



A bi-directional approach to comparing the modular structure of networks

Daniel Straulino^{1,2*} , Mattie Landman² and Neave O'Clery^{1,2,3}

*Correspondence:

d.straulino@ucl.ac.uk

¹Center of Advanced Spatial Analysis, University College London, 90 Tottenham Court Road, Bloomsbury, W1T 4TJ London, UK

²Mathematical Institute, Oxford University, Andrew Wiles Building, Radcliffe Observatory Quarter, OX2 6GG Oxford, UK

Full list of author information is available at the end of the article

Abstract

Here we propose a new method to compare the modular structure of a pair of node-aligned networks. The majority of current methods, such as normalized mutual information, compare two node partitions derived from a community detection algorithm yet ignore the respective underlying network topologies. Addressing this gap, our method deploys a community detection quality function to assess the fit of each node partition with respect to the other network's connectivity structure. Specifically, for two networks A and B, we project the node partition of B onto the connectivity structure of A. By evaluating the fit of B's partition relative to A's own partition on network A (using a standard quality function), we quantify how well network A describes the modular structure of B. Repeating this in the other direction, we obtain a two-dimensional distance measure, the bi-directional (BiDir) distance. The advantages of our methodology are three-fold. First, it is adaptable to a wide class of community detection algorithms that seek to optimize an objective function. Second, it takes into account the network structure, specifically the strength of the connections within and between communities, and can thus capture differences between networks with similar partitions but where one of them might have a more defined or robust community structure. Third, it can also identify cases in which dissimilar optimal partitions hide the fact that the underlying community structure of both networks is relatively similar. We illustrate our method for a variety of community detection algorithms, including multi-resolution approaches, and a range of both simulated and real world networks.

Keywords: Networks; Modular structure; Network comparison

1 Introduction

The last thirty years have been extremely fruitful in the study of networks, which enable us to represent and analyse the complex interconnection structure of a wide range of real world and engineered systems [40, 42, 58]. Such is the versatility and usefulness of such a representation, network analysis is increasingly popular in fields as diverse as physics, biology and sociology (see e.g., [1, 12, 25, 26, 57]). While the networks toolbox continues to expand, particularly with respect to dynamics-based approaches [2, 48] and statistical methods [5, 53], there remains a range of methodological challenges related to common applications. Seeking to fill an apparent gap, here we develop a new method to address the

© The Author(s) 2021. 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/>.

simple and widely applicable question of the comparison of the community structure of two networks with an identical set of nodes.

A node partition into communities corresponds to a division of nodes into non-overlapping sets such that each set of nodes is tightly interconnected with few connections to the rest of the network [16]. The presence of such communities is typically linked to the presence of some kind of higher-order organisation in a network, and is often related to functional or societal structure [51]. In many cases, a network may be reduced to a community-based representation which retains the fundamental structural and dynamical features of the original network [14, 34]. There exist a wide variety of methods to detect community structure, as reviewed by Schaub *et al.* [55].

Here we focus on developing a new method to compare the modular structure of networks. This is key to, for example, understanding differences in brain function in two neural networks, or understanding how partisanship has changed over time in the United States Congress. Since the presence of communities in these networks lies at the core of the question of interest, it is natural to use them as the basis of comparison. Perhaps the most widely-used community comparison technique is *normalized mutual information* (NMI) [7]. This and other similar techniques are derived from an information theoretic approach to comparing sets. By focusing exclusively on partitions, and ignoring all other features of the network (for example, the strength of the connections between different communities), these methods lose information. Furthermore, while mutual information is symmetric, there are frequent cases where one network provides useful information about the other but not vice-versa. For example, if the communities of one network are nested inside the communities of the other (meaning that one community of the second network contains several communities of the first) then the second network provides information about the structure of the first but not the other way around.

We propose a methodology that addresses both of these issues, and that is agnostic to the use of community detection algorithm within a large class of commonly used approaches. We call it the *bi-directional distance*, and it is based on a simple idea: we swap the partitions representing the communities of the two networks and evaluate whether the re-assigned communities are a good fit. In the sections below we demonstrate this method for a range of toy and synthetic networks, as well as a real-world example concerning inter-industry labour flows relevant for industrial policy. Labour flow networks illuminate the labour mobility structure of an economy, which is an important factor in the transport of industrial know-how and the localised emergence of new economic activities at the heart of regional economic growth policies. These networks were the motivation for our community comparison methodology. Since their communities have a natural interpretation as industrial clusters [46], understanding the way in which they differ provides valuable insight for industrial diversification strategies.

The paper is structured as follows. We begin by giving a short overview of both the field of community detection and current network comparison methodologies, while also illustrating the shortcomings of current community comparison techniques. We then present our new methodology and the results of experiments on toy and synthetic networks. Second, we demonstrate the flexibility and wide applicability of our method using three distinct community detection algorithms. Third, we apply our methodology to compare the inter-industry labour flow networks of 5 European countries, Germany, Sweden, the

Netherlands, the UK and Ireland. Finally, we extend our methodology to the more complex case of multi-scale community detection.

2 Literature review

2.1 Community detection

Given the diversity of community detection algorithms, it is useful to classify approaches according to the nature of the problem at hand. Schuab *et al.* [55] propose four broad categories of community detection objective: community detection as the minimization of some constraint function, community detection as clustering into densely connected groups, community detection aimed at identifying structurally similar nodes and community detection as the simplified description of dynamical flows on the network. We follow this categorization to briefly introduce some of the most popular community detection algorithms.

In the first category we find cut based approaches which aim to minimize the number or weight of the edges between communities. In other words, the constraint to minimize is calculated as the number of edges that must be deleted (or cut) to achieve the partition, e.g., “ratio-cut” [18]. Graph partitioning is heavily used in parallel computing, in the layout of circuits, and to design serial algorithms to solve partial differential equations [14].

In the second category we find, among others, modularity optimization [43] which is one of the most popular community detection algorithms. A key difference with cut based approaches is that modularity does not a priori set the number of clusters, nor constrains them to be of similar size. Although not without its limitations, this approach has seen widespread use and given rise to a large number of variations and adaptations [14]. It has been used to investigate protein interaction networks, functional brain networks, and ecological networks among other applications.

The structural category groups nodes that exhibit similar connectivity patterns, that is, they are similar if they connect to similar nodes with equal probability. One popular technique for identifying groups in this manner is the stochastic block model (SBM) [22]. This approach is based on the assumption that nodes can be divided into certain classes, and that the likelihood of a connection between two nodes is determined by their respective classes. This gives rise to a probabilistic model for the network such that, via identification of the best fitting parameters, latent groups in the network can be recovered. Originating in the social networks literature, this approach is frequently used to uncover sub-groups in a social network.

The last category deploys dynamics on networks to uncover modular structure [29, 42]. In this category we find Markov stability [8] and InfoMap [54]. Markov stability, for example, exploits random walker dynamics on a network to detect the presence of node communities. If a walker, which jumps from node to node with probability proportional to edge weight, gets ‘trapped’ in a region for a period of time this indicates a cluster of nodes with high internal connectivity. This approach has been applied to dynamical processes on networks such as consensus and synchronisation (e.g., [2, 47]).

In general, as we have seen, different types of application call for different approaches to detecting community structure. Furthermore, no particular methodology consistently outperforms others [52]. Hence, any network comparison technique that focuses on the comparison of modular structure will need to accommodate a variety of underlying community detection techniques.

2.2 Network comparison methods

There is a rich literature in network comparison [11] spanning from global methods (e.g., spectral distances) to more localised approaches (e.g., graph editing). Since we are interested in comparing labelled networks, where each node represents a particular instance (a person, a place, etc.), we focus on methods where there is a one to one correspondence between the nodes of both networks.

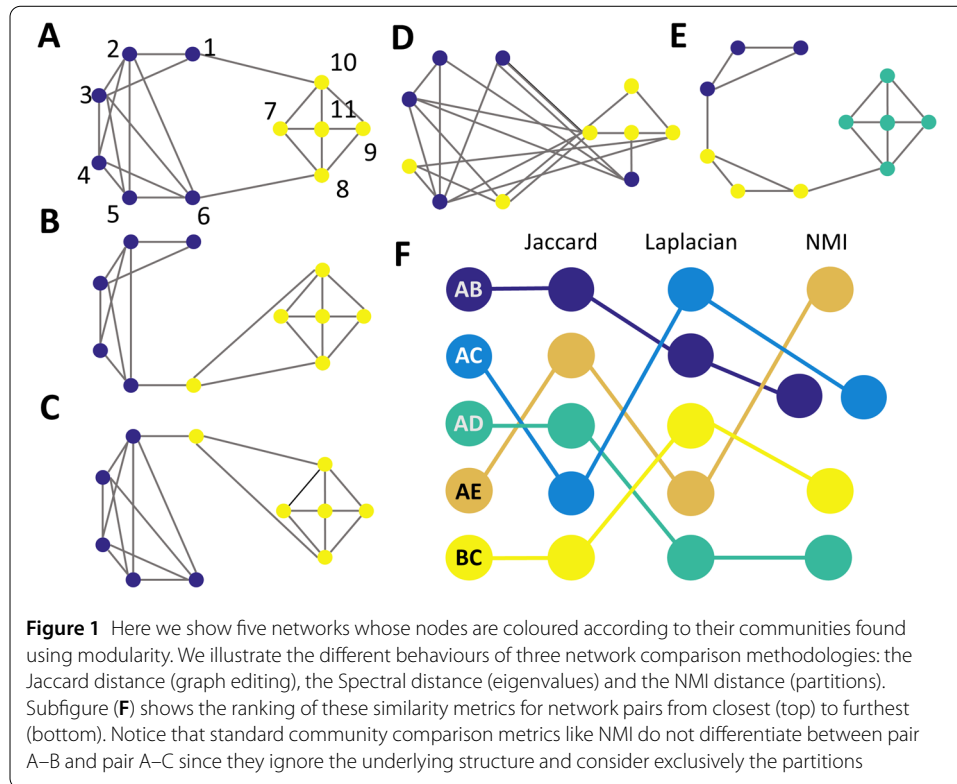
Perhaps the simplest approach is to count the proportion of edges that need to be added and deleted to transform one graph into the other, a measure which originated in the field of information theory and is known as the Hamming distance [19]. While easy to interpret, the Hamming distance is highly sensitive to the density of the graphs, a shortfall that is addressed by the Jaccard distance [23]. The simplicity of these methods has seen them become very popular, not just in the network comparison literature, but also in the fields of information theory and machine learning [11]. Nevertheless, they are very local in scope, looking only at the direct neighbours of each node, with all deletions and insertions given equal weight regardless of changes in properties such as node importance as captured by various centrality measures. Therefore, these methods are not well suited to capture similarities at medium or large scales.

