Journal of Computational Social Science
https://doi.org/10.1007/s42001-023-00207-w

RESEARCH ARTICLE

Grounding force-directed network layouts with latent space models

Felix Gaisbauer[1,2] · Armin Pournaki[2,3,4] · Sven Banisch[5] · Eckehard Olbrich[2]

Received: 28 November 2022 / Accepted: 16 April 2023

Abstract

Force-directed layout algorithms are ubiquitously used tools for network visualization. However, existing algorithms either lack clear interpretation, or they are based on techniques of dimensionality reduction which simply seek to preserve network-immanent topological features, such as geodesic distance. We propose an alternative layout algorithm. The forces of the algorithm are derived from latent space models, which assume that the probability of nodes forming a tie depends on their distance in an unobserved latent space. As opposed to previous approaches, this grounds the algorithm in a plausible interaction mechanism. The forces infer positions which maximise the likelihood of the given network under the latent space model. We implement these forces for unweighted, multi-tie, and weighted networks. We then showcase the algorithm by applying it to Facebook friendship, and Twitter follower and retweet networks; we also explore the possibility of visualizing data traditionally not seen as network data, such as survey data. Comparison to existing layout algorithms reveals that node groups are placed in similar configurations, while said algorithms show a stronger intra-cluster separation of nodes, as well as a tendency to separate clusters more strongly in multi-tie networks, such as Twitter retweet networks.

Keywords Network embeddings · Force-directed layout algorithms · Latent space models · Network visualization · Social networks

All authors have received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 732942 (www.Odycceus.eu).

✉ Felix Gaisbauer
felix.gaisbauer@weizenbaum-institut.de

[1] Weizenbaum Institute for the Networked Society, Berlin, Germany

[2] Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

[3] Laboratoire Lattice, CNRS & ENS-PSL & Université Sorbonne nouvelle, Paris, France

[4] Sciences Po, médialab, Paris, France

[5] Karlsruhe Institute for Technology, Karlsruhe, Germany

Published online: 29 May 2023

Introduction

This contribution aims to bring together two strands of research: latent space approaches to network analysis and force-directed layout algorithms (FDLs). FDLs are used ubiquitously for network exploration, illustration, and analysis in a wide variety of disciplines [2, 16, 17, 19, 24, 43, 51, 56, 61, 62]. While often producing intuitively compelling and aesthetic layouts, FDLs remain a double-edged sword. Two different families of approaches are subsumed under the term—both suffer from certain shortcomings. For FDLs that build on an electric-spring metaphor (e.g., Ref. [35]), it is unclear how to precisely interpret node positions and corresponding patterns, such as clusters (discussed recently in Refs. [34, 62]). FDLs that are instances of stress models, on the other hand, are simply directed towards the preservation of certain topological features of the network, such as geodesic distance between node pairs (see, e.g., Ref. [36]). They also tend to perform poorly on scale-free networks, a network type often found in social networks [14]. Besides, appropriate algorithm choice from the variety of FDLs available remains arbitrary in both cases.

We argue and show that explicit interpretability can be provided by latent space approaches, which have the goal of embedding a network in an underlying latent space, and where link probabilities are related to proximity in said space.

This contribution combines different approaches to embed networks in a meaningful space. It also makes the differences between several force-directed algorithms that embed networks clearer, especially for scientists who do not have an overview of the vast literature in graph drawing with FDLs. To this end, we first briefly sketch different types of FDLs in Sect. 2. We distinguish between FDLs based on electric-spring models, and stress models which are based on multidimensional scaling (MDS).[1] We sketch the approaches and show differences in terms of what their layouts mediate in low-dimensional spaces. Specifically, the former have a tight connection to modularity, while the latter usually aim at arriving at uniform edge lengths rather than conveying community structure. In Sect. 3, we introduce latent space approaches to network analysis, and subsequently show how force terms of a new type of FDL can be derived from said latent space models, where the forces move nodes towards positions and parameters which maximise the likelihood for the network under the given model. Especially for social networks, latent space models can serve as a plausible probabilistic models of social behavior and, in principle, can also be used for link prediction. We derive force equations for three types of networks (Sect. 4): unweighted networks, multi-tie networks (such as the much-studied Twitter retweet networks), and weighted networks. We present an implementation of the FDL, which we call *Leipzig Layout*, as well as a number of real-world networks spatialised with it in Sect. 5; we also show that the existing algorithms, specifically ForceAtlas2 [35], Fruchterman Reingold [23], and Yifan Hu [32], differ from the presented FDL. Section 6 concludes with a summary and an outlook.

[1] This article will not cover MDS in general, but only insofar as the layout algorithms are connected to it.

Force-directed layout algorithms

In the graph drawing literature, force-directed network layouts encase two different families of approaches. On the one hand, they include layout techniques that mimic a physical system in which forces that depend on the distance between nodes act on them until the system is at rest. This is what we will refer to as electric-spring models. On the other hand, layout algorithms have been developed that are types of MDS. While not actually simulating physical forces, they are referred to as 'force-directed' algorithms in the graph drawing community [38].[2] MDS is an approach to dimensionality reduction: information contained in a (high-dimensional) distance matrix is sought to be embedded in a low-dimensional space while keeping distances as truthful to the high-dimensional ones as possible. These approaches minimise an underlying objective function called stress function, which is why they are also referred to as stress models.[3]

Electric-spring models

Initially, force-based network visualization algorithms, both based on electric-spring and stress models, had been conceived to facilitate graph reading. They were supposed to make small networks readable in the sense that paths and nodes in the network were clearly accessible, that the edges had similar lengths and that the network was drawn as symmetric as possible [21].[4] These readability criteria were also referred to as 'aesthetic' (see, e.g., Refs. [8, 12, 20, 32, 58]). Progress in network science and the sudden availability of very large network data sets at the end of the millennium—for which a comprehension of individual node positions and paths was illusory—shifted focus: now, networks needed to be drawn such that community structure and topological features were mediated in the layout. Electric-spring models (partly having been developed already before this complex turn, notably in Refs. [21, 23]) turned out to be useful and efficient tools for this task. The algorithms have in common that all nodes repel each other, while connected nodes are additionally drawn together by their edges. The repulsive force F_r, $F_r \propto d^r$ (d is the distance between nodes), has a smaller exponent than the attractive force F_a, $F_a \propto d^a$, such that $a > r$.

Noack, in a seminal work [47], connected electric-spring models to modularity, one of the most central measures of clustering in networks in use today. Roughly

[2] Hence, one could also call electric-spring models FDLs 'proper', but we will stick to the established terminology for the rest of this article.

[3] It is worth noting already here that this distinction is permeable to a certain extent. Electric-spring models also minimise an objective function, namely the potential energy of the system. However, the forces for different algorithms have usually been chosen, because they produce aesthetically pleasing layouts, and not because the implicit energy function has certain desirable properties—such as preserving high-dimensional distances well. On the other hand, force equations governing a system can be derived from a stress function. However, typically, to find the stress function minimum, techniques such as stress majorization [25] are used.

[4] This paragraph largely follows [34] in its account of the history of (force-directed) network layout algorithms.

speaking, modularity Q compares the proportion of links connecting nodes within a group of nodes with the proportion expected if the edges in the network were randomly rewired [45]. Community detection algorithms, such as the Louvain algorithm [9], aim to find partitions of a network that maximise this value. Noack showed that, under certain constraints, modularity can be transformed into an expression that equals the energy function of electric-spring layouts. Constraints for the equivalence are as follows:

(i) Nodes can only be placed either at the same position (then, they belong to the same cluster) or at distance 1 from each other (if they are not in the same cluster).
(ii) The algorithms operate in a space of (at least) $k - 1$ dimensions, where k is the number of modularity clusters (usually, FDLs embed networks in a two-dimensional space).
(iii) The exponents of attractive and repulsive force should be non-negative. (Obviously, if the repulsive force has a negative exponent, placement of nodes at the same position would be impossible.)

For electric-spring layouts, this means that if they fulfill (ii) and (iii), their energy-minimal states are *relaxations* of modularity maximization: They make community structure in networks visible without constraint (i) of having to sort nodes into different, fixed partitions with distance 0 or 1 from each other. They can assign continuous positions in space. Or, phrased the other way round: modularity is then a special case of their energy function.

However, Noack's finding is diluted by the fact that network visualizations with FDLs are commonly restricted to two (or at most three) dimensions; and moreover, most electric-spring models in use today employ a negative exponent for the repulsive force. Noack also gave qualitative observations of which algorithms, even if they do not *exactly* fulfill the constraints above, tend to produce results that resemble modularity clusterings. Exponents in the forces should be characterized by $a \geq 0$ and the closer to 0, the better, $r \leq 0$, and $a - r \approx 1$.[5]

The connection to modularity—which, notably, had not been intended in the design of the algorithms—helped give additional credibility to FDLs in general. They were not only used for illustrative purposes [2, 16, 17], but also to explore and analyse network data [19, 43, 51, 61, 62]. It is, however, unclear what information electric-spring models add to modularity clustering by placing the nodes in a continuous space. It has been stressed that while they have been widely used, a thorough assessment of what exactly is entailed by the produced layouts has not been provided yet [35].[6] Moreover, many different types of FDLs have been developed,

[5] ForceAtlas2 ($a = 1, r = -1$) is in that sense more closely related to modularity than FruchtermanReingold ($a = 2, r = -1$) or Yifan Hu (which uses similar forces to FruchtermanReingold, but with a multilevel algorithm).

[6] [34] departs into a somewhat different direction than the work presented here by proposing certain interventions which help interpret what is visible in electric-spring models that are already in use today, but both approaches try to tackle the shortcomings that are sketched above.

several of which at least approximately subsume modularity clustering.[7] But which one of them constitutes an appropriate choice for a certain data set at hand?[8]

Stress models

A separate branch of graph drawing is constituted by the so-called stress approaches, which are a type of MDS. They are commonly also subsumed under the term 'force-directed layout algorithms'. They relate a good graph drawing to good isometry with respect to higher dimensional node distances [38]. That is, uniform edge length is aspired. Stress approaches can generally be used for any data where distances between data points are given. In the special case of networks, they seek to embed a network in a low-dimensional space (for visualization, dimensions should obviously be at most three) and distances between nodes are typically given by the geodesic shortest path between two nodes [13, 36, 38]. This is achieved via a stress function which is supposed to be minimised. A typical stress function is of the form

$$\sum_{i<j} w_{ij}\Big(\big\|\mathbf{x}_i - \mathbf{x}_j\big\| - \mathrm{dist}(i,j)\Big)^2. \tag{1}$$

$\mathbf{x}_i$ and $\mathbf{x}_j$ refer to node positions in the low-dimensional space, $\mathrm{dist}(i,j)$ is the (graph-theoretical) distance between i and j, and w_{ij} can be given by a term such as $w_{ij} = \mathrm{dist}(i,j)^{-\alpha}$ and serves as a normalization constant [38]. The stress function was introduced by Kruskal and Seery [39], and popularized by the graph drawing algorithm of Kamada and Kawai [36]. Kamada and Kawai used the Newton–Raphson method to minimise the stress function iteratively. Other, more efficient approaches, such as stress majorization, have also been conceived [25]. Electric-spring models only differentiate between connected and disconnected nodes. Stress models encode different target distances between nodes based, e.g., on the shortest path between them. Hence, the former are of lower descriptive complexity [38].

