



OPEN

Top influencers can be identified universally by combining classical centralities

Doina Bucur

Information flow, opinion, and epidemics spread over structured networks. When using node centrality indicators to predict which nodes will be among the top influencers or superspreaders, no single centrality is a consistently good ranker across networks. We show that statistical classifiers using two or more centralities are instead consistently predictive over many diverse, static real-world topologies. Certain pairs of centralities cooperate particularly well in drawing the statistical boundary between the superspreaders and the rest: a *local centrality* measuring the size of a node's neighbourhood gains from the addition of a *global centrality* such as the eigenvector centrality, closeness, or the core number. Intuitively, this is because a local centrality may rank highly nodes which are located in locally dense, but globally peripheral regions of the network. The additional global centrality indicator guides the prediction towards more central regions. The superspreaders usually jointly maximise the values of both centralities. As a result of the interplay between centrality indicators, training classifiers with seven classical indicators leads to a nearly maximum average precision function (0.995) across the networks in this study.

Social influence, news, as well as infectious diseases diffuse in society, following links drawn between participants by frequent contact, mutual interests, collaboration, communication, or transportation. The influence of a single node in such a network measures the extent to which the node, acting as the seed of a multi-hop diffusion process, will activate the rest of the network (this is the cascade size in the domain of online social networks, and the attack rate or outbreak size in epidemiology). Even assuming that the network links are known and the process of diffusion can be modelled or measured, *predicting the top influential nodes* when knowing their nodes' topological centrality indicators remains difficult, also because of the diversity and size of social contact topologies. This study shows, on real-world social networks, that in many networks the joint values of two or more (dissimilar) node centrality indicators are predictive for the influence of the node, and that good combinations are between one local centrality which measures the size of the node's neighbourhood and one global centrality: a variant of the eigenvector centrality, closeness, or the node's core number. We illustrate with examples how the addition of such a second centrality to the prediction process is beneficial on some networks, and show that simple, interpretable statistical models can be machine-learned in a supervised fashion on two or more centrality indicators, with almost universally good results across many real networks and network categories.

Most prior studies predict top influencers by a *ranking method*¹: the nodes in a network are ranked according to a single centrality, with the top assumed to be the best influencers. No single centrality is consistent in performance across realistic case studies. The degree centrality was found a weak predictor in early studies, over both simulated and measured diffusion^{2,3}. With the susceptible-infectious-recovered (SIR)⁴ diffusion model and also with measured diffusion, the top $f = 5\%$ spreaders in a small number of networks were better predicted by their core numbers than by degree or betweenness centrality^{5,6}. The predictive power of the core number was later shown to not generalise, for SIR influence at or above the epidemic threshold. In road networks, the core number correlated little with the spreading ability of a node, while in social networks the degree and core number were either equally predictive⁷, or variably predictive with f ⁸. Over a test suite of ten networks, the eigenvector centrality was on average better than the core number⁹. While refinements of classical centrality indicators were designed^{7,10–20}, also alternative ideas to combine classical centrality indicators into a predictor of influence started in 2011. A metric equal to the betweenness centrality of a node, divided by a power of its degree²¹ was used to recognise the seed of a diffusion process, but was not successful on a real-world topology. By 2020 (the time of this writing), some methods^{22–26} were not applied beyond relatively small or few networks, and also provide no explanation or intuition for the results. A scalable method based on graph neural networks²⁷ was black-box and

Department of Computer Science, University of Twente, Drienerlolaan 5, 7522 NB Enschede, The Netherlands. email: d.bucur@utwente.nl

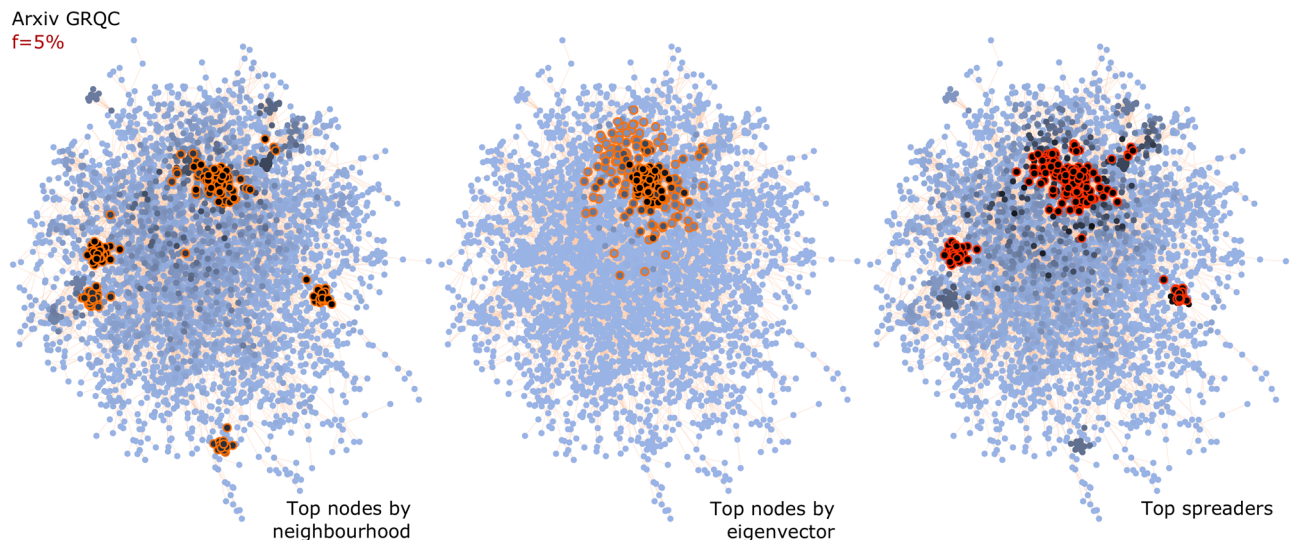


Figure 1. Comparing the location of the top nodes as ranked by (left) neighbourhood size, (centre) eigenvector centrality, and (right) SIR spread size at the epidemic threshold for the coauthorship network Arxiv GRQC. The network layout is force-directed. The colour of the nodes in each panel shows the value of that metric: darker nodes have higher centrality values or spread size. The top $f = 5\%$ of the nodes in each case are encircled.

cannot explain its decision. More interpretable approaches^{28,29} aggregated the individual rankings or values of two or more centralities, with coefficients based on the correlations between the rankings, or the information entropy of a centrality. This obtained recognition rates above 0.7 in 16 networks, with 9–18% improvement over the best single ranking in five of these networks (a lower 1–5% in the rest) and drew the conclusion that the same set of centralities suit networks with similar Laplacian spectra, but making a stronger conclusion on the connection between network topologies and centralities requires more network samples²⁸.

Two recent studies gained more detailed insight. Over all non-isomorphic small networks (up to 10 nodes), one normalized spectral centrality (PageRank or Katz centrality) together with degree (or another measure of network density) predicted well the exact expected SIR spread sizes³⁰. For the related problem of maximising collective influence, PageRank plus metrics related to the node's degree and neighbourhood brought 2–5% improvement compared to the baseline greedy heuristic in real networks³¹. Here, we aim for more general answers: are there other *good combinations* of classical centralities? can one *explain* the added value of a centrality? does the predictive power of a combination of centralities *generalise* across many topologies?

We give an early example in Fig. 1, for the 4158-node coauthorship network Arxiv GRQC. The top $f = 5\%$ of nodes by the size of their neighbourhood (the sum of degrees of nearest neighbours), encircled on the left in Fig. 1, form clusters distributed in the network. The top nodes by the eigenvector centrality (centre) are instead local to one cluster. Neither of these solutions entirely coincides with the correct set of top spreaders, but reasoning with both sets of data leads to a good prediction. The true top spreaders by the SIR diffusion model at the epidemic threshold are shown on the right: these are located in and around only that subset of the clusters with a large neighbourhood which have *also* marginally higher eigencentality values due to being in or close to the high-eigencentality cluster. (Fig. 6 will provide more detail).

We study a large and diverse set of real-world test networks of sizes between 1000 and 70,000 nodes, assuming complete knowledge of the links in the network. The predictive power of two or more centrality indicators is measured by training a supervised statistical classifier on sample nodes from each network. The ground truth for the influence of any node is estimated accurately via the simulation of the SIR diffusion model with that node as the seed of diffusion—possible here since there is one seed, unlike in studies on collective influence, where an approximate greedy heuristic must instead be used as a baseline³¹. The target of the classification is then a binary variable which shows whether the node is in the true top $f\%$ of spreaders. While the results are diverse across the set of networks, we find six universally good pairs between one local centrality which measures the density of the node's extended neighbourhood and one global centrality (eigencentality or PageRank, closeness, core number), and give an intuition for why they complement each other well. With all seven classical centralities, the average precision function is close to perfect (0.995) and the average recognition rate is 0.921.

