
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Hric, Darko; Kaski, Kimmo; Kivelä, Mikko

Stochastic block model reveals maps of citation patterns and their evolution in time

Published in:
Journal of Informetrics

DOI:
[10.1016/j.joi.2018.05.004](https://doi.org/10.1016/j.joi.2018.05.004)

Published: 01/08/2018

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Hric, D., Kaski, K., & Kivelä, M. (2018). Stochastic block model reveals maps of citation patterns and their evolution in time. *Journal of Informetrics*, 12(3), 757-783. <https://doi.org/10.1016/j.joi.2018.05.004>



Regular article

Stochastic block model reveals maps of citation patterns and their evolution in time



Darko Hric, Kimmo Kaski, Mikko Kivelä*

Department of Computer Science, Aalto University School of Science, P.O. Box 12200, FI-00076, Finland

ARTICLE INFO

Article history:

Received 2 May 2017

Received in revised form 30 May 2018

Accepted 30 May 2018

Keywords:

Web of science

Citation networks

Evolution of science

Stochastic block model

ABSTRACT

In this study we map out the large-scale structure of citation networks of science journals and follow their evolution in time by using stochastic block models (SBMs). The SBM fitting procedures are principled methods that can be used to find hierarchical grouping of journals that show similar incoming and outgoing citations patterns. These methods work directly on the citation network without the need to construct auxiliary networks based on similarity of nodes. We fit the SBMs to the networks of journals we have constructed from the data set of around 630 million citations and find a variety of different types of groups, such as communities, bridges, sources, and sinks. In addition we use a recent generalization of SBMs to determine how much a manually curated classification of journals into subfields of science is related to the group structure of the journal network and how this relationship changes in time. The SBM method tries to find a network of blocks that is the best high-level representation of the network of journals, and we illustrate how these block networks (at various levels of resolution) can be used as maps of science.

© 2018 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The process of creating scientific knowledge relies on publications that are often stored and archived, with the primary purpose of preserving and distributing the knowledge obtained through research. These archives can also be used to study the science making itself, for example, by extracting information of collaborations, citations, or keywords of the published articles. Research in this field has a fairly long and rich history with wide range of research topics, like the assessment and prediction of performance and quality of individual papers, researchers, institutions, journals, fields, and even countries (Althouse, West, Bergstrom, & Bergstrom, 2009; Lehmann, Jackson, & Lautrup, 2008; Nerur, Sikora, Mangalaraj, & Balijepally, 2005), as well as identification of various large scale structures of science (Boyack & Klavans, 2014; Carpenter & Narin, 1973; de Solla Price, 1965; Leydesdorff, Carley, & Rafols, 2013; Small, 1999; Waltman, van Eck, & Noyons, 2010), journal classification (Janssens, Zhang, Moor, & Glänzel, 2009; Leydesdorff, 2006; Wang & Waltman, 2016; Zhang, Liu, Janssens, Liang, & Glänzel, 2010), following research trends (Chen, 2013; Persson, 2010; Porter & Rafols, 2009), and recognizing the emerging fields or researchers (Cozzens et al., 2010; Lambiotte & Panzarasa, 2009; Shibata, Kajikawa, Takeda, Sakata, & Matsushima, 2011; Small, Boyack, & Klavans, 2014; Small & Greenlee, 1989).

Bibliographic databases, like Web of Science, Scopus, and Google Scholar, store metadata of scientific publications, which can be used to analyse science making at all levels, from large scale structure to performance of individual papers. The number

* Corresponding author.

E-mail addresses: darko.hric@aalto.fi (D. Hric), kimmo.kaski@aalto.fi (K. Kaski), mikko.kivela@aalto.fi (M. Kivelä).

of *entities* in the data, including articles, journals, citations, and scientists is very large and keeps growing exponentially ([Appendix A](#); Pan, Petersen, Pammolli, & Fortunato, 2016). To make sense of such massive amounts of data available about science one needs to simplify it and find its inherent patterns. This idea is not different from creating *maps* that provide a simplified description of reality, i.e. maps of science that describe the endeavour of science in a broad sense (Boyack, Klavans, & Börner, 2005; Chen, 2013; Small, 1999). Such a map needs to provide a reasonably accurate simplification of the structures it is mapping, i.e. individual elements need to be grouped (or clustered) to preserve large-scale patterns, while obscuring small and unimportant details. However, this is not a trivial task, and finding an optimal simplification accurately and reliably is becoming even more challenging as the networks under study continue to grow.

Conventional data analysis tools, such as clustering or dimension reduction methods, can be used to simplify the data about the complex relationships between the data entities. Representing the entities as vectors of their features is a common and practical abstraction that allows the use of clustering methods in the space of features, in which the most similar entities are grouped based on the similarity of the used features. These vectors can contain citation information between the entities, and one can define similarity measures, like bibliographic coupling, co-citation, distance between citation vectors (Euclidean, cosine, Jaccard, etc.), and correlation coefficients between the citation vectors or publication texts (abstracts, keywords, etc.) (Boyack et al., 2005; Carpenter & Narin, 1973; Janssens et al., 2009; Kessler, 1963; Leydesdorff & Rafols, 2012; Marshakova, 1973; Small, 1973; Wang & Koopman, 2017).

The data of scientific progress can be analysed with a variety of methods once the data has been preprocessed. The dimensionality reduction techniques project the vectors into the most significant subspaces revealing groups of correlated entities (multidimensional scaling, factor analysis, etc.) (Leydesdorff et al., 2013; Small, 1999). Classical clustering techniques, e.g. hierarchical clustering and k-means, operate on the full space of features, and provide clusters of similar entities, based on implicitly or explicitly defined similarity measure or distance (Boyack et al., 2005; Modha & Spangler, 2000; Punj & Stewart, 1983; Silva, Rodrigues, Oliveira, da, & Costa, 2013; Wang & Koopman, 2017). The factor analysis applied separately to the citing and cited direction of the complete citation matrix, enables further specialization into the types of groups it finds, since by using only one direction at a time, it detects groups based on past and future citations, separately (Leydesdorff & Rafols, 2009). The co-citation and bibliographic coupling use similarities in citations in the future and past respectively, and thus provide a separation naturally (Weinberg, 1974). The results of this type of analysis depends on the preprocessing step of constructing the data vectors and similarities, and great care is needed in interpreting the results (Boyack et al., 2005; Gläser, Glänzel, & Scharnhorst, 2017; van Eck & Waltman, 2009).

The bibliometric data can also be analysed by constructing networks—such as the citation network between journals—and directly finding structure in them using the general purpose tools for analysing the networks. The development of such methods within network science has exploded since massive amounts of data on large variety of networks—such as on social and transportation networks—have become available (Boccaletti, Latora, Moreno, Chavez, & Hwang, 2006; Newman, 2003). A prominent way of finding structure in citation networks using these methods is to investigate network clusters or communities (Fortunato, 2010; Fortunato & Hric, 2016; Porter, Onnela, & Mucha, 2009), which are subnetworks that have a large number of links inside them (Chen & Redner, 2010; Lambiotte & Panzarasa, 2009; Lancichinetti & Fortunato, 2012; Radicchi, Fortunato, & Vespignani, 2012; Rosvall & Bergstrom, 2008). The assumption with most of these methods is that the network is constructed from densely connected cores of nodes or journals that have a relatively small number of citations to the rest of the network. This is in contrast to the methods based on similarity of journals that can find groups with a strong preference for receiving or giving citations from a certain subset of journals, for instance work of applied research can cite theoretical works, without being cited back.

Even if one would accept the premise that the community-like structures are relevant in citation networks, many community detection methods are besieged with intrinsic problems. Very often they detect structures even in case of random networks by mistaking noise for data, they might be very sensitive to small perturbations (noise), and posses a “resolution limit”, i.e. suffering from the inability to identify communities below a certain size that depends on the total size of the network (Fortunato & Barthélémy, 2007; Guimerà, Sales-Pardo, & Amaral, 2004). The performance, reliability, and even the results to some extent depend on the choice of a method from the large set of currently available methods.

The problems with community detection methods are well-known in the network science literature, and the need to find the richer structure in networks than those obtained by partitioning nodes to communities has been acknowledged for many types of networks (Leskovec, Lang, Dasgupta, & Mahoney, 2009; Palla, Derényi, Farkas, & Vicsek, 2005; Rombach, Porter, Fowler, & Mucha, 2014; Wang & Hopcroft, 2010; Xie, Kelley, & Szymanski, 2013). Very recently, as a solution to this problem, the old idea of using stochastic block models (SBMs) as models of network structure (Holland, Laskey, & Leinhardt, 1983; Lorrain & White, 1971; Wasserman & Anderson, 1987) has received renewed attention, because of the theoretical and algorithmic advances that enabled their use in a reliable and scalable way (Bianconi, 2009; Karrer & Newman, 2011; Peixoto, 2012a). SBM is a model in which nodes belong to *blocks* (the name for groups in the SBM paradigm) and edges are created between (and within) the blocks with some fixed probabilities for each pairs of blocks. The methods based on SBMs work by finding the model which best explains the network data. The best explanation is not necessarily the model that would have most likely produced the data, but the simplicity of the model must also be taken into account, and the principled and powerful ideas from statistical inference literature are used to avoid such overfitting. One can consider the blocks as “super nodes” that are connected with weighted edges, and SBM methods then—by definition—try to find the “super network” that is the best simplification of the original network.

Here we take the advantage of the recent advances in SBM methods found in the network science literature and apply them to large scale citation networks between journals. We use journal citation networks from Thomson Reuters Citation Index® for the years ranging from 1900 to 2013 which contains hundreds of millions of citations. Many studies concentrated on small subsets of the citation network (An, Janssen, & Milios, 2004; Grossman, 2002; Nerur et al., 2005; Pieters, Baumgartner, Vermunt, & Bijmolt, 1999; Porter & Rafols, 2009; Shibata et al., 2011; Zhang et al., 2010), while others were interested in large-scale patterns (Boyack et al., 2005; de Moya-Anegón et al., 2007; Leydesdorff & Rafols, 2009). We focus on the large scale citation networks that are constructed using all articles in this bibliographic dataset. First we divide the full time period into the time windows of 5 or 10 years and use the articles in those windows to construct networks of the journals active in each window. That is, we take snapshots of the contemporary science at different points of time and track the important developments by fitting them with hierarchical SBMs. We visualize the resulting block structure at multiple levels of hierarchy, and illustrate the presence of blocks that are not community-like by categorizing them as sources, sinks, bridges, and communities. Moreover, we follow the evolution of these block categories in 16 largest fields of science in time and report the large-scale changes in them over more than a hundred-year observation period.

The citation networks can be studied in isolation but they can also be augmented and compared with many other data sources such as journal categorizations, article keywords, and author information. Previous studies have, for example, compared predetermined journal categories to network clusters (Boyack et al., 2005; Janssens et al., 2009) or to factors from factor analysis (Leydesdorff & Rafols, 2009). They have also constructed networks using categories as nodes (Zhang et al., 2010) and evaluated the quality of categorisations using criteria that favour community-like categories (Wang & Waltman, 2016). Here we will utilize a recently developed generalization of the SBM method that allows the inclusion of any “tag” information about the nodes (Hric, Peixoto, & Fortunato, 2016) and use it to analyse how much information the predetermined journal categorizations carry about the group structure of the citation networks. This approach does not assume that the journal classifications are the ground truth, but determines the suitability of subject categories for describing citation structure by asking how much better we can do in estimating the citation flows with the classifications than we can do without the knowledge of the classifications. The construction of contemporary citation networks allows us to track the congruity of the subject categories with citation patterns throughout the last century.

The paper is organized as follows. The process of building annotated journal networks from raw citation data is described in Section 2. The stochastic block models are introduced and described in Section 3. Then the visualization of the citation networks is described and a selection of results is presented in Section 4. More detailed analysis of journal groups properties is done in Section 5, while Section 6 deals with their evolution in time. Next a comparison between the subject categories and citation structure is developed and presented in Section 7. Conclusions are made in 8. Some basic properties of the data and additional results are presented in Appendices A and D.

2. Data

All the networks constructed in this paper are based on data on articles and citations extracted from three Thomson Reuters Citation Index® datasets (*Science Citation Index Expanded™*, *Social Sciences Citation Index®*, and *Arts & Humanities Citation Index®*). This database contains information about the publishing year and the venue (journal, proceeding, conference, etc.) of articles, and each venue (from now on called *journal*) is assigned to none, one, or several *subfields*. We join the subfields into larger *fields* similar to Parolo et al. (2015). The data set spans from year 1900 to 2013 and contains about 76,000 journals, approximately 5.5M articles, and about 630M citations in total. A more detailed description of the data can be found in Appendix A.

As the full data set spans for more than a hundred years, it includes information needed to track development of modern science. We aim to investigate how the citation patterns have evolved during this time period and to that end we split the data into multiple time windows, each of which is then used to construct a contemporary network of journals. The total volume of publications and citations is growing exponentially in time (Pan et al., 2016), and because of this we set the time window length to ten years before 1970s and to five years afterwards.

2.1. Network construction

In each time window, a node corresponds to an active journal that has publications in the given time period. The connections between the journals are constructed using outgoing citations from these journals such that there is a directed link from journal *a* to journal *b* if an article in journal *a* cites an article in journal *b*, and the weight of this link is taken to be the number of such citations. For each time window we only count the contemporary citations satisfying the following two criteria: (1) the cited article is published in a journal that is active in the time window, and (2) the time difference between the citing and the cited article is shorter than the length of the window. This procedure ensures that all articles in the time window contribute equally (with their citations) to network links.

We have also tested a different approach for selecting the contemporary citations where both the citing and the cited article were required to be within the fixed time window. The more strict filtering of contemporary citations brings imbalance to incoming and outgoing citations of articles depending whether they are published at the beginning or towards the end of the window: those at the beginning have larger pool of articles they can receive citations from than the pool they can cite,

and the opposite for articles towards the end of the window. We replicated all the results reported in this article using the networks created with this method, and they were very similar to the ones reported here (see [Appendix C](#)).

2.2. Simplified networks

All of our results with the exception of the ones described in Section 7 are obtained using the weighted and directed networks of journals described above. Although all methods used in this manuscript are fully capable of dealing with non-simple networks, since the calculation of probabilities for directed and weighted networks was not implemented in the version of `graph-tool` we used (2.19dev), in Section 7 we had to use simplified networks (undirected, unweighted, and without self-loops). A naive method of discarding link directions and weights, and removing self-loops, leaves the networks very dense, and is a poor approximation because it regards all links equally important, irrespective of their direction or weight. A usual approach is to set a global threshold on the link weights that keeps only the strongest links, or to use only the links that form a maximum spanning tree ([Kruskal, 1956](#); [Macdonald, Almaas, & Barabási, 2005](#)). Both of these are global methods, meaning that the decision on whether a link will be kept or removed depends on the weight distribution and the structure of the full network. We use a local thresholding method, in which statistical significance of weights of links of every node are calculated based on a null model defined for each node separately ([Serrano, Boguñá, & Vespignani, 2009](#)). This significance measure is the probability that link weight is compatible with the null hypothesis, in statistical inference known as the p value, but here denoted with α . By keeping only the links that have α value lower than a certain threshold, we are dismissing all links that do not significantly differ from ones created randomly, while those that are kept can be considered significant (not random), and this “significance” is controlled with the value of the threshold. We tested the range of thresholds and found that the results are robust against the change of the threshold value α (see Section 7). The results are shown for α in the range 0.05, ..., 0.25 that preserve about 6% to about 21% of the most important links (representing about 23% to about 45% of citations) and about 51% to about 99% of nodes, respectively.