As has been stated already, MDS approaches to network visualization are types of dimensionality reduction. They simply try to preserve graph-theoretic distances as well as possible. This is why the literature on FDLs appears to have split in two different directions: Electric-spring approaches mediate tightly connected clusters and are evaluated on that basis [34, 47, 62]. The graph drawing community has oriented towards MDS models, since they have an explicit objective function and there are powerful techniques allowing efficient low-dimensional embeddings [13, 14, 25, 38].

Scientists might seek to interpret their layout based on a reasonable interaction mechanism between nodes. For example, what is often sought by social scientists

[7] We note here that modularity clustering is not without significant weaknesses, such as its resolution limit [22] or strong degeneracies of high-scoring solutions [26].

[8] Certain quality measures to *compare* network layouts have been proposed, such as the normalized edge length [46] corresponding to the total geometric length of the edges of a network divided by the graph density and the total geometric distance between nodes. However, these do not give meaning to the produced layout beyond network-immanent topological features.

is a plausible account of human behavior which forms the basis of the embedding. Stress models do not provide this. They *do* optimize an explicit objective function which has a clear interpretation.[9] However, their objective function aims at preserving a specific topological feature of a network which has been translated into a distance matrix.[10] Moreover, stress models tend to yield "poor results on certain classes of graphs, which include small worlds and other graphs with many shortcuts or low diameter, and scale-free graphs with highly skewed degree distributions, large 1-shells, or other forms of structural imbalance" [14, p. 228]. Since, for example, social networks often have scale-free properties, this is rather problematic. All in all, it is clear what network embeddings with stress models mean, but this meaning does not point beyond network-immanent topological features.

Latent space models

FDLs are often implemented in easily accessible tools, such as Gephi [7].[11] Not all researchers using the tools possess the methodological training to assess the mechanics behind them. However, the problems sketched above give an additional explanation for the fact that the limits or benefits of a chosen FDL are usually not discussed and "tools such as Gephi [are often treated as] as black boxes" [15]. As we have seen above, while being heavily used, layouts produced with FDLs are either ambiguous (in the case of electric-spring models) or aim at enforcing uniform edge length.

We propose to base a new type of FDL grounded in latent space models, with which network layouts can be interpreted explicitly in terms of a probabilistic model of node interactions. The approach is general, but as we show, it is especially fruitful for social networks. The distances between nodes can then be given a rigorous interpretation, and forces can be chosen that are suitable to the network data one wants to analyse.

Latent space approaches to network analysis have been developed to infer latent positions from their interactions. They represent a class of models based on the assumption that the probability that two actors establish a relation depends on their positions in an unobserved social space [28, 31]. The social space can be constituted by a continuous space, such as an Euclidean space, or a discrete latent space, where each node is in one of several latent classes [42]. Models of this type have also been introduced under the name of spatially embedded random networks [6]. Under this umbrella term, the Waxman model [65] and random geometric graphs [18, 50] were recognized as specific examples. Recently, latent space models have been employed in the estimation of continuous one-dimensional ideological positions from social

[9] As opposed to electric-spring models, where interpretation is ambiguous and reasons for algorithm choice are only given ex post [34, 46].

[10] This is not always straightforward for social networks, especially multi-tie networks. What, for example, would be the distance between two nodes i and j if i had retweeted j once, but j had set off 10 tweets?

[11] Gephi incorporates mostly electric-spring models, but the stress model algorithm by Kamada and Kawai [36] is implemented, e.g., in the Python networkx [27] package.

media data [4, 5, 33], specifically from Twitter follower networks. These works covered large quantities of users and showed good agreement with, e.g., party registration records in the United States [5]. The estimation of positions in the latent space was achieved with correspondence analysis in [5], while in [4], the posterior density of the parameters was inferred via Markov Chain Monte Carlo methods. Latent space models have also been used to embed dynamic networks [53, 54], and have recently been extended to additionally include attribute information or multiple relationships in one shared embedding [64, 66].

This is where the present work intersects: We attempt to take an alternative route to arrive at a specific form of force equations for FDLs. We obtain the forces on the basis of latent space models. The positions of the nodes in an assumed latent space influence the probability of ties between them—the closer their positions, the more probable it is that they form a tie. We derive an FDL as a maximum-likelihood estimator of such a model. This approach clarifies the underlying assumptions of our layout algorithm and makes the resulting layout *interpretable*. We derive three different forces for three different types of networks, specifically adapted to the task of embedding them in a political space: unweighted, multi-tie, and weighted networks. Moreover, alternative interaction models can in principle be used to develop force-directed layouts in a completely analogous way. For this, the present work can serve as a blueprint.

While some might claim that the visualization of a network only serves illustrative purposes, their wide-spread use, not only for exploration and illustration, but also visual analysis of networks [17, 19, 43, 61, 62] underscores the necessity of this enterprise: Exploration and interpretation are, in practice, guided by force-directed layouts for many researchers from a variety of disciplines. And especially in the (computational) social sciences, where network data are abundant and exploration and analysis are tied closely to readily implemented algorithms, algorithms that make their assumptions explicit and are suitable to the data that is supposed to be analysed are needed.

From latent space models to force equations

In this section, we show how force terms in a force-directed layout algorithm can be derived from latent space models of node interactions. Central to this procedure is the assumption that nodes tend to form ties to others that are close to them in a latent social space. The closer two nodes, the higher the probability that one forms a tie to the other. Since none of the positions (as well as none of the additional parameters of the statistical model which will be introduced in the corresponding subsections) are directly observed, the statistical problem posed here is their inference. Given the latent space model, one can determine the likelihood function $L(G)$ for any observed network. The positions and parameters are then inferred via maximum-likelihood estimation. In our approach, this is done by treating the negative log-likelihood as a potential energy. The minima of this potential energy are the local maximisers of the likelihood. Its derivatives with respect to the positions and parameters of the nodes

can be considered as forces that move the nodes towards positions that maximise the likelihood.

We will cover three different types of directed networks: unweighted networks, such as the follower networks covered by Barberá and colleagues [4, 5], multi-tie networks (which include Twitter retweet networks), and weighted networks. Undirected networks are implicitly included as a special case where $a_{ij} = a_{ji}$ and each node only has one additional parameter α. We will present the derivation of the forces for the unweighted case in detail. The complete derivations for the other two cases are given in Appendices 7 and 8.

Unweighted networks

Consider an unweighted graph $G = (V, E)$ with nodes $i \in V$ and edges $(i, j) \in E$. The graph can be described by an adjacency matrix $A = \{a_{ij} | a_{ij} = 1 \text{ if } (i, j) \in E\}$. Now, let us assume that the nodes are represented by vectors $\mathbf{x}_i \in \mathbb{R}^n$ and $d_{ij} = d(\mathbf{x}_i, \mathbf{x}_j)$ denotes the Euclidean distance between $\mathbf{x}_i$ and $\mathbf{x}_j$.[12] For the probability of a tie between two nodes, we choose (see [31])

$$p(a_{ij} = 1) = \text{logit}^{-1}\left(\alpha_i + \beta_j - d_{ij}^2\right) = \frac{\exp\left(\alpha_i + \beta_j - d_{ij}^2\right)}{1 + \exp\left(\alpha_i + \beta_j - d_{ij}^2\right)}. \tag{2}$$

The probability is dependent on the squared Euclidean distance between the two node positions.[13] That the probability is dependent on the squared distance is also assumed in [4], while in [5, 31], the linear distance is used. α and β are additional parameters that also influence the probability of a tie. α_i can be interpreted as an activity parameter related to the out-degree of node i: The higher α_i, the higher the probability of a tie from i to others. β_j influences the probability of ties *to* j, i.e., the in-degree of node j. The parameters allow nodes that occupy the same position in space to have different degrees—as an example, there might be people with roughly the same political position as, say, a state leader, but it is generally unreasonable to expect that these users have the same amount of followers on social media. On the other hand, some users might simply be more active than others, hence forming more ties, while sharing a political position.

The model introduced here has already been well established in the literature [4, 5, 31]. However, alternative models can be posited that might be more fitting for certain network types. In that case, the derivation lined out below can be carried out analogously.

[12] The question of dimensionality is an open one both for latent space models and for FDLs. We are concerned with network *visualization* in this contribution, and hence restrict ourselves to two dimensions. Especially with respect to political embeddings, already two dimensions (left–right plus an additional one) which already explain much of the structure in these networks (see [48]). Nevertheless, it is highly desirable to investigate dimensionality effects in a systematic way both for latent space models and FDLs.

[13] This is a convenient choice, since it yields, as we will see shortly, linear forces.

For a given graph G, the likelihood function $L(G)$ can be written as the product of the probability of an edge *if* there exists an edge between two nodes, and the probability of there not being an edge if not

$$\begin{aligned} L(G) &= \prod_{(i,j)\in E} p\big(a_{ij}=1\big) \prod_{(i,j)\notin E} \big(1-p\big(a_{ij}=1\big)\big) \\ &= \frac{\prod_{(i,j)\in E} \exp\left(\alpha_i+\beta_j-d_{ij}^2\right)}{\prod\limits_{\substack{i,j \\ i\neq j}} \left(1+\exp\left(\alpha_i+\beta_j-d_{ij}^2\right)\right)}. \end{aligned} \tag{3}$$

The logarithm of the likelihood is given by

$$\begin{aligned} LL(G) :&= \log L(G) \\ &= \sum_{(i,j)\in E} \left(\alpha_i+\beta_j-d_{ij}^2\right) - \sum_{\substack{i,j \\ i\neq j}} \log\left(1+\exp\left(\alpha_i+\beta_j-d_{ij}^2\right)\right). \end{aligned} \tag{4}$$

If we consider the *negative* log-likelihood as a potential energy, its minimas are the local maximisers of the likelihood. Its (negative, once again) derivatives with respect to the positions $\mathbf{x}_i$ of the nodes can be considered as *forces* that move the nodes towards positions that maximise the likelihood. For a concrete node i', an attractive force is generated by node j' if i' establishes a tie to j'

$$F^{j'}_{\text{att},i'} = \frac{\partial}{\partial \mathbf{x}_{i'}}\left(\alpha_{i'}+\beta_{j'}-d_{i'j'}^2\right) = -2\big(\mathbf{x}_{i'}-\mathbf{x}_{j'}\big). \tag{5}$$

If j' also establishes a tie to i', the same attractive force is applied again. On the other hand, a rejecting force is always present for each possible tie[14]

$$F^{j'}_{\text{rej},i'} = -\frac{\partial}{\partial \mathbf{x}_{i'}} \log\left(1+\exp\left(\alpha_{i'}+\beta_{j'}-d_{i'j'}^2\right)\right) = \frac{1}{1+\exp\left(-\alpha_{i'}-\beta_{j'}+d_{i'j'}^2\right)} 2\big(\mathbf{x}_{i'}-\mathbf{x}_{j'}\big). \tag{6}$$

Another repulsive force on i' appears for this node pair for the potential tie from j' to i'.