The practical use of these results is twofold. The method of supervised classification can be ported to any new network where the assumption of complete knowledge about the links is satisfied. For a more realistic estimation of node influence, empirical diffusion data^{3,6}, when available, can replace the mathematical model of diffusion. More importantly, the basic principles of centrality pairing can help with the design of more effective centrality indicators or ranking algorithms, and can improve the understanding of diffusion outcomes in social networks.

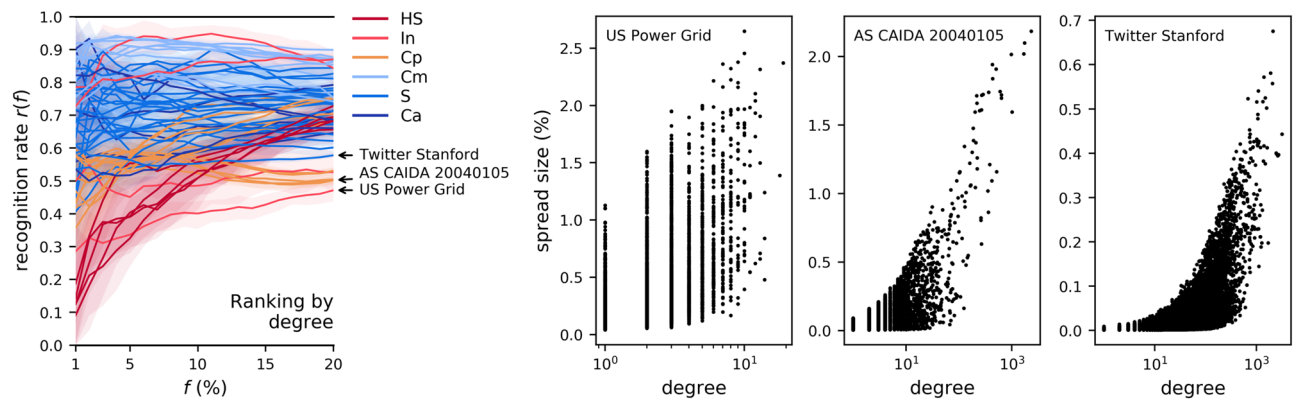


Figure 2. (left) The recognition rate by degree, across f , for all networks, at the SIR epidemic threshold. Each data line corresponds to a network, with the 95% confidence interval shown as a shaded, partly transparent area. The network categories are *Ca* (Coauthorship, 6 networks), *Cm* (Communication between people, 11 networks), *Cp* (Computer, 11 networks), *HS* (offline Human Social, 5 networks), *In* (Infrastructure, 4 networks), *S* (online Social, 23 networks). (right) Degree-influence scatterplots for three of the worst-case networks.

Results

We run an empirical study over 60 real-world examples of static network topologies (listed in Table 1 in Methods). The networks are directed, unweighted, and fall into six categories: human social networks (separately, online or offline), human networks formed by professional coauthorship or online communication, computer networks, and physical infrastructure. The influence of a node is the SIR spread size when the node is the seed of diffusion, estimated via Monte Carlo simulation (see Methods). Analyses are shown in this section for the SIR influence at the epidemic threshold λ_c for every network; they hold also above the epidemic threshold, at $1.5 \cdot \lambda_c$ (with numerical results for these shown in the Supplementary Information).

We study seven classical centrality indicators and their combinations, as follows.

- Local metrics, simple to compute, reflect the density of a node's neighbourhood: the *degree*, *neighbourhood* (the sum of the degrees of direct neighbours), and *two-hop neighbourhood* (the sum of the degrees of neighbours exactly two hops away).
- The *core number* results from k-shell decomposition.
- Distance-based centralities, such as *closeness* and *betweenness*, reflect the importance of nodes by their link distances in the network. Of these two popular centralities, in prior studies on the SIR model, betweenness showed weak predictiveness both as a ranker of nodes in large networks^{5,7} and also in combinations with other centralities on small networks³⁰. We thus study here the closeness centrality.
- Normalised spectral centralities: *PageRank* and *eigenvector* centrality.

The predictive power of single centralities is inconsistent across networks. We first show that the ability of any one centrality indicator to predict the top spreaders across a large number of network cases is too variable to be of universal practical use. Take a network of N nodes, f a fraction, and the task of selecting the best fN spreaders in the network. The standard ranking method has each centrality rank the nodes in this network; the top fN nodes by this ranking are put forward as best spreaders^{5–9} (see Methods). The predictive power of the degree centrality is shown in Fig. 2, across all networks, at the epidemic threshold. This is measured via the *recognition rate* (also called recall) $r(f)$: the fraction of correctly identified top spreaders (Eq. 1 in Methods); the 95% confidence interval around $r(f)$ is shown as a shaded area. In Fig. 2, for each of the three categories of networks with lowest recognition rates at $f = 20\%$, the worst-case network is named. The degree-influence scatterplots, also in Fig. 2, show the reason: a correlation between degree and influence does exist even in these worst cases, but with too wide a variance of influence per degree for accurate ranking.

Compared to the degree, the performance of the core number as a ranker is much less consistent across networks (Fig. 3). The same cause holds for the three worst-case networks marked in the figure: all have few k-shells (between 1 and 5), so the core number by itself is not a discriminative variable for a ranking task. In the very worst case (as in the case of Gnutella25), the network has a single k-shell, so predicting the top spreaders by ranking the nodes in the network is the same as doing a random draw. In Fig. 3, three more networks are marked, for which ranking by core number gives good recognition rates at $f = 20\%$, but poor rates when $f < 5\%$. The scatterplots between core number and influence show the cause. The nodes with the highest core number in the Twitter Stanford network are poor spreaders; a topological reason for this was found in a prior study focused on the core number⁸: the most effective core in the network depends not only on its core number, but also on its connectivity to other cores. Even in other topologies, in which high core numbers do correlate with wide spreading (as is the case for Twitch ES and US Airports), the highest core contains many nodes of very variable influence, so the core number alone is not a sufficiently discriminative variable when f is low.

Neither the degree centrality nor the core number are universally better than the other across the network space. If the core number can be a more accurate ranker in some cases (Fig. 3 shows values of $r(f)$ closer to 1

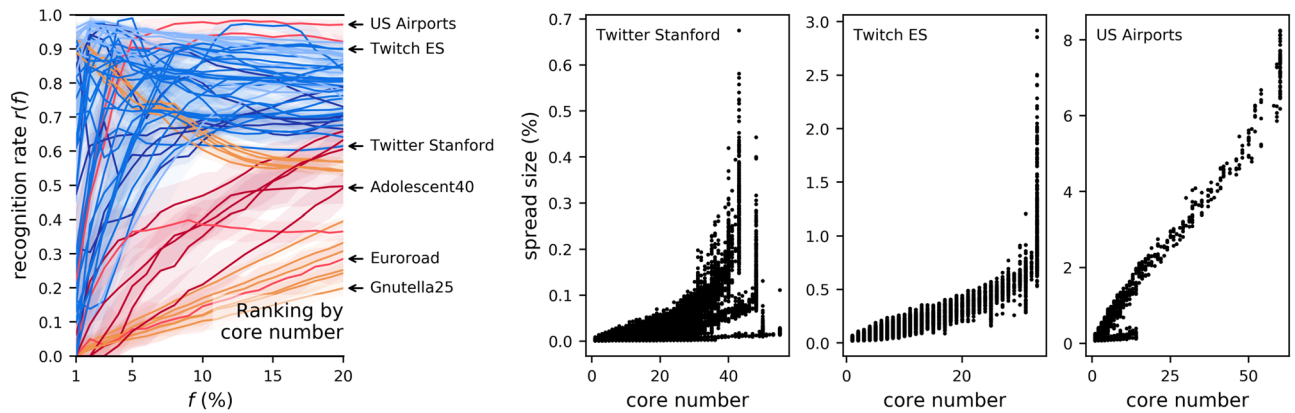


Figure 3. As Fig. 2, but with the core number as the ranker.