3. Stochastic block model

Networks and graphs can be measured and summarized at many levels of granularity, starting from global or macroscopic measures—such as the total number of links or diameter—to local or microscopic measures such as node degree or the clustering coefficient ([Newman, 2010](#)). Here we concentrate on describing networks in a mesoscopic scale that is between these two extremes. Network analysis methods that work at this level of granularity almost exclusively deal with sets of nodes and links called *communities* or depending on the field of research, *clusters*, *groups*, *modules*, etc. ([Boccaletti et al., 2006](#); [Fortunato, 2010](#); [Schaeffer, 2007](#); [Wasserman & Faust, 1994](#)). There is not a single, precise definition of community, but most often it is described as a set of nodes with more connections between them than to the rest of the network ([Fortunato, 2010](#); [Porter et al., 2009](#)). The community paradigm assumes that a network can be described as a collection of tightly knit sets of nodes, which are loosely connected to each other.

Stochastic block model relaxes the assumption about the nature of constituent sets of nodes such that they only need to be equivalent in the way they connect to other groups (called *blocks* in the SBM paradigm), which in effect allows for a description beyond the community structure, like bipartite, core-periphery, etc. ([Barucca & Lillo, 2016](#); [Karrer & Newman, 2011](#)). SBM is a generative model, meaning that it assumes a model of the underlying structure and prescribes a procedure for building networks that have this structure in common. The model is defined by assigning all nodes to disjoint sets¹ and setting the number of links between and within blocks. Obeying the above described constraints, a network is generated by randomly placing links between nodes. An alternative description is to set the probabilities for placing a link between any two blocks, but we used the link counts following the approach laid out in [Peixoto \(2014b\)](#).

Nodes within blocks share the probabilities for links towards the nodes in other blocks but also including their own block. In journal citation networks this means that all the journals in a block have the same citation patterns to other blocks. They can, for instance, receive most of their citations from one set of journals, and give them out to another set, or have higher than average probability to exchange citations with some blocks and lower-than-average probability with other blocks. Two blocks could also have identical citation patterns to other blocks, but different number of internal citations. All this tells us that this model groups nodes (journals) into classes by their role in the network, whichever those are.

Once the model is known, building networks with the prescribed block structure is straightforward. However, the more common situation is opposite to this: only a single realization of the model of the empirical network at hand is known, and parameters of the model that most likely produce this network, need to be inferred. Finding the most likely parameters is a highly non-trivial task, and many approaches to solve it have been used ([Wasserman & Anderson, 1987](#)). All approaches use an objective function, in one form or the other, that measures the probability of the given parameters to be the ones that produced the observed network. The problem with this naive approach is that the best fitting model will be too detailed and will reproduce the observed network with very high accuracy, which goes against the purpose of the models in providing a good simplification of the reality. The cause for this is the fact that the simple approach uses all available data, including noise, for fitting the parameters. In the extreme case a highly detailed model ends up putting all the nodes in their own

¹ The assumption about blocks being disjoint sets of nodes can be relaxed ([Peixoto, 2015](#)).

blocks, since this reproduces the network perfectly. A simple solution to prevent this from happening is to introduce the number of blocks as a constraint in the fitting procedure (Karrer & Newman, 2011). This works fine in cases where the number of blocks is known, otherwise it needs to be inferred from data, for instance by finding a balance between the model description length and its goodness-of-fit to data. Some of the approaches in this direction are listed in Latouche, Birmele, and Ambroise (2012). In this work we are taking advantage of the recent progress with this topic that estimates the number of blocks by including the information necessary for describing the model parameters into the total information amount being minimized such that it penalizes a too lengthy description (Peixoto, 2013, 2015).

One of the main issue with the classical SBMs, in addition to overfitting, is that they assume Poisson-distributed node degrees, so any deviation from the expected structure is considered a feature of the data and the fitting algorithm tries to find a model that would adequately describe it. Fitting this model with a network with a different degree distribution would yield blocks that represent classes of nodes based on their degrees, not just on which other nodes they connect to. So, for instance, highly cited journals would be separated from journals publishing review articles, that presumably cite in large volumes but do not get cited as much. Also, journals that play a central role in their respective communities would be put together, even if communities might not be otherwise related. The *degree-corrected stochastic block model* is “blind” to this kind of structure: it separates nodes into classes based on their degree, which removes unfair competition for links between nodes of largely different degrees (Karrer & Newman, 2011). This effectively includes the degree sequence into the model that does increase the information needed to describe the model, but the benefit of a simpler model for describing the data lowers the total description length in most cases.

The SBM and its degree-corrected variant can be expanded to describe hierarchical structure of blocks (Peixoto, 2014b). In this case, each level constitutes a network of blocks and the best fit of SBM of the level below it, starting from the network itself at the bottom level, to the trivial one-block level on top. Fitting all the levels is done simultaneously to obtain the minimum description length of the whole structure. This approach shows several advantages, like allowing for even smaller blocks that can reliably be inferred and providing a multi-resolution view of the network. Hierarchical SBM can be viewed as a stack of progressively simpler weighted networks—each level in the hierarchy is a zoomed-out version of the level below it. This means that we can inspect connectivity patterns of blocks of nodes at the desired level of detail/block sizes. The existence of the smallest identifiable blocks is in community detection literature known as the *resolution limit* and it is defined as the smallest group that a method is able to identify (Fortunato & Barthélémy, 2007). One of the most commonly used objective functions is modularity (Newman & Girvan, 2004), for which it has been shown that it is not able to separate small groups, even in obvious cases, if the number of links inside them is too small compared to the number of links in the rest of the network (Fortunato & Barthélémy, 2007). It should be noted that SBMs are not completely immune to resolution limit problems (Choi, Wolfe, & Airola, 2012; Peixoto, 2013), which presents a problem in the comparative analysis of networks spanning two orders of magnitude, as is the case with our time slices. The smallest detectable blocks scale with the network size, which means that high-resolution levels for large networks would remain hidden from us. Fitting the blocks at all resolutions for the hierarchical SBM at the same time reduces the limit, because each lower level uses blocks from the level above as a constraint, so block inference at the lower levels is in effect done locally in each higher-level block (Peixoto, 2014b).

The criteria and algorithms described above are implemented in `graph-tool` Python module, which we use to do all SBM fitting in this work (Peixoto, 2014a).

4. Visualizing citation flows between blocks

Hierarchical SBM levels can be visualized as networks that provide us with a multi-resolution map of the citation network. Network visualization is a powerful tool for visual inspection of complex network data, but its readability and thus usefulness depends on the level of details presented and the total size of the network. In addition, if the network is dense, i.e. average node degree is large, it is even harder to make a clear picture of it. Lower levels contain a lot of information in fine detail, but they are often vast and too dense for visualization to be readable. But not all weighted links are equally important if we are interested in large-scale patterns. This means we can depict only the most important ones, making the network significantly less dense. This is again done by the procedure outlined in Section 2, which keeps only the links that form the backbone of the network (Serrano et al., 2009). For the purpose of this analysis the directions of links are preserved.

Three examples of networks of SBM blocks for the time periods of 1920s, 1960s, and 1995–2000, at the levels most similar, i.e. the closest matching number of blocks to subfields and fields are depicted in Fig. 1.

Clustering of blocks of similar fields is quite visible in all six networks. Medicine forms the biggest cluster, followed by Biology, Chemistry, and Physics. Large-scale structure remains similar for all three time periods: Medicine is tightly connected to Biology, which is then connected to Chemistry and Physics. This structure is in accordance with the previously published maps (Leydesdorff & Rafols, 2009; Rosvall & Bergstrom, 2008). The figures also include interesting small-scale details that vary between the resolutions and time windows. For instance, the interdisciplinary blocks are often located at the boundaries between other fields, but this effect is more visible in the subfield resolution level.

The fact that there are multiple blocks with the same dominant field in levels having the most similar number of blocks to the number of assigned fields, is a consequence of large heterogeneity of the field sizes and importance. Large fields also have rich internal structure that overshadows small, more simple fields in process of inferring the best blocks at each hierarchical level. As a first approximation, we have chosen the number of fields as a guide for choosing the most appropriate

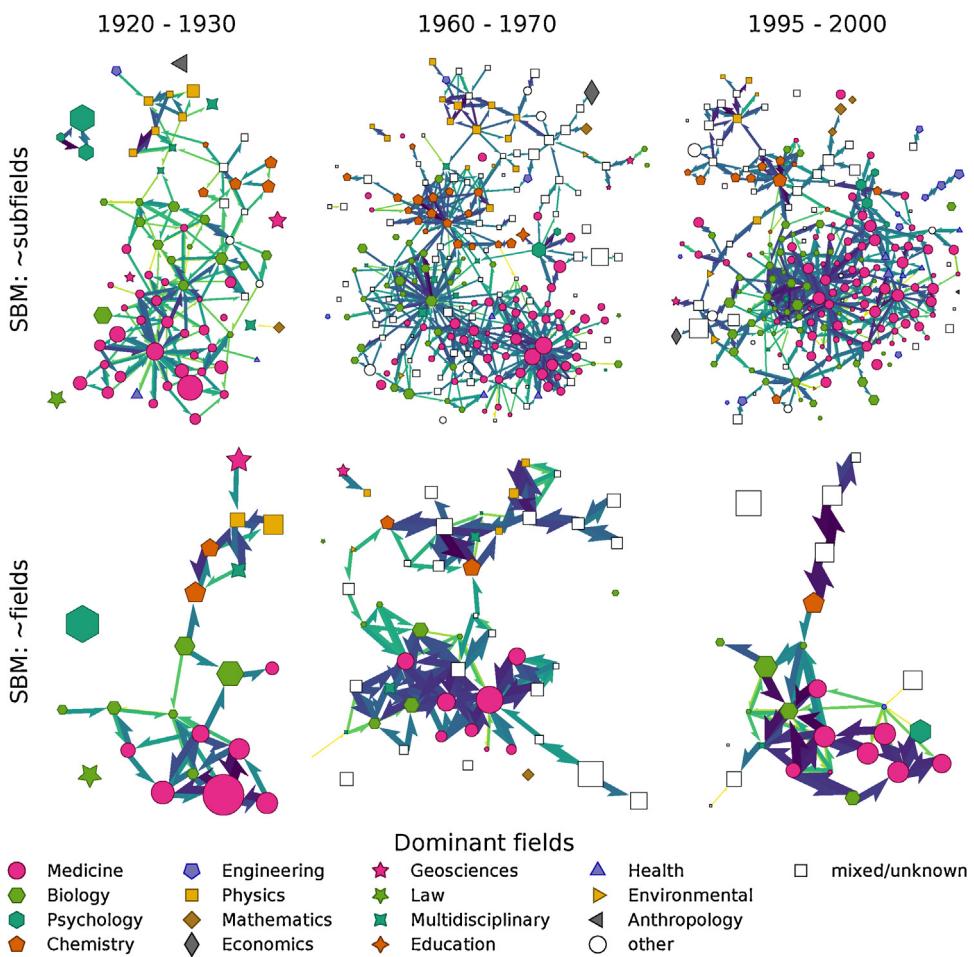


Fig. 1. Networks of SBM blocks for the time periods of 1920s, 1960s, and 1995–2000. The top row shows resolution levels most similar to subfields, bottom to fields. The node shapes and colours represent the dominant field in that block and the node sizes are proportional to the number of articles. Directed links are coloured and sized logarithmically, according to the number of citations they carry. Node and link sizes are normalized per pairs of networks from the same time window, so that these sizes can be compared between the two resolution levels but not across the time windows. The link colours are normalized for each network separately. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

level as it simple and intuitive, while another, more complex approach comparing the full size distributions did not provide significantly different results (see Appendix B).

5. Connectivity patterns of blocks

In contrast to traditional community detection, blocks in SBM are not limited to “dense subgroups” where the citations stay inside the blocks, but a block can also represent a structure where the citations flow out of or into the block, as long as all journals in the block behave in similar way. We summarize the type of block in terms of citation flows by counting the number of citations entering the block s_{in} (articles in the block are being cited), leaving the block s_{out} (they cite articles in other blocks) and internal citations s_{int} (citing articles in the same block)².

² One can view the system of flows and blocks as a weighted network. In the literature of weighted complex networks (Barrat, Barthélémy, Pastor-Satorras, & Vespignani, 2004), weighted sum of a node's links is called *strength*, and for the directed networks it can be in-, out- and internal strength: s_{in} , s_{out} and s_{int} .

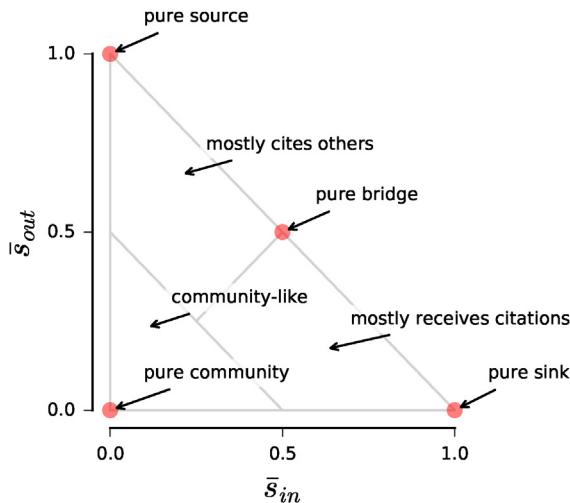


Fig. 2. Properties of blocks depending on their location on the connectivity pattern plot. \bar{s}_{in} and \bar{s}_{out} are relative in- and out- flows of citations to/from the block. Due to relation $\bar{s}_{in} + \bar{s}_{out} \leq 1$ points are confined to regions outlined by grey lines. Regions, as well as red points at special locations, are marked with arrows and annotated with descriptions of the properties of blocks at those locations. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Note that the sum of the three flow counts represents the total activity of the journals in the block. Here we are not interested in the total activities but in the type of flows. We investigate these types by separating the total flow from our flow measures and concentrate on relative in-, out-, and internal flows. These normalized flows are defined as:

$$\begin{aligned}\bar{s}_{in} &= s_{in}/s_{tot}, \\ \bar{s}_{out} &= s_{out}/s_{tot}, \quad \text{where } s_{tot} = s_{in} + s_{out} + s_{int}. \\ \bar{s}_{int} &= s_{int}/s_{tot},\end{aligned}\tag{1}$$

By normalising the flows by the total activity we reduce the number of free parameters needed to describe blocks' flow patterns from three to two. That is, the sum of the three relative flows equals to one, and knowing two of them is enough as the third one can always be calculated based on them. This means that we can report the two out of the three flow measures that are the most convenient for us.

For visualizing the connectivity patterns of blocks, the choice of \bar{s}_{in} and \bar{s}_{out} as x and y coordinates makes it easy to visually assess the block's properties from its position on the plot as illustrated in Fig. 2. Since the sum $\bar{s}_{in} + \bar{s}_{out}$ must be ≤ 1 , points can lie only in the triangle bounded by the diagonal $(0, 1) - (1, 0)$. Proximity to the origin tells us how much "self-centred" or "community-like" the block is, while the distances from the axes signify the balance between receiving and giving out citations. It helps to consider four extremal cases for a block, marked with red points in Fig. 2: **(0,0)** *Pure community*. Journals in this block are isolated from the rest of the network. Articles published in these journals only cite and get cited by articles in journals from this block. **(1,0)** *Pure sink*. Journals in this block only receive citations and do not cite at all. **(0,1)** *Pure source*. Journals in this block do not get cited, but cite others. **(0.5,0.5)** *Pure bridge*. Journals in this block cite and get cited equally, but there are no citations within the block. In most cases, the values lie somewhere between these extremes. The triangular space can be divided into three regions outlined by grey lines in Fig. 2 that contain blocks with the following properties: **Inner triangle** is community-like ("a community in a weak sense" [Radicchi, Castellano, Ceconi, Loreto, & Parisi, 2004]). More than half of the citations pertaining to this block stay within it. **Upper wing** mostly cites others. **Lower wing** mostly receives citations.