The derivative of Eq. (4) with respect to $\alpha_{i'}$ and $\beta_{i'}$ gives us the forces on the parameters of node i', such that

$$F^{j'}_{\alpha_{i'}} = a_{i'j'} - \frac{1}{1+\exp\left(-\alpha_{i'}-\beta_{j'}+d_{i'j'}^2\right)} = a_{i'j'} - p\big(a_{i'j'}=1\big) \tag{7}$$

and

[14] Note the sign reversal in the exponent of the exponential function in the denominator in the last equivalence, which stems from $\exp : (\alpha_{i'}+\beta_{j'}-d_{i'j'}^2)/\left(1+\exp\left(\alpha_{i'}+\beta_{j'}-d_{i'j'}^2\right)\right) = 1/\left(1+\exp\left(-\alpha_{i'}-\beta_{j'}+d_{i'j'}^2\right)\right)$.

$$F^{j'}_{\beta_{i'}} = a_{j'i'} - \frac{1}{1+\exp\left(-\alpha_{j'} - \beta_{i'} + d^2_{i'j'}\right)} = a_{j'i'} - p\left(a_{j'i'} = 1\right). \tag{8}$$

The sum over all forces on the single parameters yields the difference between the actual and the expected in-/out-degree. In the equilibrium state, where this sum yields 0, the observed in-/out-degree of nodes equals the one expected under the model of Eq. (2)

$$\sum_{j'} F^{j'}_{\alpha_{i'}} = d^{\text{out}}_i - \left\langle d^{\text{out}}_i \right\rangle, \tag{9}$$

$$\sum_{j'} F^{j'}_{\beta_{i'}} = d^{\text{in}}_i - \left\langle d^{\text{in}}_i \right\rangle. \tag{10}$$

Multi-tie networks

Force equations can also be derived for networks which are constituted by a number of binary signals between nodes—for example, when users of an online platform create several posts, each of which can be taken up by others (e.g., through liking or sharing the respective post). A much-studied case are Twitter retweet networks, which are frequently employed to investigate opinion groups on the platform [16, 17].

We consider, for each action k initiated by a node j (e.g., a tweet), an unweighted graph $G^k_j = (V, E^k_j)$ with nodes $i \in V$ and edges $(i,j) \in E^k_j$, where an edge means that i has formed a tie to j upon action k. The graph for each k can be described by an adjacency matrix $A^k_j = \{a_{ij} | a_{ij} = 1 \text{ if } (i,j) \in E^k_j\}$. This constitutes an m-star graph with $m = |E^k_j|$.[15]

Analogously to the unweighted case, we assume the probability of establishing *a single tie upon action k from user i to user j* with

$$p(a^k_{ij} = 1) = \frac{1}{1+\exp\left(-\alpha_i - \beta_{jk} + d^2_{ij}\right)}, \tag{11}$$

where each action k of j has its own parameter β_{jk} which affects the in-degree of j.

The log-likelihood for the multi-tie network can be written as

[15] In general, multi-tie networks are also constituted by networks with an arbitrary number of edges between nodes. In our model here, we are concerned with networks with a maximum number of edges from node i to node j which is determined by the number of actions m_j of j that i can form a tie on.

$$LL(G) = \sum_{j} \sum_{k=1}^{m_j} \left(\sum_{(i,j) \in E_{jk}} \left(\alpha_i + \beta_{jk} - d_{ij}^2 \right) - \sum_{\substack{i \\ i \neq j}} \log\left(1 + e^{-\alpha_i - \beta_{jk} + d_{ij}^2} \right) \right). \quad (12)$$

Here, in addition to user pairs, we sum over all actions k. The derivation of forces for this case is largely analogous to the unweighted case and can be found in Appendix 7, along with the concrete force equations.

Weighted networks

So far, we have assumed a binary signal between node pairs—e.g., whether an individual follows another one or not, or whether someone shares certain content of another individual or not. Extending the model to the non-binary case can be achieved by exchanging the ordinary logit model [Eq. (2)] with an ordered logit or proportional odds model.

There, a response variable has levels 0, 1, ..., n (e.g., people rate their relationships to others on a scale from 0 to 6, or similar). We consider the general case of weighted networks with finite weights that can be transformed into natural numbers (including 0), i.e., an adjacency matrix $A = \{a_{ij} = k | k \in \mathbb{N}_0\}$. The probability of the variable being greater than or equal to a certain level k is given by [29, 37]

$$p(a_{ij} \geq k) = \frac{1}{1 + \exp\left(-c_k - \alpha_i - \beta_j + d_{ij}^2 \right)}, \quad (13)$$

where $k = 0, 1, ..., n$. ($c_0 = \infty$, $c_{n+1} = -\infty$.) The probability of a_{ij} equal to a certain k is given by

$$p\left(a_{ij} = k\right) = P\left(a_{ij} \geq k\right) - P\left(a_{ij} \geq k + 1\right). \quad (14)$$

The likelihood $L(G)$ is given by

$$L(G) = \prod_{i,j} \left(\frac{1}{1 + \exp\left(-c_{a_{ij}} - \alpha_i - \beta_j + d_{ij}^2 \right)} - \frac{1}{1 + \exp\left(-c_{a_{ij}+1} - \alpha_i - \beta_j + d_{ij}^2 \right)} \right), \quad (15)$$

and the log-likelihood by

$$LL(G) = \sum_{\substack{i,j \\ i \neq j}} \log\left(\frac{1}{1 + \exp\left(-c_{a_{ij}} - \alpha_i - \beta_j + d_{ij}^2 \right)} - \frac{1}{1 + \exp\left(-c_{a_{ij}+1} - \alpha_i - \beta_j + d_{ij}^2 \right)} \right). \quad (16)$$

The force equations derived from the log-likelihood (for the case of a three-point scale) can be found in Appendix 8. Potential applications of a visualization with these forces are manifold. In smaller data sets, a non-binary signal between nodes, e.g., rating of the relationships between individuals of a social group, might be given between all node pairs. But often, there might be cases for which a subset of individuals (say, politicians, public figures, etc.) or items (e.g., the importance of political goals, technologies, etc.) are rated by others, such that they form a bi-partite network. Then, only the rating individuals receive an α, while the rated just have a β-parameter. The interpretation of the parameters α and β might need adjustment: They now rather refer to the tendency of individuals to give/receive rather high/low ratings.

Real-world networks

The force-directed layout algorithm was implemented in JavaScript, building upon the d3-force library [10]. There, force equations are simulated using a velocity Verlet integrator [57, 63]. A ready-to-use implementation, which we call *Leipzig Layout*, is available under https://github.com/pournaki/leipzig-layout.[16] It builds upon the force-graph library [3] to interactively display the graph and the evolution of node positions in the simulation of forces. Note that at the current state, this layout tool works reasonably fast for networks below 10,000 links. We include a performance benchmark for convergence speed in Appendix 9. Validation is included in Appendix 10.

We use Leipzig Layout to spatialise several real-world networks: Undirected Facebook friendship networks, the directed Twitter follower network of the German parliament, the Twitter retweet network of a debate surrounding the publication of a letter on free speech by Harper's magazine, and a survey on different types of energy-generating technologies.

Facebook100: Haverford & Caltech

The Facebook100 data set consists of online social networks collected from the Facebook social media platform when the platform was only open to 100 universities in the US [59]. The data set contains social networks of students of particular universities with quite rich metadata (e.g., gender, year, residence, or major). We analyse friendship networks—undirected networks where a tie between users represents that both have agreed to connect with each other as 'friends' on the platform. We spatialise the friendship network of Haverford University in Fig. 1 A, B (links have been omitted for better accessibility). On the left, it is visible that students are spatially layered according to their year by the layout algorithm. The first-year students are visually separated from the others. The layout becomes denser for students who have been at the university for a longer time. The layers are ordered

[16] For the moment, this implementation is restricted to unweighted graphs. An extension for weighted and multi-tie graphs will be published on the same repository.

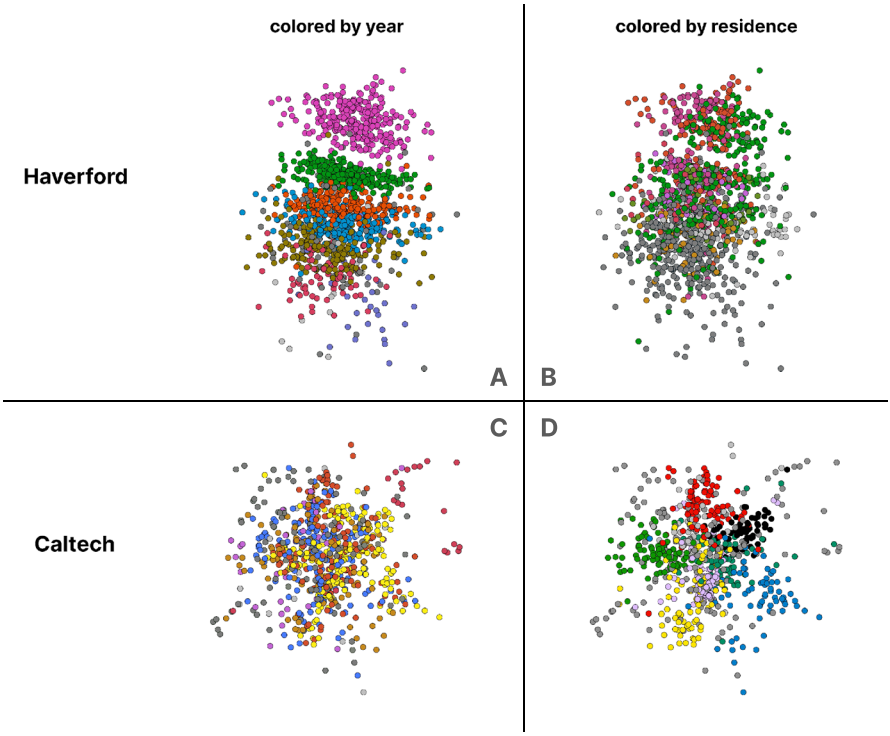

Fig. 1 Friendship network of students of Haverford University (top) and California Institute of Technology (bottom), colored by year (left), and residence (right). The spatialisation of the former layers students by year (**A**, chronologically ordered from top to bottom, with first-year students colored pink, second-year students colored green, etc.; dark grey nodes correspond to students whose year is unknown). First-year students are visually separated from the others, while the layout becomes denser if students have been at university for a longer time. In (**B**), it is also visible that first-year students show a higher tendency to mix with others they share residency with (dark grey: dorm unknown). For Caltech, the network out of the Facebook100 data set with the highest assortativity with respect to residence, nodes are visibly placed according to dorm membership (**D**, dark grey: year/dorm unknown), and less so with respect to year (**C**)

chronologically. It seems that if students form cross-year ties, they tend to connect to others from adjacent cohorts. The local assortativity distribution with respect to residence of the students of Haverford has been analysed in detail in [49]. There, it was found that first-year students tend to form ties to other students from their dorms, while students from higher years show less of a tendency to mix only with others they share residence with. This behavioral pattern can also be discerned in the spatialisation (Fig. 1B): For the first-year students, the students sharing a dorm tend to be placed rather close to each other, while for students from higher years, this is not the case. As a complement, we spatialise the network for the university with the highest overall assortativity with respect to dormitory in the data set: Caltech. There, the students' years do not influence Facebook friendship to a large extent; rather, students' friendships are more strongly guided by their residences [52]. This