Motif based methods, which compare the frequency of certain sub-graphs (motifs) such as triangles or stars, have proved very popular in applications to protein interaction networks [32]. Despite being relatively local in scope, they are able to capture patterns beyond immediate neighbourhoods. This technique is particularly useful when the motifs represent functional components of the network. However, just as for graph editing distances, these methods fail to capture similarities in large scale connectivity or organisation patterns across networks.

On the other end of the spectrum we have spectral distances. These are global measures based on the eigenvalues of either the adjacency matrix or some version of its Laplacian. Since the eigenvalues are related to topological properties of the network, such as connectivity, they capture global features of the network. Nevertheless, these approaches are permutation invariant, and thus ignore the labels of the nodes which renders them impractical for many applications (consider a network of friendships, the spectral distance does not differentiate between the original network and one obtained by randomly exchanging individuals).

There are other approaches, both global and local, beyond those that we have outlined above but there is comparatively little work comparing networks at a mesoscopic or modular scale. By far the most widely-used community-based comparison technique is *Normalized Mutual Information* (NMI) [7]. NMI, which has been borrowed from the field of information theory, defines the similarity of two partitions as the mutual information of the two partitions, normalized by a combination of the entropy of each partition. It therefore measures how easily one can infer one partition, given the other. Modifications of NMI have been proposed to address some of its limitations, namely that it can do very poorly when the number of communities in two networks differ [44, 49].

Figure 1 illustrates some key network comparison methodologies. We have chosen three popular methodologies: the Jaccard distance, the Spectral distance, and NMI. While the Jaccard distance considers B to be the most similar to A, the Laplacian ranks C as more similar, and NMI indicates that E is the closest network to A. It is clear that these metrics capture different features of the networks. We are interested in comparing the community



structure, which is shown via the colouring of the nodes in the figure. We observe that E is mostly embedded in A and therefore quite similar. This is picked up by NMI. Comparing B and C to A, it appears that C is ‘closer’ since the node that changes community is relatively peripheral to both the blue and yellow clusters. However, since NMI only uses information on partitions—and ignores the strength of community as well as the role of the nodes—it does not differentiate when comparing A to B or C.

Although the above methodologies for comparing networks based on community structure have proven to be useful, the fact remains that there are a number of shortcomings. Perhaps the most prominent issue is that NMI and its modifications rely exclusively on partitions of network, and ignore the quality of these partitions. By reducing the community structure to a partition, these methodologies discard important information about the network [50]. There have been some efforts to address this issue [4], but they rely on correction factors (usually weighting the nodes by their degree). Furthermore, we described earlier how different community detection algorithms differ in their objectives, but partition based methodologies do not account for these differences.

In this paper we propose an alternative framework to comparing networks: partition swapping. This approach presents three marked advantages: it is simple to understand and compute, it incorporates information beyond the partition of the network, and it is flexible, adapting to different community detection algorithms. Its simplicity comes from the underlying idea: for a pair of networks, we assess the fit of each node partition with respect to the other network’s connectivity structure. Since the quality function takes into account the edge weights in order to evaluate the fit, we include information about the network topology. And by changing the quality function we immediately adapt to the application at hand.

2.3 Inter-industry labour flow networks

One of the initial motivations for the development of a new modular network comparison method was to compare the community structure of several industry networks from different countries. The edges of these industry networks are derived from worker data on inter-industry job switches. The extent to which workers jump between industry pairs corresponds to the degree of skill-overlap, a key metric for the ‘relatedness’ between industries which underlies network-based modelling of regional growth paths [15, 36]. This type of network was first introduced by Neffke *et al.* [35] and is known as a skill-relatedness network (SRN) in the literature. A key feature of these networks is their modular structure [46], which constrains regional diversification paths and knowledge diffusion. One might expect that these networks exhibit a near-universal structure, with skill-based industry clusters conserved across countries. On the other hand, differences in structure may illuminate potential growth paths and opportunities unseen in models based solely on data from a single country. Here we investigate these questions, seeking to deploy our new methodology to illuminate both similarities and differences across networks.

Industry networks where nodes are connected based on some kind of skill or capacity overlap have emerged as a key modelling tool in evolutionary economic geography. The key idea is that regions are constrained by the capabilities and know-how embedded in their workforce and industrial structure [15, 20, 38]. In order to diversify and grow, regions tend to move into new industries that share similar or ‘related’ capabilities in an incremental path dependent fashion [21, 35, 36]. In order to model this process, a range of different types of industry networks capturing the ‘relatedness’ between industry pairs have been proposed. These include, amongst others, geographic clustering or co-location of industries as a proxy for general capability overlap [20], occupational similarity as a proxy for labour sharing [13], and collaboration on patents as a proxy for knowledge sharing [24]. Here we focus on ‘skill-relatedness’ which was proposed by Neffke *et al.* [35]. Intuitively, if many workers move from one industry to another, then it is likely that these industries share a high degree of skill-similarity. This metric is relatively less noisy and has more even industry coverage relative to other alternatives.

Network-based models for regional growth and diversification paths broadly work as follows. We can think of the ‘location’ of a region in an industry network in terms of the collection (subgraph) of industries it currently has. Potential new industries share existing skills and capabilities, and correspond to neighbouring nodes in the network. Hence, if a region has many industries on the periphery of the network, it has few related industries into which it is able to diversify. However, if a region has industries which are more central and connected, it has many more potential diversification paths. Skill-relatedness networks have been applied to model and predict the diversification paths of regions in a large number of contexts [35, 37, 46]. They have also been used in various other applications, such as models for urban formality growth [45, 47], employment resilience and adaptability [10] and knowledge spillovers from MNEs [6].

Recently, O’Clery *et al.* [46] showed that these networks have a distinct modular structure. Hence, they contain communities (i.e., groups of industries) within which workers switch much more easily (relative to switches between industries in different communities). The authors refer to these communities as ‘skills-basins’ in reference to the degree of skill and knowledge sharing within industry groupings. The presence of modular structure is a key topological feature of these networks that constrains the industrial diversification

opportunities of regions, and can be used to improve network-based models for regional diversification and growth patterns.

In this study, we are interested in comparing the modular structure of various countries' SRNs. Within the related diversification literature, the SRN has been assumed to be identical across both space and time. Neffke *et al.* [37] showed that there was little variation between the SRN when constructed for East and West of Germany, as well as for Germany compared to Sweden [35]. The authors estimated the similarity between the two networks by calculating the Spearman correlation between edges. Based on these results, various studies have used the SRN of a different country within their analysis when the SRN of the country under consideration was not available.

We believe, however, that these networks may be different, particularly on a mesoscopic-scale (i.e., with respect to their modular structure). As an SRN is constructed from all inter-industry labour flows within a country, the network reflects intricacies of the structure of the local labour market. We expect that the degree of inter-industry labour flows is highly dependent on the historical economic progress of a country and its institutional labour market structures. We are interested in both similarities and differences in the modular structure of these networks. Similarities uncover universal structure in inter-industry skill-sharing, and portability of networks across contexts. Differences may provide insight into hidden growth opportunities: linkages and clustering patterns present in one context may suggest potential unseen paths in another.

3 Results

3.1 A bi-directional distance metric

A large class of community detection algorithms are based on optimizing an objective function F that measures the goodness of fit of a partition according to some desired property, whether structural (for example modularity from [41]), dynamic (see [8]) or other (see [54]). We propose to compare the modular structure of two networks, say A and B, by computing the ratio of A's F -score under B's optimal partition to its F -score under its own optimal partition, and vice versa. In a sense, our two dimensional method tells us how well network A describes modular structure of B as well as the other way around.

More formally, let $A, B \in M^{n \times n}$ be the adjacency matrices of two networks of size n , where there is a one to one correspondence of the nodes in each network given by the identity function. In other words, the labelling of the nodes matter and we assume that the labels of both networks are identical. Let \mathcal{P} be the space of partitions of $[n]$, the set of integers up to n . Without loss of generality, let $F : M^{n \times n} \times \mathcal{P} \mapsto [0, 1]$ be an objective function.¹ We then denote

$$P_A := \arg \max_{P \in \mathcal{P}} F(A, P),$$

$$P_B := \arg \max_{P \in \mathcal{P}} F(B, P),$$

¹Any bounded quality function can be linearly transformed to match this range. For unbounded functions, it is sufficient to shift them to $[0, \infty)$ as the ratio will remain in $[0, 1]$

the partitions that maximize the objective function for each network. We are now ready to define our (non-symmetric) distance score:

$$d(A, B) = 1 - \frac{F(A, P_B)}{F(A, P_A)}. \quad (1)$$

The ratio on the right tells us how well network A describes modular structure of B—a high value indicates that the optimal partition of B is indeed also a good partition of network A. By construction, $d(A, B)$ ranges over $[0, 1]$, and is 0 if and only if $F(A, P_A) = F(A, P_B)$. By swapping A and B we can obtain a second distance score. In this way, for every pair of networks, we have a pair of distance scores that reflect how well their respective partitions capture each others' community structure. We propose a two dimensional distance, the partition swapping score *BiDir*:

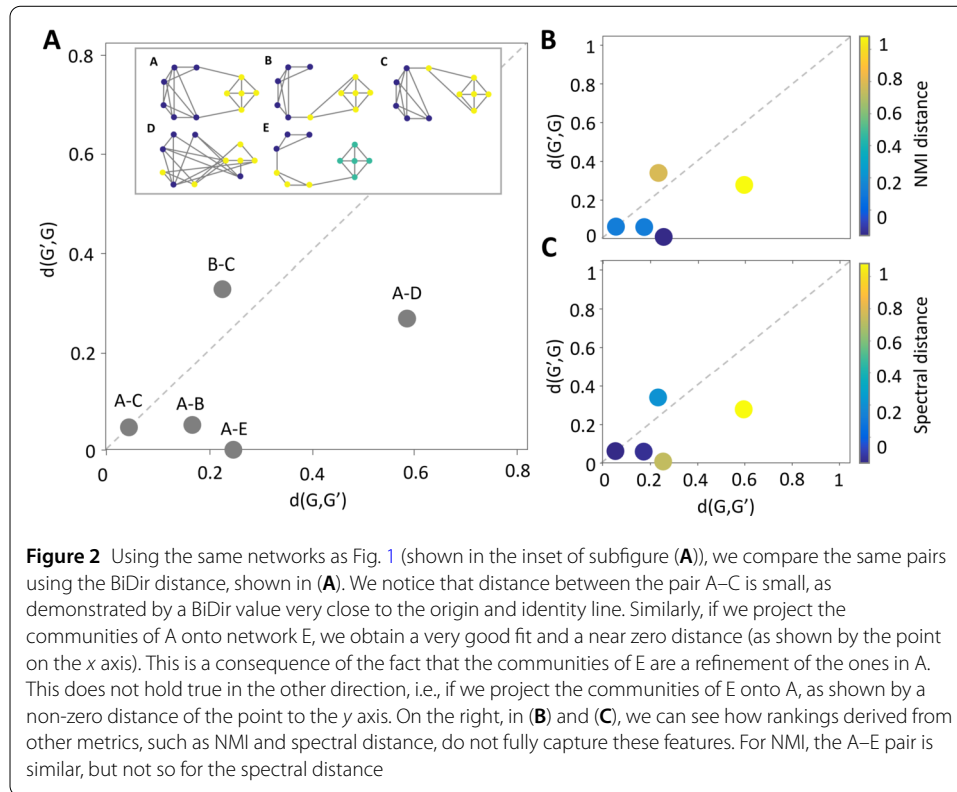
$$\text{BiDir}(A, B) = (d(A, B), d(B, A)). \quad (2)$$

Hence, for any pair of networks A and B, we compute a two dimensional or bi-directional distance score in order to compare their modular structure.

When we calculate each of the distance coordinates, we compare the value of the quality function on an alternative partition to its value on the optimal partition. This ratio will therefore depend on the shape of the quality function but notice that both the numerator and the denominator are calculated on the same network. So even if networks A and B have very different densities (which in some cases might affect the range of the values that the quality function takes), $F(A, P_B)$ and $F(A, P_A)$ depend only on network A and can thus be compared. Nevertheless, notice that if $F(A, P_A)$ is very close to 0, indicating that network A is not modular, the ratio will not be very informative, as arbitrary partitions will produce similar scores. Therefore, our framework will be best suited to networks with relative high scores of the quality function. Furthermore, it is necessary that the quality function properly discriminates good partitions from bad ones, and thus the cases in which it fails to do so will result in non-informative *BiDir* scores.

This method has three distinct advantages over previous community or partition comparison methods. First, it can be adapted to any community detection algorithm that is based on optimizing an objective function. Second, it goes beyond simply comparing the resulting partition arising from a community detection algorithm but accounts for the underlying network structure, specifically the strength of the connections within and between communities. This means, for example, that our method can capture differences between networks with similar partitions but where one of them might have a more defined or robust community structure. Third, it can also identify cases in which very dissimilar partitions hide the fact that the community structure of both networks is relatively similar. A network might be well suited to multiple distinct partitions—for example, when there are nested communities—all of which will score highly in the objective function. Unlike other methods, our framework will pick up this signal.

In Fig. 2 we illustrate the key features of the *BiDir* distance using the toy networks from Fig. 1. We use modularity optimization as the community detection algorithm. As outlined above, we obtain not one but two values for each pair of networks, which are plotted against each other in two-dimensional space in subfigure (A).



The first thing to notice is that the two components of our distance convey distinct pieces of information. Points that are positioned away from identity line (diagonal) correspond to asymmetric behaviour, whereby one of the swapped partitions scores higher than the other. We can interpret this a case where the optimal partition of the first network is well-described by the modular structure of the second. The optimal partition of the second network, however, is not as compatible with the first. For example, we can see that the point corresponding to pair A–E is very close to the x axis. This indicates that if we project the communities of A onto network E, we obtain a very good fit. In fact, since the distance is near zero, the partition found on E is virtually as good a description of A’s community structure as the communities found on A. Looking at the other direction, whereby we project the communities of E onto network A, do we not find as good a fit. By eyeballing the relevant networks, shown in the inset, we notice that the communities of E are a refinement of the communities of A, and thus we can think of them as being nested in the communities of A. This explains why the distance $d(E, A)$ is so small.

The added information provided by the two components of the BiDir distance comes at a price. While single value distances are relatively easy to interpret, the two-dimensional score we obtain requires one to carefully consider the directionality of each dimension. Therefore, while BiDir is well suited to investigate hierarchical or nested similarities, it does not provide a unique way of ranking the networks in terms of similarity. Different choices on how to collapse the two dimensions into one, i.e. geometric or arithmetic mean, will produce different results, and the application of interest should guide this choice.

The second key observation is that our approach implicitly takes into account the network’s connectivity patterns, and not just the partition, by making use of the objective

or quality function. Notice that, unlike NMI (or like any other purely partition-based methodology), we differentiate between pairs A–B and A–C.

They both are equally distant from the x axis, and so the optimal partition of A, when projected onto B or C, results in more or less the same modularity score. The community structure of C, however, scores better than B when projected onto network A. Again, looking at two networks we can understand why this is the case.

In the former case, when we project the partition of A onto network B and C, the partition misplaces a single node (relative to the optimal partition of B and C). In the latter case, when we project the partition of B and C onto network A, we see that while both the partitions derived from B and C differ from A by a single node in the first case a node that is central to the blue community in A changes groups, while in the second it is the most peripheral node that does. Hence, the distance between A and C is smaller.

This simple example shows how BiDir, by exploiting the objective function of the community detection algorithm, captures information on connectivity structure underlying the community structure of both networks, and is therefore capable of distinguishing cases which NMI and related methodologies would not be able to discriminate.

At the same time, this dependence on the objective function comes with some trade-offs with respect to metrics that are agnostic on the community detection algorithm. While NMI and other similar metrics can be used to compare the performance of different community detection algorithms, by comparing them on benchmark datasets, BiDir uses a different function to calculate the distance for each community detection algorithm, and thus results for each algorithm are not directly comparable.

While the motivation for the BiDir distance was investigating the differences between skill-relatedness networks, the framework can be applied more generally. The key assumptions of our method are:

- The networks must be node-aligned (this implies that both networks have the same size).
- For each network we have a partition that represents its modular structure.
- This partition is optimal for a given quality function.²

We emphasize that by modular structure we mean that it can be divided into subgraphs, where each subgraph is more tightly connected within itself than it is connected to the rest of the graph. Different community detection algorithms operationalize this concept in different ways, which is why they all define different quality functions. The quality function itself can add some restrictions to the applicability of our framework, for example the Markov stability algorithm does not work on disconnected graphs, modularity doesn't work on directed networks. But as long as the networks being compared are node aligned and feature modular structure, our framework can be applied.

3.2 Synthetic networks

Having introduced the BiDir distance, we will now construct several families of networks drawn from a stochastic block model (SBM) to further illustrate its features and test its performance. Specifically, we probe the effect of the strength of the community structure, and nestedness and overlap of communities, on the BiDir distance.

²While many quality functions are optimized with greedy algorithms that might not identify the optimal partition, in practice it is enough that the partition used is close to optimal.

Stochastic block models provide a good generating model for networks with a clearly defined community structure. First introduced to study the structure of social networks [22], they have been widely adopted because they provide a straightforward way of generating modular networks as well as a simple benchmark for testing the statistical significance of communities. In a classical stochastic block model, n nodes are partitioned into a set of k communities $C = \{C_1, \dots, C_k\}$. A matrix P of $k \times k$ encodes inter-community connection probabilities (the probability that nodes $v \in C_i$ and $u \in C_j$ share an edge is given by P_{ij}). Hence, we can directly specify the number of communities as well as the likelihood of the connection within and across communities.