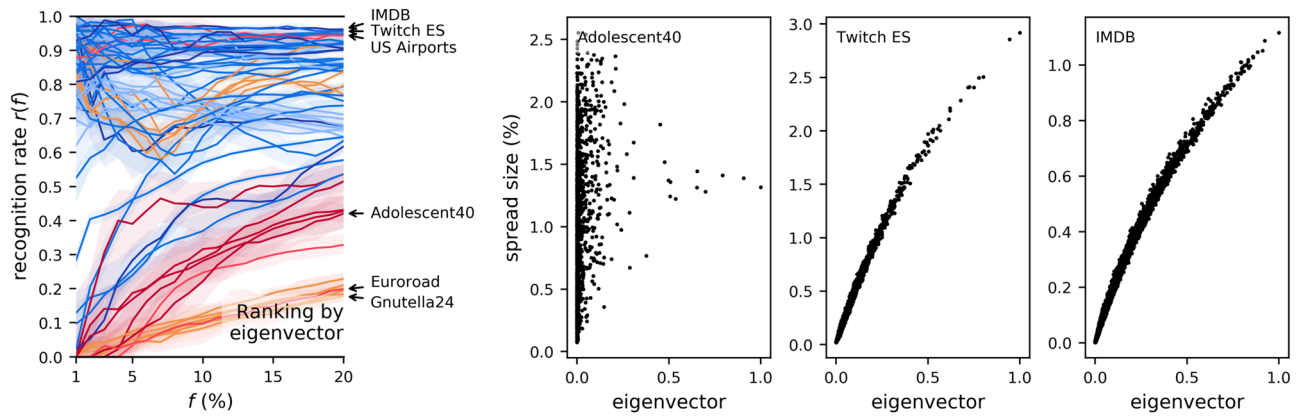


Figure 4. As Fig. 2, but with the eigenvector centrality as the ranker.

for the core number, as was also found in prior studies on selected topologies^{5,6}, it is also a poor predictor in absolute terms when $f < 5\%$ for many networks, and also across all f values when the network doesn't have a strong core structure. For online human networks (categories *Ca*, *Cm*, and *S* in this study), and with $f > 5\%$, Figs. 2 and 3 show the two centralities to be comparable, with the core number marginally better. In general, as recognised before^{7–9}, the predictive power of the core number is not consistently better than the degree centrality for SIR influence.

Another popular ranker, the eigenvector centrality was previously found (on average across a set of networks) more predictive than the core number⁹. By the summary in Fig. 4, this is the case for low values of f , but there is still a wide variance between networks. In some cases (such as Gnutella24 and Euroroad, marked in the figure), the distribution of centrality values is such that ranking is not better than a random draw; in others, such as Adolescent40, there is little correlation between the centrality and influence, so the ranking remains poor. In the best of cases (for two of which scatterplots are shown in the figure), this correlation is strong, which explains why the eigenvector centrality can be a very good predictor across the range of f .

A second performance metric is also of interest: the *precision function* $p(f)$ (Eq. 1 in Methods), which compares the SIR influence of the predicted nodes with the SIR influence of the correct top spreaders. A $p(f)$ value close to 1 for a prediction task means that, regardless whether or not the exact top spreaders were identified, the influence of the nodes which were identified is close to that of the set of top spreaders—so $p(f)$ does not penalise node substitutions, if the substitutes are similar in terms of influence. For ranking by single centralities, the results for both the recognition rate and the precision function are shown in Fig. 5. Each data point marks the performance of a ranking task, over a given network, for a value of f in 1, 2, . . . 20%. (To make the data points visible despite many partial overlaps, each data point is a horizontal line; this line does not denote the uncertainty of the data, but is of fixed size.) The centroid of each data cloud summarises the performance of that centrality over this set of networks. Overall, the neighbourhood centrality makes for the best single ranker, with an average recognition rate of 0.804 and an average precision function of 0.962. The two-hop neighbourhood (not shown in the figure) is only slightly worse (on average 0.781 and 0.942, respectively). PageRank is the least accurate, with an average recognition rate of 0.487, and an average precision function of 0.727. This latter result is not entirely surprising: although widely used for ranking nodes in network structures³², PageRank was found before to not be a competitive predictor for measured diffusion in various networks^{6,9}.

Next, we show that certain pairs of centrality indicators have, together, sufficient topological information about network nodes to improve the accuracy of the prediction tasks.

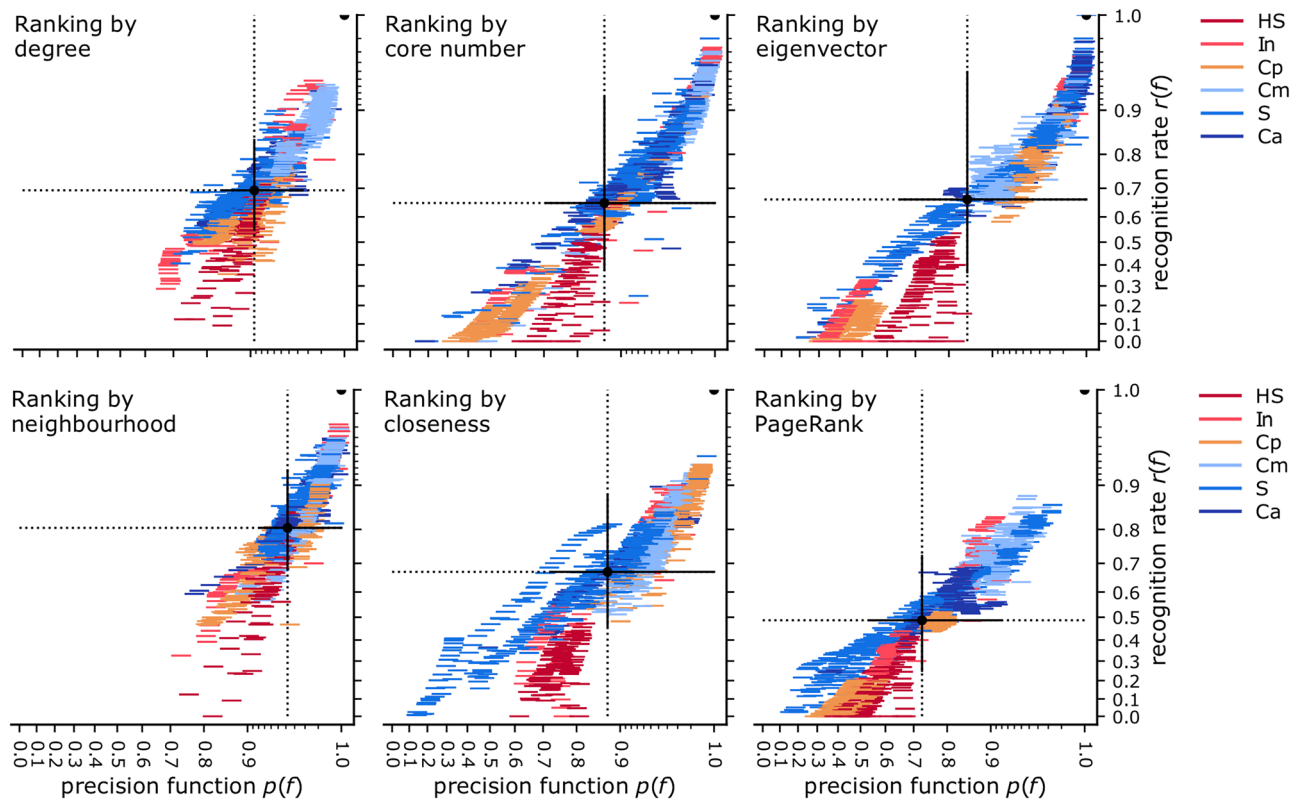


Figure 5. The success of single-centrality ranking at predicting spreaders, across all networks and values of f , at the SIR epidemic threshold. The scales are quadratic. Each data point (a horizontal line of fixed size) denotes a prediction task, and the colour shows the category of the network (listed in Table 1 in Methods). The centroid of the point cluster and the standard deviation on both axes are marked with a solid dot and lines. The point of perfect scores (1,1) is also marked with a half circle. The neighbourhood centrality is the best overall single ranker, with an average precision function of 0.962 and an average recognition rate of 0.804.

Pairs of centralities combine into better predictors. A statistical classifier is now trained with multivariate data from part of the nodes in each network. The result is one trained classifier per network and fraction f . For training, a centrality is one input feature. The target variable (or class) is binary, and it shows whether or not a node is in the top fraction f in the network by spread size. The two performance metrics for the classifiers are the same as for ranking tasks, with the difference that the recall $r(f)$ is now improved as the F1 score, which is the harmonic mean between the precision of classification and the recall (for motivation, see Methods, Eq. 2).

Parsimonious statistical models are beneficial to gain clear intuition about the results. We report here the most *interpretable* statistical models which have good performance: support-vector machine (SVM) with second-degree polynomials as kernels (see Methods), whose *decision boundaries* between classes are simple to understand. We verified that other, higher-variance statistical models based on decision trees have similar performance (with numerical results for Random Forests shown in the Supplementary Information). We start with training SVM classifiers with two centralities, and show that, for certain network examples, certain pairs of centralities build on each other's strengths and obtain predictive models that are significantly better than either centrality alone.

Combinations with the eigenvector centrality. We show four network examples in Fig. 6. For each network, the left panel maps the distribution of the spread size at the epidemic threshold for all the nodes in the network, against the pairing of the eigencentality with a neighbourhood indicator. The right panel notes a value for f , and colours the nodes according to their true class: the red nodes are the top f by spread size. Also in the right panel, two dotted lines show the decision boundaries made by the corresponding single-centrality rankers. If $f = 1\%$, these boundaries are the 99th percentiles for either centrality; a ranker will predict as top spreaders all nodes above this boundary. These ranking boundaries are improved upon by the classifier, whose decision boundary is shown as the transition between background colours, with a blue (or darker) background showing the centrality space where the top spreaders are predicted to be. (Note that only part of this centrality space may be occupied by nodes; in other words, not every combination of centrality values may be physically possible.) The optimal decision boundary would leave no nodes misclassified and would lead to values of 1 for both the precision function and the recall or F1 score.