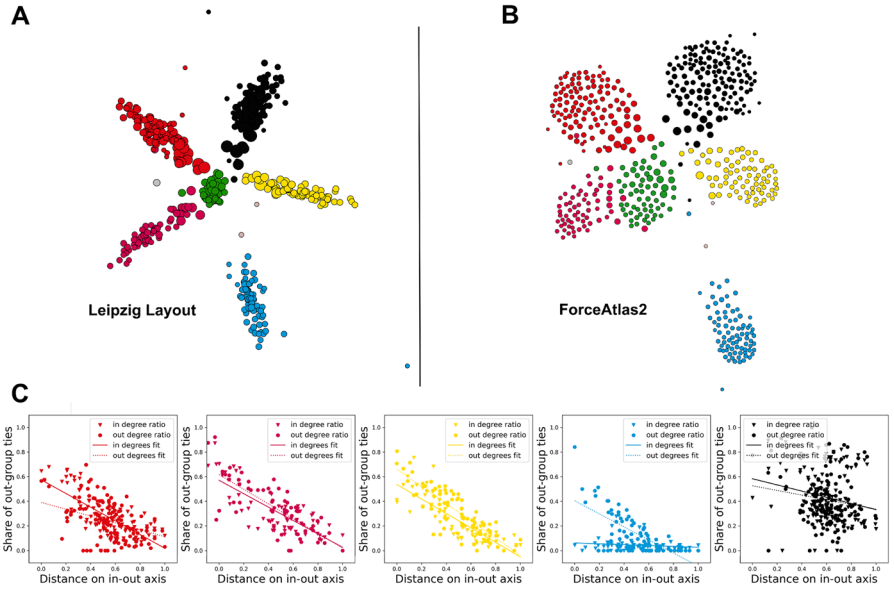

Fig. 2 Leipzig layout of the follower network of all German deputies that have a Twitter account (**A**). Members are colored according to their party and node size corresponds to overall node degree. Clear division between parties, as well as a stronger division between the right-wing party *AfD* and the other parties is visible. All parties except the *Greens* are arranged on a one-dimensional axis. This is explained by a difference in cross-party ties between politicians of the same party: The further out a member on the party-internal axis, the fewer cross-party ties to and from them have been established (except for the *AfD*, which does not receive many ties from other parties no matter where the users are placed) (C, colored according to parties, linear fits included). ForceAtlas2, in comparison, has a stronger separation of nodes within party clusters due to its rejecting force being proportional to d^{-1} (**B**)

is reproduced by Leipzig Layout: Fig. 1 D shows that students that share residence are visibly placed close to each other. On the other hand, students are less strongly grouped according to their university year (C).

German parliament: Twitter follower network

While [4, 5] aim for the estimation of one-dimensional ideological positions of politicians (and their followers), the FDL proposed here embeds nodes in a two-dimensional space. We spatialise the Twitter follower network of all then-members of the German parliament that had an active Twitter account in Fig. 2 as of July 2021. The parties (members colored according to their typical party color) are quite visibly separated. They are located along a circle that quite accurately mirrors the political constellation in federal German politics. The center-left to center-right parties (*SPD*, *Bündnis 90/Die Grünen* (Green party), *CDU/CSU*) are positioned between *Die Linke* (left party) and the market-liberal *FDP*. The *AfD* (blue), a right-wing populist party with which collaboration has been ruled out by all other parties, accordingly occupies a secluded area. Interestingly, within

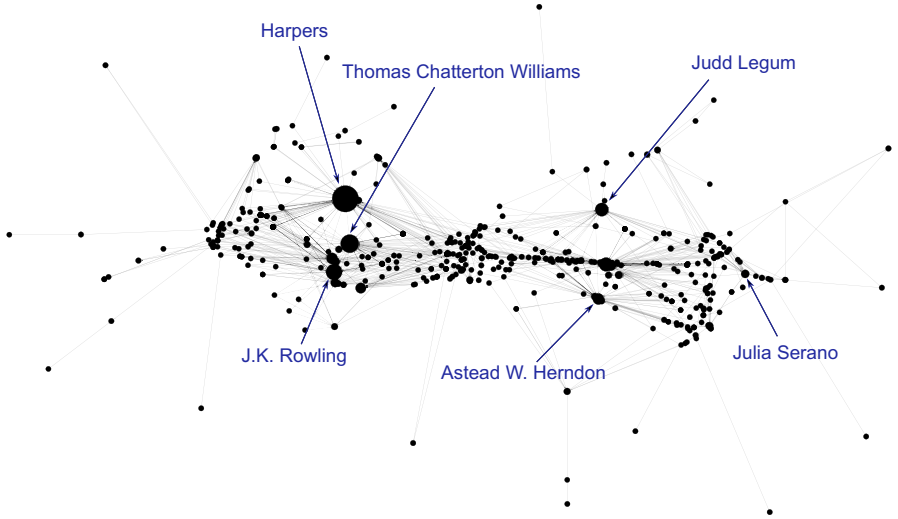

Fig. 3 Retweet network of Twitter debate about a letter on free speech published by Harper's magazine (node size proportional to in-degree). A two-camp division is visible, where the left pole includes the magazine as well as prominent signees, while the right pole contains critics

parties, a one-dimensional arrangement is visible (except for the *Greens*). This mirrors the amount of cross-party ties, as well as the users' activity on Twitter: The further out on an axis between the innermost and outermost party member users are placed (outliers excluded), the smaller their share of ties to other parties (Fig. 2C). The densely packed placement of the Greens can be explained by the fact that they tend to use Twitter quite homogeneously (see Appendix 12) in the sense that there are no users which are inactive or lack followers: Each of them (except one) has an in-/out-degree of at least 50. Moreover, the party members are followed by and follow all other parties (except for the *AfD*) in a well-balanced fashion. α and β are generally correlated with out and in-degree of the nodes (again, see Appendix 12). They exhibit a very pronounced linear correlation for the Greens.

This layout also illustrates the difference between ForceAtlas2 (see (B) in Fig. 2) and the layout algorithm at hand here: ForceAtlas2 incorporates a rejecting force between node pairs proportional to d_{ij}^{-1}, which leads to a stronger separation of nodes within clusters. Hence, while the overall arrangement of parties is similar to Leipzig Layout, parties themselves are more strongly spaced out. A comparison to spatialisations with the algorithms Yifan Hu and FruchtermanReingold can also be found in Appendix 12.

Moreover, we observe that several minima are inferred by both Leipzig Layout and the other FDLs—depending on the initial positions of the nodes. While the existence of different local minima is a general problem of FDLs, one can simply select the outcome with the highest likelihood with the present approach. A different, but less likely local minimum inferred with Leipzig Layout is also presented in Appendix 12. The more likely minimum, which is displayed in Fig. 2, is also the

politically more plausible one: In Appendix 12, *SPD* is placed closer to *FDP* than *CDU/CSU*, while the latter two parties have more commonalities (especially when it comes to economic policy).

Retweet network: Harper's letter

In July 2020, Harper's magazine published an open letter signed by 153 public figures defending free speech which they saw endangered by 'forces of illiberalism.' Not only Donald Trump was denounced as contributing to illiberalism, but also some groups who advancing "racial and political justice," who had "intensified a new set of moral attitudes and political commitments that tend to weaken our norms of open debate and toleration of differences in favor of ideological conformity" [1]. On Twitter, the letter was controversially discussed (see also https://blog.twitterexplorer.org/post/harpers_letter/). The layout of the retweet network reproduces a division between critics and supporters of the letter: On the left side of Fig. 3, the account of Harper's magazine as well as prominent signees such as Thomas Chatterton Williams and Joanne K. Rowling are visible, while the right pole includes critics of the letter and its signees, such as Judd Legum, Astead W. Herndon, and Julia Serano. Serano, a transgender activist, criticized that what the signees referred to as 'free speech' has prevented marginalized groups from speaking out, and accused Rowling of having spread disinformation about trans children. That she was voicing rather specific criticism which aimed towards certain signatories of the letter is mirrored in her position close to the margin of the inferred space. Legum and Herndon are placed closer to the center: Legum noted in a relatively nuanced critique that the signees of the letter are not silenced in any way, while Herndon published several ironical tweets about the letter. Interestingly, the division of clusters visible in the layout is not as pronounced as in the spatialisation of the network with ForceAtlas2 and Yifan Hu (see Fig. 11 in Appendix 12), a finding that calls for further systematic investigation.

Survey data

With the weighted layout, not only generic network data can be spatialised, but also surveys: There, evaluated items as well as respondents are nodes, and forces only exist between items and individuals.

In Fig. 4, we visualize a survey where respondents were asked about their attitude towards six different energy-generating technologies [55]. The responses represent the initial attitudes of respondents with respect to the technologies before being confronted with several pro and counter arguments. Responses were initially given on a nine-point scale, which was aggregated to a three-point scale for visualization. The presented technologies were gas and coal power stations, onshore and offshore wind stations, biomass power stations, and open-space photovoltaics (which we refer to as solar in Fig. 4).

The distribution of respondents over the inferred space, given by a density plot (the lighter the color, the more respondents lie in a region of the layout), shows that

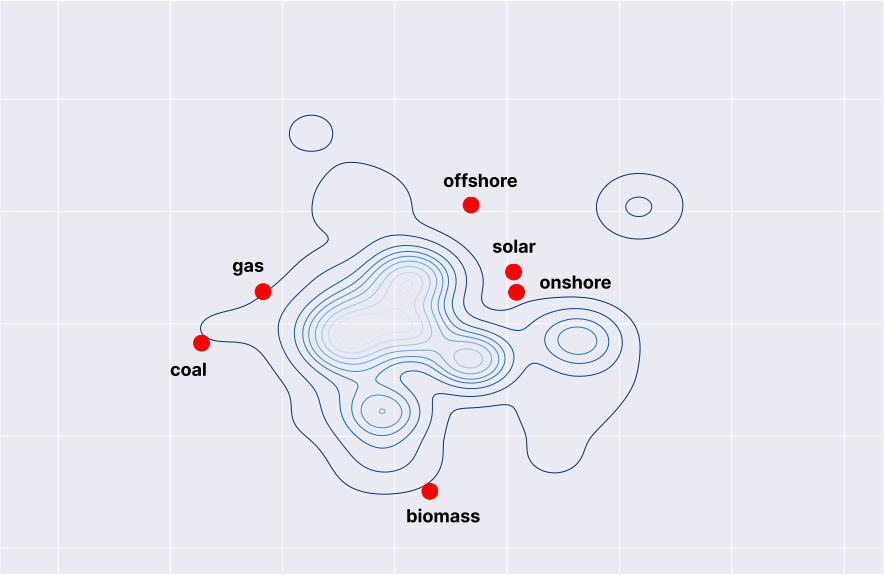

Fig. 4 Visualization of a survey on six different energy-generating technologies. The distribution of respondents is plotted as a density in the background (the lighter, the denser they are distributed in an area). Respondents are distributed close to gas, renewable energy-generating technologies, and between them. Two technological axes are visible: One from coal and gas to the renewables, and one among technologies using renewable sources of energy, with onshore and solar occupying central positions, while offshore and biomass are located opposite of each other

the vast majority of respondents are located between gas and onshore, solar and offshore energy-generating technologies, while coal is placed far away from most respondents.