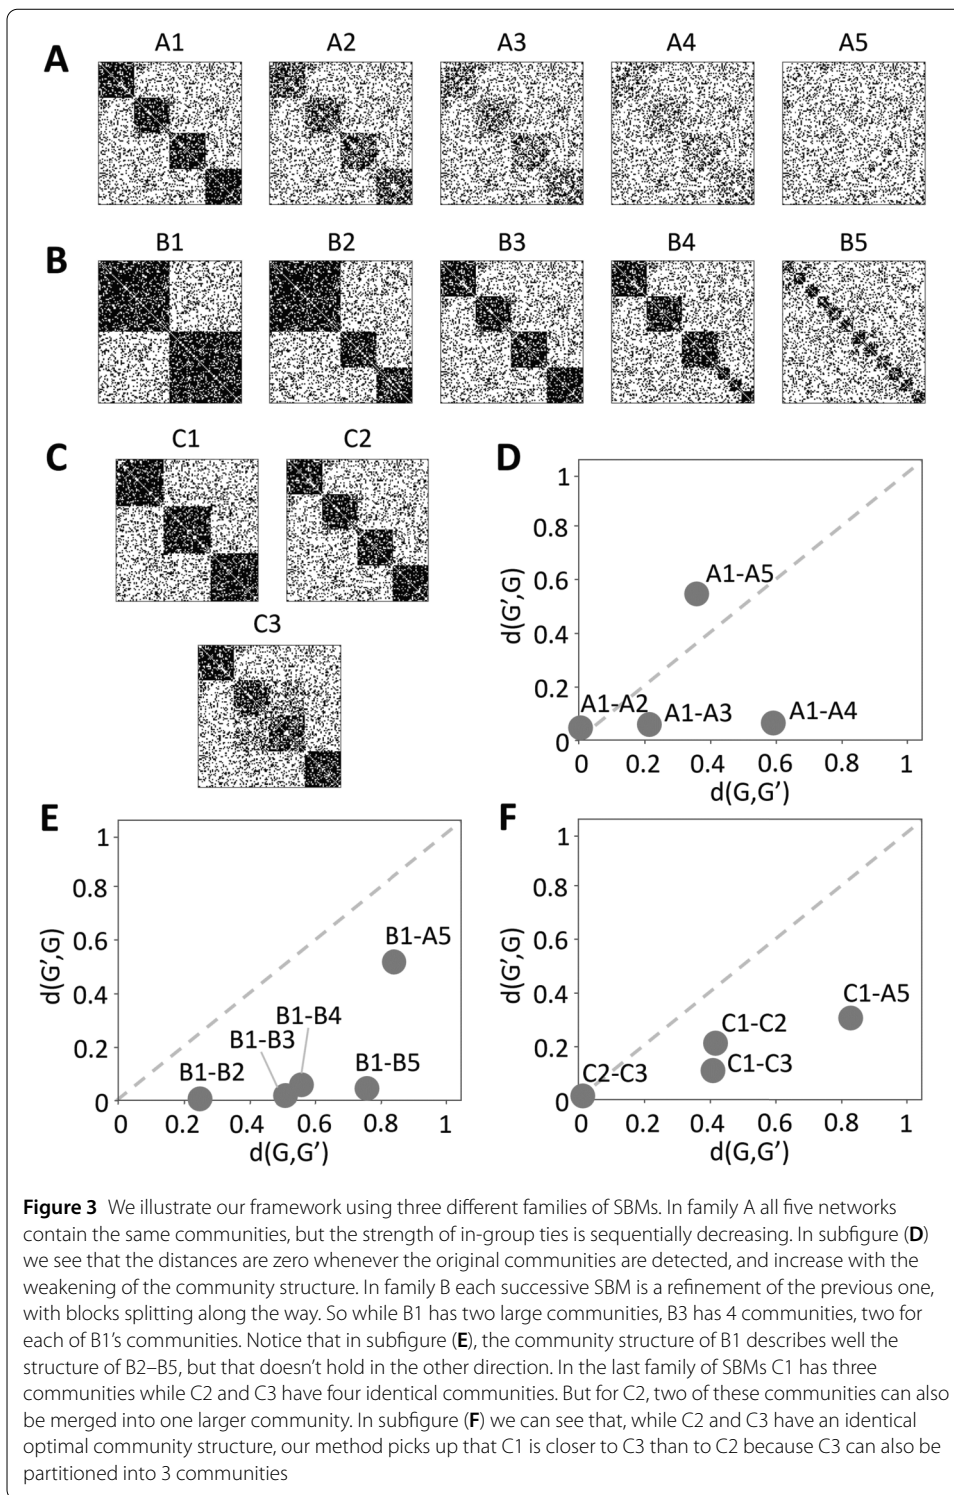
It is possible to use SBMs to generate networks that exhibit a nested community structure. In this case, smaller communities are nested inside large communities. This is type of a structure that we observe in many real world networks. For example consider scientific collaboration networks: at a small scale we might identify individuals that belong or belonged to the same institution or research group, but these groups might also work together within a country or region, creating two levels of modular structure.

We generate three families of SBM, as illustrated in Fig. 3. Full details on the parameters used to generate these ensembles is provided in the SI. As before, modularity is used as the community detection algorithm. The distances shown correspond average BiDir distances over 1000 instances of each SBM.

The first family has four SBMs, all with the same community structure but with successively decreasing probability of in-group ties. As long as the communities can be recovered, their structure remains intact and thus the distance between them will be 0. Once the communities can no longer be properly identified by our community detection algorithm, the distances increase. We see in subfigure (D) that when we project the community structure of A1 onto A2, A3 and A4, we obtain values close to zero (i.e., the points lie along the x axis). As expected, A1's community structure is very close to being optimal for A2, A3 and A4 as well. However, if we instead project the community structure detected for A2, A3 and A4 onto A1, the distance progressively increases as the community detection algorithm struggles to find the underlying or generating structure. The final SBM A5, where inter and intra community connection probabilities are equal, does not have a community structure and is thus most distant from A1.

The second family is composed of five SBMs, each a refinement of the previous one. In other words, B1 has two communities, B2 has three communities, obtained by splitting one of the communities of B1 into two, B3 has four communities and so on. We compare the first SBM B1 to the other four. We observe, as we would expect, that as we split the communities further, we increase the distance from the original network. More specifically, if we project the community structure of B1 onto B2–B5, we again obtain values close to zero (i.e., the points lie along the x axis). Hence the original partition B1 remains informative even after a few splits since the new communities are nested inside the original one. Looking in the other direction, however, the more dis-aggregate partitions describe progressively less well the structure of B1.

Our last family of SBMs is slightly more complicated. C1 has three communities, but none of them maps onto the communities in C2 or C3. C2 and C3 are both split into four identical communities, but for C3 two of these communities are also strongly connected between them, and thus C3 can also be reasonably partitioned into three communities. In subfigure (F) we can see that the $D(C2, C3) = (0, 0)$ —since the partition of C3 into four



communities has higher modularity than the partition into three, C2 and C3 the same optimal community structure. But when comparing them to C1, we see a different picture. C3 is closer to C1 because the middle community of C1 is close to the combined middle block of C3. C2 is further from C1 as the middle community of C1 overlaps two very distinct blocks in C2. We see that BiDir distance rightly identifies that although C2 and

C3 have the same optimal partitions, their underlying community structure is markedly distinct.

Hence, even when the optimal partitions of two networks appear to exhibit little overlap, our method will generate a low distance score when there exists an alternative and almost optimal partition that is similar. It is well known [14] that an exhaustive optimization of modularity is impossible due to the high number of possible partitions, and furthermore, very different partitions might obtain similar near optimal scores. Therefore looking exclusively at the resulting partitions is likely to miss important structural similarities that are nevertheless captured by the BiDir distance whenever these two distinct partitions produce similar scores for the optimization function.

3.3 Application to inter-industry labour flow networks

In this section, we compare the modular structure of five European Skill-Relatedness Networks (SRN) using the BiDir distance. Recall that, although it has previously been argued that the SRN is universal, we expect to find variation across the modular structure of different countries' SRN. Understanding this variation will unveil how structural characteristics of labour markets impact on industrial diversification paths.

An SRN is constructed as follows. Two industries are considered skill-related if there is a larger number of job switches between them compared to what would be expected at random. Formally, if ϕ_{ij} is the number of job switches between industry i and j within a given time period, then the skill-relatedness is given by

$$S_{ij} = \frac{\phi_{ij} / \sum_k \phi_{ik}}{\sum_l \phi_{lj} / \sum_{lk} \phi_{lk}}. \quad (3)$$

The matrix is then made symmetric by averaging with its transpose and re-scaling the values so that they range from -1 to 1 to give the adjacency matrix:

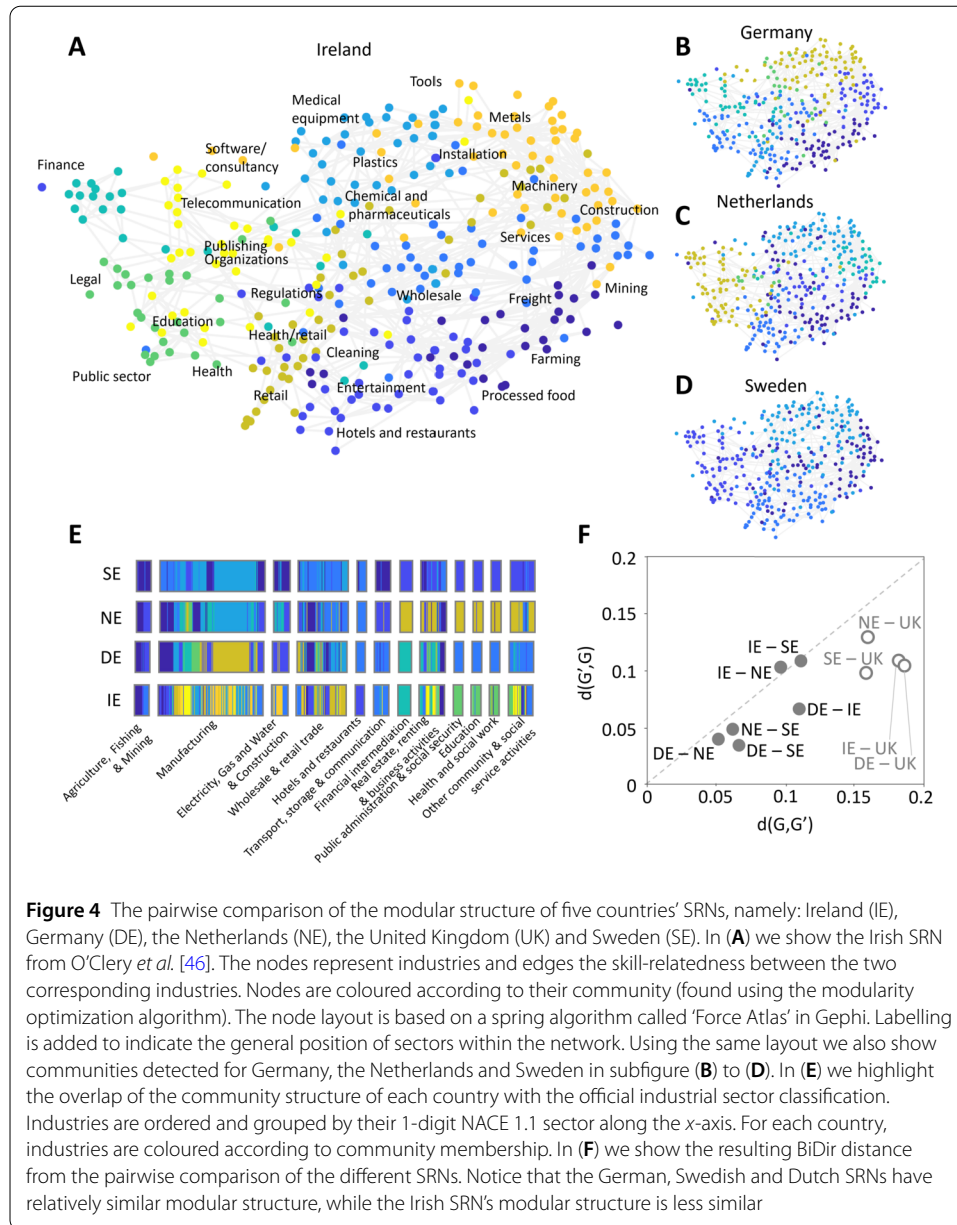
$$A_{ij} = \frac{S_{ij} + S_{ji} - 2}{S_{ij} + S_{ji} + 2}. \quad (4)$$

Note that values that are greater than 0 indicate that the number of job switches is greater than what would be expected at random under a null model (specifically, the configuration model [33]). We therefore conserve only positive values of this matrix. The reader is referred to Neffke *et al.* [35] for a detailed discussion of these methodological considerations.