In Fig. 3 we visualize the types of connectivity patterns of blocks found in the citation networks. We display the data for three time periods (1920s, 1960s, and 1995–2000) and for three levels of resolution. This gives us an overview of the types of mesoscale structures one can find in the citation networks. A large number of blocks detected by the SBM method fall outside the inner triangle, and are thus not community-like structures even in the weak sense (Radicchi et al., 2004). This is especially evident for high-resolution levels where the vast majority of inferred blocks are not communities. Note that the inclination towards non-community-like structures is a feature of the data as the SBM fitting method we use does not have a preference for any particular block structure that is plausible within the SBM framework.

Note that, in the hierarchical structure of blocks, the level that contains larger blocks cannot have less community-like blocks than a level that contains smaller blocks, and for the one block at the top of the hierarchy all citations are internal. This can be illustrated by considering the merger of two blocks: their internal citations remain internal, but part of their external citations that go between them become internal to the merged block (c.f. Appendix E). In consequence, the average

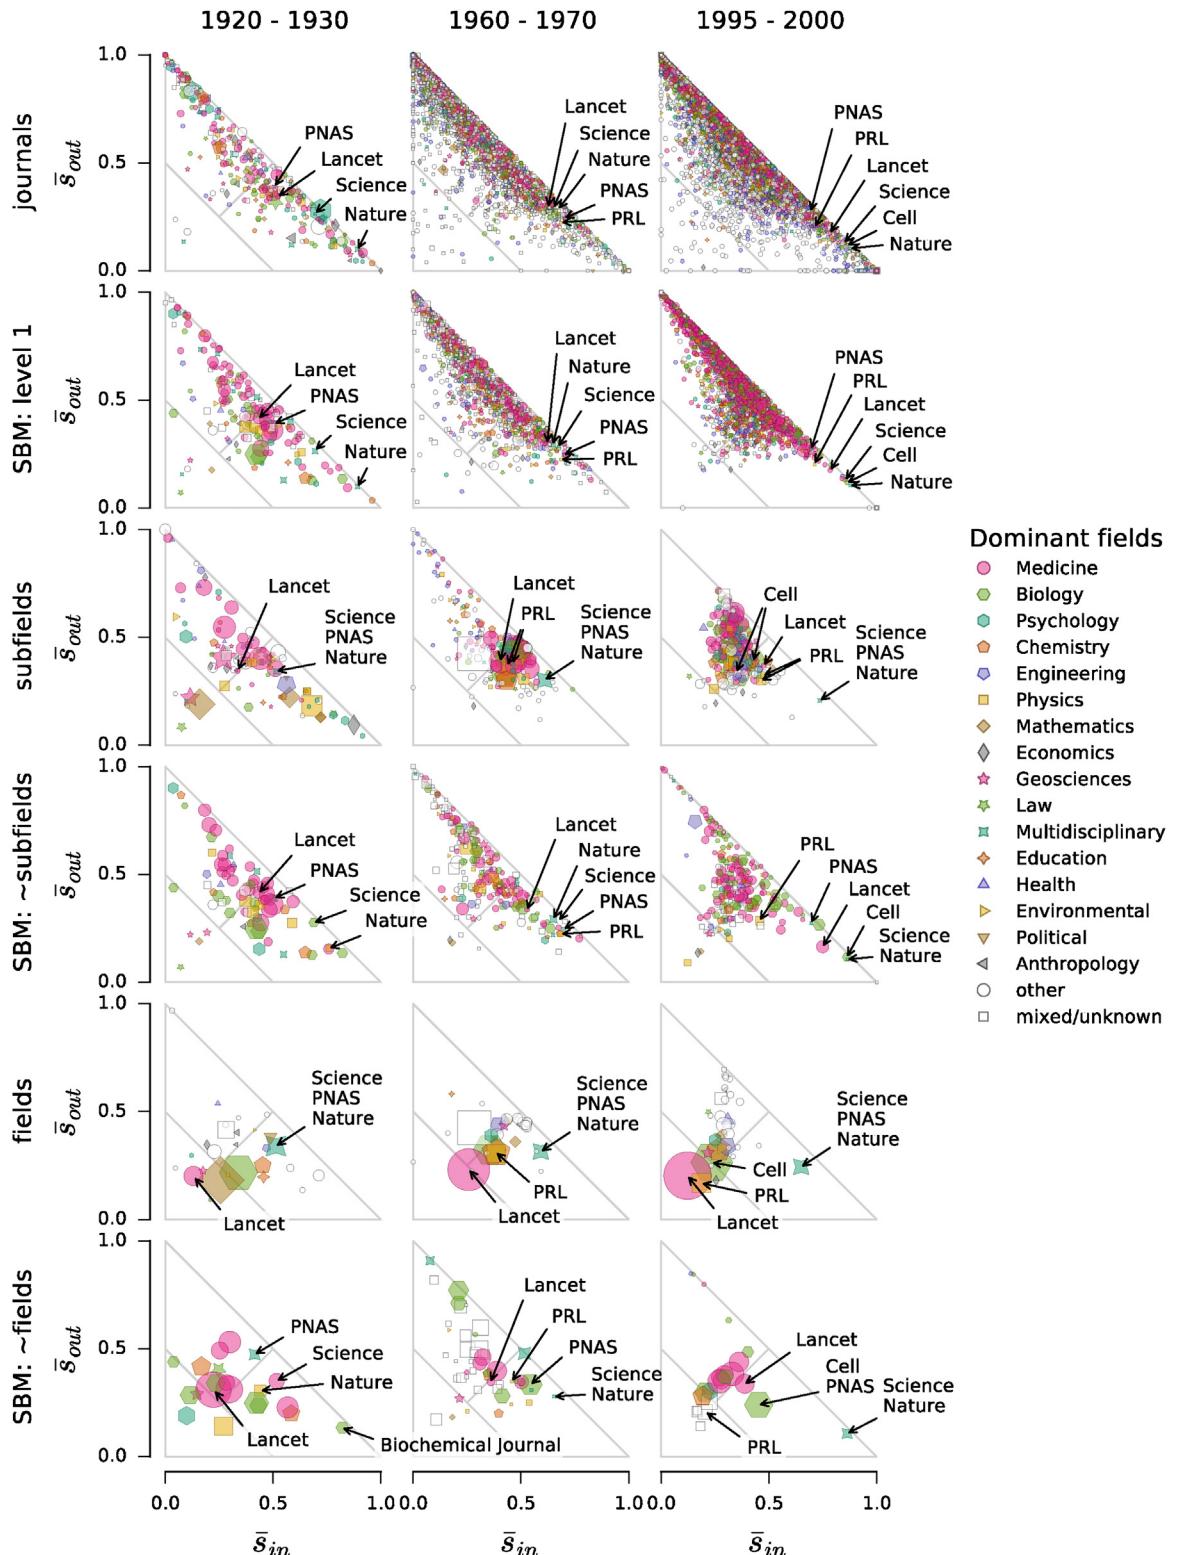


Fig. 3. Connectivity patterns of structural and classification blocks, for the time periods of 1900–1910, 1950–1960, and 1995–2000. Each row corresponds to a different block type: journals themselves; the first SBM level; subfields classification; SBM level with the closest matching number of blocks to the number of subfields (SBM: ~subfields); fields classification; and SBM level with the closest matching number of blocks to the number of fields (SBM: ~fields). Blocks are represented as points (coloured and shaped according to the dominant field in a block) with coordinates being the incoming and outgoing citations as fractions of total citations for each block (\bar{s}_{in} and \bar{s}_{out} , Eq. (1)). Point areas are proportional to the total citations (strength s_{tot}) pertinent to the block. Blocks with fewer than 10 total citations are not shown. Blocks containing a selected set of journals are annotated. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

fraction of internal citations in the block at the higher level of hierarchy can only be equal to or higher than the weighted average of the two blocks at the lower level of hierarchy.

The procedure outlined in this section can be used with any kind of partition of the journals into “blocks”, not just ones inferred for SBM. The journals are explicitly partitioned in the data by subfield and field classification, and we want to compare these partitionings to ones given by the blocks found by the SBM method. For our purposes it is useful to view both of these partitions as block structures, but to avoid confusion we name the blocks as determined by the classification data *classification blocks*, and the ones inferred by SBM from the citation patterns’ *structural blocks*. Journals themselves form elementary blocks, which can be viewed as the “zeroth” level of hierarchical SBM or any other hierarchical block structure. These zeroth level blocks give us properties of individual journals. Journals are assigned to subfields (in the dataset), which are in turn grouped into fields. One would expect this classification to be reflected in citation patterns, since it should group similar journals together. We are able to test this assumption by comparing the properties of artificial blocks, defined by subfields and fields, with blocks inferred from citations. Hierarchical SBM provides us with many levels at different resolutions, and in Fig. 3 we compare the level with the most similar number of blocks to the number of subfields and fields in that network.

Individual journals span the whole space of in- and outflow balance (there are many strong sinks as well as strong sources), while the overwhelming majority do not predominantly cite themselves. There are annotated journals of higher prominence that are found in the lower wing, which is to be expected since they receive more citations than give out.

The first SBM level determines the smallest non-trivial structural blocks, and it is the smallest grouping of journals that cite, and are cited, in a similar way as they have similar citation patterns towards the rest of the network. Compared with the level of journals the spread of points is reduced in all time periods, although to a varying degree. The spread is reduced with time: there is almost no change in 1920s, some change in 1960s and a significant constriction of values for 1995–2000. This could be a consequence of the resolution limit, where smaller details are increasingly harder to capture as the size of the network increases, or the outlying journals carry less information in later years, so they are combined with more moderate ones.

At the level of subfields (the 3rd and 4th rows in Fig. 3), both structural blocks and classification blocks display a pattern where multidisciplinary blocks are located mostly on the cited side. This is expected given that these blocks are dominated by high impact journals publishing high quality articles from a wide spectrum of fields. Comparing the distribution of points for classification blocks versus structural blocks we see that the latter ones are more evenly spread out in 1960s and 1995–2000, while this is not the case in 1920s. There are also more community-like blocks in the structural case, in particular in the fields of Economics, Physics, and Mathematics and Geosciences in 1920s. This means that SBM captures a broader spectrum of block types, while classification blocks tend to be more similar to each other or they have a preference for certain properties.

The highest level of hierarchy we focus on—for both classification blocks and structural blocks—is the level of fields (5th and 6th rows in Fig. 3). Constriction in classification blocks is again present, albeit not so strong as for the subfields. With time, Multidisciplinary field strongly separates from the bulk that remains more elongated in the community-bridge direction and slightly leaned toward source-like behaviour. Similar to the level of subfields, Multidisciplinary block tends to become more cited with time for both classification and structural blocks, with this behaviour being more pronounced for structural blocks and the last time period. Medicine and Physics get separated into a community-like block, meaning that the most citations remain within their field, which was not the case for the subfields. The SBM finds multiple blocks of mixed fields that are source- and community-like for 1960s, and only community-like for 1995–2000. Most of the mixed blocks in 1960s are comprised of unclassified journals, number of which rises in the second half of the century (see Appendix A for details). For classification blocks these journals are all collected under the “mixed/unknown” field (large white square).

Note that one needs to be careful when comparing blocks across different panels in Fig. 3, as there is no guarantee that blocks with similar qualities in different panels comprise of the same journals. Deeper analysis in this direction would require listing all journals in a block, or annotating journals of interest (which is done for a few journals in Fig. 3).

6. Evolution of block connectivity patterns in time

In the previous section we summarized the types of citation flows of individual blocks. We will next build on this summarization method, and quantify the evolution of citation flows of specific fields. To be more precise, we ask the question: what is the expected type of citation flow of the block where a randomly chosen article of a given field belongs to?

The expected flows for a journal in a field can be calculated by collecting all journals belonging to a field and taking a weighted average of the flows of the blocks they belong to. To calculate average flows for articles, averaging needs to be weighted by a fraction of journals’ publications out of the total number in the field:

$$\bar{s}_x^f = \frac{\sum_j a_j \bar{s}_x}{\sum_j a_j} , \quad (2)$$

where \bar{s}_x is any of the three average flows (in-, out-, and internal) from Eq. (1), a_j is the number of articles published by the journal j , and the sums go over all journals of the field f . If we do this for networks in all time windows, in addition to differentiating fields among themselves, we can follow the evolution of citation flow patterns of individual fields over time.

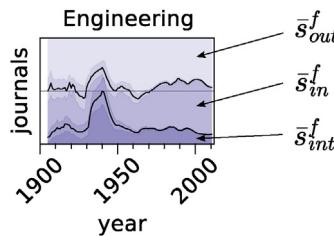


Fig. 4. Plot of the evolution of the average citation flows, for the example field Engineering. Time (in years) is on the horizontal axis. Relative citation flows of blocks containing journals from the respective field [see Eq. (2)] are shown as shaded regions: bottom part (the darkest) is for internal flow \bar{s}_{int}^f , middle (lighter) for incoming flow \bar{s}_{in}^f , and top part (the lightest) for outgoing flow \bar{s}_{out}^f . The average value of each flow is marked with a black line, and the error of the mean is shown as a transitional shade between regions. The central horizontal line marks 50% of the flow.

Similar to the previous section, we repeat the analysis for structural blocks inferred by the SBM and classification blocks given as classifications of the data (fields and subfields). That is, the average flows \bar{s}_x^f (in-, out- and internal) of a field f are calculated over all the field's journals, and the values \bar{s}_x are taken from blocks, either the structural blocks or the classification blocks, the journals belong to.

We illustrate the approach used to plot the evolution of the average citation flows of block connectivities in Fig. 4. Time (in years) is on the horizontal axis, while vertical axis is split by relative citation flows (Eq. (2)) of blocks containing journals from the respective field: bottom part (the darkest) represents internal flow \bar{s}_{int}^f , middle (lighter) incoming flow \bar{s}_{in}^f , and top part (the lightest) outgoing flow \bar{s}_{out}^f . The average value of each flow is marked with a black line, and the error of the mean is shown as a transitional shade between regions. The central horizontal line marks 50% of the flow.

The evolution of the average citation flows or block connectivities for 16 most prominent fields is shown in Fig. 5. Fields are ordered by the surface area under the curve gained by taking logarithms of numbers of journals in each field. Logarithmic scale is used to mitigate the impact of the exponential increase of the number of journals (c.f. Appendix A).

For half of the 16 fields, the internal flows of structural blocks are larger than the internal flows of classification blocks. Mathematics, Economics, Geosciences, Law, Education, Political Science, and Anthropology journals are found to reside in structural blocks that are more community-like than the corresponding classification blocks. This could mean that SBM is able to find their “natural communities”—ones that capture the most of the citation flows—while the classifications alone are not able to achieve. The strength of this effect varies, with the most striking examples being Mathematics, Law, and Political Science. Classification blocks of Economics, and partially Geosciences and Law are themselves quite community-like. The opposite effect is present for Medicine and Biology, meaning that the structural blocks are less community-like than classification ones. This can be explained by the fact that these fields have a large number of both subfields and journals, so they contain a rich internal structure of blocks with large flow of citations between them. Their internal citation structure might also be different from other fields, as they contain highly cited papers describing methods and procedures (Small & Griffith, 1974).