Several density peaks exist which correspond to respondents with similar response profiles: one between gas and biomass, one placed rather centrally between gas, offshore, solar, onshore and biomass, two between biomass and solar/onshore, and one at the margin of the space, but closest to offshore and solar/onshore technologies. Even more interesting is the arrangement of technologies themselves, since it shows that collectively, response profiles of individuals create two orthogonal axes along which technologies are placed: One axis is visible from renewables towards technologies relying on fossil sources of energy (gas and coal). On the other hand, renewable sources of energy are distributed along an perpendicular axis. Onshore and solar occupy central positions there, while offshore and biomass are located opposite of each other. The respondents' distribution and the arrangement of technologies are in line with the average ratings and rating correlations between individuals (see Appendix 13). Average ratings for coal are reported to be significantly lower than for any other technology in [55], gas receives a neutral rating, and renewable technologies (offshore, onshore, solar, and biomass) are rated positively on average. Biomass receives the lowest average rating of the renewables, which is reflected in the distribution of respondents. Biomass has, among the renewables, the

weakest correlations with the other renewables. On the other hand, ratings are not negatively correlated with coal or gas. This is mirrored by its placement in Fig. 4.

Outlook

FDLs are frequently employed for network visualization across a variety of scientific disciplines. However, the algorithms are either based on forces whose explicit interpretation is still unclear or on techniques dimensionality reduction which seek to preserve network-immanent topological features.

We have presented a path towards interpretable FDLs based on latent space models. We have derived force equations for Leipzig Layout, a FDL that serves as a maximum-likelihood estimator of said models. We have posited three variants of the FDL, which are applicable to unweighted, multi-tie, and weighted networks, respectively. Exemplary spatialisations of several real-world networks show that important properties of the networks (assessed through different network measures) are reflected by node placement. Moreover, commonalities with, but also differences to the existing FDLs have been pointed out: The latter tend to exhibit a stronger separation of nodes within tightly connected clusters, and, for retweet networks, also between each other.

The new type of FDL presented here makes the assumptions it is based on—the underlying latent space model—explicit, and hence constitutes an attempt to put FDLs on a plausible model of node interactions. Latent space models are well established in the estimation of ideological positions on the basis of (social) networks. In most cases, the ideology estimates have been carried out on a one-dimensional axis. Leipzig Layout infers a two-dimensional latent space.

The model chosen here can, if found necessary for certain network types, be replaced by alternative interaction models. For this purpose, the derivation of forces above can serve as a blueprint. The present approach might also be used to motivate parameter choices for the existing FDLs, such as ForceAtlas2: The degree of influence of edge weights there, for example, can be chosen freely. However, the choice could be guided by agreement with the weighted case of the algorithm implemented here.

The spatialisation of survey data presented above points beyond traditional usage of FDLs, an avenue which should be explored further. Moreover, embeddings produced on latent space models can be used for link prediction, for example, if data are incomplete or if future behavior is supposed to be prognosticated.

Recently, approaches have been developed which assist or replace the force-directed algorithm with neural networks, especially graph neural networks (GNNs) [11, 40]. This might potentially increase convergence speed considerably and might form a possible route to make a layout with the energy (i.e., loss) function above more efficient. We note here that recent work in machine learning also used force-directed layouts that were derived as gradients of an objective function: stochastic neighbor embedding (SNE) [30] and t-SNE [44, 60] used this technique to embed a graph in a lower dimensional space. In these cases, the

objective function was not the likelihood of a statistical model but the KL-divergence between two probability distributions.

Limitations remain: The convergence of FDLs to local minima is a problem that cannot be solved by the present approach. The follower network presented in Fig. 2, for instance, possesses several equilibria for which the parties were allocated in different order, both for Leipzig Layout as well as the three algorithms it was compared with. Nevertheless, the underlying model of Leipzig layout allows a comparison of the log-likelihood of several equilibria, out of which the most likely can then be chosen. Minor design options which serve readability could be included in the algorithm in the future. For example, allowing nodes to overlap allows a precise representation of the inferred space. However, if individual nodes are supposed to be discerned, an option that prevents node overlap could be introduced.

Moreover, the role of dimensionality for the outcomes of latent space inference in general has not been studied systematically [42]. For network visualization, an extension to a three-dimensional latent space would be of interest in comparison to the two-dimensional case studied here.

Appendix A: force derivation—multi-tie networks

For multi-tie networks, we stipulate the probability of establishing *a single tie upon action k from user i to user j* by

$$p\left(a_{ij}^{k}=1\right)=\frac{1}{1+\exp\left(-\alpha_i-\beta_{jk}+d_{ij}^{2}\right)}. \tag{A1}$$

The derivation of the forces for this case is analogous to the unweighted case, except that we sum over tweets instead of user pairs.

The log-likelihood for a given network can be written as

$$LL(G)=\sum_{j}\sum_{k=1}^{m_j}\left(\sum_{(i,j)\in E_{jk}}\left(\alpha_i+\beta_{jk}-d_{ij}^{2}\right)-\sum_{\substack{i\\ i\neq j}}\log\left(1+e^{-\alpha_i-\beta_{jk}+d_{ij}^{2}}\right)\right). \tag{A2}$$

The force on the position of node i' exerted by node j' is given by

$$\begin{aligned}\partial_{x_{i'}}^{j'}LL(G)=2(\mathbf{x}_{i'}-\mathbf{x}_j)\Bigg[&\sum_{k=1}^{m_{j'}}\left(-a_{i'j'}+\frac{1}{1+\exp(-\alpha_{i'}-\beta_{j'k}+d_{i'j'}^{2})}\right)+\\ &\sum_{k=1}^{m_{i'}}\left(-a_{j'i'}+\frac{1}{1+\exp(-\alpha_{j'}-\beta_{i'k}+d_{i'j'}^{2})}\right)\Bigg].\end{aligned} \tag{A3}$$

The force on $\alpha_{i'}$ exerted by j' is given by

$$\partial^{j'}_{\alpha_{i'}} LL(G) = \sum_{k=1}^{m_{j'}} \left(a_{i'j'} - \frac{1}{1 + \exp\left(-\alpha_{i'} - \beta_{j'k} + d^2_{i'j'}\right)} \right); \tag{A4}$$

the force on each $\beta_{i'k}$ due to j' by

$$\partial^{j'}_{\beta_{i'k}} LL(G) = a_{j'i'} - \frac{1}{1 + \exp\left(-\alpha_{j'} - \beta_{i'k} + d^2_{i'j'}\right)}. \tag{A5}$$

Appendix B: Force derivation—weighted networks

A possibility of extending the model to weighted graphs, i.e., an adjacency matrix $A = \{a_{ij} = k | k \in \mathbb{N}_0\}$, is the ordered logit or proportional odds model. There, the response variable has levels 0, 1, ..., n (e.g., people rate their relationships to others on a scale from 0 to 6). The probability of the variable being greater than or equal to a certain level k is given by

$$p(a_{ij} \geq k) = \frac{1}{1 + \exp\left(-c_k - \alpha_i - \beta_j + d^2_{ij}\right)} \tag{B1}$$

($c_0 = \infty, c_{n+1} = -\infty$). The probability of a_{ij} equal to a certain k is given by

$$p(a_{ij} = k) = p\left(a_{ij} \geq k\right) - p\left(a_{ij} \geq k + 1\right). \tag{B2}$$

The log-likelihood of the graph is then given by

$$LL(G) = \sum_{\substack{i,j \\ i \neq j}} \log\left(\frac{1}{1 + \exp\left(-c_{a_{ij}} - \alpha_i - \beta_j + d^2_{ij}\right)} - \frac{1}{1 + \exp\left(-c_{a_{ij}+1} - \alpha_i - \beta_j + d^2_{ij}\right)} \right). \tag{B3}$$

As the simplest example, we turn to the case with three levels, where the probabilities are given by[17]

$$p(a_{ij} = 0) = 1 - \frac{1}{1 + \exp\left(-c_1 - \alpha_i - \beta_j + d^2_{ij}\right)} \tag{B4}$$

[17] In practice, one might often encounter data sets where individuals have rated a limited set of others, e.g., politicians, which themselves might not participate in the rating. This yields a bi-partite graph for which β only applies to the rat*ed* individuals, while α-values are present only for the rat*ing* individuals.

$$p(a_{ij}=1)=\frac{1}{1+\exp\left(-c_1-\alpha_i-\beta_j+d_{ij}^2\right)}-\frac{1}{1+\exp\left(-c_2-\alpha_i-\beta_j+d_{ij}^2\right)} \tag{B5}$$

$$p\left(a_{ij}=2\right)=\frac{1}{1+\exp\left(-c_2-\alpha_i-\beta_j+d_{ij}^2\right)}. \tag{B6}$$

The derivative of Eq. (B3) with respect to $\mathbf{x}_i$, α_i and β_i gives us the forces on the position and the parameters of node i. Additionally, the derivative with respect to c_1 and c_2 estimates the cut points. We introduce the following abbreviations:

$$C_1=-c_1-\alpha_{i'}-\beta_{j'}+d_{i'j'}^2$$

and

$$C_2=-c_2-\alpha_{i'}-\beta_{j'}+d_{i'j'}^2.$$

The first part of the force on the position of node i' exerted by node j' is given by

$$\partial_{x_{i'}}\log\left(p(a_{i'j'}=0)\right)=2(\mathbf{x}_{i'}-\mathbf{x}_{j'})\frac{e^{C_1}}{\left(1+e^{C_1}\right)^2-\left(1+e^{C_1}\right)}=2\left(\mathbf{x}_{i'}-\mathbf{x}_{j'}\right)\frac{1}{1+e^{C_1}} \tag{B7}$$

if $a_{i'j'}=0$. If $a_{i'j'}=1$, the force is given by

$$\partial_{x_{i'}}\log\left(p(a_{i'j'}=1)\right)=2\left(\mathbf{x}_{i'}-\mathbf{x}_{j'}\right)\frac{\left(1+e^{C_2}\right)^{-2}e^{C_2}-\left(1+e^{C_1}\right)^{-2}e^{C_1}}{\left(1+e^{C_1}\right)^{-1}-\left(1+e^{C_2}\right)^{-1}}. \tag{B8}$$

Or, if $a_{i'j'}=2$

$$\partial_{x_{i'}}\log\left(p(a_{i'j'}=2)\right)=-2\left(\mathbf{x}_{i'}-\mathbf{x}_{j'}\right)\frac{1}{1+e^{-C_2}}. \tag{B9}$$

For the second part of the force on i', caused by $a_{j'i'}$, one simply needs to take the appropriate term out of the three above and switch i' and j' for α and β.