Here we compare the SRNs for Germany [37], Ireland [46], the Netherlands [10], the United Kingdom and Sweden [35]. In all cases, the industry classification corresponds to the 4-digit NACE 1.1 industry classification.³ The graph intercept of these networks (with the exception of the UK, see footnote below) is used to ensure that all networks have the same size and are node-aligned—a requirement of the BiDir metric.

In Fig. 4(A) we show a visualisation of the Irish SRN from O'Clery *et al.* [46]. Each node in the network represents an industry and each edge the skill-relatedness between its two corresponding industries. The node layout is based on a spring algorithm called 'Force

³Due to data availability, the United Kingdom SRN was constructed using the 4-digit NACE 2 classification and then converted to the NACE 1.1 classification. Readers are referred to the SI for a more detailed discussion regarding this conversion.



'Atlas' in Gephi, where more skill-related industries are positioned closer together. We have added labelling to indicate the general position of sectors in the network. We observe that service-orientated industries and government activities tend to be located on the left hand side of the network, whereas heavy goods, construction, manufacturing and agricultural sectors dominate the right. Retail (bottom) and business (top) activities lie in between.

The modular structure of the network is shown via the node colouring. The communities were detected using modularity [41]. Specifically, the partition maximizing the modularity function from 10,000 iterations of the Louvain algorithm is adopted. Using the same layout, we visualise the communities found for Germany, the Netherlands and Sweden on the right.

Figure 4(E) highlights the overlap of the community structure of each country with the official industrial sector classification. Industries are ordered and grouped by their 1-digit

NACE 1.1 sector along the x -axis. For each country, industries are coloured according to community membership, corresponding to the node colouring in subfigure (A) to (D). Hence, blocks of colour within a sector indicates a close correspondence between a community and a sector. Conversely, a diversity of colours within a sector indicates a community structure that is quite distinct from the official classification.

This figure provides insight into how the modular structure varies across the different SRNs. First, we observe that the Irish SRN has more communities than the other 3 networks. The Irish SRN has 9 communities, the Netherlands and Germany have 7 communities, while Sweden only has 5 communities. We can also see that industries within certain sectors are always clustered together across all four networks (as indicated by solid colour blocks), such as financial intermediation, public administration and social security, agriculture and the hotel and restaurant sectors. These communities show universal structure in inter-industry skill-sharing. Furthermore, insights such as the increased subdivision of the manufacturing sector in Ireland compared to other SRNs can also be observed.

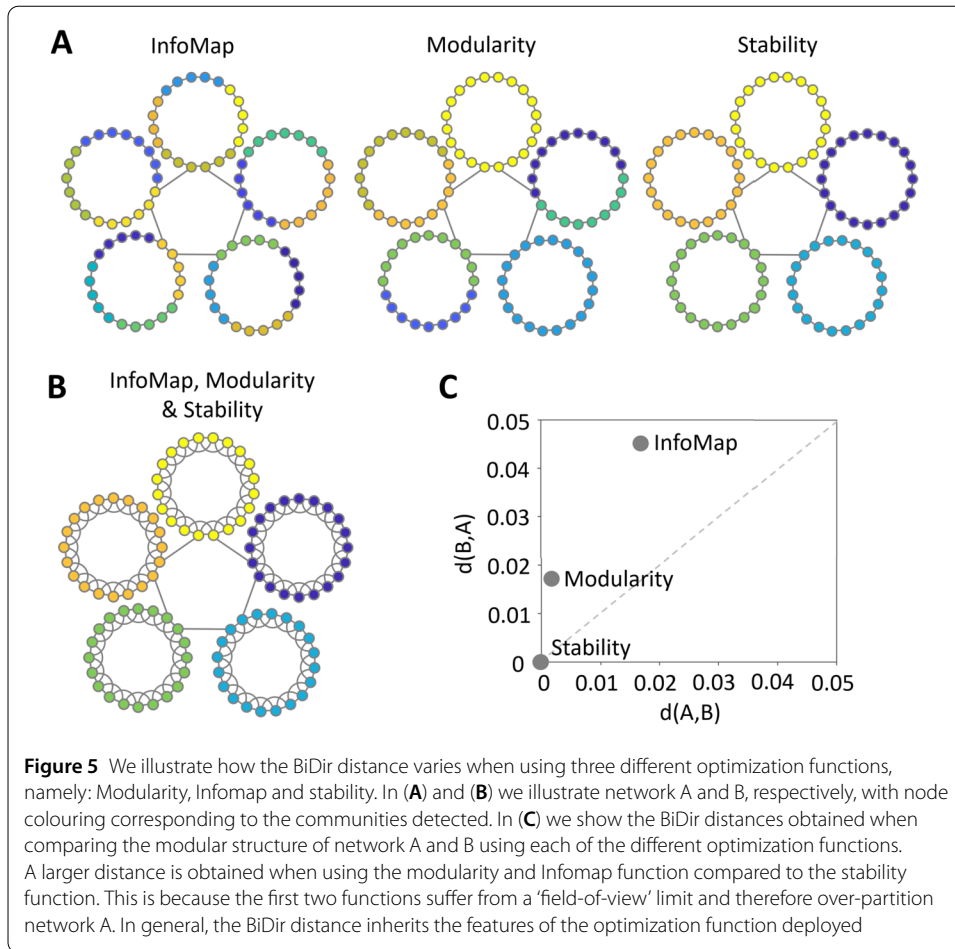
We now quantify the distance between these networks using our BiDir distance. The pairwise comparison of the various SRNs is shown in Fig. 4(F).⁴ We observe that Germany and the Netherlands are the closest in terms of modular structure with the smallest distance, and both are similar to Sweden. Ireland appears to have a very different modular structure compared to all other countries. Comparison of the Irish SRN and the German SRN uncover an asymmetric relationship. If we project the German communities onto the Irish SRN, we obtain a good fit—but not vice versa. Visual inspection of subfigures (A), (B) and (E) confirms that the communities of Ireland are somewhat nested inside the communities of Germany. For example, the manufacturing and social service sectors consists of many small communities in Ireland compared to just a few larger communities in Germany. Hence, inter-industry flows, and thus knowledge diffusion and skill-sharing that constrains development paths, are more fragmented in Ireland. From an Irish policy perspective, the German SRN might be informative in the design of smart interventions. For example, policies could be developed to facilitate and encourage worker mobility between sectors in, for example, manufacturing that are highly interconnected in Germany, but not yet in Ireland.

While it is generally accepted [37] that the skill-relatedness networks vary little across space and time, our analysis shows that there are important differences across the European countries we analyse. While the modular structures of the Northern countries (Germany, the Netherlands and Sweden) are relatively similar, those of Ireland and the UK are less so. Furthermore, the directionality of our distance metric enables us to pick up nuanced differences between SRNs (specifically in terms of nested community structure, like in the case of Germany and Ireland), and uncover unseen potential linkages, which are key to policy efforts to generate regional industrial growth and diversification potential.

3.4 Different optimization functions

The flexibility of BiDir, in terms of its adaptability to a class of community detection algorithms (with the use of an objective function), is advantageous in that it facilitates the

⁴The United Kingdom's SRN has fewer industries compared to the other four countries, and also much more sparse. This is because it was constructed from a longitudinal survey including just 1 per cent of UK workers. Hence, when we compare the UK to the other countries, we compare the subgraphs induced by the set of overlapping nodes (and corresponding edges) in each case. A more detailed comparison of the UK to the other countries is included in the SI.



comparison of networks across a wide variety of applications. This is because different optimization functions define communities differently and will therefore often detect different partitions even on a single network. Thus far, we have illustrated the use of BiDir when using the modularity optimization function [39] as our objective function. Here, we illustrate BiDir distances obtained when comparing the modular structure of two networks using three different optimization functions. Specifically, we consider InfoMap [54], Markov stability [8] and modularity functions. Specifically, we illustrate the degree to which properties of an optimization function are inherited by and reflected in our distance measure.

We compare the modular structure of networks A and B illustrated in Fig. 5(A) and (B), respectively. Network A is a ring-of-rings and is an example of a network with non-clique-like communities (communities that have a high average diameter). The network is taken from Schaub *et al.* [56]. Network B is a ring-of-circulant sub-graphs $C_n(1, 2)$.⁵ It has a similar community structure as network A, however the clique-like structure of each community is enhanced (the average diameter of each community is decreased) by the addition of intra-community edges.

For modularity and InfoMap the partition is obtained by applying the Louvain- and internal InfoMap heuristic 1000 times respectively, and choosing the partition which either

⁵A circulant graph is a graph of order n in which the i^{th} vertex is adjacent to the $(i+j)^{\text{th}}$ and $(i-j)^{\text{th}}$ graph vertices for each j in a so-called connection set. A circulant graph of order n and with connection set $\{1, 2\}$ is denoted as $C_n(1, 2)$.

maximises the modularity or minimises the InfoMap function. Similarly, in the case of the multi-resolution stability algorithm (where time is the intrinsic resolution parameter), for each time a resulting partition is obtained by applying the Louvain heuristic 1000 times. For each time resolution, we identify the partition that maximizes the stability function. In order to choose the partition corresponding to the optimal resolution, we calculate the mean variation of information⁶ (VI) [31] at each resolution, and choose the partition corresponding to the resolution with lowest VI. We consider the full range of partitions obtained at different resolutions in the next section.

The resulting partitions are visualised via node colouring. As noted by Schaub *et al.* [56], we observe that both InfoMap and modularity over-partition network A (each ring is partitioned into multiple communities). This can be explained by the ‘field-of-view’ limit of which both of these functions suffer [17]. The ‘field-of-view limit’ is an upper limit on the effective diameter of the communities that can be detected with one-step dynamical community detection techniques [56]. As the rings in network A have high diameters (non-clique like structure) both these algorithms over partition the rings, creating communities with smaller diameters. However, this over-partitioning is not observed in the case of network B. This is because each of the circulants in network B have a smaller diameters and the optimization functions are able to partition them into their own community. As the Markov stability function is a multi-step dynamical community detection algorithm it does not suffer from the field-of-view limit [56], and obtains a community structure consistent with what we would expect for both network A and B.