The in-flows \bar{s}_{in}^f and out-flows \bar{s}_{out}^f are quite balanced for most fields, with a few notable exceptions. Multidisciplinary field has noticeably more incoming citations than outgoing, for all hierarchy levels and for both types of blocks. This does not necessarily mean that all journals in Multidisciplinary field attract citations, but it can be due to the fact that it contains several high profile journals,³ like Nature, Science, and PNAS. Individual journals in Political Science field receive more citations than they give out, while this difference vanishes for the blocks they belong to. The opposite is true for Medicine, Health, and Environmental (they give out more citations than they receive) and this behaviour remains present for blocks in higher levels.

In time domain, fields exhibit a wealth of behaviours. Some have relatively stable patterns (Medicine, Chemistry, Economics, Multidisciplinary, and to some degree Health and Political Science), most of the others have large changes around the World War II, while some have uniform and steady shifts (Education and Anthropology are the most notable).

The most striking feature is the sudden rise in share of outgoing flows at the time of the WWII, mostly for classification blocks, and to some degree for journals of some fields. The largest rises are in decreasing order for: Law, Geosciences, Mathematics, Anthropology, Physics, Biology, and Environmental. In Biology and Chemistry a faint effect is also visible at the level of journals, while in case of Environmental it is mostly a drop in internal flow. Given that this effect is almost invisible for structural blocks, the observed changes most likely do not arise from the change in journals' citation patterns, but in the way they are classified. This sudden change correlates with the large increase of the number of subfields (c.f. Appendix A).

Some fields exhibit slow but steady change over time, predominantly in the structural blocks. Anthropology journals are citing more external literature and less themselves with time, which is also visible in the smallest structural blocks, to lesser extent in the structural blocks of size of subfields, and not at all in the structural blocks of size of fields. This means that considerable amount of the citation flow that is increasingly going out of the journals is nevertheless retained inside the

³ For this reason some authors have in similar analyses excluded these journals altogether (Zhang et al., 2010).

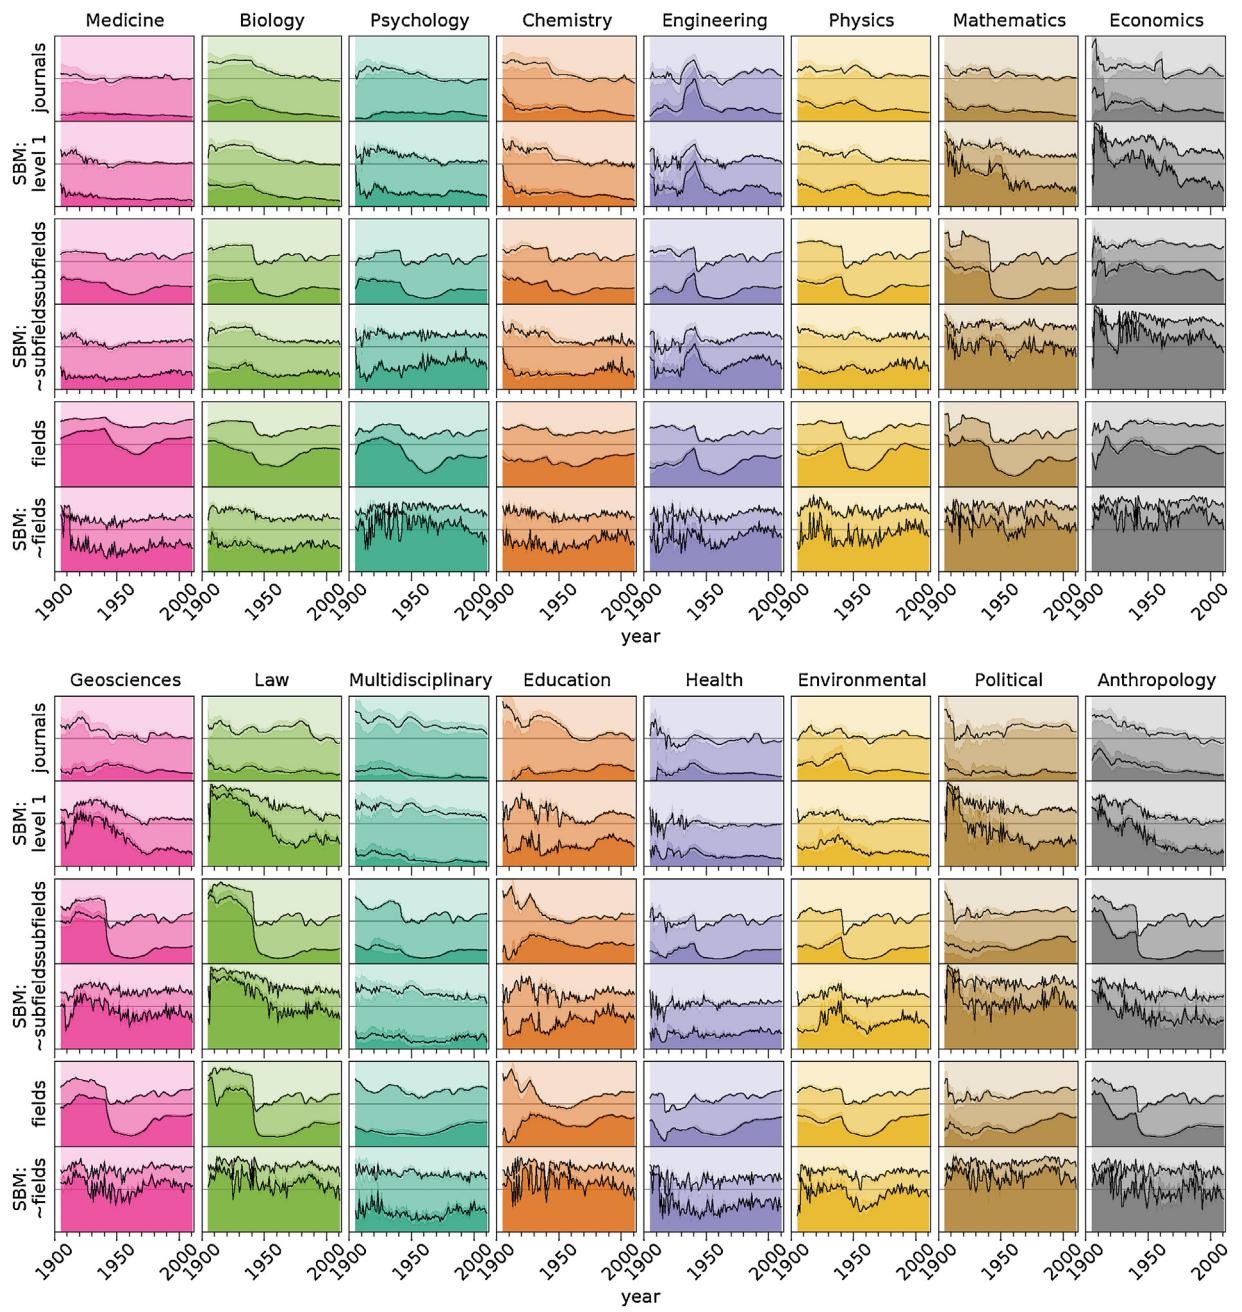


Fig. 5. Evolution of structural and classification block connectivities in time, for 16 largest fields. Fields are grouped in columns (colours are the same as in Fig. 3) and each row is for different type of blocks, from top: journals themselves; the first SBM level; subfields classification; SBM level with the closest matching number of blocks to the number of subfields (SBM: ~subfields); fields classification; and SBM level with the closest matching number of blocks to the number of fields (SBM: ~fields). Time (in years) is on the horizontal axis, while vertical axis is split by relative citation flows of blocks containing journals from the respective field: outgoing, incoming and internal flows (from top to bottom, respectively). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

structural blocks. Biology, Chemistry, Multidisciplinary, and Education journals show similar behaviour, but their outgoing flows remain stable also in the structural blocks of size of subfields. A plausible explanation is that as a new journal appears in the field, it “steals” some of the citation flows from the old journals while mimicking their citation pattern towards the rest of the network, which means that they will all be nevertheless put into the same block by SBM. Proving such explanations would require a more detailed analysis.

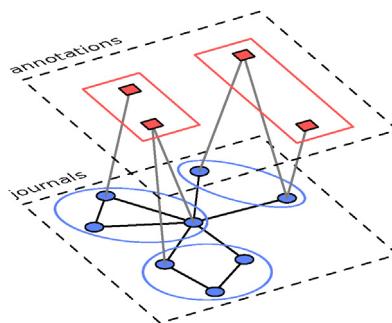


Fig. 6. Schematic representation of joint journal-annotations model. Journals with citations connecting them (blue circles and black lines) are augmented with annotations (red squares and grey lines). SBM is fitted onto the whole network such that blocks of journals are separate from blocks of annotations (blue circles and red squares respectively). Note that a journal can have multiple annotations, or it can be unannotated. Here we use subfield and field classifications as the annotations of journals, but any other data on journals can be used. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

7. Predictive power of subject categorisations

Citation networks can be augmented with a wealth of information about articles, journals, and authors. These include subfields, tags, keywords, author affiliations, etc. In this work we use classification of journals into subject categories (sub-fields) provided in the dataset, and we are interested in how much does this classification correspond to the structural blocks found in the citation patterns.

The comparison of two partitions—such as classifications, clusterings, or block structures—is in general often done using some comparison measure, such as Jaccard index, Omega index, and Variation of Information (Meilă, 2007). It is typical to compare partitions arising from a classification given in the data and the groups arising from the network structure as returned by some community detection method (Bommarito, Katza, & Zelnick, 2010; Chen & Redner, 2010; Hric, Darst, & Fortunato, 2014; Lancichinetti & Fortunato, 2009; Yang & Leskovec, 2015). This is a viable option in our case as well, but we would have to make a choice of a comparison measure. Because the question of how similar two partitions are is ill-defined, each comparison measure realises it differently and can even return different results (Fortunato & Hric, 2016; Meilă, 2007; Traud, Kelsic, Mucha, & Porter, 2011).

Instead of asking how similar the two partitions are, we ask the question: *what can we know about the citations of a journal from its classification?* Exactly this question is answered by including node annotations into SBM as it is done by Hric et al. (2016), which is based on the notion that annotations on nodes are just meta-information one has about the network—there is no principled difference between the data about connections between two nodes (links) and between a node and its annotations. In literature dealing with the community detection in networks this distinction between data and annotations is often made explicit, either by treating annotations as a sort of “ground truth” for groups (Yang & Leskovec, 2012a, 2012b, 2015), or as features that need to be learned by the model (Newman & Clauset, 2016). Here instead, annotations are treated as nodes of a bipartite network consisting of “data” nodes (journals in the citation network) and “annotation” nodes (subfields or fields of the journals), and connection exists between data node and all of its annotations (there can be any number of them, including zero).

Fig. 6 illustrates the resulting combined network that consists of two kinds of nodes (journals and annotations) and two kinds of links (citations and journal-annotation assignment). SBM is then fitted with a constraint that each inferred block must contain only one kind of nodes. The benefits of this procedure is that the node annotations contribute to the inferred node blocks, and annotations are also grouped into blocks of “equivalence”.

Working in this framework, the question from the beginning of this section can be formulated as: *how much information gain does one get about links of a single node, after learning the node’s annotations?* To answer it the following procedure is used. A small fraction of nodes is removed from the network (5% or 100, whichever is smaller), turning them into “extra nodes”—the nodes we are missing the information on, and would like to know our chances in correctly guessing where their links connect to. Then, the blocks are inferred for both the original network (without annotations), and one described above (with annotations included in an additional layer). The probabilities for node’s links to connect to nodes that belong to existing blocks are defined only by the block the node belongs to. Without annotations to tell us which block does the extra node belong to, the only thing we know about the node is its degree, and thus our best guess for the probability of this node to belong to a block is to use the size distribution of the blocks, as it is the only information we have. In case we do have the node’s annotations, we know its links in the annotations layer which narrows our choice of blocks it can belong to and thus raises the probabilities of guessing the correct links. If we denote the probability for guessing all links of node i without knowing annotations with P_i and the same by using annotations as $P_i(\text{ann})$ we can quantify the relative improvement with the *predictive likelihood ratio* λ_i :

$$\lambda_i = \frac{P_i(\text{ann})}{P_i + P_i(\text{ann})}. \quad (3)$$

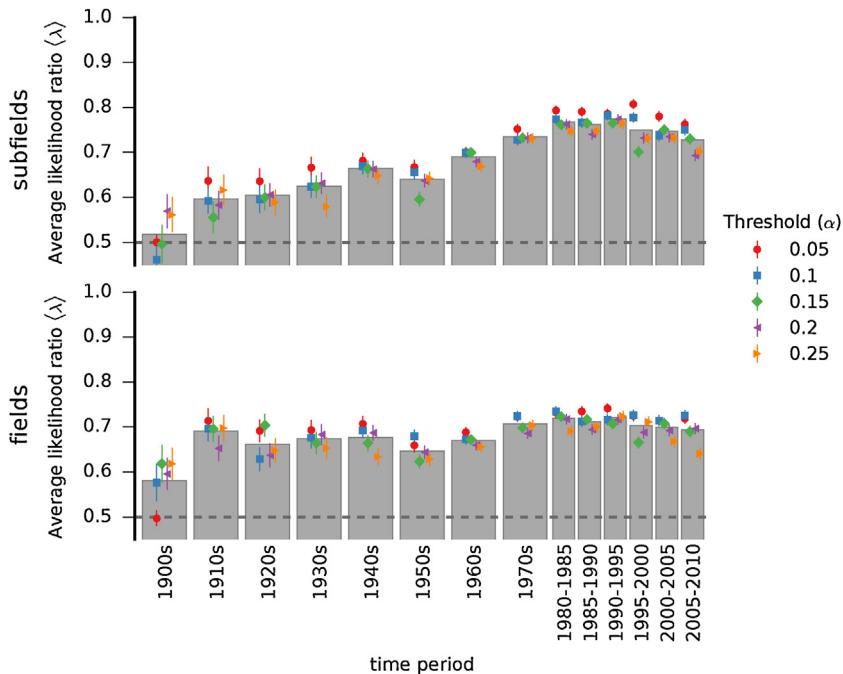


Fig. 7. Node prediction performance, measured by the average predictive likelihood ratio $\langle \lambda \rangle$ for subfields and fields [see Eq. (3)]. The $\langle \lambda \rangle$ values are calculated for simplified networks (see Section 2.2) for 14 time slices used previously, with five threshold values $\alpha: 0.05, \dots, 0.25$. Each bar corresponds to a single time window, with the height of the bar being the average over α values shown as dots on top of each bar. For each α value, the average and standard deviation over ten samples is shown. Each sample is formed by randomly removing 5% or 100 nodes, whichever is smaller.

The predictive likelihood ratio λ_i takes values from $[0, 1]$ and is above 0.5 if annotations improve link prediction power, around 0.5 if they do not change it, and below 0.5 if they decrease it. The average λ_i of all sample nodes is the *average predictive likelihood ratio* $\langle \lambda \rangle$ for a dataset.