There are clear commonalities among the improved decision boundaries in Fig. 6: for Facebook Artists, Brightkite, and Arxiv GRQC, the *joint* increase in the values of both centralities in the pair is what determines

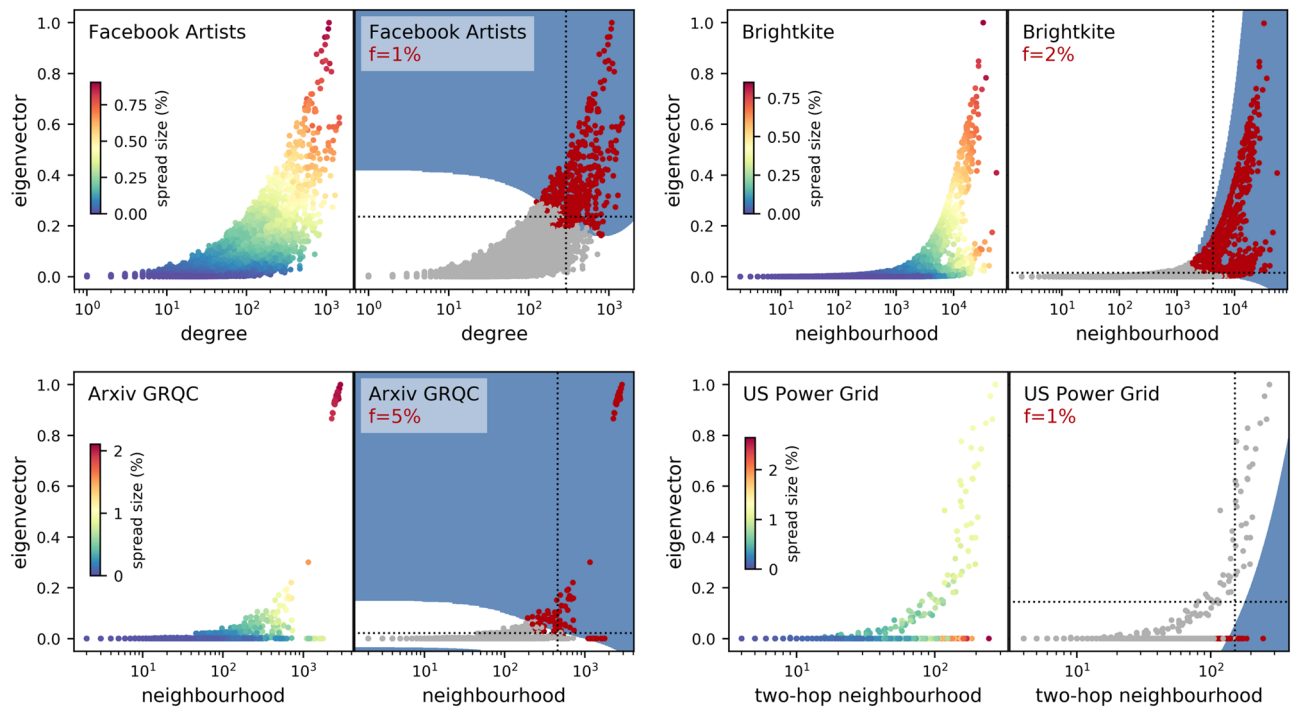


Figure 6. Network examples for which eigenvector centrality combined with another centrality improves the predictions of single-centrality rankers. In every left panel, a scatterplot of node centralities versus spread size. In every right panel, the top spreaders are coloured in red (or darker), the decision boundaries for rankers using either centrality are dotted lines, and the background colour shows the decision boundaries for the classifiers: a blue (or darker) background denotes the area predicted for top spreaders.

an effective spreader. For Facebook Artists and Brightkite (both relatively large networks of over 50,000 nodes), ranking the nodes by only one centrality would place some nodes in the wrong class; unlike this, the two-centrality classifier (F1 scores of 0.920 and 0.924, respectively) draws a decision boundary that is much closer to optimal. We illustrated the intuition behind the Arxiv GRQC result (F1 score 0.900) in Fig. 1: the size of the local neighbourhood does affect the spreading ability of nodes, but proximity to the ‘hub’ of high eigencentrality also helps.

There are also exceptions from this. The US Power Grid network (4941 nodes) shown in the same figure has an outlying cluster of low-eigencentrality nodes as top spreaders, while the lesser spreaders instead follow the expected trend described above. Supplementary Figure S1 shows the cause: a small hub of high eigencentrality values lies at a periphery of the network, while a larger region of nodes with large neighbourhoods (but low eigencentrality) is located far apart. It is the latter, larger region which enables the top 1% of the spreaders, and the classifier is able to learn this pattern slightly better, with a 0.162 increase (F1 score 0.509) compared to the $r(f)$ of ranking by the two-hop neighbourhood alone.

Combinations with the core number. A similar intuition holds when pairing the core number with eigenvector centrality, and also with neighbourhood centralities. (Other pairings with the core number are less effective.) We show two examples in Fig. 7. Again it is the joint increase in both centralities which enables superspreading. For Facebook Politicians (F1 score 0.894), Fig. 7 (bottom) also illustrates the intuition. A number of dense cores are distributed in the network, with the highest core numbers not in close proximity, but isolated by regions of low density. On the other hand, a single region of high eigencentrality exists, and the top 5% of spreaders are located exactly in those cores of highest eigencentrality. Interestingly, pairing the core number with a neighbourhood centrality (GooglePlus, F1 score 0.968) also shows that not all the nodes in dense cores are equally good spreaders, and that their neighbourhood size can help to make a selection.

Combinations with closeness. Closeness also plays a role similar to the eigencentrality—that of guiding the selection of nodes away from more peripheral nodes with dense neighbourhoods, towards the centre of the network, with an increase in performance. Figure 8 shows two examples. In the Adolescent41 offline social network (1,640 nodes), the best ranker is that by neighbourhood ($r(f) = 0.469$), but when considering also closeness, the F1 score rises to 0.598. On the topology of the network (at the bottom of the same figure), closeness values identify only very few of the top spreaders, while the neighbourhood size identifies more; the correct top spreaders, however, again lie in a region where both centralities jointly have high values. In the Gnutella05 computer network, for a similar reason, the best ranker is instead closeness ($r(f) = 0.594$), but when considering also the two-hop neighbourhood, the F1 score rises to 0.725.