In Fig. 5(C), we illustrate the BiDir distances (comparing network A and B) obtained when adopting each of these optimization functions. As expected, Infomap and modularity result in the largest distances. Both of these distances are also highly asymmetric. We find that if we project the communities of B onto network A, we obtain a good fit—but not vice versa. The BiDir distance therefore recognises the subdivided communities found using modularity and Infomap. As InfoMap results in the most severe over-partitioning of the network, it results in the largest and most asymmetric distance. As the Markov stability algorithm does not suffer from the ‘field-of-view limit’ it obtains the same community structure for network A and B, and results in a distance of (0,0). The different distances obtained using the different optimization functions illustrates that the BiDir distance inherits the properties of the community detection algorithm used.

3.5 Multi-scale optimization functions

Although most well-known community detection algorithms seek to obtain a single node partition, in many cases it is more natural to analyse a range of partitions—from many small communities to a few large communities—when investigating the modular structure of a network. Such hierarchical structure can be informative, revealing layers of organisation. For example, consider a university friendship network. Larger communities may reveal institutional structures such as departments or colleges, while smaller communities might capture friendship or social circles.

⁶The variation in information is a metric used to measure the robustness of our partitions. The variance of information calculates the amount of information (in an information-theoretical sense) two partitions share. Two partitions that are similar will exhibit a low variation of information. For each resolution a pairwise variance of information is calculated for each partition obtained from the Louvain algorithm. The average of these is then the mean variation of information at the resolution. If this value is low, it shows that the detected partitions are similar, and therefore more robust [31].

Hence, when comparing the modular structure of two networks, one may want to consider not a single partition but a range of partitions for each network. This approach can reveal, for example, distinct distances corresponding to particular scales, and consequently uncover the resolution at which the two networks are most similar. For example, two networks may be quite similar in terms of their modular structure at a coarse scale, but distinct at a finer or more dis-aggregate scale. Taking further advantage of the flexibility of our method, here we illustrate the use of the BiDir distance for a multi-resolution optimization function.

To illustrate how we adapt the BiDir distance to this case, again consider two networks, A and B. We compute the ratio of A's F -score under B's optimal partition found at resolution β to its F -score under its own optimal partition found at resolution α , and vice versa. Formally,

$$d(A_\alpha, B_\beta) = 1 - \frac{F(A, P_{B_\beta})}{F(A, P_{A_\alpha})}, \quad (5)$$

where P_{A_α} is the optimal partition of network A obtained at resolution α and P_{B_β} is the optimal partition of network B obtained at resolution β . Hence, the final BiDir distance corresponds to:

$$D(A_\alpha, B_\beta) = (d(A_\alpha, B_\beta), d(B_\beta, A_\alpha)). \quad (6)$$

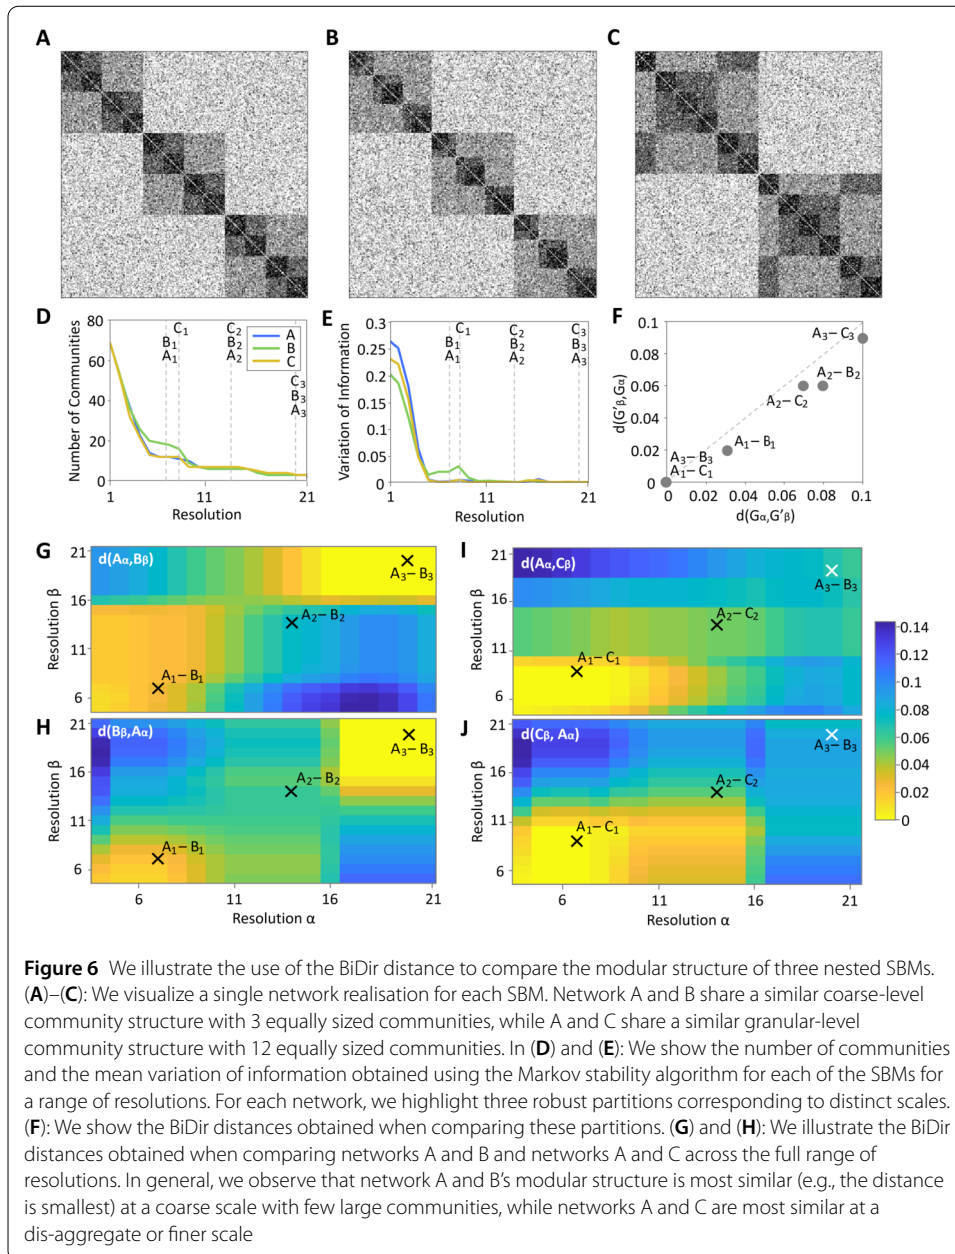
In practice, we calculate this distance for a range of α and β values such that small values of α or β correspond to fine partitions with many small communities, while larger values correspond to fewer larger communities.⁷

To illustrate this approach, we compare three families of networks generated by hierarchical SBMs that share modular structure at three different scales. All three SBMs consist of 300 nodes (see SI for more details). In Fig. 6(A)–(C), we visualise the adjacency matrix for a single network generated by each SBM. Note that A and B share a coarse-level modular structure of three equally size communities. However, at a more granular-level, their community partitions are more dissimilar. On the other hand, network A and C share a granular-level modular structure of 12 equally sized communities.

Next, we use the multi-resolution Markov stability community detection algorithm to detect a range of partitions for each of the networks. Recall, that this algorithm is based on a simple random walk model [8, 27, 28]. The key idea behind this method is that it sets a walker to roam on a network—jumping from node to node with probability proportional to the edge weight. If the walker gets trapped in a region of the network (a group of nodes) for a prolonged period this corresponds to a group of densely connected nodes which form a community. Here, the Markov ‘time’ is used as the resolution parameter. Intuitively, if we let a walker roam for longer periods on the network, the walker will detect larger and larger communities. Therefore, by varying the Markov times we detect communities on a range of scales: from many small communities to few large communities. More detail on this algorithm can be found in Lambiotte *et al.* [28].

We adopt the Louvain algorithm [3] as our search algorithm to find the optimal node partition with respect to our objective function (the stability function [8]) at each resolu-

⁷The range of α and β is network specific and chosen by the user.



tion. We use variance of information (across a large number of realisations of the Louvain algorithm) to assess which of these partitions is the most robust. In subfigure (D) and (E), we display the number of communities and the mean variation of information for each resolution.⁸ We highlight three resolutions that are both robust, with low variation of information, and correspond to three distinct scales of community structure.

First, we use the BiDir distance to compare these chosen partitions in the standard manner, as shown in (F). We clearly observe that the distance varies depending on the resolutions chosen. For example, partitions A_3 and B_3 , corresponding to larger scale coarse

⁸Note that for all results, we calculate the average values obtained from generating 100,000 network instances of each SBM and running the Markov stability algorithm with 1000 iterations of the Louvain heuristic on each network instance generated.

partitions of both A and B, are much more similar than A_1 and B_1 (fine) or A_2 and B_1 (moderate).

Next, we compare the modular structures of network A and B across the full range of resolutions. Figure 6 ((G) and (H)) shows values obtained for different (α, β) resolution pairs. As our distance is two-dimensional, (G) corresponds to direction $d(A(\alpha), B(\beta))$ and (H) corresponds to direction $d(B(\beta), A(\alpha))$. Consistent with our example above, we observe that A and B's modular structure is most similar (e.g., the distance is smallest) at a coarse scale with few large communities. As we increase the granularity of the community structure, the modular structure of the networks becomes more distinct (at a moderate scale) and then again more similar (at a fine scale).

We also observe an asymmetric relationship with respect to modular distance. For small α and moderate β we observe that $d(A(\alpha), B(\beta)) < d(B(\beta), A(\alpha))$. Hence, if we project the communities of network B onto network A, we obtain a good fit—but not vice versa. In this region network B has fewer and larger communities (9 communities) than network A (which has 12 communities). Thus, our metric picks up that the stability of both partitions are similar with regard to network A.

Conversely, when we compare networks A and C in (I and J), we observe that the smallest distance is obtained at a fine resolution corresponding to detection of smaller communities. As the resolution increases, and larger communities are detected, the modular structure of the networks becomes less similar. Here we observe a stronger symmetric relationship between the two dimensions of the BiDir distance.