The force on α caused by j' is given by

$$\partial_{\alpha_{i'}}\log\left(p(a_{i'j'}=0)\right)=-\frac{1}{1+e^{C_1}} \tag{B10}$$

or

$$\partial_{\alpha_{i'}}\log\left(p(a_{i'j'}=1)\right)=\frac{\left(1+e^{C_1}\right)^{-2}e^{C_1}-\left(1+e^{C_2}\right)^{-2}e^{C_2}}{\left(1+e^{C_1}\right)^{-1}-\left(1+e^{C_2}\right)^{-1}} \tag{B11}$$

or

Table 1 Performance benchmarks for networks of different sizes

Network name	Dir./undir	N_{nodes}	N_{edges}	Convergence time (s)
Zachary's Karate Club	Undirected	34	156	2
Lazega Lawyers	Undirected	71	798	3
German parliamentarians	Directed	538	48260	5
Caltech FB Friendships	Undirected	1538	33312	60
Haverford FB Friendships	Undirected	2892	119178	90
Harper's letter (multi-tie)	Directed	540	1017	3600

$$\partial_{\alpha_{i'}} \log\big(p(a_{i'j'} = 2)\big) = \frac{1}{1 + e^{-C_2}}. \tag{B12}$$

Now, there is no second force part—this is the only contribution by the pair i' and j'. The force on $\beta_{i'}$ is given by the analogous term where i' and j' are, again, switched for α and β, and the level of $a_{j'i'}$ is considered.

The forces on c_1 and c_2 by the pair are given by

$$\partial_{c_1} \log\big(p(a_{i'j'} = 0)\big) = \partial_{\alpha_{i'}} \log\big(p(a_{i'j'} = 0)\big), \tag{B13}$$

$$\partial_{c_1} \log\big(p(a_{i'j'} = 1)\big) = \frac{e^{C_1}\left(1 + e^{C_1}\right)^{-2}}{\left(1 + e^{C_1}\right)^{-1} - \left(1 + e^{C_2}\right)^{-1}}, \tag{B14}$$

$$\partial_{c_1} \log(\mathrm{p}(a_{i'j'} = 2)) = 0. \tag{B15}$$

(For the second part, switch i' and j' for α and β and consider $a_{j'i'}$.)

$$\partial_{c_2} \log\big(p(a_{i'j'} = 0)\big) = 0, \tag{B16}$$

$$\partial_{c_2} \log\big(p(a_{i'j'} = 1)\big) = -\frac{e^{C_2}\left(1 + e^{C_2}\right)^{-2}}{\left(1 + e^{C_1}\right)^{-1} - \left(1 + e^{C_2}\right)^{-1}}, \tag{B17}$$

$$\partial_{c_2} \log\big(p(a_{i'j'} = 2)\big) = \partial_{\alpha_{i'}} \log\big(p(a_{i'j'} = 2)\big). \tag{B18}$$

(For the second part, switch i' and j' for α and β and consider $a_{j'i'}$.)

Appendix C: Performance benchmark

The convergence times in Table 1 were measured on an Apple Silicon M1 Max machine. Note that, especially for sharing networks, our first implementation of the algorithm is still very slow compared to existing layout algorithms such as FA2.

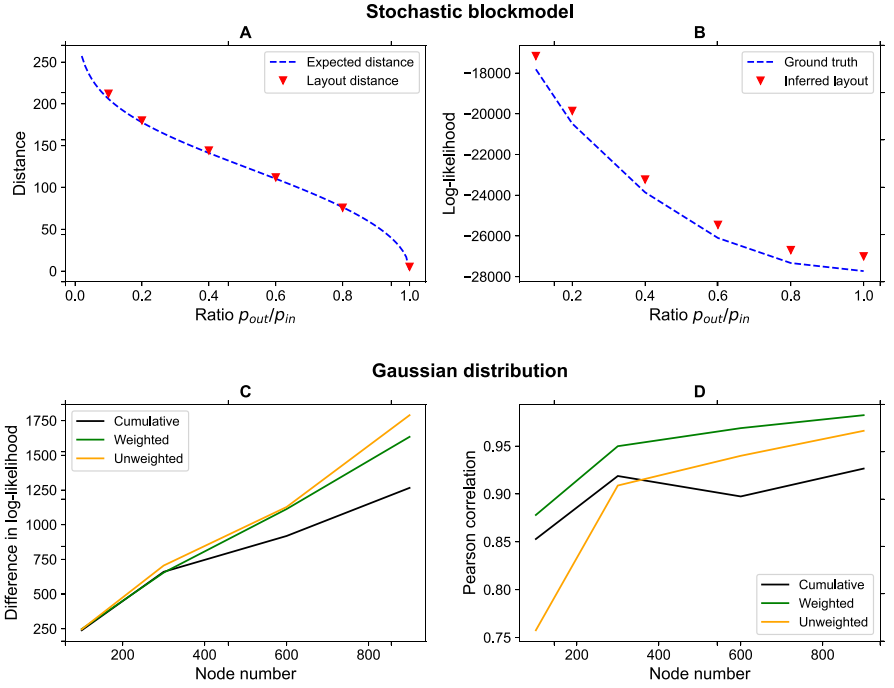

Fig. 5 Expected distance of an SBM (two blocks, 100 nodes each) with varying p_{out} compared to the distance between the center of mass of the clusters in the proposed layout algorithm, averaged over 5 runs (**A**). Not only is the inferred distance by the force-directed layout algorithm nearly identical to the expected one throughout, but the log-likelihood of the inferred latent space surpasses the ground truth in all cases (**B**). The difference between inferred log-likelihood and the log-likelihood of the ground truth for a Gaussian distribution of two groups of nodes ($\sigma = 1/12$, $d = 5/6$, averaged over 3 runs) is given in (**C**). In all cases, the log-likelihood of the inferred latent space surpasses the ground truth (i.e., them difference in log-likelihood is positive). Still, similarity between ground truth and inferred distances between nodes is high (increasing with number of nodes), as is visible in the average Pearson correlation between distance matrices (**D**)

Further work is required to accelerate the performance for larger networks, which could be achieved using approximations such as Barnes Hut, or by randomly sampling nodes for the computation of repulsive forces instead of computing them between all pairs.

Appendix D: Validation

Validation for the unweighted case was performed by testing the agreement with the expected distance of a stochastic blockmodel (SBM) of two blocks with varying p_{out} and $p_{\text{in}} = 0.5$. In the above model, the expected distance can be computed by placing the nodes of each block on the same point in space, and then choosing the distance d between the two blocks, so that $p_{\text{out}} = \frac{1}{1+\exp(-d^2)}$ (α and β are set to 0, so that the expected degree for all nodes in one block is the same). With this underlying

Table 2 Average z-scores for a Mantel test for layout of a network generated from a Gaussian distribution of two groups of nodes with a sigma of 1/12 and a distance of 5/6 between the groups with varying node number (averaged over three runs)

Node number	100	300	600	900
Unweighted	36.65	77.49	90.56	94.91
multi-tie	43.07	77.04	84.58	92.12
Weighted	45.60	83.32	92.84	95.69

latent space, one can draw a network according to the given probabilities and let the layout algorithm infer the latent space again. Averaged over five runs and for 100 nodes per block, we observe that the inferred distance of the centers of mass of the blocks are nearly identical to the expected one (Fig. 5 A), and the log-likelihood of the inferreds latent space surpasses the one of the actually drawn one in all cases (Fig. 5 B).

Moreover, we compare the log-likelihood of the inferred latent space with the one of the actually drawn network from a Gaussian distribution of two groups of nodes with a σ of 1/12 and a distance of 5/6 between the groups with varying node number, averaged over three runs. In all cases, the log-likelihood of the inferred latent space is higher than the ground truth (Fig. 5C). Still, similarity to the ground truth distances between nodes was high throughout, which we assessed with a Mantel test [41]. Pearson correlation between distance matrices can be inspected in Fig. 5D. Table 2 gives the average z-scores for the Mantel test.

Appendix E: Bayesian correction of force term

For the inference of the model parameters, we have only one sample. Although our model is a probabilistic model, our empirical distribution contains either $p(a_{ij} = 1) = 1$ if there is an edge, or $p(a_{ij} = 1) = 0$, if not. This can lead to divergences of the model parameters in the maximum-likelihood solution. If one node is connected to all the other nodes, for instance, i.e., the graph contains a (N-1)-star subgraph, α_i will diverge in the maximum-likelihood solution ($\alpha_i \to \infty$). One way to avoid this problem is to formulate is as a Bayesian inference problem. This is a well-defined problem, even with a single data point. Lets denote the entries of the adjacency matrix of our graph G as a_{ij} and the corresponding random variables A_{ij} or $\mathbf{A}$, respectively. Moreover, lets call the parameter vector of our model $\boldsymbol{\theta} = \{\boldsymbol{\alpha}, \boldsymbol{\beta}, \mathbf{x}\}$, with $\theta_i = \{\alpha_i, \beta_i, \mathbf{x}_i\}$ and the corresponding random variable X. Then, the posterior distribution of the parameters is given as

$$p(X|\mathbf{A}) = \frac{p(\mathbf{A}|X)p(X)}{p(\mathbf{A})}. \tag{E1}$$

Here, $p(\mathbf{A}|X)$ is the likelihood (see Eq. (3)), $p(X)$ comprises all prior distributions, and $p(\mathbf{A})$ is the marginal likelihood of the data. Instead of asking for the parameters that maximise the likelihood, we can now ask for the parameters that maximise the

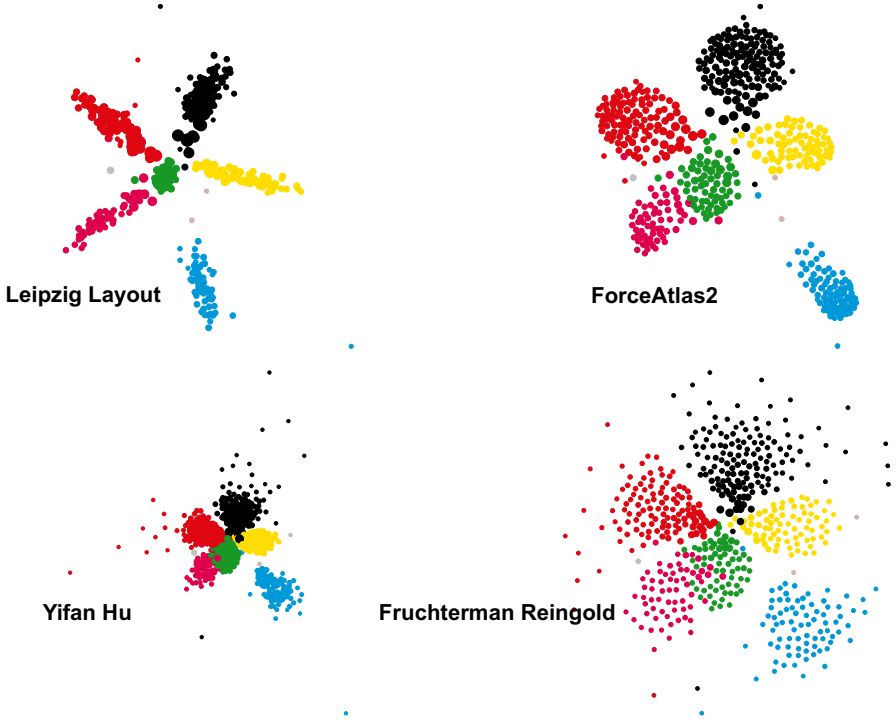

Fig. 6 Bundestag follower network, comparison of Leipzig Layout (top left), ForceAtlas2 (top right), Yifan Hu (lower left), and Fruchterman Reingold (lower right). Overall placement of parties is similar, but Leipzig layout, which allows closer placement of nodes, arranges all parties (except the *Greens*) along a one-dimensional axis

posterior $p(X|\mathbf{A})$ or, equivalently, its logarithm. If one considers the gradients of the logarithm of the posterior again as forces, the likelihood term produces the same forces as in the maximum-likelihood case. However, we may get additional forces from the prior term, which, for instance, can prevent the activity parameters from diverging.[18]