This measure is not sensitive to the total number of blocks or annotations, as it depends only on the “power” of those annotations to predict the block structure. The only important thing is how aligned the annotations are to the structural blocks.

We use the average predictive likelihood ratio $\langle \lambda \rangle$ to measure the ability of subfield and field classifications to predict journals’ citations in the simplified networks (see Section 2.2). Predicting the links of the simplified networks is equivalent to predicting which journals are the most important sources and destinations for the citations of the extra nodes, because these networks only include the most important links for each node and do not include the actual citation counts for the links. The values for $\langle \lambda \rangle$ are presented in Fig. 7 for each individual time slice, and for five threshold levels $\alpha: 0.05, 0.1, 0.15, 0.2$, and 0.25 .

Overall, both subfield and field classification correlate positively with the citation structure. The only exception is the low threshold networks for 1900s, in which knowing the journal’s (sub)field does not help in predicting its citations. Low threshold values in this already small network caused the loss of large fraction of links and nodes, which lowered the quality of the approximation by the simplified and thresholded network. Higher threshold values do not have this problem.

Subfields are more predictive with time, although there is a slight decline for the last 15 years, with the possible reason being a some sort of over-specialization of subfields, which does not necessarily correspond to the citation patterns of the journals being classified. Fields, on the other hand, remain a good proxy for large-scale citation structure throughout the whole time period.

7.1. Predictability of individual fields

The method described in the previous section answers the question of how much information we gain about journal’s citations if we know what subfields it is classified into. We will next divide this question into smaller parts, and ask how much information do we gain by knowing that a journal belongs to a specific subfield.

Using the model from Hric et al. (2016) it is possible to calculate how much information gain (for guessing node’s links) does a single annotation provide, in comparison to a case where annotations are assigned randomly. Information gain relative to the random case is defined as predictiveness μ_a , it is defined per annotation block a , and it is not affected by the total number of blocks or annotations. Further details and formulas can be found in Hric et al. (2016).

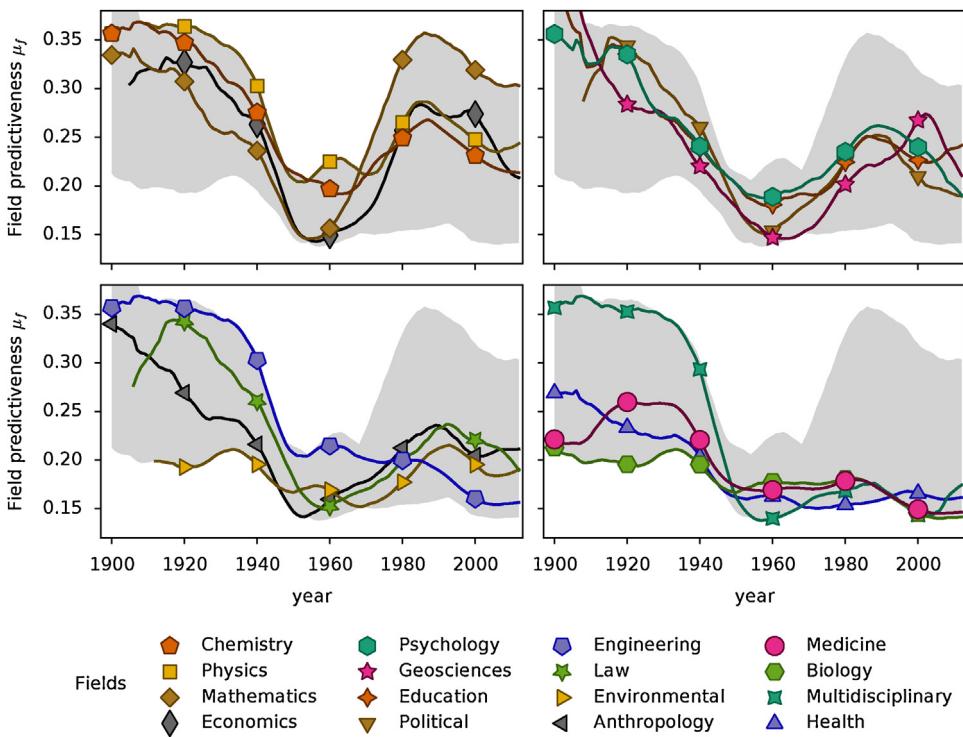


Fig. 8. Predictiveness measured by μ of the top 16 fields, for classification into fields, for all available years. The values are ten-year sliding averages of average μ for all time windows using that year. Each panel contains a group of four fields in the order of decreasing μ after 1960. Shaded region in the background is the total span of values.

Here we again consider subfields and fields as annotations of journals. After fitting the SBM onto the two-layered network of citations and annotations, the predictiveness of each block of subfields (or fields) is calculated. Fields in the same block inherit predictiveness of the block, which follows from the SBM assumption that all annotations in a block are equivalent. By calculating the field predictivenesses for all time windows, on top of being able to compare fields to each other, their change in time can also be tracked both relatively and absolutely.

Because the implementation of SBM fitting algorithm is probabilistic, variations in results are to be expected with small differences in the networks (for instance two 10-year windows with 9-year overlap), and even in different runs of fitting function on the same network. These variations cause μ_a values for the network from adjacent time windows to vary considerably, obscuring more general trends. Taking an average μ_a of all time windows that a year belongs to, we are able to overcome these fluctuations. Additional clarity is achieved by ten-year sliding averages of these values.

Here we present the predictivenesses of individual fields, for the case where classification into fields was used. The results for the case where subfields were used instead, are presented in [Appendix D](#).

Based on field predictiveness in [Fig. 8](#), the time can be split into three periods: before 1940s, the transition period, and after 1970s. Before 1940s the fields have on average higher predictiveness, but since the data is quite scarce for this period, one needs to be careful when drawing conclusions. In the transition period all the fields have very poor predictiveness, experiencing a rebound after 1970 for all but handful of fields in the last panels (Engineering, Environmental, Medicine, Biology, Multidisciplinary, and Health). This can be a sign of major changes science has gone through after the WWII. Mathematics has the highest predictiveness in the third period by a visible margin, while before 1940s the best scoring fields are Engineering and Multidisciplinary. It has risen sharply from the bottom in the transition period to the top in just 25 years, i.e. from 1955 to 1980. This means that citation patterns of Mathematics papers became more characteristic after 1970s, which is picked up by SBM and Mathematics journals end up in a small number of exclusive blocks. For Engineering it is the opposite: it fared very high before 1940s, did not suffer hard in the transition, but never recovered. The same can be said about Multidisciplinary field: it had even sharper drop and did not really recover.

On the other side of the spectrum are often large fields (Engineering, Medicine, Biology), or related to a large field (Environmental, Anthropology, and Health). Their large size means that they contain rich structure within themselves, which gets detected by the SBM as large number of blocks. Hence knowing just the field label tells little about the small blocks within the field.

8. Conclusions and discussion

The tools developed in network science are routinely used to analyse the citation networks, and in particular network clustering methods are often used to identify large-scale patterns in these networks. Recently, there has been significant advances in stochastic block modelling methods within the network science literature, but large-scale analysis of citation patterns using these methods have been missing until now. The main difference between the SBM methods and the conventional network clustering is that the SBM methods can be used to detect a variety of mesoscale structures—not just dense subgraphs. Therefore, the SBM methods lead to compelling advantages over traditional network clustering only for networks in which the dominant structure is not a collection of dense communities that are sparsely connected to each other. In this article we have shown that this is the case for citation networks by fitting the SBMs to them: the blocks that best explain these networks are only rarely even weak communities. This observation is true across multiple resolution levels and through the whole history of modern science.

The idea of finding sets of journals with similar citation patterns and similar roles is not new, but the SBM formalises it in a way that is novel to citation network analysis. Nodes in a block do not necessarily share the same links to all other nodes, but they do share similar links to all other blocks. Previously similar structures were found by calculating similarity matrices between journals (or other elements) and applying data clustering methods that work on the similarities. This type of multi-stage method requires researcher to first choose a similarity measure and then a clustering method and the type of clusters found reflect the unique combination of these two choices. The SBM method is more transparent: there is a single easy-to-interpret model that is fitted to the data. Model fitting has its own pitfalls and problems, but many of these difficulties have been solved for SBM in the recent literature leading to a very robust and efficient fitting methods. This point is demonstrated by the fact that the SBM method works consistently with minimal preprocessing for both relatively small networks representing citation patterns in the 1900s and in the recent large-scale networks.

Even though the citations between journals are consistently organized in a way that the SBMs, best explaining their creation, are mostly not community-like even in the weak sense (Radicchi et al., 2004), there might be other scientometric data displaying stronger clustering structure. Furthermore, there might be reasons for trying to find community-like structure that originate from research questions that differ from the one laid out here. For example, graph-clustering methods based on stationary random walkers, such as the InfoMap method (Rosvall & Bergstrom, 2008), have been used to analyse citation networks (Velden, Yan, & Lagoze, 2017). In this case the question to be answered by the study, is about how a random walker perceives the network, not necessarily how the network was created. For instance, the random walker could model a researcher following citations in the attempt to build a mental image of the network structure. It is interesting and important to study the differences in the groups found from these different starting points (Gläser et al., 2017), and to do so one needs to keep in mind and justify the fundamental assumptions of each method and not apply them blindly. In order to facilitate such comparison we have made the detailed results of this study public (<http://doi.org/10.5281/zenodo.1255837>). (Unfortunately, the paper-level data is proprietary and cannot be published, instead we can only publish the aggregated results of our analysis including the block hierarchies, citation counts between the blocks, and journal citation networks.)

The journals are classified manually into subfields and fields, and these classes can be viewed as human-curated blocks. Such classification block structures should correspond to the structure of the citation networks, and it is interesting to study how these two are related. A typical—although often flawed (Hric et al., 2014)—approach for doing this is to consider the classification blocks as “ground truth” and compare them to structural blocks using some subset of the many measures designed to quantify similarity or dissimilarity of partitions. Here, we have taken a more direct and well-grounded approach to ask the question of how much can the field classifications help in predicting the citation patterns. We used the SBMs to test this, and found that overall the subfield classifications have become more helpful. The ability to predict citations based on categorization was further split for different time periods and for different fields revealing a rich variety of temporal patterns.

We have demonstrated that the groups of journals found with the SBM method are often not communities in the sense that they have less citations going inside the group than in and out of the group. This result points towards the explanation that the citation networks are not only organized in communities but contain more rich structures. However, in all statistical inference the results depend on the validity of the family of models that were fitted to the data. That is, we have found out that the hypotheses about a process that has generated the network structure that contain mixed types of blocks are better than the ones with mostly community-like blocks (or only a single block), but many other models could be conceived and they might turn out to explain the citation network structure better than the SBMs we used. Alternative models, such as SBMs with overlapping blocks or some completely different kinds of generative models, could be compared to the models we used by utilizing the minimum description length principle also used here to find the SBMs that best explain the data.

The methods based on SBMs are elegant and principled answers to the need present in the literature to find groups of journals and other entities sharing similar citation patterns and creating maps of science, and therefore the SBM methods have the potential of becoming one of the standard tools in analysing citation networks in the future. The work presented here only lays a basis for such future analysis, even though we have analysed modern science in its full length and width. There are several immediate research directions that are still interesting and open. Obviously, case studies can be focused on subsets of journals, at various levels of details, and for any time period, not just the ones considered here. This approach would allow one to zoom into the individual blocks and investigate their composition and characteristics in detail, and possibly shed more light on the mechanisms creating the block structures we have observed in this study. Further, as the

large-scale networks constructed here are based on only contemporary citations, one can explore the role of older literature by including the long-range citations that were omitted here to the analysis.

The possibility to include auxiliary data to the process of inferring blocks in citation network also opens up new types of research and new avenues. One can study how the rich annotation data of scientific articles is related to the structure of the citation networks. For example, the relationships of keywords, authors, institutional affiliations, and countries to the structure of the networks can be analysed. The SBM method used here partitions annotations—in addition to journals—into blocks of equivalence, and these annotation blocks deserve to receive detailed attention. For instance, one can find blocks of authors, institutions or countries that have the same role in the structure of the citation networks. Lastly, rich hierarchical structures returned by SBM are ripe for comparison to existing ontologies that is left for future work.

Author contributions

Conceived and designed the analysis: Darko Hric and Mikko Kivela.

Contributed data or analysis tools: Darko Hric and Mikko Kivela.

Performed the analysis: Darko Hric.

Wrote the paper: Darko Hric, Kimmo Kaski and Mikko Kivela.

Acknowledgements

We thank Tiago P. Peixoto for providing technical help regarding `graph-tool` Python module, and useful comments on the manuscript. This work was supported by the European Community's H2020 Program under the scheme "INFRAIA-1-2014-2015: Research Infrastructures", Grant agreement No. 654024 "SoBigData: Social Mining & Big Data Ecosystem" (<http://www.sobigdata.eu>). D.H. and K.K. gratefully acknowledge MULTIPLEX, grant number 317532 of the European Commission. We acknowledge the computational resources provided by the Aalto Science-IT project. Certain information included herein are results of analysis using the Science Citation Index Expanded, Social Science Citation Index and Arts & Humanities Citation Index, prepared by Clarivate Analytics®, Philadelphia, Pennsylvania, USA, © Copyright Clarivate Analytics® 2018. All rights reserved.

Appendix A. Dataset: basic statistics

The total number of journals, articles and citations in the dataset is 76418, 55199417 (of which 38212193 have citations flowing in or out) and 632340116, respectively.

The number of journals, articles, in- and out- citations for every year are presented in Fig. 9, left. Exponential growth of science in all aspects is clearly visible. On the right are plotted numbers of journals and citations in networks for each time window. They are also growing exponentially which prompted us to use shorted time windows after 1970s. Due to using only the references not older than the length of the time window, only a certain fraction of citations is used, which is plotted with grey lines on the right side of Fig. 9. Ten-year windows retain around 70% of the citations, while 5-year ones around 35%.

In contrast to all other measures, number of subject categories grows only linearly in time (Fig. 10, left). It needs to be pointed out that these values do not necessarily represent the reality truthfully, since it is assumed that classifications of a journal remain the same throughout its lifespan, for simplicity reasons. Fractions of classified journals depicted on the right side of Fig. 10 show that most journals are classified to at least one subfield, with the lowest fraction of classified journals being around year 1970.

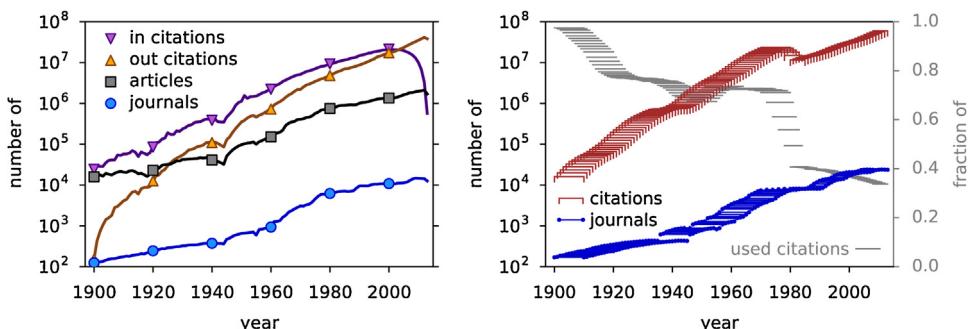


Fig. 9. Number of journals, articles, and citations. Left: values for every year in the dataset. Right: values for every network made from slicing the data in time windows, including the fraction of used citations out of all citations pertinent on journals in each time window.

