. Rev. E* **80**, 035101(R) (2009).
- McInnes, L. et al. UMAP: Uniform Manifold Approximation and Projection. *J. Open Source Softw.* **3**, 861 (2018).
- Belkin, M. & Niyogi, P. Laplacian eigenmaps and spectral techniques for embedding and clustering. In *Advances in Neural Information Processing Systems*, 585–591 (2002).
- Pei, H., Wei, B., Chang, K. C.-C., Lei, Y. & Yang, B. Geom-gcn: Geometric graph convolutional networks. *Int. Conf. Learn. Represent.* (2020).
- Lü, L. & Zhou, T. Link prediction in complex networks: a survey. *Phys. A Stat. Mech. Appl.* **390**, 1150–1170 (2011).
- Kitsak, M., Voitalov, I. & Krioukov, D. Link prediction with hyperbolic geometry. *Phys. Rev. Res.* **2**, 043113 (2020).
- Lü, L., Pan, L., Zhou, T., Zhang, Y.-C. & Stanley, H. E. Toward link predictability of complex networks. *Proc. Natl Acad. Sci.* **112**, 2325–2330 (2015).
- Sun, J. et al. Revealing the predictability of intrinsic structure in complex networks. *Nat. Commun.* **11**, 574 (2020).
- García-Pérez, G., Aliakbarisani, R., Ghasemi, A. & Serrano, M. Á. Precision as a measure of predictability of missing links in real networks. *Phys. Rev. E* **101**, 052318 (2020).

36. Blondel, V. D., Guillaume, J.-L., Lambiotte, R. & Lefebvre, E. Fast unfolding of communities in large networks. *J. Stat. Mech. Theory Exp.* **2008**, P10008 (2008).
 37. Pedregosa, F. et al. Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
 38. Sokal, R. R. & Michener, C. D. A statistical method for evaluating systematic relationships. *Univ. Kans. Sci. Bull.* **38**, 1409–1438 (1958).
 39. Gholamy, A., Kreinovich, V. & Kosheleva, O. *Why 70/30 or 80/20 Relation Between Training and Testing Sets: A Pedagogical Explanation*. Departmental Technical Reports (CS) (2018).
 40. Hébert-Dufresne, L., Grochow, J. A. & Allard, A. Multi-scale structure and topological anomaly detection via a new network statistic: the onion decomposition. *Sci. Rep.* **6**, 1–9 (2016).
 41. Hao, Y. et al. Integrated analysis of multimodal single-cell data. *Cell* **184**, 3573–3587.e29 (2021).
 42. Litviňuková, M. et al. Cells of the adult human heart. *Nature* **588**, 466–472 (2020).
 43. Giles, C. L., Bollacker, K. D. & Lawrence, S. Citeseer: an automatic citation indexing system. In *Proceedings of the Third ACM Conference on Digital Libraries*, 89–98 (1998).
 44. McCallum, A. K., Nigam, K., Rennie, J. & Seymore, K. Automating the construction of internet portals with machine learning. *Inf. Retr.* **3**, 127–163 (2000).
 45. Wang, X. et al. Heterogeneous graph attention network. In *The world Wide Web Conference, 2022–2032* (2019).
 46. Shchur, O., Mumme, M., Bojchevski, A. & Günnemann, S. Pitfalls of graph neural network evaluation. *Relational Representation Learning Workshop, NeurIPS 2018* (2018).
 47. Craven, M. et al. Learning to extract symbolic knowledge from the world wide web. In *Proceedings of the Fifteenth National/Tenth Conference on Artificial Intelligence/Innovative Applications of Artificial Intelligence, AAAI '98/IAAI '98*, 509–516 (American Association for Artificial Intelligence, 1998).
 48. Rozemberczki, B. & Sarkar, R. Characteristic functions on graphs: birds of a feather, from statistical descriptors to parametric models. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, 1325–1334 (ACM, 2020).
 49. Rozemberczki, B., Allen, C. & Sarkar, R. Multi-scale attributed node embedding. *J. Complex Netw.* **9**, cnab014 (2021).
- ationEU/PRTR”; Grant PID2022-137505NB-C22 funded by MCIU/AEI/10.13039/501100011033 and by ERDF/EU; Generalitat de Catalunya grant number 2021SGR00856. R.J. acknowledges support from the fellowship FISSDUR funded by Generalitat de Catalunya. P.H. acknowledges support from the NSF AccelNet-MultiNet program under Grant No. 1927425 at North-eastern University. M.B. acknowledges the ICREA Academia award, funded by the Generalitat de Catalunya.

Author contributions

R.J. implemented the methods. M.A.S. supervised the work and wrote the first draft of the manuscript. R.J., P.H., M.B., and M.A.S. contributed to the research design, the analysis of the results, and the writing of the paper.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at

<https://doi.org/10.1038/s44260-024-00013-z>.

Correspondence and requests for materials should be addressed to M. Ángeles Serrano.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024

Acknowledgements

We acknowledge support from: Grant TED2021-129791B-I00 funded by MCIU/AEI/10.13039/501100011033 and the “European Union NextGener-