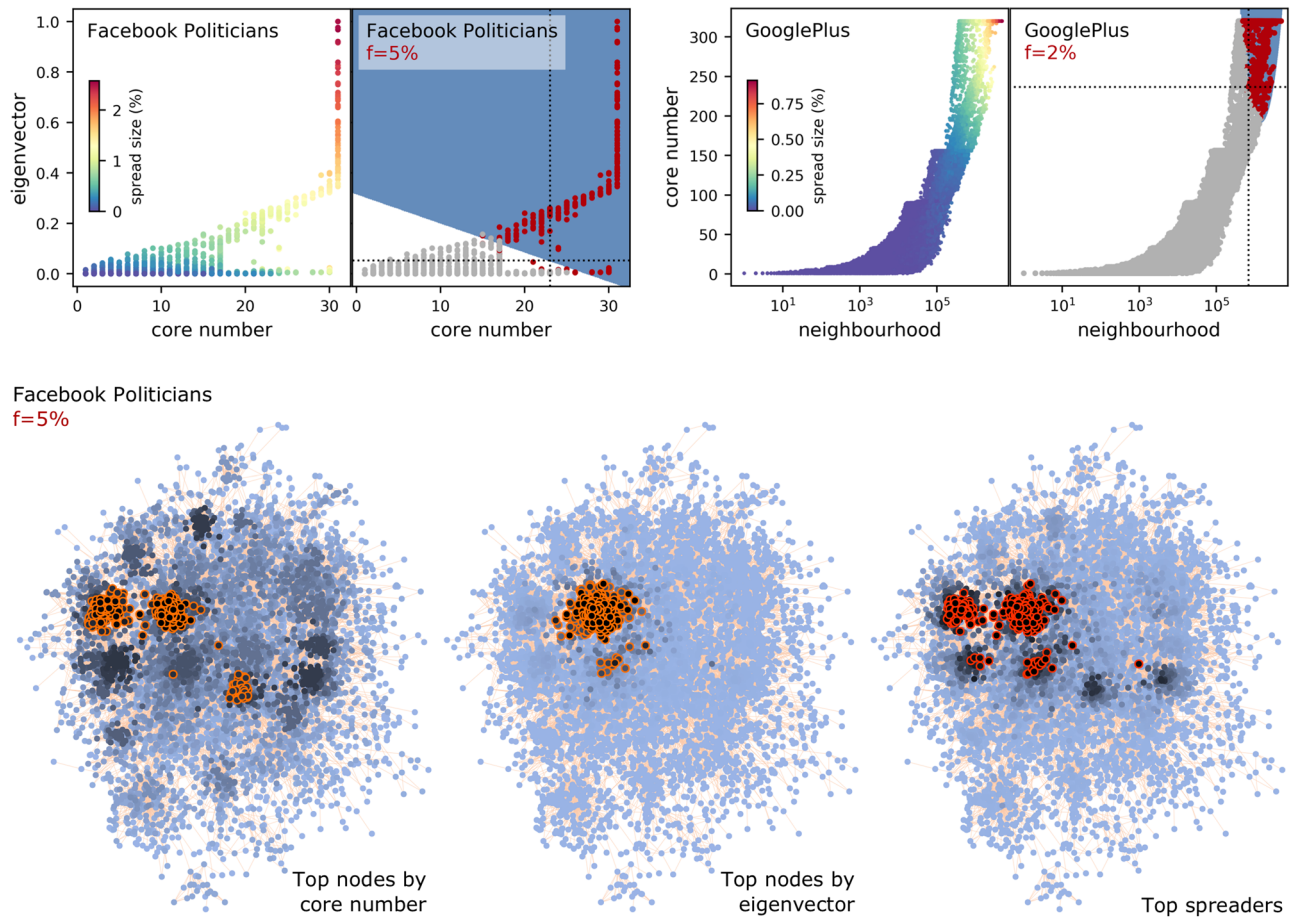


Figure 7. As Fig. 6, for core number combined with another centrality.

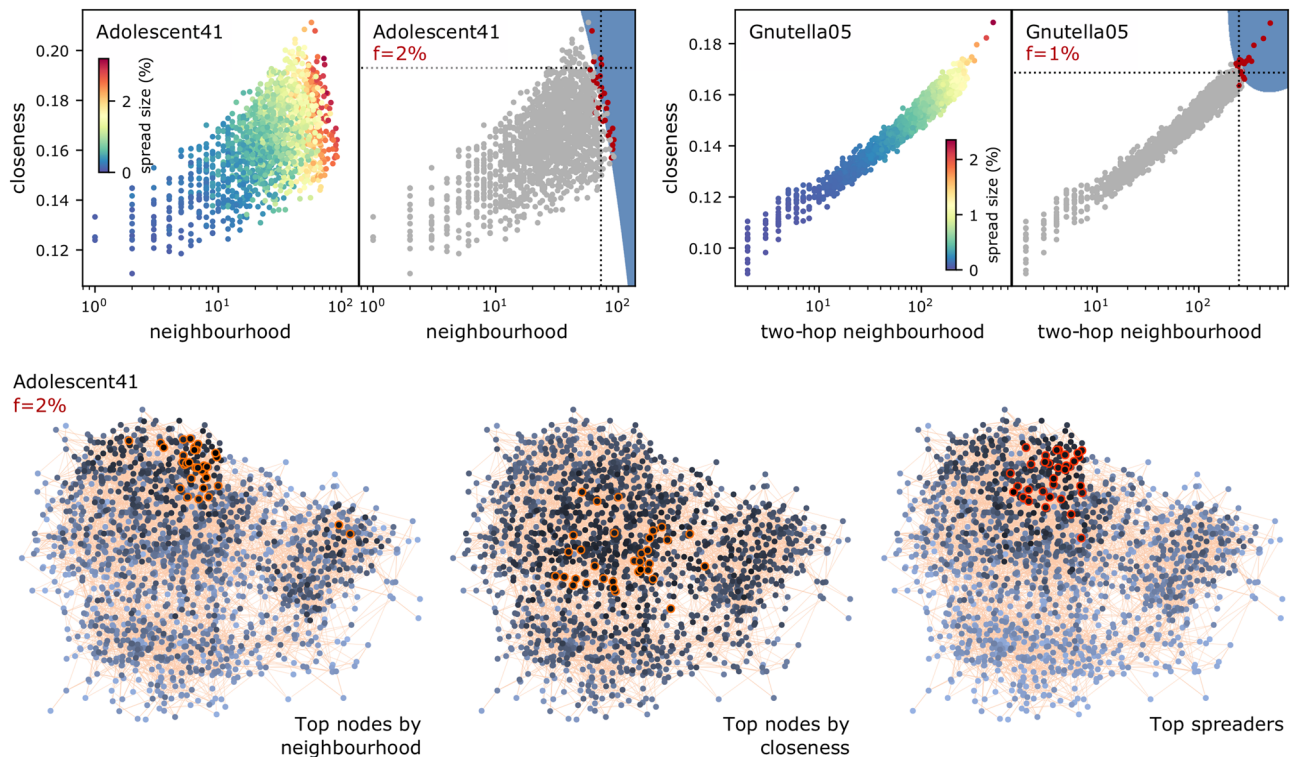


Figure 8. As Fig. 6, for closeness combined with another centrality.

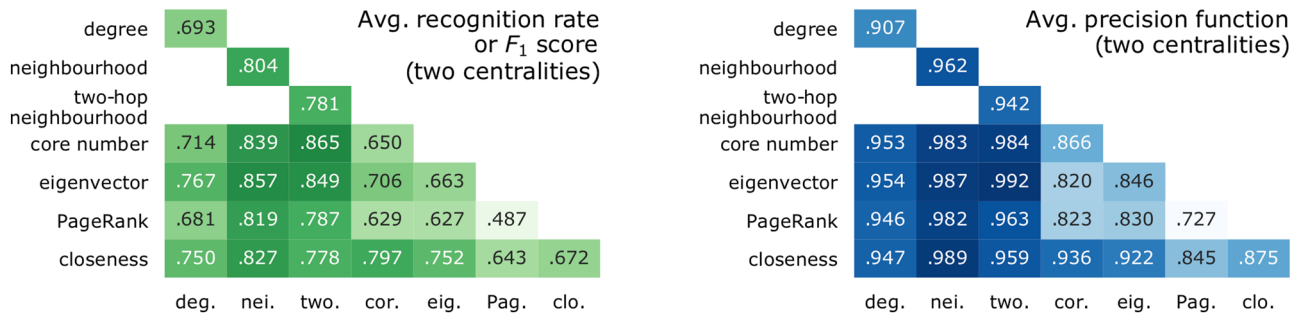


Figure 9. The success of single and pairs of centralities at predicting spreaders: for each pair of centralities, the average performance score across all networks and values of f . The diagonal is the result of ranking by a single centrality and it is scored by the recognition rate and the precision function. The rest of the matrix is the result of classification by two centralities and is scored by the F1 score and the precision function.

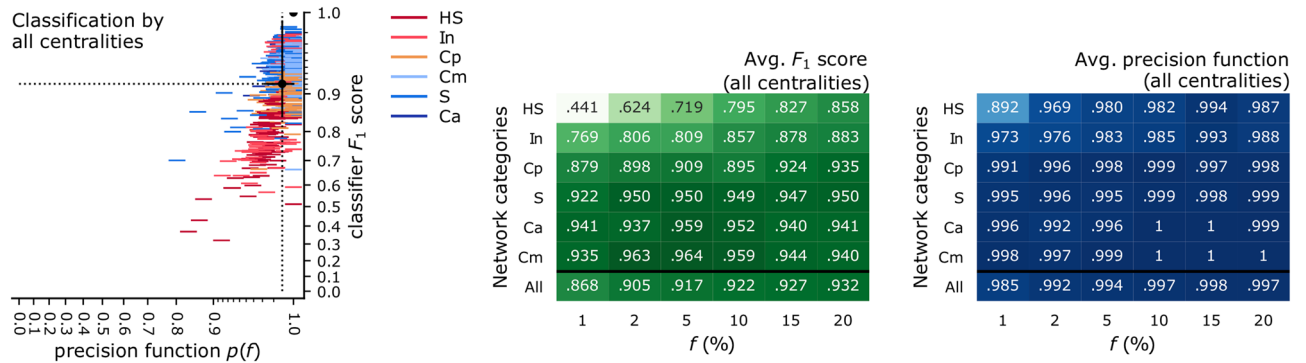


Figure 10. The success of classifiers using all centralities at predicting spreaders, across all networks and values of f , at the SIR epidemic threshold. (left) Each data point denotes a prediction task, and the colour shows the category of the network (listed in Table 1). The centroid of the point cluster and the standard deviation on both axes are marked (counterpart to Fig. 5). (right) The average performance scores across all networks in one of six network categories, and across all values of f (counterpart to Fig. 9).

In the examples from Figs. 6, 7 and 8, each classifier’s decision boundary improves upon the decision boundary of the best ranker such that $r(f)$ is raised by between 0.090 and 0.213. Among our 60 test cases, we also found other examples of networks, combined with certain values for f , for which the single-centrality rankers could not be improved by any classifier. For example, only when $f = 1\%$, none of the five Adolescent networks is resolved any better by using two centralities—but also there the performance improves when f increases.