Appendix F: Real-world networks and comparison to other layout algorithms

German parliament In the main text, the Bundestag follower network layout was only compared to ForceAtlas2. In Fig. 6, we also show the follower network spatialised with Yifan Hu [32] (lower left) and Fruchterman Reingold [23] (lower

[18] [4] assumes normal priors for all parameters of the model. Nevertheless, it is noted in the main text that flat priors are used except for α and the positions $\mathbf{x}$. The mean of the prior distribution of α is set to 0 and the prior distribution for the positions is $\mathcal{N}(0, 1)$.

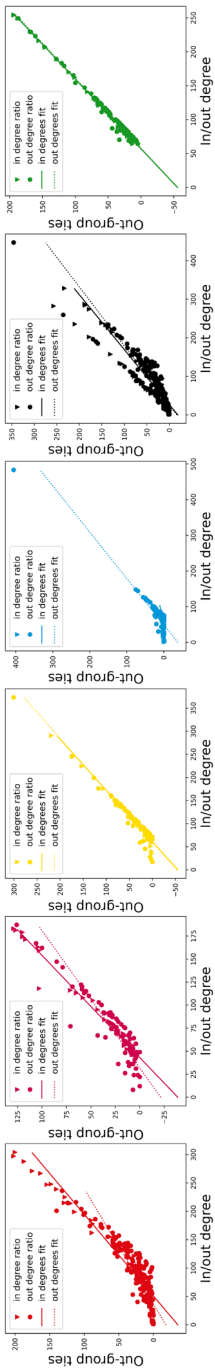

Fig. 7 Scatter plot of in/out-degree and number of out-group ties for each party. No user from the *Green* party has in- or out-degree of less than 50. They use Twitter quite homogeneously, which explains their central placement in the force-directed layout (as well as the homogeneously distributed incoming and outgoing links among parties, see Fig. 7)

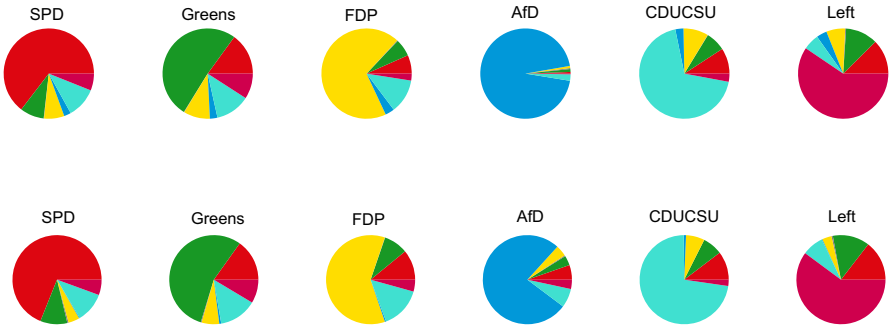

Fig. 8 Incoming and outgoing links from/to different parties by party. *Greens* are followed and follow all other parties (except *AfD*) quite uniformly, which explains their central placement

right). All layout algorithms roughly reproduce party divisions. Fruchterman Reingold tends to distribute nodes homogeneously in spaces. *Die Linke* and the *Greens* have some overlap in this layout. A less pronounced overlap is also visible for Yifan Hu, which produces a comparably dense spatialisation.

The central placement of the *Greens* in the Leipzig Layout (as well as in the other layout algorithms) can be explained by their homogeneous usage of the platform: Each user of the party has an in/out-degree of at least 50 (which is not the case for the other parties; see Fig. 7). Moreover, the party members are followed and follow all other parties in a quite uniform fashion (see Fig. 8).

Moreover, α and β are correlated with node out- and in-degree for each party; see Fig. 9. In the case of the Greens, the correlation appears to be (except for one outlier) strongly linear.

Different (local) minima exist for this network, one of which is displayed in Fig. 10. There, the *FDP* is placed between *SPD* and *Die Linke*. The inference of local minima is a general problem of FDLs—nevertheless, in the present framework, one can compare the log-likelihood of different equilibria and take the spatialisation with the highest likelihood (i.e., the lowest negative log-likelihood). The log-likelihood of Fig. 10 is around 61,700, while it is roughly 57,700 for Fig. 2. The minimum with the higher likelihood is also the politically more plausible one: *SPD* and *Die Linke* are, especially when it comes to economic policy, oftentimes strongly opposed to the market-liberal *FDP*. On the other hand, *FDP* and *CDU/CSU* have often stressed that they are parties that have many things in common, such that Fig. 2 seems to be closer to political reality.

Harper's letter

Figure 11 shows the Harper's letter retweet network spatialised with the four layout algorithms. Fruchterman Reingold, again, arranges the nodes rather uniformly in space. Interestingly, Yifan Hu and ForceAtlas2 produce a much more polarized spatialisation than Leipzig Layout.

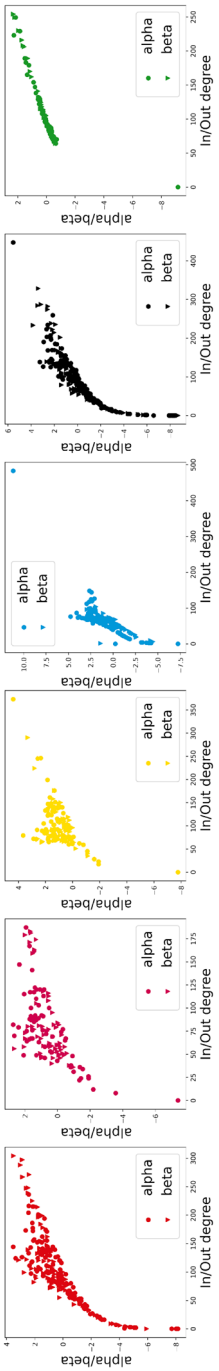

Fig. 9 Scatter plot of in/out-degree and α and β for each party

Fig. 10 Local minimum of the follower network of German deputies. Parties are still visibly separated, but *FDP* is now placed between *SPD* and *Die Linke*

Appendix G: Survey on energy-generating technologies—correlations

Pairwise Pearson correlation coefficients between ratings of different energy-generating technologies (for the aggregated three-point scale) can be found in Fig. 12. Solar and onshore technologies exhibit the strongest correlation. Gas and coal, as well as offshore and onshore technologies are also correlated relatively strongly. Biomass has, among the renewables, the weakest correlations with the other renewables. On the other hand, ratings are not negatively correlated with coal or gas. This is mirrored by its placement in Fig. 4 at a certain distance from solar and onshore, and also offshore.

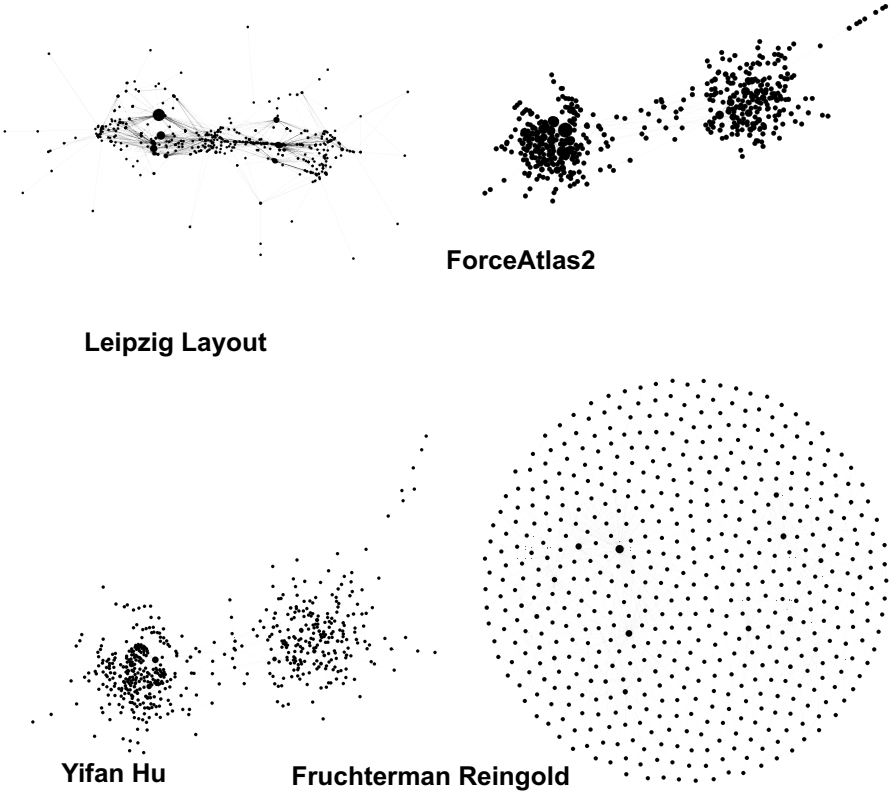

Fig. 11 Harper's letter retweet network comparison of Leipzig Layout (multi-tie case, top left), ForceAtlas2 (top right), Yifan Hu (lower left), and Fruchterman Reingold (lower right)

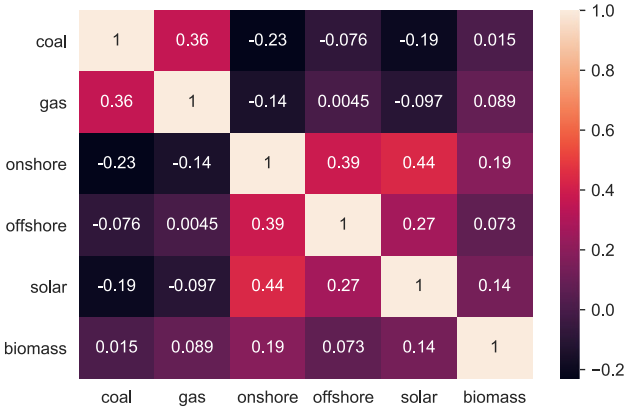

Fig. 12 Correlations between ratings of different energy-generating technologies for the aggregated 3-point scale

Funding Open Access funding enabled and organized by Projekt DEAL.