4 Discussion

Comparing community structure across networks remains a challenging task. The majority of current methodologies equate the community structure to the graph partition, but this approach ignores key topological information. While we also compare two partitions, by assessing the fit of one partition to the topology of the other network (and vice versa), the bi-directional distance we propose incorporates information on the network structure. This approach enables us to pick up similarities and differences in the underlying network topologies not captured in the partitions alone. Agnostic about the community detection quality function used within a broad class, this approach can be deployed across a wide range of applications.

The BiDir distance assumes that the partition assigned to each network is optimal (under some chosen community detection quality function). This might appear like a strong assumption but the BiDir distance is well suited to handle cases where there exist alternative partitions that are close to optimal, since they all are assigned similar scores by the quality function. This feature enables our method to detect topological similarities that are not captured by (potentially dissimilar) partitions. Nevertheless, for networks where the mesoscale structure is not modular (for example if they exhibit a core-periphery structure, or disassortative behaviour), our methodology will be less applicable. While it will distinguish networks with high modularity from networks with a different structure, it will not be useful to compare between the latter.

The multi-dimensionality of BiDir means that it is neither a metric nor a quasi-metric in a formal sense. While the directionality provides more information on the similarities and differences in modular structure, there are cases where it will be ambiguous. If we are interested in identifying which networks are the most similar (or different) an extra step

is required to reduce the BiDir to a single score, which can be done by, for example, taking the sum of the squares. Even then, it is important to keep in mind that there can be no clustering comparison methodology that can be generally used in all applications, as no single metric will satisfy all possible desirable axioms [30].

By construction, the BiDir distance inherits the properties of the quality function. Therefore, it is well suited for comparing the community structure of networks when the application of interest is related to said function. For example, in the inter-industry flows, we are interested in flows within the network, which are captured by the stability algorithm. But it also constrains its use. Unlike NMI, it cannot be used to benchmark different clustering algorithms using ground truth data, since each algorithm has a different quality function.

A key aspect of community detection is that the specific algorithm is chosen according to whether structural, dynamic or other features are of interest. Since we use the quality function in our method, we retain this focus on the features of interest. However, our method is limited to algorithms that optimise an objective function. Extensions to other algorithms could be explored in future work. For example, if one was interested in comparing two networks using an SBM model, one could compare the statistical significance of two alternative SBM models (e.g., using a posterior odds ratio) instead of using an f -score.

Finally, our measure identifies the presence of similar modular structure between network pairs. However, when it comes to applications, often more information is required to understand ‘where’ in these networks similarities (or differences) occur. This was clearly illustrated in the comparison of the inter-industry labour flow networks, where a finer investigation of community overlap was used to shed light on sectors with either conserved or distinct patterns of labour flows across countries. The adaption of our bi-directional distance to a finer community-level distance is a potentially fruitful future research endeavour.

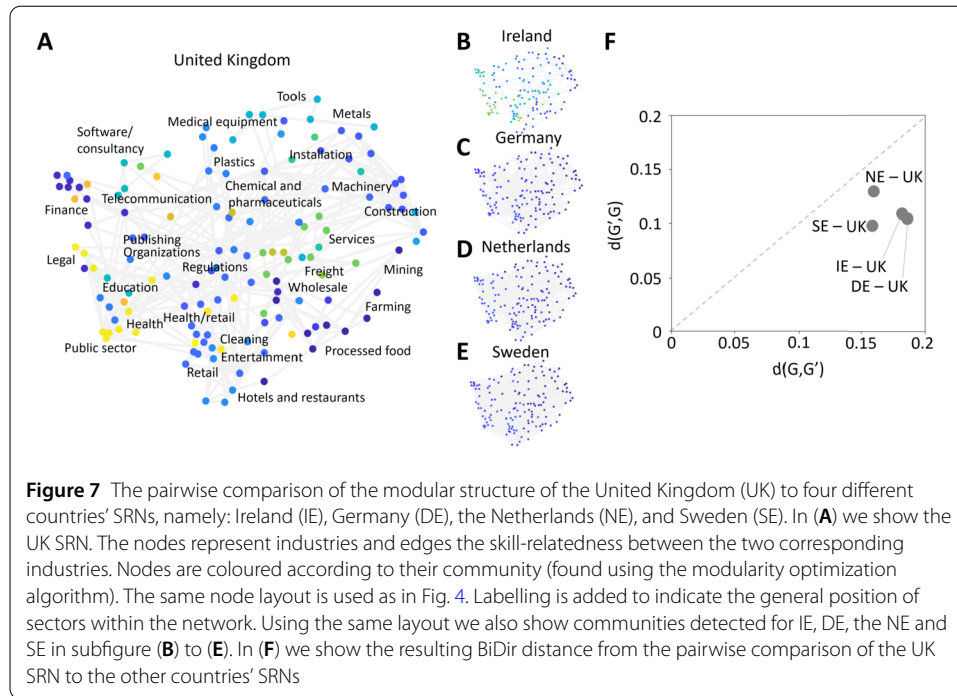
Appendix A: Inter-industry labour flow networks

Here, we further elaborate on the comparison of the modular structure of the United Kingdom Skill-Relatedness Network (SRN) to the Irish, German, Dutch and Swedish SRNs using the BiDir distance.

Recall, the BiDir distance requires the networks being compared to be node-aligned (*i.e.* all the networks need to have the same node set). To ensure that this is the case, we first convert the UK network so that industries are defined according to the NACE 1 industry classification. We then take the node intercept of all the networks being compared. Next, we elaborate on each of these steps.

Although the UK SRN has been constructed using the methodology of Neffke *et al.* [35], due to data availability it was constructed using the 4-digit NACE 2 industry classification. Using the official concordance table provided by Eurostat, we convert the industries from a 4-digit NACE 2 to a 4-digit NACE 1 industry classification. We adopt the network-based convergence methodology of Diodato [9].

Once all the networks are in the same industry classification, we take the node intercept of all the networks being compared. We observe that the UK’s SRN has fewer industries compared to the other four countries. It is also much more sparse. This is because it was constructed from a longitudinal survey including just 1 per cent of UK workers (ASHE).



On the other hand, all other networks were constructed using administrative data containing all workers within the formal sector. Hence, when comparing the UK SRN to those of other countries, only the subgraph induced by the set of overlapping nodes (and their corresponding edges) are compared in each case.

In Fig. 7(A) we show a visualisation of the UK SRN. Each node in the network represents an industry and each edge the skill-relatedness between its two corresponding industries. The same node layout and node labelling is used as in Fig. 4(A). The modular structure of the network is shown via the node colouring. The communities were detected using modularity [41]. Specifically, the partition maximizing the modularity function from 10,000 iterations of the Louvain algorithm is adopted. Using the same layout, we visualise the communities found for Irish, Germany, the Netherlands and Sweden on the right. Note that for all the networks only the induced sub-graph of the SRNs are shown.

We now quantify the distance between the UK SRN and the other four countries' SRN using our BiDir distance. The pairwise comparison is shown in Fig. 7(F). We observe that the modular structure of the UK network is very different to those of the other countries. We also observe that all the distances show an asymmetric relationship. For all four countries a smaller distance is obtained when projecting the community structure onto the UK network than vice versa. As the UK network is particularly sparse, the modularity algorithm unveils a more fine-grained community structure. Our measure is therefore picking up this difference in the scale at which the communities are most robust.

Appendix B: Technical details regarding SBM construction

In Fig. 3 we can see the performance of our framework in three different families of SBMs. We now summarize how these three different families were generated. All networks consist of 120 nodes and are generated according to the following parameters:

- All networks in family *A* consist of 4 identical blocks of size 30, and the SBM parameters are $P_{ij} = 0.20$ and $P_{ii} = 0.80, 0.60, 0.45, 0.35, 0.20$ for each successive network. Since for the last network, *A5*, both probabilities are the same, we do not have any a priori community structure, and we can use this network as a point of comparison.
- For all networks in family *B*, $P_{ij} = 0.20$ and $P_{ii} = 0.80$. But each successive networks splits some of the blocks of the previous one. So *B2* split one of the original blocks into two, *B3* splits the other original block to obtain 4 identical medium blocks. *B4* splits one of these 4 blocks into 3 smaller ones, and *B5* repeats that procedure on the other 3 medium blocks.
- *C1* and *C2* are produced by again using $P_{ij} = 0.20$ and $P_{ii} = 0.80$, with the first one split into 3 blocks and the second one into 4 blocks. *C3* has the same blocks as *C2*, but $P_{2,3} = P_{3,2} = 0.40$.

Furthermore, in Fig. 6 we illustrate the use of the BiDir distance using a multi-resolution community detection algorithm. Here, we provide more detail regarding how the three nested SBMs were constructed. All three of the networks consists of 300 nodes and have implanted communities at 3-levels. As illustrated in the adjacency matrices of each network, each network contains regions with one of 4 different shades of grey. These regions correspond to 4 different values of P_{ij} . From darkest to lightest the regions correspond to the following four values of $P_{ij} \in \{0.8, 0.53, 0.27, 0.15\}$. The size of the different regions for each network are given as follows:

- Network A consists of 12 identical blocks of size 25. Two of these blocks are then merged together to form 6 identical blocks of size 50. Similarly, two of these blocks are then grouped together to form 3 identical blocks of size 100.
- Network B, consists of 18 blocks of size 20, 10, 10, 25, 10, 25, 20, 10, 10, 25, 10, 25, 20, 10, 10, 25, 10, 25, respectively. Two consecutive blocks are then joined together to create 9 blocks of size 30, 35, 35, 30, 35, 35, 30, 35, 35, respectively. Finally, three consecutive blocks are joined together to create three identical blocks of size 100—similar to Network A.
- Network C, consists of 12 identical blocks of size 25, similar to Network A. These blocks are merged together to form four communities containing three blocks of size 75. Finally, two 75-sized blocks are joined together to form two communities of size 150.

Acknowledgements

We would like to thank Ioan Alexandru Puiu and Renaud Lambiotte for their insightful suggestions and comments. Furthermore, we thank the anonymous reviewers whose suggestions have helped improve and clarify our manuscript.

Funding

This project (NOC and DS) was funded by a Turing-HSBC-ONS Economic Data Science Award 'Network modelling of the UK's urban skill base'. ML received support from the Skye Foundation Trust and the Oppenheimer Memorial Trust.