From all pairs of centralities, the combination of two-hop neighbourhood and core number has the best average F1 score (0.865) across all the network cases in this study, and across the range of f . On the other hand, the combination of two-hop neighbourhood and eigenvector has the best average precision function (0.992). Figure 9 is a summary for the averages of both performance scores across all single centralities (on the diagonal) and pairs of centralities (the rest of the matrix). All possible pairs of centralities are studied, except for the redundant combinations between degree and neighbourhood, and between the two types of neighbourhood centralities. The six pairs which improve significantly on the most predictive ranker are all composed of one of the neighbourhood centralities, and one of: core number, eigenvector centrality, closeness, or PageRank. These six pairs improve on both recall and precision function.

Multi-centrality predictors and summary of results. While the previous subsection demonstrated that centrality indicators can play on each others’ strengths and improve the prediction of top spreaders by the SIR diffusion model at the critical threshold, we now show that classifiers using all seven centralities as features give near-perfect prediction on most network examples. One exception is that of offline human social networks (the HS network category) and only at very low fractions f . This category contains networks that are not structurally unusual, but are some of the smallest networks in the study, which leads to very few training data points, thus lower classification performance.

We train a seven-centrality SVM classifier for each prediction task, and summarise the results in Fig. 10. The centroid of all prediction scores (Fig. 10, left) is an average recognition rate of 0.921, and an average precision function of 0.995. While the precision function was almost as high (0.992) when training the classifier using only the eigenvector centrality and the two-hop neighbourhood as features (Fig. 9), the average recognition rate is now further improved by adding more features to the statistical model. Not all six network categories are equal: a breakdown of the scores by network category and by the value of the fraction f (Fig. 10, right) shows that

recognising the top 1% of spreaders in the Adolescent networks (the HS network category) remains difficult. All other prediction tasks are resolved well, particularly when performance is measured by the precision function, which ranges between 0.969 and 1.

These conclusion hold also above the epidemic threshold, at $1.5 \cdot \lambda_c$; numerical results showing very similar prediction scores are in Supplementary Fig. S2. They are also not an artefact of the type of statistical model used in the classifier. When training nonlinear Random Forest classifiers, which are high-variance so—in general—are able to obtain better performance than the polynomial SVM, a similar conclusion emerges (Supplementary Fig. S3), so there no significant advantage to using higher-variance classifiers.

Discussion

Insights gained. We showed that two or more classical centrality indicators can contain sufficient statistical information about the nodes in a real-world network to train an accurate supervised predictor of SIR influence, and outperform node rankers. The decision boundaries between the two classes, as learnt by classifiers, demonstrate where the advantage of multi-variate prediction comes from: certain centrality indicators are particularly good complements to others. Notably, there are multiple answers to the question: what is a good pair of centralities? For the degree centrality, the best complement is the eigenvector centrality. For the neighbourhood centrality (the best overall single ranker), three other centralities make good complements: the eigenvector centrality, closeness, and core number (with PageRank also close). For those network cases where multi-variate prediction has an advantage, the joint distribution of the centralities and the SIR influence is such that one centrality (or, a one-dimensional decision boundary) is insufficient to classify the nodes accurately, but a multi-dimensional decision boundary is able to refine the decision in the most important region of centrality values. When the entire set of classical centralities are used, the prediction performance is close to optimal (to an average recognition rate of 0.921, and an average precision function of 0.995).

We showed the topological intuition behind this improvement in the prediction of superspreaders. Often, when a subset of the top nodes by local centrality indicators are located in more peripheral regions of the network, global centrality indicators step in and act as a selector and guide towards the effective centre of the network, so that the nodes selected jointly maximise the values of both centralities. In exceptional topologies, when the global centrality has high values at a peripheral location (such as US Power Grid, in Supplementary Fig. S1), the roles reverse: the local centrality becomes the selector, and the statistical model learns that high global centrality values are not beneficial.

Practical use, assumptions, and limitations. The basic insight of jointly maximising the values of two or more centralities can help improve existing, unsupervised node ranking methods. The advantage of ranking algorithms is that they are unsupervised, i.e., require no ground truth; their disadvantage is lower recall and precision.

Network practitioners can also use supervised classification as presented here, and train a new classifier on a new network. While this method delivers good predictions, it assumes (a) complete knowledge of the network links, and (b) means to estimate the spread size for a fraction of the network nodes. If historical diffusion data is available (such as the number of retweets on Twitter), this data replaces the need to simulate a theoretical diffusion model in order to obtain ground truth for the spread size. Only a fraction of nodes need ground truth data, since the statistical classifier is trained on a random sample of the nodes in the network, and will predict the class for the others. The size of the training data necessary to obtain good predictions depends on the network and on the distributions of centrality and influence values, but is expected to be small. In Supplementary Fig. S4, we measure the required training set size from the learning curves of three of the largest networks in this study. These show that, to obtain maximum performance, some networks only require a training data size of 1% of the network size, while others need around 10%. The set of centralities to use as features can be tailored to the computational budget available. The type of statistical model can also be tailored with the network size: heuristic training algorithms, such as those training Random Forest classifiers, scale better with large networks.

Future work. There are follow-ups to explore as continuations of this study, at the intersection between real-world network dynamics and machine learning. A method to train a single statistical model for predicting superspreaders across networks is desirable, as long as its performance remains good; this was previously achieved only for small networks³⁰. An unsupervised or semi-supervised learning method (for example, based on clustering nodes using the same centrality indicators as features, such as in the related work³³ from the domain of natural-language processing) would lower the computational load required to estimate the spread size of many nodes. Other directions include the prediction of other measures of node influence (such as the measured diffusion of information in large online social networks⁶) and of node importance (such as the ability of a node to block the diffusion of information), and the study of other types of networks (such as different network categories, networks with node and link attributes, and networks with dynamic structure).

Methods

Networks, centrality indicators, and the estimation of node influence. Most of our network case studies (see Table 1 for the overview) model entire communities at a specific point in time. This is the case for the high-school friendships in the Adolescent networks, the daily Gnutella peer-to-peer file sharing networks, the five sets of institutional email exchanges, or the networks of mutual likes between verified Facebook pages. A minority of the networks (such as the Facebook Stanford friendships, collected from survey participants) are instead bounded samples from a larger community. All are (transformed into) directed, strongly connected, and unweighted networks; when the original version in the repository had timestamp, attribute, or weight annota-

Cat.	Rep.	Network	Size	λ_c
HS	K	Adolescent36	1671	0.29
HS	K	Adolescent40	1679	0.23
HS	K	Adolescent41	1640	0.26
HS	K	Adolescent49	1149	0.25
HS	K	Adolescent50	2155	0.25
S	K	Advogato	3140	0.041
Ca	S	Arxiv Astro	17,903	0.0155
Ca	S	Arxiv CondMat	21,363	0.045
Ca	S	Arxiv GRQC	4158	0.080
Ca	S	Arxiv HEPPh	11,204	0.008
Ca	S	Arxiv HEPTh	8638	0.0925
Cp	S	AS CAIDA 20040105	16,301	0.033
Cp	S	AS CAIDA 20041206	18,501	0.028
Cp	S	AS CAIDA 20051205	20,889	0.028
Cp	S	AS CAIDA 20061225	23,918	0.027
Cp	S	AS CAIDA 20071112	26,389	0.030
S	S	Brightkite	56,739	0.0185
Cm	S	Email Enron	33,696	0.012
Cm	S	Email EU	34,203	0.022
Cm	K	Email Linux	18,531	0.0075
Cm	M	Email UCL	12,625	0.035
Cm	K	Email URV	1133	0.070
S	S	Epinions	32,223	0.0135
In	K	Euroroad	1039	1.3
S	S	Facebook Artists	50,515	0.007
S	S	Facebook Athletes	13,866	0.030
S	S	Facebook Companies	14,113	0.057
S	S	Facebook Government	7057	0.014
S	M	Facebook New Orleans	63,392	0.0098
S	S	Facebook Politicians	5908	0.031
S	S	Facebook Public figures	11,565	0.020
S	S	Facebook Stanford	4039	0.011
S	S	Facebook TV shows	3892	0.049
S	S	GitHub	37,700	0.0105
Cp	S	Gnutella04	4317	0.29
Cp	S	Gnutella05	3234	0.32
Cp	S	Gnutella24	6352	0.39
Cp	S	Gnutella25	5153	0.42
Cp	S	Gnutella30	8490	0.35
Cp	S	Gnutella31	14,149	0.38
S	S	GooglePlus	69,501	0.0019
S	K	Hamsterster	2000	0.029
Ca	M	IMDB	47,719	0.003
In	K	OpenFlights	3354	0.024
S	K	PGP	10,680	0.065
S	S	Twitch DE	9498	0.0085
S	S	Twitch EN	7126	0.033
S	S	Twitch ES	4648	0.014
S	S	Twitch FR	6549	0.0098
S	S	Twitch RU	4385	0.0185
S	S	Twitch PT	1912	0.013
S	S	Twitter Stanford	68,413	0.0115
In	K	US Airports	1402	0.020
In	K	US Power Grid	4941	0.87
Cm	K	WikiTalk AR	8797	0.018
Cm	K	WikiTalk IT	36,356	0.008