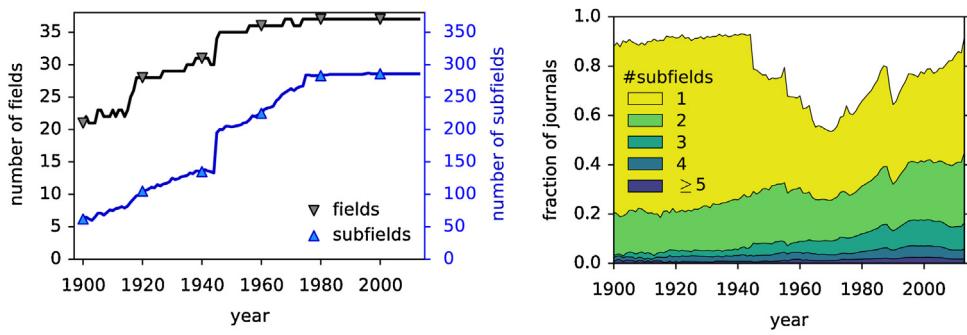


Fig. 10. Number and distribution of subject categories. Left: number of fields and subfields for every year. Right: fractions of journals classified in different number of subfields.

Appendix B. Number of blocks vs the number of (sub)fields

B.1 Distribution of block sizes

Distributions of block sizes for the time periods of 1920s, 1960s, and 1995–2000, are shown in Fig. 11. Block sizes span several orders of magnitude, and only the lowest level have a distribution that resemble scale-free, while higher ones start having characteristic scale—block sizes are clearly grouped around a certain value. This is particularly evident for fields and subfields levels.

The evolution of the number of blocks is shown in Fig. 12. Since the sizes of the networks of each slice grows exponentially (Fig. 9), the super-linear growth of the number of blocks of all levels is to be expected. Number of fields and subfields grow much slower (Appendix A).

B.2 Finding the best matching hierarchy levels

SBM finds multiple partitions of nodes into blocks (levels), organized hierarchically: blocks in one level are elements of blocks from the level above. Field and subfield assignments represent another hierarchy of levels. In this paper we compared the partitions found by the SBM with field and subfield assignments. In order to do this, and given that there are usually more than two levels in the SBM hierarchy, we needed to find which of the SBM levels correspond the best to the classification levels.

Since we take no a priori assumption about the similarity of the partitions, in order to find the best matching levels from two hierarchies, we look only at the “resolution” or “granularity” of each level. This information is fully contained in the distribution of block sizes, so the best matching levels are those that have the most similar distribution of block sizes.

Given that block size distributions have similar shapes (Appendix B.1), the total number of blocks (or equivalently, the average block size, because the total number of nodes is the same in the comparison) is a good proxy for similarity. To test this assumption, we compared the results of this simple approach with results obtained from comparing the full distributions using Jensen–Shannon divergence (JSD). Table 1 contains detailed values for the time periods of 1920s, 1960s, and 1995–2000, and a summary statistics for all 104 time slices. In the three time slices, JSD mostly suggests levels that have the most similar number of blocks (column B_I). When this is not the case, the best matching level by the number of blocks has the second smallest value of JSD. This is consistent across the whole span of slices, which can be seen from the last table that shows the number of times JSD suggested a level that is different from the most-similar-number-of-blocks level by the value “offset”. JSD is off by ± 1 , but most often it is the same as the number-of-blocks case.

Given these results, there is no strong argument to use the full block size distributions, so we decided to use levels with the most similar number of blocks.

B.3 Note on logarithmic binning

Due to the fact that distributions of group sizes have a fat-tail, using discrete distribution similarity measures directly is not the best approach, because our distributions have most values equal to zero, and thus do not contribute to the score. Most values being zero is a consequence of “fat-tail” distribution of values, where many values are small, and just a few are very large. In other words, the probability of any group having a certain size is not uniform across the range of sizes. Large groups have smaller chance of being matched (the chance that there are two large groups of the exactly same sizes, one in each partition being compared), so in effect small groups will be overestimated. This problem can be circumvented by using

Table 1

The best matching hierarchical levels according to the number of blocks and Jensen–Shannon divergence of the block size distributions. Tables in the first row contain values for field and subfield levels on one side, and all hierarchical levels on the other, for the time periods of 1920s, 1960s, and 1995–2000. The number of fields and subfields (B) are given in brackets. Columns are: the hierarchical level, number of blocks (B_l), and Jensen–Shannon divergence between block size distributions (JSD). The most similar level according to each criteria is marked in bold: the most similar number of blocks, and the smallest value of JSD. The bottom table contains the number of occurrences of a specific offset, defined as the difference between the best-matching level according to the number of blocks, and the JSD.

1920s		
Level	B_l	JSD
Fields ($B=29$)		
1	177	0.2651
2	79	0.1498
3	39	0.0443
4	21	0.0506
5	8	0.2751
6	3	0.6359
7	2	0.6292
Subfields ($B=119$)		
1	177	0.0525
2	79	0.0190
3	39	0.0980
4	21	0.1852
5	8	0.4457
6	3	0.6475
7	2	0.6426
1960s		
Fields ($B=38$)		
1	1266	0.4333
2	466	0.2952
3	229	0.1800
4	108	0.0579
5	51	0.0828
6	21	0.2077
7	10	0.3302
8	4	0.4188
9	2	0.6062
Subfields ($B=264$)		
1	1266	0.2319
2	466	0.0919
3	229	0.0257
4	108	0.0794
5	51	0.1832
6	21	0.3561
7	10	0.5863
8	4	0.5599
9	2	0.6931
1995–2000		
Fields ($B=39$)		
1	2780	0.5948
2	885	0.4954
3	457	0.4112
4	227	0.2794
5	124	0.1830
6	60	0.1266
7	26	0.1641
8	11	0.3091
9	5	0.4332
10	2	0.4510
Subfields ($B=289$)		
1	2780	0.4696
2	885	0.2947
3	457	0.1855
4	227	0.1088
5	124	0.1171
6	60	0.2303
7	26	0.4721
8	11	0.4191
9	5	0.4960
10	2	0.5185
All slices (104 in total)		
Offset	Count	
fields		
-1	4	
0	54	
1	46	
Subfields		
-2	1	
-1	39	
0	64	

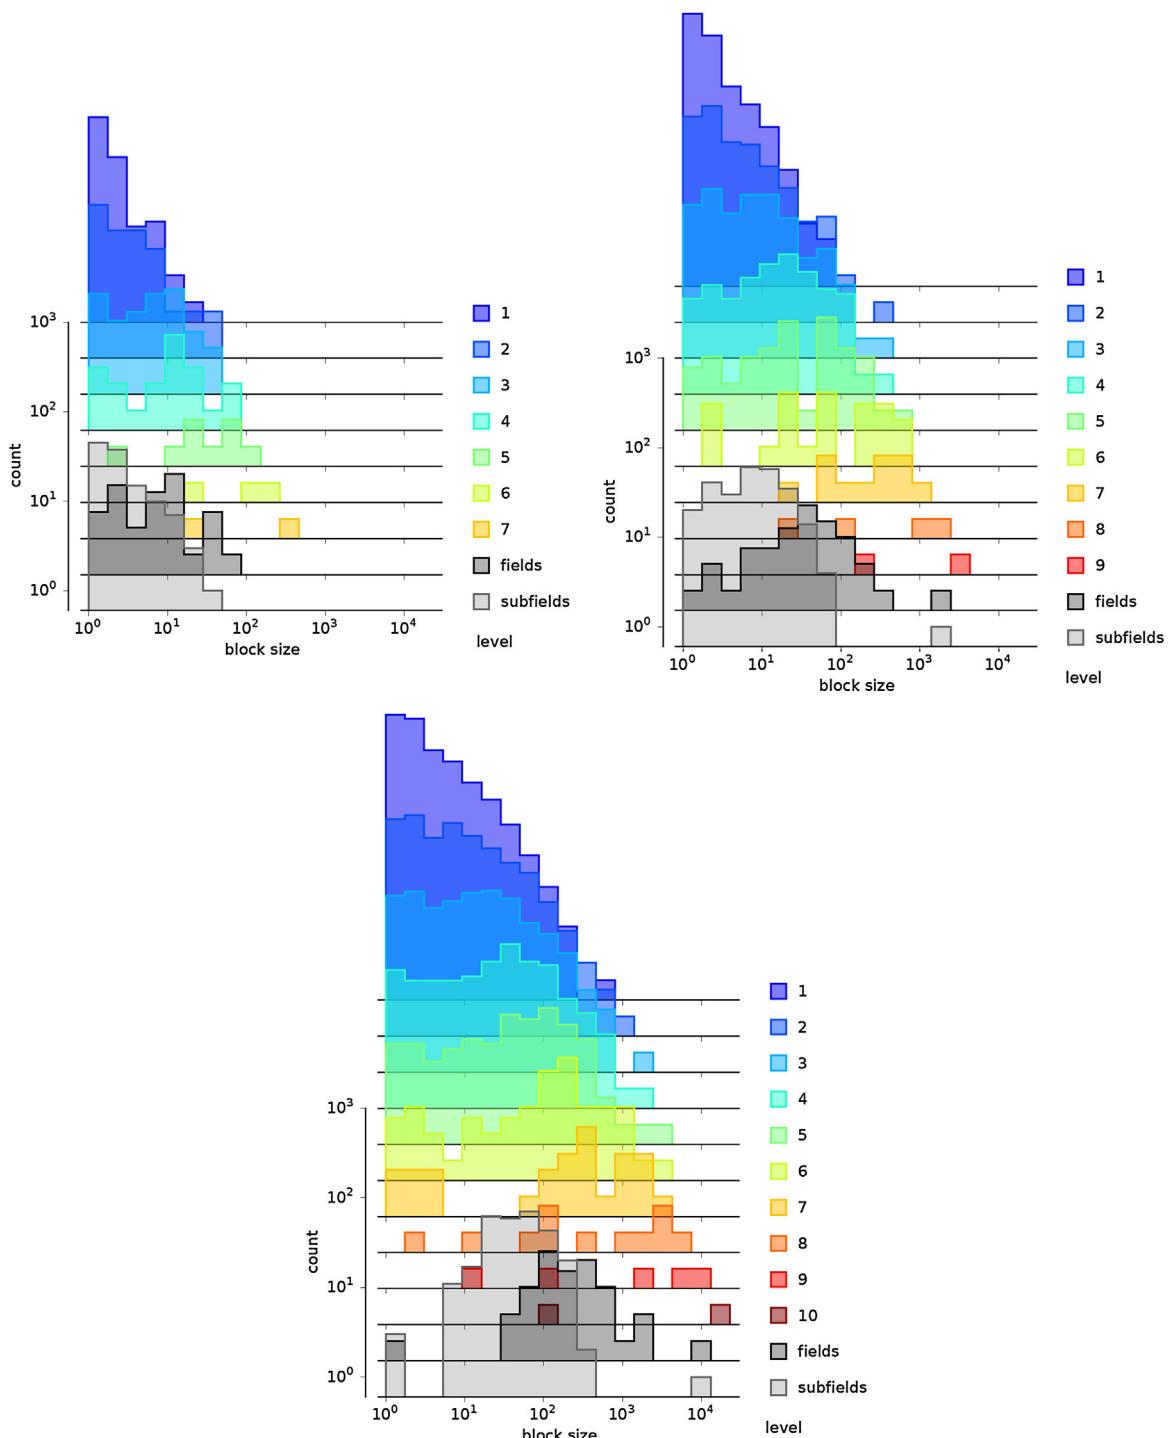


Fig. 11. The distributing of block sizes per hierarchical level, fields and subfields for the time periods of 1920s, 1960s, and 1995–2000.

logarithmic binning, which we used for calculating Jensen–Shannon divergence, and it is also why we used logarithmic scale in Fig. 11, and in many other places. Logarithmic binning introduces additional arbitrariness in the form of the number and size of the bins. In our tests, this choice had a small effect on the final results, but since we chose the simpler approach, it did not receive further attention.

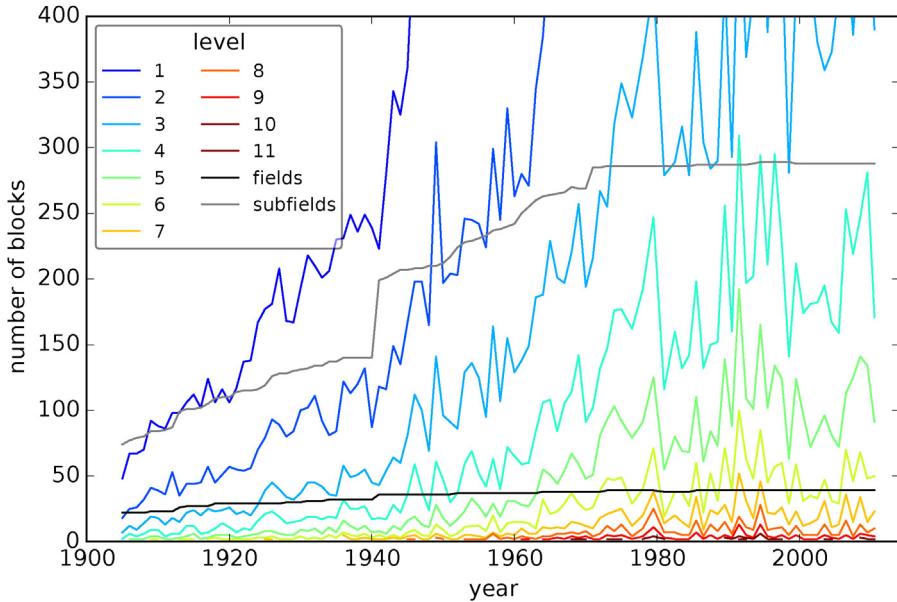


Fig. 12. The evolution of the number of blocks per hierarchical levels, fields and subfields.

Appendix C. Alternative approach for selecting contemporary citations

The simplest approach to selecting the contemporary citations is to divide the data into strict time windows. This means that both citing and cited paper have to be published in the same time window. In our work we have used a more forgiving approach that does not require both citing and cited papers to be within a given time window, one of them is enough. Additional requirements are that the time difference between them is less than the length of the window, and that the paper that is not inside the window is published in a journal that was active (it published papers) during the observed time window. Here we present the results obtained using the more strict approach.

Fig. 13 shows the evolution of block connectivities in time for 16 largest fields. In comparison to Fig. 5, most fields behave very similarly, with some notable differences. The fraction of outgoing citations are decreasing instead of increasing in time for most levels in education, and for journals level in Chemistry. Perhaps the largest difference is the shrinkage of incoming citations for multidisciplinary fields, in all levels and during the whole time period.

Fig. 14 shows average predictive ratio for unknown nodes, that measures how well can node's links be predicted knowing only its annotations. The average values are similar to the ones in Fig. 7, with larger variability across both values of the threshold parameter α and within multiple samples for the same α .