Abbreviations

BiDir, bidirectional distance; NMI, normalized mutual information; SBM, stochastic block model; SRN, Skill Relatedness Network.

Availability of data and materials

The methodology to reproduce the synthetic networks is outlined in the Supplementary Materials. The SRNs for Germany and Sweden were constructed by Frank Neffke [35] and are available at <http://www.frankneffke.com/relatedness.html>, and the SRN for the Netherlands was constructed by Dario Diodato [10] and is available at <http://dariodiodato.com>. The SRNs for Ireland and the UK are available upon request.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

NO, DS and ML designed the study and wrote the paper. DS and ML carried out the analysis. All authors read and approved the final manuscript.

Author details

¹Center of Advanced Spatial Analysis, University College London, 90 Tottenham Court Road, Bloomsbury, W1T 4TJ London, UK.

²Mathematical Institute, Oxford University, Andrew Wiles Building, Radcliffe Observatory Quarter, OX2 6GG Oxford, UK.

³Alan Turing Institute, British Library, 96 Euston Road, NW1 2DB London, UK.

Publisher's Note

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

Received: 14 October 2020 Accepted: 25 February 2021 Published online: 10 March 2021

References

1. Barabási A-L, Albert R, Jeong H (1999) Mean-field theory for scale-free random networks. *Phys A, Stat Mech Appl* 272(1–2):173–187
2. Barahona M, Pecora LM (2002) Synchronization in small-world systems. *Phys Rev Lett* 89(5):54101
3. Blondel VD, Hendrickx JM, Olshevsky A, Tsitsiklis JN (2005) Convergence in multiagent coordination, consensus, and flocking. In: 44th IEEE conference on decision and control, pp 2996–3000. <https://doi.org/10.1109/CDC.2005.1582620>
4. Chakraborty T, Dalmia A, Mukherjee A, Ganguly N (2017) Metrics for community analysis: a survey. *ACM Comput Surv* 50(4):1–37
5. Cranmer SJ, Desmarais BA (2011) Inferential network analysis with exponential random graph models. *Polit Anal* 19(1):66–86
6. Csáfordi Z, László L, Lengyel B, Kiss KM (2016) The effect of labor flows, ownership and skill-relatedness on firm productivity. In: Proceedings of international academic conferences. International Institute of Social and Economic Sciences, number 4006263
7. Danon L, Diaz-Guilera A, Duch J, Arenas A (2005) Comparing community structure identification. *J Stat Mech Theory Exp* 2005(09):P09008
8. Delvenne J-C, Yaliraki SN, Barahona M (2010) Stability of graph communities across time scales. *Proc Natl Acad Sci* 107(29):12755–12760
9. Diodato D (2018) A network-based method to harmonize data classifications. Technical report, Utrecht University, Department of Human Geography and Spatial Planning
10. Diodato D, Weterings AB (2014) The resilience of regional labour markets to economic shocks: exploring the role of interactions among firms and workers. *J Econ Geogr* 15(4):723–742
11. Donnat C, Holmes S (2018) Tracking network dynamics: a survey of distances and similarity metrics. arXiv preprint, [arXiv:1801.07351](https://arxiv.org/abs/1801.07351)
12. Dunne JA, Williams RJ, Martinez ND (2002) Network structure and biodiversity loss in food webs: robustness increases with connectance. *Ecol Lett* 5(4):558–567
13. Farjoun M (1994) Beyond industry boundaries: human expertise, diversification and resource-related industry groups. *Organ Sci* 5(2):185–199
14. Fortunato S (2010) Community detection in graphs. *Phys Rep* 486(3–5):75–174
15. Frenken R, Van Oort F, Verburg T (2007) Related variety, unrelated variety and regional economic growth. *Reg Stud* 41(5):685–697
16. Girvan M, Newman ME (2002) Community structure in social and biological networks. *Proc Natl Acad Sci* 99(12):7821–7826
17. Good BH, De Montjoye YA, Clauset A (2010) Performance of modularity maximization in practical contexts. *Phys Rev E* 81(4):046106
18. Hagen L, Kahng AB (1992) New spectral methods for ratio cut partitioning and clustering. *IEEE Trans Comput-Aided Des Integr Circuits Syst* 11(9):1074–1085
19. Hamming RW (1950) Error detecting and error correcting codes. *Bell Syst Tech J* 29(2):147–160
20. Hausmann R, Hwang J, Rodrik D (2007) What you export matters. *J Econ Growth* 12(1):1–25
21. Hidalgo CA, Klinger B, Barabási A-L, Hausmann R (2007) The product space conditions the development of nations. *Science* 317(5837):482–487
22. Holland PW, Laskey KB, Leinhardt S (1983) Stochastic blockmodels: first steps. *Soc Netw* 5(2):109–137
23. Jaccard P (1912) The distribution of the flora in the Alpine zone. 1. *New Phytol* 11(2):37–50
24. Jaffe AB (1989) Characterizing the “technological position” of firms, with application to quantifying technological opportunity and research spillovers. *Res Policy* 18(2):87–97
25. Jeong H, Tombor B, Albert R, Oltvai ZN, Barabási A-L (2000) The large-scale organization of metabolic networks. *Nature* 407(6804):651–654
26. Krause U (1997) Soziale dynamiken mit vielen interagierenden Akteuren. Eine problemskizze. In: Modellierung und Simulation von Dynamiken mit vielen interagierenden Akteuren p 3751
27. Lambiotte R, Delvenne JC, Barahona M (2008) Dynamics and modular structure in networks. arXiv preprint, [arXiv:0812.1770](https://arxiv.org/abs/0812.1770)
28. Lambiotte R, Sinatra R, Delvenne J-C, Evans TS, Barahona M, Latora V (2011) Flow graphs: interweaving dynamics and structure. *Phys Rev E* 84(1):017102
29. Masuda N, Porter MA, Lambiotte R (2017) Random walks and diffusion on networks. *Phys Rep* 716:1–58

30. Meilä M (2003) Comparing clusterings by the variation of information. In: Learning theory and kernel machines. Springer, Berlin, pp 173–187
31. Meilä M (2007) Comparing clusterings—an information based distance. *J Multivar Anal* 98(5):873–895
32. Milo R, Shen-Orr S, Itzkovitz S, Kashtan N, Chklovskii D, Alon U (2002) Network motifs: simple building blocks of complex networks. *Science* 298(5594):824–827
33. Molloy M, Reed B (1995) A critical point for random graphs with a given degree sequence. *Random Struct Algorithms* 6(2–3):161–180
34. Mucha PJ, Onnela J, Porter MA (2009) Communities in networks. *Not Am Math Soc* 56:1082–1097
35. Neffke F, Henning M (2013) Skill relatedness and firm diversification. *Strateg Manag J* 34(3):297–316
36. Neffke F, Henning M, Boschma R (2011) How do regions diversify over time? Industry relatedness and the development of new growth paths in regions. *Econ Geogr* 87(3):237–265
37. Neffke F, Henning M, Boschma R, Lundquist K-J, Olander L-O (2011) The dynamics of agglomeration externalities along the life cycle of industries. *Reg Stud* 45(1):49–65
38. Nelson RR, Winter SG (1982) An evolutionary theory of economic change. The Belknap Press of Harvard University Press, Cambridge
39. Newman M (2008) The physics of networks. *Phys Today* 61(11):33–38
40. Newman ME (2003) The structure and function of complex networks. *SIAM Rev* 45(2):167–256
41. Newman ME (2006) Modularity and community structure in networks. *Proc Natl Acad Sci* 103(23):8577–8582
42. Newman ME, Barabási A-LE, Watts DJ (2006) The structure and dynamics of networks. Princeton University Press, Princeton
43. Newman ME, Girvan M (2004) Finding and evaluating community structure in networks. *Phys Rev E* 69(2):026113
44. Newman MEJ, Cantwell GT, Young J-G (2020) Improved mutual information measure for clustering, classification, and community detection. *Phys Rev E* 101(4):042304
45. O’Clery N, Curriel RP, Lora E (2019) Commuting times and the mobilisation of skills in emergent cities. *Appl Netw Sci* 4(1):118
46. O’Clery N, Flaherty E, Kinsella S (2019) Modular structure in labour networks reveals skill basins. arXiv preprint, [arXiv:1909.03379](https://arxiv.org/abs/1909.03379)
47. O’Clery N, Gomez A, Lora E (2016) The path to labour formality: urban agglomeration and the emergence of complex industries. Harvard Center for International Development Working Paper
48. O’Clery N, Yuan Y, Stan G-B, Barahona M (2013) Observability and coarse graining of consensus dynamics through the external equitable partition. *Phys Rev E* 88(4):042805
49. Onnela J-P, Fenn DJ, Reid S, Porter MA, Mucha PJ, Fricker MD, Jones NS (2012) Taxonomies of networks from community structure. *Phys Rev E* 86(3):036104
50. Orman GK, Labatut V, Cherifi H (2012) Comparative evaluation of community detection algorithms: a topological approach. *J Stat Mech Theory Exp* 2012(08):P08001
51. Palla G, Derényi I, Farkas I, Vicsek T (2005) Uncovering the overlapping community structure of complex networks in nature and society. *Nature* 435(7043):814–818
52. Peel L, Larremore DB, Clauset A (2017) The ground truth about metadata and community detection in networks. *Sci Adv* 3(5):e1602548
53. Robins G, Snijders T, Wang P, Handcock M, Pattison P (2007) Recent developments in exponential random graph (p^*) models for social networks. *Soc Netw* 29(2):192–215
54. Rosvall M, Axelsson D, Bergstrom CT (2009) The map equation. *Eur Phys J Spec Top* 178(1):13–23
55. Schaub MT, Delvenne J-C, Rosvall M, Lambiotte R (2017) The many facets of community detection in complex networks. *Appl Netw Sci* 2(1):1–13
56. Schaub MT, Delvenne JC, Yaliraki SN, Barahona M (2012) Markov dynamics as a zooming lens for multiscale community detection: non clique-like communities and the field-of-view limit. *PLoS ONE* 7(2):e32210
57. Sood V, Redner S (2005) Voter model on heterogeneous graphs. *Phys Rev Lett* 94(17):178701
58. Strogatz SH (2001) Exploring complex networks. *Nature* 410(6825):268–276

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