Continued

Cat.	Rep.	Network	Size	λ_c
<i>Cm</i>	<i>K</i>	WikiTalk NL	18,598	0.012
<i>Cm</i>	<i>K</i>	WikiTalk PT	21,747	0.009
<i>Cm</i>	<i>K</i>	WikiTalk RU	22,664	0.011
<i>Cm</i>	<i>K</i>	WikiTalk ZH	10,831	0.013

Table 1. The 60 case studies. We use the largest strongly connected component, whose size (node count) is reported here. **Rep.** denotes the source repository: *K* (KONECT^{35,36}), *M* (H. Makse³⁷), or *S* (SNAP³⁸). **Cat.** is the KONECT category most suited to the case study: *Ca* (Coauthorship), *Cm* (Communication between people), *Cp* (Computer), *HS* (Human Social, offline), *In* (Infrastructure), *S* (Social, online). λ_c denotes the epidemic threshold, estimated numerically.

tions, these were removed. The direction of the edges is reversed when needed, to model information flow—so the degree centrality of interest is the out-degree. To be able to study the closeness centrality³⁴ which computes the lengths of shortest paths, only the largest strongly connected component (SCC) was kept. These networks were selected from public repositories such that (a) they fit into these six categories, and (b) have the size of their SCC above 1,000 nodes. The upper bound on network size is simply imposed by finite computing resources.

The following centrality indicators were computed for every node in every network: its *degree*, *neighbourhood* (i.e., the sum of the degrees of the nearest neighbours, previously denoted k_{sum} and found to be a competitive predictor in a previous study⁶), *two-hop neighbourhood* (as before⁶ for nearest neighbours exactly two hops away and previously denoted k_{2sum}), *PageRank*³⁴ with a 0.85 damping factor, *eigenvector centrality*³⁴, *closeness centrality*³⁴, and *core number*⁵. An additional set of indicators that we tried, the *link strength* of a node towards upper, equal, or lower shells⁸, denoted r^u , r^e , or r^l , did not provide notable results.

The ultimate influence of a node in a network is estimated numerically, as the average among 10^4 runs of the susceptible-infectious-recovered (SIR)⁴ diffusion model for infectious diseases. In SIR, an infectious node infects a susceptible neighbour at a *rate* β (meaning the number of infection events per time unit, so can be higher than 1). An infectious node recovers at a rate μ . The effective transmission rate is $\lambda = \beta/\mu$. Here, we take $\mu = 1$ and study the normalized rate λ .

As λ increases in SIR simulations, the size of the outbreaks increase from an infinitesimal fraction to a finite fraction of the network size. The regime of interest is neither very low λ values (in which case, the diffusion remains localised to the neighbourhood of the seed node) nor very high (in which case, all nodes should reach a large fraction of the network). Since our test cases are both finite in size, and diverse (a scenario studied previously³⁹), we estimate the *epidemic threshold* λ_c numerically by identifying it with the variability measure³⁹

$\Delta = \frac{\sqrt{(\rho^2) - (\rho)^2}}{(\rho)}$. Here, ρ denotes the random variable of outbreak size from different seed nodes, and $\langle \cdot \rangle$ denotes the mean. Given a value for λ , Δ is estimated by setting seed nodes from a random sample of 10^4 of the nodes in a network (or the entire network size, if this is smaller). After estimating Δ for a range of λ values at regularly spaced intervals, we take λ_c to be the position of the peak of Δ . The resulting values are noted in Table 1. The maximum spread size (influence) at λ_c in any network is between 0.7% and 6% of the network size (with two exceptions among the smallest infrastructure networks, where this reaches 8% and 11%).

Ranking by a single centrality. *Method.* We first predict superspreaders using the single-centrality ranking method common in prior studies^{5–9}, and also carry forward the performance metrics defined in these studies. This ranking method builds the assumption that higher centrality values for a node will also indicate higher node influence. Given a centrality C , first all the nodes have their values for C computed. The top fraction f of spreaders is then predicted to be the fraction f of nodes with the highest values for C . At ties between nodes (which occur for discrete-valued centralities such as degree and core number) a random subset of the tied nodes are selected. This random sampling is then repeated 10^2 times for a bootstrap technique (described below), which averages among the scores of these individual random choices.

Performance metrics. In prior studies, this ranking is evaluated via two metrics. Denote by I_f the set of the top fraction f of nodes as ranked by their SIR influence, and by C_f the set of top fraction f of nodes as ranked by their centrality values; the sizes of these sets are equal for a given f , $|I_f| = |C_f|$. Also denote by ρ_i the spread size when setting node i as seed. The *recognition rate* $r(f)$ measures the extent to which the identities of the predicted superspreaders match the true identities⁶. A synonym for the recognition rate is *recall*. The *precision function* $p(f)$ is a weaker, but more practically useful performance measure comparing the spread of the predicted superspreaders to that of the true top spreaders:

$$r(f) = \frac{|I_f \cap C_f|}{|I_f|} \quad \text{and} \quad p(f) = \frac{\text{avg}_{i \in C_f} \rho_i}{\text{avg}_{i \in I_f} \rho_i} \quad (1)$$

Both metrics take values in the interval $[0, 1]$. An imprecision function $\epsilon(f)$ was defined previously⁵, such that lower values of $\epsilon(f)$ are better. Here, to present the two metrics in a unified fashion, we use instead $p(f) = 1 - \epsilon(f)$, such that higher values are better for both $r(f)$ and $p(f)$. A confidence interval was originally provided for $r(f)$ by bootstrap⁶. Here, we apply a bootstrap technique when estimating both metrics. Given a

network of N nodes, 10^2 times, we draw a random sample of the N nodes uniformly with replacement. Among these nodes, the ranking method is applied and a prediction is made and evaluated via either $r(f)$ or $p(f)$, as needed. The final value for each performance metric is the average, together with the 95% confidence interval among these samples.

Classification by a combination of centralities. *Method.* A multi-centrality method learns a *discriminative statistical model* able to classify network nodes into superspreaders or not. For this, a dataset is formed for every network; a record describes a node via its centrality values (the predictors). When training the model to recognise the top fraction f of the nodes, the nodes are ranked by their true SIR spread size, and each node is assigned one of two target classes based on whether or not they are in the top fraction f . The model is trained and tuned on a training fraction $t = 0.5$ of the nodes (sampled randomly without replacement), and tested on the remaining nodes.

A binary statistical classifier learns a decision boundary between the classes. We use a *support-vector machine* (SVM)⁴⁰, which learns optimal separating hyperplanes in the multi-dimensional predictor space, including in cases where the classes overlap in this space. Here, the optimal decision boundary is that which leaves the largest margin in space between the classes, with still allowing some data points to fall on the wrong side of the boundary. SVMs have advantages: (a) they are optimal learners rather than heuristics, and (b) the kernel function K and the regularisation parameter C , which ultimately give the shape and variance of the boundary⁴¹, are tunable hyperparameters.

We aim to obtain the simplest, most interpretable classifier with good performance; higher-variance classifiers bring little performance advantages for this problem, and may lose in interpretability. The results presented are for second-degree polynomials K (which gives a low-variance model, less prone to overfitting), C tuned in the range $[1, 100]$ with five-fold cross-validation, and a fixed tolerance for the stopping criterion⁴² of $5e-4$. No class weights are added to balance the classes artificially. (We tested other, higher-variance statistical models: SVMs with third-degree polynomials for K , and nonlinear models based on decision trees, either boosted or in ensembles⁴³; since they had similar performance to the SVM with a second-degree polynomial for kernel, we retain and present the results for the latter.) We show the decision boundaries learnt by two-centrality models via plotting them in the predictor space.

Performance metrics. For a network of size N and the fraction f , a classifier produces a guess for the class of each network node in the test set. We port the same notation C_f to mean here the set of nodes classified as top spreaders. The number of superspreaders predicted in this way is decided by the classifier, and may not equal fN . We measure the overlap between the classifier prediction and the ground truth with metrics similar to Eq. 1. In binary classification, the measure $r(f)$ as defined in Eq. 1 is called *recall* or *sensitivity*. It is a useful metric, but insufficient to characterise the classifier: alongside making many correct choices (giving a high true positive rate, $|I_f \cap C_f|$), the classifier may also add many false positives. The *precision* metric helps to quantify the false positives, and a classical metric is the combination of recall and precision in their harmonic mean, the *F1 score*⁴⁴:

$$\text{recall}(f) = r(f) = \frac{|I_f \cap C_f|}{|I_f|} \quad \text{precision}(f) = \frac{|I_f \cap C_f|}{|C_f|} \quad \text{F1 score}(f) = \frac{2}{\text{recall}(f)^{-1} + \text{precision}(f)^{-1}} \quad (2)$$

Note that *precision* is an established name in the area of Information Retrieval⁴⁴, while the *imprecision function* $\epsilon(f)$ which gave the *precision function* $p(f)$ was defined recently⁵ for analysing networks. Although the names are unfortunately too similar, their meaning is different and should not be confused.