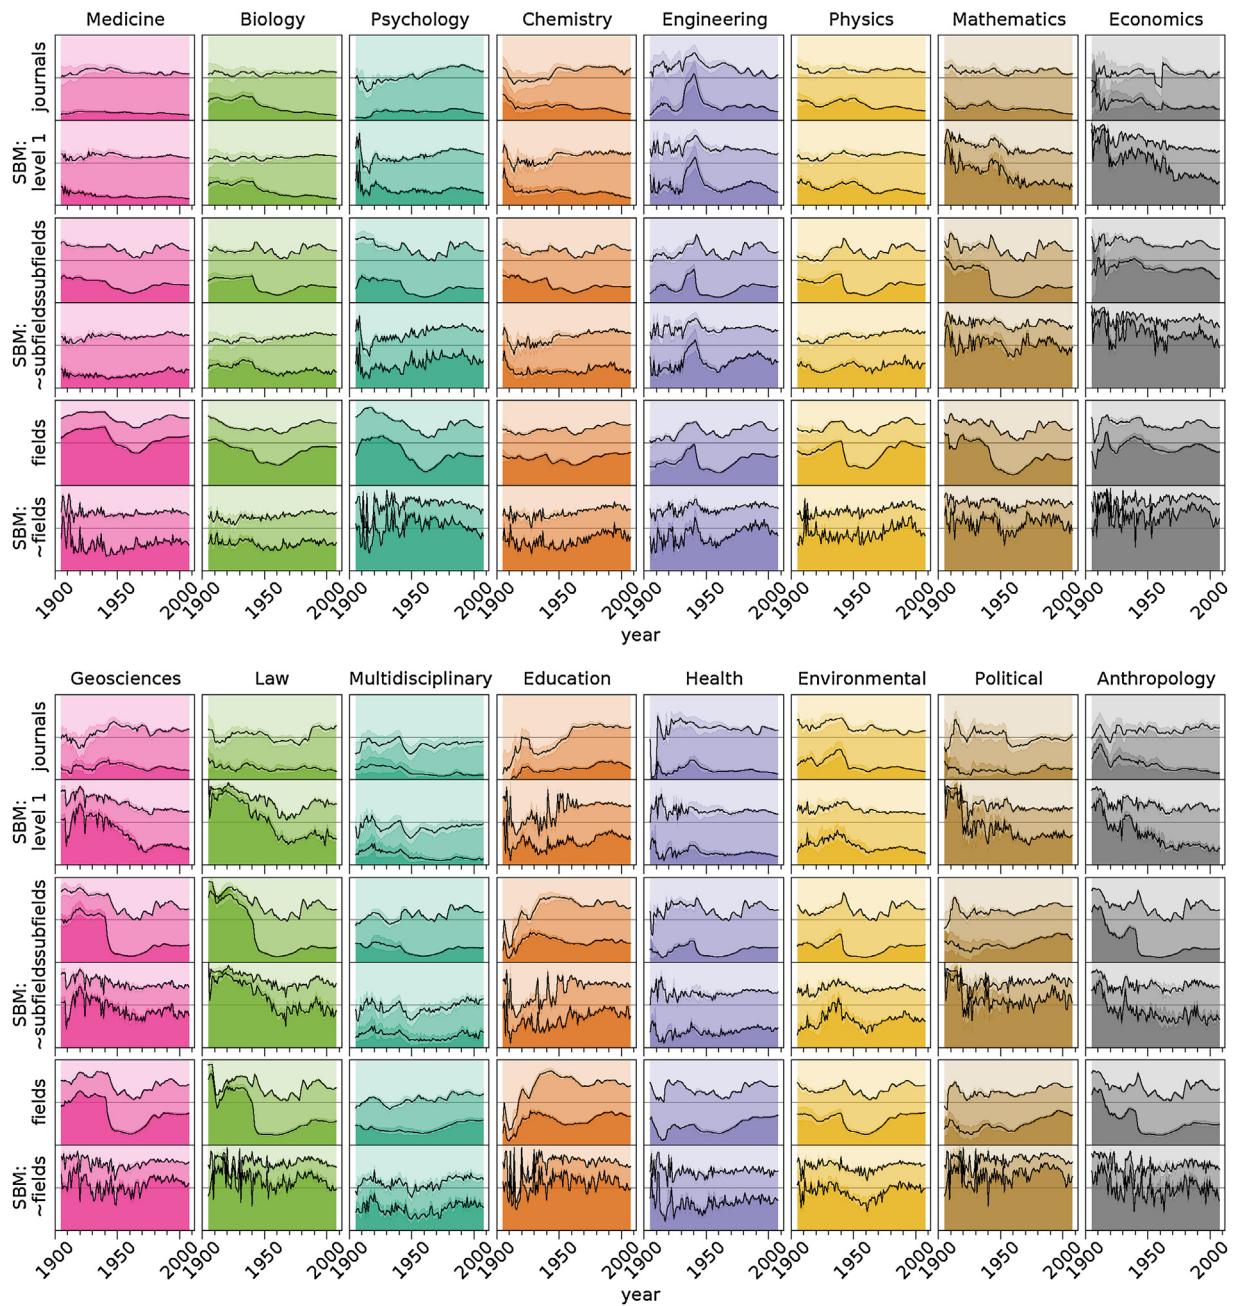


Fig. 13. Evolution of structural and classification block connectivities in time, for 16 largest fields, using an alternative citation selection criteria. Fields are grouped in columns (colours are the same as in Fig. Fig. 3) and each row is for different type of blocks, from top: journals themselves; the first SBM level; subfields classification; SBM level with the closest matching number of blocks to the number of subfields (SBM: ~subfields); fields classification; and SBM level with the closest matching number of blocks to the number of fields (SBM: ~fields). Time (in years) is on the horizontal axis, while vertical axis is split by relative citation flows of blocks containing journals from the respective field: outgoing, incoming and internal flows (from top to bottom, respectively). The results are very similar to Fig. 5. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

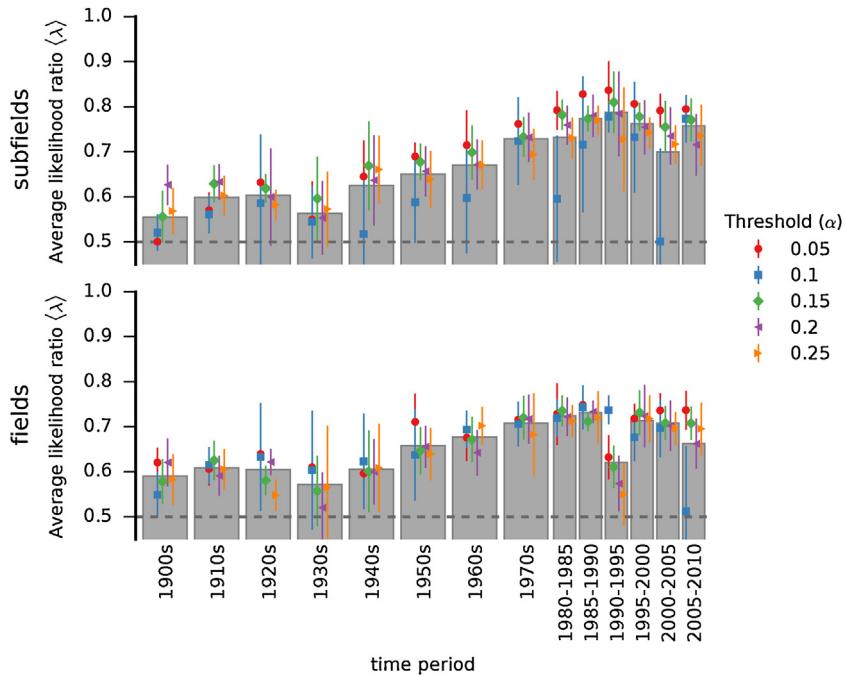


Fig. 14. Node prediction performance, measured by the average predictive likelihood ratio ($\langle \lambda \rangle$) for subfields and fields [see Eq. (3)], using an alternative citation selection criteria. The ($\langle \lambda \rangle$) values are calculated for simplified networks (see Section 2.2) for 14 time slices used previously, with five threshold values $\alpha: 0.05, \dots, 0.25$. Each bar corresponds to a single time window, with the height of the bar being the average over α values shown as dots on top of each bar. For each α value, the average and standard deviation over ten samples is shown. Each sample is formed by randomly removing 5% or 100 nodes, whichever is smaller. Averages are similar to the ones in Fig. 7 with much higher standard deviations.

Appendix D. Predictiveness of individual fields

Field predictivenesses (μ_f) from Section 7.1 are calculated using field categories as “meta-information”. The same can be calculated for the case when subfields are used instead, the results of which are presented in Fig. 15. Although μ values are calculated for subfields, values for fields are calculated as averages of all subfields belonging to a field. In addition to averages, errors of the mean are shown as shaded regions.

In general, subfields are slightly better predictors of journal’s citations than fields. This is true for all subfields, regardless of the field they belong to, which in consequence prevents clear separation of fields based on this criteria. Nevertheless, the order of fields by predictiveness of their subfields is roughly the same as for the predictiveness of fields themselves (c.f. Fig. 8 in Section 7.1).

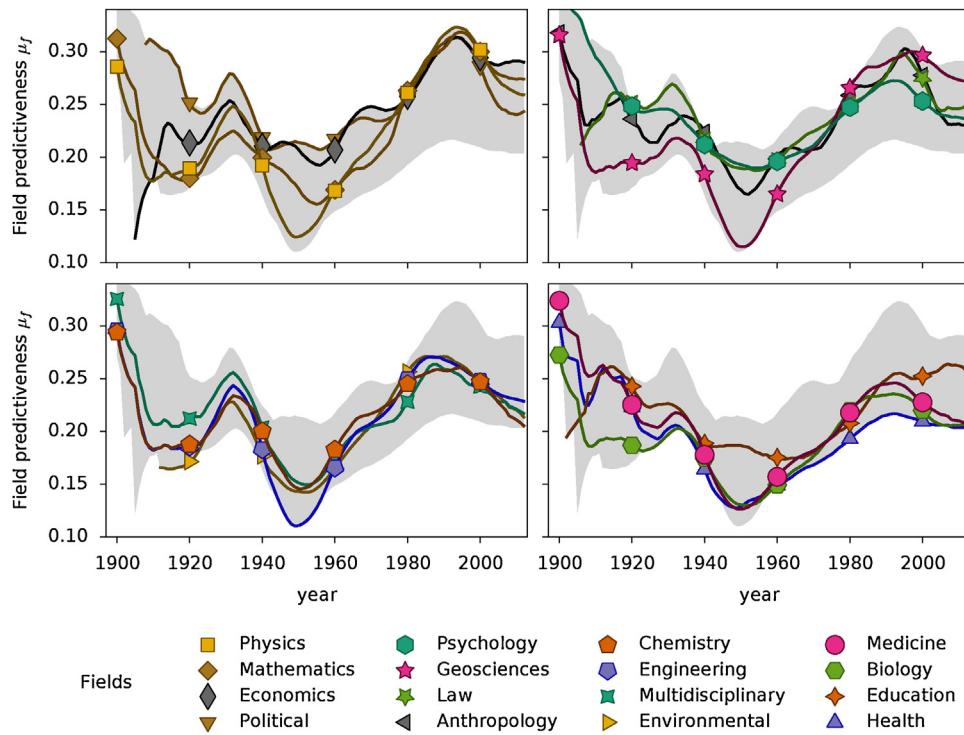


Fig. 15. Predictivenesses measured by μ of the top 16 fields, for classification into subfields, for all available years. The values are the averages of overlapping time windows, additionally smoothed by 10-year sliding average. Each panel contains a group of four fields in the order of decreasing μ after 1960. Total span of values is drawn as shaded region in the background.

Appendix E. Formulas for connectivity patterns of merged blocks

Each level in a hierarchy takes the blocks from the level below it as nodes of the network that it models. In this process, the properties of the small blocks (\bar{s}_{int} , \bar{s}_{in} , and \bar{s}_{out}) are propagated in non-trivial way into the properties of the merged block.

Here we consider the properties of a block that is formed when two blocks at the lower level are merged (see Fig. 16). Internal flow of the merged block (S_{int}) contains all flows inside a dashed rectangle in Fig. 16, while in- and out- flows (S_{in} , S_{out}) are sums of flows of the constituent blocks:

$$\begin{aligned} S_{int} &= s_{int}^1 + s_{int}^2 + s_{12} + s_{21}, \\ S_{in} &= s'_{in}^1 + s'_{in}^2, \\ S_{out} &= s'_{out}^1 + s'_{out}^2, \end{aligned} \tag{4}$$

where s_{int}^b is the internal flow of block b , $s_{b1,b2}$ is the flow between blocks $b1$ and $b2$, and $s'_{in/out}^b$ are in- and out- flows of block b that are not internal to the resulting block. These flows relate to the flows of the small blocks in the following way:

$$\begin{aligned} \bar{s}_{int} &= s_{int}, \\ \bar{s}_{in}^{b1} &= s_{b2,b1} + s'_{in}^{b1}, \quad \text{where } b1, b2 \in \{(1, 2), (2, 1)\}. \\ \bar{s}_{out}^{b1} &= s_{b1,b2} + s'_{out}^{b1}, \end{aligned} \tag{5}$$

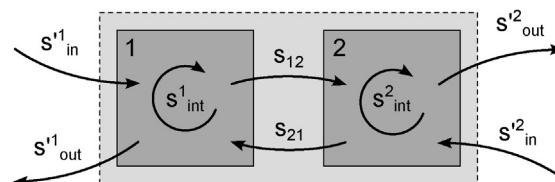


Fig. 16. Schematic representation of citation flows in the merge of two blocks.

Normalized flows of the merged block are obtained by dividing the absolute flows by the total flows of the block exactly as also described in Eq. (1):

$$\begin{aligned}\bar{s}_{\text{int}} &= S_{\text{int}} / S_{\text{tot}}, \\ \bar{s}_{\text{in}} &= S_{\text{in}} / S_{\text{tot}}, \quad \text{where } S_{\text{tot}} = S_{\text{int}} + S_{\text{in}} + S_{\text{out}}. \\ \bar{s}_{\text{out}} &= S_{\text{out}} / S_{\text{tot}}.\end{aligned}\tag{6}$$

Expressing the normalized flows of the merged block with the normalized flows of the constituent blocks, we get:

$$\begin{aligned}\bar{s}_{\text{int}} &= \frac{s_{\text{tot}}^1}{S_{\text{tot}}} \bar{s}_{\text{int}}^1 + \frac{s_{\text{tot}}^2}{S_{\text{tot}}} \bar{s}_{\text{int}}^2 + \frac{s_{12} + s_{21}}{S_{\text{tot}}}, \\ \bar{s}_{\text{in}} &= \frac{s_{\text{tot}}^1}{S_{\text{tot}}} \bar{s}_{\text{in}}^1 + \frac{s_{\text{tot}}^2}{S_{\text{tot}}} \bar{s}_{\text{in}}^2 - \frac{s_{12} + s_{21}}{S_{\text{tot}}}, \\ \bar{s}_{\text{out}} &= \frac{s_{\text{tot}}^1}{S_{\text{tot}}} \bar{s}_{\text{out}}^1 + \frac{s_{\text{tot}}^2}{S_{\text{tot}}} \bar{s}_{\text{out}}^2 - \frac{s_{12} + s_{21}}{S_{\text{tot}}},\end{aligned}\tag{7}$$

where

$$s_{\text{tot}}^b = s_{\text{int}}^b + s_{\text{in}}^b + s_{\text{out}}^b + s_{12} + s_{21}, \quad \text{for } b \in \{1, 2\}.\tag{8}$$

The three formulas from Eq. (7) have the same form:

$$\bar{s}_f = \frac{s_{\text{tot}}^1}{S_{\text{tot}}} \bar{s}_f^1 + \frac{s_{\text{tot}}^2}{S_{\text{tot}}} \bar{s}_f^2 \pm \frac{s_{12} + s_{21}}{S_{\text{tot}}} = w_1 \bar{s}_f^1 + w_2 \bar{s}_f^2 \pm w_{12}.\tag{9}$$

where w_1, w_2 , are the fractions of flows pertinent to blocks 1 and 2, and w_{12} is the fraction of flow that goes between blocks 1 and 2, out of total flow S_{tot} .