Data availability statement The implementation of Leipzig Layout is available under https://github.com/pournaki/leipzig-layout.

Declarations

Conflicts of interest The authors declare that no conflict of interest exists.

References

1. (2020). A letter on justice and open debate. https://harpers.org/a-letter-on-justice-and-open-debate/. Accessed 1 Aug 2021
2. Adamic, L. A., & Glance, N. (2005). The political blogosphere and the 2004 us election: divided they blog. In: *Proceedings of the 3rd International Workshop on Link Discovery*, pp. 36–43.
3. Asturiano, V. (2018). Force-graph. https://github.com/vasturiano/force-graph. Accessed 19 Oct 2021.
4. Barberá, P. (2015). Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. *Political Analysis, 23*(1), 76–91.
5. Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right. *Psychological Science, 26*(10), 1531–1542.
6. Barnett, L., Di Paolo, E., & Bullock, S. (2007). Spatially embedded random networks. *Physical Review E, 76*, 056115.
7. Bastian, M., Heymann, S., & Jacomy, M. (2009). Gephi: An open source software for exploring and manipulating networks. In: *Third International AAAI Conference on weblogs and social media*.
8. Batini, C., Furlani, L., & Nardelli, E. (1985). What is a good diagram? A pragmatic approach. In: *Proceedings of the Fourth International Conference on Entity-Relationship Approach*, pp. 312–319.
9. Blondel, V. D., Guillaume, J.-L., Lambiotte, R., & Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment, 2008*(10), P10008.
10. Bostock, M. (2015). d3-force. https://github.com/d3/d3-force. Accessed 19 Oct 2021.
11. Both, C., Dehmamy, N., Yu, R., & Barabási, A.-L. (2023). Accelerating network layouts using graph neural networks. *Nature Communications, 14*(1), 1560.
12. Brandes, U. (2001). Drawing on physical analogies. In: M. Kaufmann & D. Wagner (Eds.), *Drawing graphs: methods and models* (pp. 71–86). Springer.
13. Brandes, U., & Pich, C. (2006). Eigensolver methods for progressive multidimensional scaling of large data. In: *International Symposium on Graph Drawing*, pp. 42–53. Springer.
14. Brandes, U., & Pich, C. (2008). An experimental study on distance-based graph drawing. In: *International Symposium on Graph Drawing*, pp. 218–229. Springer.
15. Bruns, A. (2013). Faster than the speed of print: Reconciling 'big data' social media analysis and academic scholarship. *First Monday, 18*(10), 1–5.
16. Conover, M. D., Goncalves, B., Ratkiewicz, J., Flammini, A., & Menczer, F. (2011). Predicting the political alignment of twitter users. In: *2011 IEEE Third International Conference on Privacy, Security, Risk and Trust and 2011 IEEE Third International Conference on Social Computing*, pp. 192–199.

17. Conover, M. D., Ratkiewicz, J., Francisco, M. R., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on twitter. *Icwsm, 133*(26), 89–96.
18. Dall, J., & Christensen, M. (2002). Random geometric graphs. *Physical Review E, 66*(1), 016121.
19. Decuypere, M. (2020). Visual network analysis: a qualitative method for researching sociomaterial practice. *Qualitative Research, 20*(1), 73–90.
20. Di Battista, G., Eades, P., Tamassia, R., & Tollis, I. G. (1994). Algorithms for drawing graphs: An annotated bibliography. *Computational Geometry, 4*(5), 235–282.
21. Eades, P. (1984). A heuristic for graph drawing. *Congressus Numerantium, 42*, 149–160.
22. Fortunato, S., & Barthelemy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences, 104*(1), 36–41.
23. Fruchterman, T. M., & Reingold, E. M. (1991). Graph drawing by force-directed placement. *Software: Practice and Experience, 21*(11), 1129–1164.
24. Gaisbauer, F., Pournaki, A., Banisch, S., & Olbrich, E. (2021). Ideological differences in engagement in public debate on twitter. *PLoS One, 16*(3), e0249241.
25. Gansner, E. R., Koren, Y., & North, S. (2004). Graph drawing by stress majorization. In: *International Symposium on Graph Drawing*, pp. 239–250. Springer.
26. Good, B. H., De Montjoye, Y.-A., & Clauset, A. (2010). Performance of modularity maximization in practical contexts. *Physical Review E, 81*(4), 046106.
27. Hagberg, A., Swart, P., and S Chult, D. (2008). Exploring network structure, dynamics, and function using networkx. Technical report, Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
28. Handcock, M. S., Raftery, A. E., & Tantrum, J. M. (2007). Model-based clustering for social networks. *Journal of the Royal Statistical Society: Series A (Statistics in Society), 170*(2), 301–354.
29. Harrell, F. E., Jr. (2015). *Regression modeling strategies: With applications to linear models, logistic and ordinal regression, and survival analysis*. Springer.
30. Hinton, G., & Roweis, S. T. (2002). Stochastic neighbor embedding. *NIPS, 15*, 833–840.
31. Hoff, P. D., Raftery, A. E., & Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association, 97*(460), 1090–1098.
32. Hu, Y. (2005). Efficient, high-quality force-directed graph drawing. *Mathematica Journal, 10*(1), 37–71.
33. Imai, K., Lo, J., Olmsted, J., et al. (2016). Fast estimation of ideal points with massive data. *American Political Science Review, 110*(4), 631–656.
34. Jacomy, M. (2021). Situating Visual Network Analysis. PhD thesis.
35. Jacomy, M., Venturini, T., Heymann, S., & Bastian, M. (2014). Forceatlas2, a continuous graph layout algorithm for handy network visualization designed for the gephi software. *PLoS One, 9*(6), e98679.
36. Kamada, T., Kawai, S., et al. (1989). An algorithm for drawing general undirected graphs. *Information Processing Letters, 31*(1), 7–15.
37. Kleinbaum, D. G., Dietz, K., Gail, M., Klein, M., & Klein, M. (2002). *Logistic regression*. Springer.
38. Koren, Y., & Civril, A. (2008). The binary stress model for graph drawing. In: *International Symposium on Graph Drawing*, pp. 193–205. Springer.
39. Kruskal, J. B. (1980). Designing network diagrams. In: *Proc. 1st General Conf. On Social Graphics*, 1980, pages 22–50. US Dept. of the Census.
40. Liu, Z., Wang, Y., Bernard, J., & Munzner, T. (2022). Visualizing graph neural networks with corgie: Corresponding a graph to its embedding. *IEEE Transactions on Visualization and Computer Graphics, 28*(6), 2500–2516.
41. Mantel, N. (1967). The detection of disease clustering and a generalized regression approach. *Cancer Research, 27*(2 Part 1), 209–220.
42. Matias, C., & Robin, S. (2014). Modeling heterogeneity in random graphs through latent space models: A selective review. *ESAIM: Proceedings and Surveys, 47*, 55–74.
43. McGinn, D., Birch, D., Akroyd, D., Molina-Solana, M., Guo, Y., & Knottenbelt, W. J. (2016). Visualizing dynamic bitcoin transaction patterns. *Big Data, 4*(2), 109–119.
44. McInnes, L., Healy, J., Saul, N., & Großberger, L. (2018). UMAP: uniform manifold approximation and projection. *Journal of Open Source Software, 3*(29), 861.
45. Newman, M. E., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physical Review E, 69*(2), 026113.
46. Noack, A. (2007). Unified quality measures for clusterings, layouts, and orderings of graphs, and their application as software design criteria. Doctoral dissertation, BTU Cottbus-Senftenberg.

47. Noack, A. (2009). Modularity clustering is force-directed layout. *Physical Review E, 79*(2), 026102.
48. Olbrich, E., & Banisch, S. (2021). The rise of populism and the reconfiguration of the German political space. *Frontiers in Big Data, 4*, 731349.
49. Peel, L., Delvenne, J.-C., & Lambiotte, R. (2018). Multiscale mixing patterns in networks. *Proceedings of the National Academy of Sciences, 115*(16), 4057–4062.
50. Penrose, M. (2003). *Random geometric graphs. Number 5*. Oxford University Press.
51. Pournaki, A., Gaisbauer, F., Banisch, S., & Olbrich, E. (2021). The twitter explorer: A framework for observing twitter through interactive networks. *Journal of Digital Social Research, 3*(1), 106–118.
52. Red, V., Kelsic, E. D., Mucha, P. J., & Porter, M. A. (2011). Comparing community structure to characteristics in online collegiate social networks. *SIAM Review, 53*(3), 526–543.
53. Sarkar, P., & Moore, A. W. (2005). Dynamic social network analysis using latent space models. *ACM SIGKDD Explorations Newsletter, 7*(2), 31–40.
54. Sewell, D. K., & Chen, Y. (2015). Latent space models for dynamic networks. *Journal of the American Statistical Association, 110*(512), 1646–1657. https://doi.org/10.1080/01621459.2014.988214
55. Shamon, H., Schumann, D., Fischer, W., Vögele, S., Heinrichs, H. U., & Kuckshinrichs, W. (2019). Changing attitudes and conflicting arguments: Reviewing stakeholder communication on electricity technologies in Germany. *Energy Research & Social Science, 55*, 106–121.
56. Steinberg, B., & Ostermeier, M. (2016). Environmental changes bridge evolutionary valleys. *Science Advances, 2*(1), e1500921.
57. Swope, W. C., Andersen, H. C., Berens, P. H., & Wilson, K. R. (1982). A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters. *The Journal of Chemical Physics, 76*(1), 637–649.
58. Tamassia, R., Di Battista, G., & Batini, C. (1988). Automatic graph drawing and readability of diagrams. *IEEE Transactions on Systems, Man, and Cybernetics, 18*(1), 61–79.
59. Traud, A. L., Mucha, P. J., & Porter, M. A. (2012). Social structure of Facebook networks. *Physica A: Statistical Mechanics and its Applications, 391*(16), 4165–4180.
60. Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research, 9*, 2579–2605.
61. van Vliet, L., Törnberg, P., & Uitermark, J. (2020). The twitter parliamentarian database: Analyzing twitter politics across 26 countries. *PLoS One, 15*(9), e0237073.
62. Venturini, T., Jacomy, M., & Jensen, P. (2021). What do we see when we look at networks: Visual network analysis, relational ambiguity, and force-directed layouts. *Big Data & Society, 8*(1), 20539517211018490.
63. Verlet, L. (1967). Computer "Experiments'' on classical fluids. I. Thermodynamical properties of Lennard-Jones molecules. *Physical Review, 159*(1), 98–103.
64. Wang, S. S., Paul, S., & De Boeck, P. (2019). Joint latent space model for social networks with multivariate attributes. arXiv preprint arXiv:1910.12128.
65. Waxman, B. (1988). Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications, 6*(9), 1617–1622.
66. Zhang, X., Xue, S., & Zhu, J. (2020). A flexible latent space model for multilayer networks. In: *International Conference on Machine Learning*, pp. 11288–11297. PMLR.