The F1 score takes values in the interval $[0, 1]$. We apply to the classifier the second metric, the precision function $p(f)$, exactly as it is defined in Eq. 1. Its values can exceed 1.0, in cases when the classifier predicts fewer than fN superspreaders, and they are on average better than the true fN superspreaders; we cap higher values to 1.0. We estimate both F1 score and $p(f)$ by randomly drawing different training sets for the classifier (the same training fractions t of the nodes) 10^2 times, then training and testing the classifier on each draw. The final value for each performance metric is the average of the individual scores.

Received: 16 June 2020; Accepted: 9 November 2020

Published online: 25 November 2020

References

- Mariani, M. S. & Lü, L. Network-based ranking in social systems: three challenges. *J. Phys.: Complex.* **1**, 011001 (2020).
- Watts, D. J. & Dodds, P. S. Influentials, networks, and public opinion formation. *J. Consum. Res.* **34**, 441–458 (2007).
- Cha, M., Haddadi, H., Benevenuto, F. & Gummadi, K. P. Measuring user influence in Twitter: the million follower fallacy, in *Fourth International AAAI Conference on Weblogs and Social Media* (2010).
- Anderson, R. M. & May, R. M. Population biology of infectious diseases: part I. *Nature* **280**, 361–367 (1979).
- Kitsak, M. *et al.* Identification of influential spreaders in complex networks. *Nat. Phys.* **6**, 888–893 (2010).
- Pei, S., Muchnik, L., Andrade Jr, J. S., Zheng, Z. & Makse, H. A. Searching for superspreaders of information in real-world social media. *Sci. Rep.* **4**, 5547 (2014).
- De Arruda, G. F. *et al.* Role of centrality for the identification of influential spreaders in complex networks. *Phys. Rev. E* **90**, 032812 (2014).
- Liu, Y., Tang, M., Zhou, T. & Do, Y. Core-like groups result in invalidation of identifying super-spreader by k-shell decomposition. *Sci. Rep.* **5**, 9602 (2015).
- Macdonald, B., Shakarian, P., Howard, N. & Moores, G. Spreaders in the network SIR model: an empirical study. Preprint at <https://arxiv.org/abs/1208.4269> (2012).

10. Lü, L., Zhang, Y.-C., Yeung, C. H. & Zhou, T. Leaders in social networks, the Delicious case. *PLoS ONE* **6**, e21202 (2011).
11. Garas, A., Schweitzer, F. & Havlin, S. A k-shell decomposition method for weighted networks. *New J. Phys.* **14**, 083030 (2012).
12. Chen, D.-B., Gao, H., Lü, L. & Zhou, T. Identifying influential nodes in large-scale directed networks: the role of clustering. *PLoS ONE* **8**, e77455 (2013).
13. Zeng, A. & Zhang, C.-J. Ranking spreaders by decomposing complex networks. *Phys. Lett. A* **377**, 1031–1035 (2013).
14. Liu, J.-G., Ren, Z.-M. & Guo, Q. Ranking the spreading influence in complex networks. *Phys. A: Stat. Mech. Appl.* **392**, 4154–4159 (2013).
15. Chen, D.-B., Xiao, R., Zeng, A. & Zhang, Y.-C. Path diversity improves the identification of influential spreaders. *EPL* **104**, 68006 (2014).
16. Liu, Y., Tang, M., Zhou, T. & Do, Y. Improving the accuracy of the k-shell method by removing redundant links: from a perspective of spreading dynamics. *Sci. Rep.* **5**, 13172 (2015).
17. Liu, Y., Tang, M., Zhou, T. & Do, Y. Identify influential spreaders in complex networks, the role of neighborhood. *Phys. A: Stat. Mech. Appl.* **452**, 289–298 (2016).
18. Radicchi, F. & Castellano, C. Leveraging percolation theory to single out influential spreaders in networks. *Phys. Rev. E* **93**, 062314 (2016).
19. Wang, Z., Du, C., Fan, J. & Xing, Y. Ranking influential nodes in social networks based on node position and neighborhood. *Neurocomputing* **260**, 466–477 (2017).
20. Li, C., Wang, L., Sun, S. & Xia, C. Identification of influential spreaders based on classified neighbors in real-world complex networks. *Appl. Math. Comput.* **320**, 512–523 (2018).
21. Comin, C. H. & da Fontoura Costa, L. Identifying the starting point of a spreading process in complex networks. *Phys. Rev. E* **84**, 056105 (2011).
22. Mo, H., Gao, C. & Deng, Y. Evidential method to identify influential nodes in complex networks. *J. Syst. Eng. Electron.* **26**, 381–387 (2015).
23. Liu, Z., Jiang, C., Wang, J. & Yu, H. The node importance in actual complex networks based on a multi-attribute ranking method. *Knowl.-Based Syst.* **84**, 56–66 (2015).
24. Bian, T., Hu, J. & Deng, Y. Identifying influential nodes in complex networks based on AHP. *Phys. A: Stat. Mech. Appl.* **479**, 422–436 (2017).
25. Rodrigues, F. A., Peron, T., Connaughton, C., Kurths, J. & Moreno, Y. A machine learning approach to predicting dynamical observables from network structure. Preprint at <https://arxiv.org/abs/1910.00544> (2019).
26. Zhao, G., Jia, P., Huang, C., Zhou, A. & Fang, Y. A machine learning based framework for identifying influential nodes in complex networks. *IEEE Access* **8**, 65462–65471 (2020).
27. Fan, C., Zeng, L., Sun, Y. & Liu, Y.-Y. Finding key players in complex networks through deep reinforcement learning. *Nat. Mach. Intell.* **2**, 1–8 (2020).
28. Madotto, A. & Liu, J. Super-spreader identification using meta-centrality. *Sci. Rep.* **6**, 38994 (2016).
29. Ibnoulouafi, A., El Haziti, M. & Cherifi, H. M-Centrality: identifying key nodes based on global position and local degree variation. *J. Stat. Mech.* **2018**, 073407 (2018).
30. Bucur, D. & Holme, P. Beyond ranking nodes: Predicting epidemic outbreak sizes by network centralities. *PLoS Comput. Biol.* **16**, 1–20 (2020). <https://doi.org/10.1371/journal.pcbi.1008052>.
31. Erkol, Ş., Castellano, C. & Radicchi, F. Systematic comparison between methods for the detection of influential spreaders in complex networks. *Sci. Rep.* **9**, 1–11 (2019).
32. Lü, L. *et al.* Vital nodes identification in complex networks. *Phys. Rep.* **650**, 1–63 (2016).
33. Vega-Oliveros, D. A., Gomes, P. S., Milios, E. E. & Berton, L. A multi-centrality index for graph-based keyword extraction. *Inf. Process. Manag.* **56**, 102063 (2019).
34. Newman, M. *Networks* (Oxford University Press, Oxford, 2018).
35. Kunegis, J. KONECT, the Koblenz network collection. <http://konect.uni-koblenz.de/>. Accessed May 2020.
36. Kunegis, J. KONECT: the Koblenz network collection, in *Proceedings of the 22nd International Conference on World Wide Web*, 1343–1350 (2013).
37. Makse, H. Software and data. <https://hmkse.cny.cuny.edu/software-and-data/>. Accessed May 2020.
38. Leskovec, J. & Krevl, A. SNAP Datasets: Stanford large network dataset collection. <http://snap.stanford.edu/data>. Accessed May 2020.
39. Shu, P., Wang, W., Tang, M. & Do, Y. Numerical identification of epidemic thresholds for susceptible-infected-recovered model on finite-size networks. *Chaos Interdiscip. J. Nonlinear Sci.* **25**, 063104 (2015).
40. Ben-Hur, A., Horn, D., Siegelmann, H. T. & Vapnik, V. Support vector clustering. *J. Mach. Learn. Res.* **2**, 125–137 (2001).
41. Hastie, T., Tibshirani, R. & Friedman, J. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. (Springer, New York, NY 2009).
42. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
43. Breiman, L. *et al.* Statistical modeling: the two cultures (with comments and a rejoinder by the author). *Stat. Sci.* **16**, 199–231 (2001).
44. Van Rijsbergen, C. J. *Information Retrieval* (Butterworth-Heinemann, Oxford, 1979).

Author contributions

D.B. is the sole author, and completed all steps of the work.

Competing interests

The author declares no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-020-77536-7>.

Correspondence and requests for materials should be addressed to D.B.

Reprints and permissions information is available at 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) 2020