Appendix F. Stochastic blockmodel: basic formulas

Here we present some basic formulas and equations underpinning the stochastic blockmodel. A proper theoretical introduction to SBM would require a lot more space, so the interested reader is encouraged to consult the cited literature.

The probability of observing a specific realization of adjacency matrix A given the parameters $\theta = \{B, p_r, w_{rs}\}$ is (Zhang, Krzakala, Reichardt, & Zdeborová, 2012):

$$P(A, \mathbf{b} \mid \theta) = \prod_{i < j} \left[w_{b_i b_j}^{a_{ij}} (1 - w_{b_i b_j})^{(1-a_{ij})} \right] \prod_{n=1}^N b_n,\tag{10}$$

where a_{ij} is an element of the adjacency matrix A , node i belongs to block b_i , $w_{b_i b_j}$ is the expected number of edges between blocks b_i and b_j , N is the total number of nodes, and \mathbf{b} is a specific realization of node-block assignments. This equation assumes the nodes are assigned to B blocks randomly by a multinomial distribution $P(b_i=r)=p_r, \forall i$.

The formula in Eq. (10) tells us how probable is a specific realization of a network given the known parameters of an SBM, but we are more often interested in finding the parameters that are most likely to produce a specific realization of the network. This task is known as *statistical inference*, which finds the most probable parameters of the model by, for instance, expectation maximization or variational analysis.

In case of an SBM the number of blocks B is usually an unknown parameter, and it is the one that defines the model itself: in order to find B *model selection* needs to be used. A principled approach to model selection is *minimum description length*, which seeks to describe the system with the least amount of information possible.

If we assume a microcanonical formulation of the SBM, each network in the ensemble has the same probability of realization: $P = 1/\Omega$, where Ω is the number of graphs in the ensemble (Peixoto, 2012b). In this case the log-likelihood has the form:

$$L = \ln P = -\ln \Omega = -S,\tag{11}$$

where S is the entropy of the ensemble. In order to avoid “overfitting” (describing the system with an overly complicated model) instead of minimizing S one should minimize the sum of S and “information necessary to describe the model itself” M (Peixoto, 2013):

$$\Sigma = S + M.\tag{12}$$

S and M can be expressed in terms of the parameters of the SBM as:

$$\begin{aligned} S &\cong E - \frac{1}{2} \sum_{rs} e_{rs} \ln \left(\frac{e_{rs}}{n_r n_s} \right), \\ M &\cong Eh \left(\frac{B(B+1)}{2E} \right) + N \ln B, \end{aligned} \quad (13)$$

where E is the number of edges in the network, e_{rs} is the number of edges between blocks r and s , n_r is the number of nodes in block r , $h(x) = (1+x)\ln(1+x) - x\ln x$ is the binary entropy function, and $E \gg 1$.

For degree-corrected SBM entropy from Eq. (13) has the form (Peixoto, 2012b, 2013):

$$S \cong -E - \sum_k N_k \ln k! - \frac{1}{2} \sum_{rs} e_{rs} \ln \left(\frac{e_{rs}}{e_r e_s} \right), \quad (14)$$

where N_k is the number of nodes having degree k , and $e_r = \sum_s e_{rs}$ is the number of half-edges (stubs) incident on the block r .

For the hierarchical SBM the formulas in Eq. (13) have the form:

$$\begin{aligned} S &= S_0 + \sum_{l=1}^L S_l, \\ M &= \sum_{l=1}^L [S_l + M_{l-1}], \end{aligned} \quad (15)$$

where L is the total number of levels, S_l is the entropy of layer l , and M_l is S of the level l . Level-specific elements in Eq. (15) are (Peixoto, 2014b):

$$\begin{aligned} S_l &= \sum_{r>s} \ln \binom{n_r n_s}{e_{rs}} + \sum_r \ln \binom{\binom{n_r}{2}}{e_{rr}/2}, \\ M_l &= \ln \binom{B_l}{B_{l-1}} + \ln B_{l-1}! - \sum_r \ln n_r^l!, \end{aligned} \quad (16)$$

where $\binom{n}{m} = \binom{n+m-1}{m}$ is the number of m -combination with repetitions from a set of size n , B_l is the number of blocks in the level l , and n_r^l is the number of nodes in the block r at the level l , and the rest are the same as before.

References

- Aithouse, B. M., West, J. D., Bergstrom, C. T., & Bergstrom, T. (2009). Differences in impact factor across fields and over time. *Journal of the American Society for Information Science and Technology*, 60(1), 27–34.
- An, Y., Janssen, J., & Milios, E. E. (2004). Characterizing and mining the citation graph of the computer science literature. *Knowledge and Information Systems*, 6(6), 664–678.
- Barrat, A., Barthélémy, M., Pastor-Satorras, R., & Vespignani, A. (2004). The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11), 3747–3752.
- Barucca, P., & Lillo, F. (2016). Disentangling bipartite and core-periphery structure in financial networks. *Chaos, Solitons and Fractals*, 88, 244–253.
- Bianconi, G. (2009). Entropy of network ensembles. *Physical Review E – Statistical, Nonlinear, and Soft Matter Physics*, 79(3), 1–10.
- Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., & Hwang, D. (2006). Complex networks: Structure and dynamics. *Physics Reports*, 424(4–5), 175–308.
- Bommarito, M. J., Katza, D. M., & Zelnard, J. L. (2010). On the stability of community detection algorithms on longitudinal citation data. *Procedia – Social and Behavioral Sciences*, 4, 26–37.
- Boyack, K. W., & Klavans, R. (2014). Creation of a highly detailed, dynamic, global model and map of science. *Journal of the Association for Information Science and Technology*, 65(4), 670–685.
- Boyack, K. W., Klavans, R., & Börner, K. (2005). Mapping the backbone of science. *Scientometrics*, 64(3), 351–374.
- Carpenter, M. P., & Narin, F. (1973). Clustering of scientific journals. *Journal of the American Society for Information Science*, 24(6), 425–436.
- Chen, C. (2013). *Mapping scientific Frontiers*. London: Springer London.
- Chen, P., & Redner, S. (2010). Community structure of the physical review citation network. *Journal of Informetrics*, 4(3), 278–290.
- Choi, D. S., Wolfe, P. J., & Airolidi, E. M. (2012). Stochastic blockmodels with a growing number of classes. *Biometrika*, 99(2), 273.
- Cozzens, S., Gatchair, S., Kang, J., Kim, K. S., Lee, H. J., Ordóñez, G., et al. (2010). Emerging technologies: Quantitative identification and measurement. *Technology Analysis & Strategic Management*, 22(3), 361–376.
- de Solla Price, D. J. (1965). Networks of scientific papers. *Science*, 149(3683), 510–515.
- van Eck, N. J., & Waltman, L. (2009). How to normalize co-occurrence data? An analysis of some well-known similarity measures. *Journal of the American Society for Information Science and Technology*, 60(8), 1635–1651.
- Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, 486(3–5), 75–174.
- Fortunato, S., & Barthélémy, M. (2007). Resolution limit in community detection. *Proceedings of the National Academy of Sciences of the United States of America*, 104(1), 36–41.
- Fortunato, S., & Hric, D. (2016). Community detection in networks: A user guide. *Physics Reports*, 659, 1–44.
- Gläser, J., Glänzel, W., & Scharnhorst, A. (2017). Same data–different results? Towards a comparative approach to the identification of thematic structures in science. *Scientometrics*, 111(2), 981–998.

- Grossman, J. W. (2002, January). The evolution of the mathematical research collaboration graph. *Congressus Numerantium*, 201–212.
- Guimerà, R., Sales-Pardo, M., & Amaral, L. A. N. (2004). Modularity from fluctuations in random graphs and complex networks. *Physics Review E*, 70, 025101.
- Holland, P. W., Laskey, K. B., & Leinhardt, S. (1983). Stochastic blockmodels: First steps. *Social Networks*, 5(2), 109–137.
- Hric, D., Darst, R. K., & Fortunato, S. (2014). Community detection in networks: Structural communities versus ground truth. *Physics Review E*, 90, 062805.
- Hric, D., Peixoto, T. P., & Fortunato, S. (2016). Network structure, metadata and the prediction of missing nodes. *Physical Review X*, 031038(2160), 13.
- Janssens, F., Zhang, L., Moor, B. D., & Glänzel, W. (2009). Hybrid clustering for validation and improvement of subject-classification schemes. *Information Processing and Management*, 45(6), 683–702.
- Karrer, B., & Newman, M. E. J. (2011). Stochastic blockmodels and community structure in networks. *Physical Review E – Statistical, Nonlinear, and Soft Matter Physics*, 83(1), 1–10.
- Kessler, M. M. (1963). Bibliographic coupling between scientific papers. *American Documentation*, 14(1), 10–25.
- Kruskal, J. B. (1956). On the shortest spanning subtree of a graph and the traveling salesman problem. *Proceedings of the American Mathematical Society*, 7(1), 48–50.
- Lambiotte, R., & Panzarasa, P. (2009). Communities, knowledge creation, and information diffusion. *Journal of Informetrics*, 3(3), 180–190.
- Lancichinetti, A., & Fortunato, S. (2009). Community detection algorithms: A comparative analysis. *Physical Review E*, 80(5), 1–11.
- Lancichinetti, A., & Fortunato, S. (2012). Consensus clustering in complex networks. *Scientific Reports*, 2, 336.
- Latouche, P., Birmele, E., & Ambroise, C. (2012). Variational Bayesian inference and complexity control for stochastic block models. *Statistical Modelling*, 12(1), 93–115.
- Lehmann, S., Jackson, A. D., & Lautrup, B. E. (2008). A quantitative analysis of indicators of scientific performance. *Scientometrics*, 76(2), 369–390.
- Leskovec, J., Lang, K. J., Dasgupta, A., & Mahoney, M. W. (2009). Community structure in large networks: Natural cluster sizes and the absence of large well-defined clusters. *Internet Mathematics*, 6(1), 29–123.
- Leydesdorff, L. (2006). Can scientific journals be classified in terms of aggregated journal-journal citation relations using the Journal Citation Reports? *Journal of the American Society for Information Science and Technology*, 57(5), 601–613.
- Leydesdorff, L., & Rafols, I. (2009). A global map of science based on the ISI subject categories. *Journal of the American Society for Information Science and Technology*, 60(2), 348–362.
- Leydesdorff, L., & Rafols, I. (2012). Interactive overlays: A new method for generating global journal maps from Web-of-Science data. *Journal of Informetrics*, 6(2), 318–332.
- Leydesdorff, L., Carley, S., & Rafols, I. (2013). Global maps of science based on the new Web-of-Science categories. *Scientometrics*, 94(2), 589–593.
- Lorrain, F., & White, H. C. (1971). Structural equivalence of individuals in social networks. *Journal of Mathematical Sociology*, 1(1), 49–80.
- Macdonald, P. J., Almaas, E., & Barabási, A. L. (2005). Minimum spanning trees of weighted scale-free networks. *EPL (Europhysics Letters)*, 72(2), 308.
- Marshakova, I. V. (1973). System of document connections based on references. *Naucho-Tekhnicheskaya Informatsiya Seriya 2-Informatsionnye Protsessy i Sistemy*, 6, 3–8.
- Meilă, M. (2007). Comparing clusterings – An information based distance. *Journal of Multivariate Analysis*, 98(5), 873–895.
- Modha, D. S., & Spangler, W. S. (2000). Clustering hypertext with applications to web searching. In *Proceedings of the eleventh ACM on hypertext and hypermedia – HYPERTEXT '00* (pp. 143–152).
- de Moya-Anegón, S. G. F., Vargas-Quesada, B., Chinchilla-Rodríguez, Z., Corera-Álvarez, E., Muñoz-Fernández, F. J., & Herrero-Solana, V. (2007). Visualizing the marrow of science. *Journal of the American Society for Information Science and Technology*, 58(14), 2167–2179.
- Nerur, S., Sikora, R., Mangalaraj, G., & Balijepally, V. (2005). Assessing the relative influence of journals in a citation network. *Communications of the ACM*, 48(11), 71–74.
- Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2), 167–256.
- Newman, M. E. J. (2010). *Networks: An introduction*. Oxford University Press.
- Newman, M. E. J., & Clauset, A. (2016). Structure and inference in annotated networks. *Nature Communication*, 7.
- Newman, M. E. J., & Girvan, M. (2004). Finding and evaluating community structure in networks. *Physics Review E*, 69(2), 026113.
- 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.
- Pan, R. K., Petersen, A. M., Pammolli, F., & Fortunato, S. (2016). The memory of science: Inflation, myopia, and the knowledge network. ArXiv e-prints.
- Parolo, P. D. B., Pan, R. K., Ghosh, R., Huberman, B. A., Kaski, K., & Fortunato, S. (2015). Attention decay in science. *Journal of Informetrics*, 9(4), 734–745.
- Peixoto, T. P. (2012a). Entropy of stochastic blockmodel ensembles. *Physical Review E*, 85(5), 056122.
- Peixoto, T. P. (2012b). Entropy of stochastic blockmodel ensembles. *Physical Review E*, 85(5), 056122.
- Peixoto, T. P. (2013). Parsimonious module inference in large networks. *Physics Review Letters*, 110(14), 148701.
- Peixoto, T. P. (2014a). The graph-tool python library. Figsshare.
- Peixoto, T. P. (2014b). Hierarchical block structures and high-resolution model selection in large networks. *Physics Review X*, 4(1), 011047.
- Peixoto, T. P. (2015). Model selection and hypothesis testing for large-scale network models with overlapping groups. *Physics Review X*, 5, 011033.
- Persson, O. (2010). Identifying research themes with weighted direct citation links. *Journal of Informetrics*, 4(3), 415–422.
- Pieters, R., Baumgartner, H., Vermunt, J., & Bijmolt, T. (1999). Importance and similarity in the evolving citation network of the International Journal of Research in Marketing. *International Journal of Research in Marketing*, 16(2), 113–127.
- Porter, A. L., & Rafols, I. (2009). Is science becoming more interdisciplinary? Measuring and mapping six research fields over time. *Scientometrics*, 81(3), 719–745.
- Porter, M. A., Onnela, J. P., & Mucha, P. J. (2009). Communities in networks. *Notices of the AMS*, 56(9), 1082–1097.
- Punj, G., & Stewart, D. W. (1983). Cluster analysis in marketing research: Review and suggestions for application. *Journal of Marketing Research*, 20(2), 134.
- Radicchi, F., Castellano, C., Cecconi, F., Loreto, V., & Parisi, D. (2004). Defining and identifying communities in networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(9), 2658–2663.
- Radicchi, F., Fortunato, S., & Vespignani, A. (2012). Citation networks. In A. Scharnhorst, K. Börner, & P. van den Besselaar (Eds.), *Models of science dynamics: Encounters between complexity theory and information sciences* (pp. 233–257). Berlin, Heidelberg: Springer Berlin Heidelberg.
- Rombach, M. P., Porter, M. A., Fowler, J. H., & Mucha, P. J. (2014). Core-periphery structure in networks. *SIAM Journal on Applied mathematics*, 74(1), 167–190.
- Rosvall, M., & Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4), 1118–1123.
- Schaeffer, S. E. (2007). Graph clustering. *Computer Science Review*, 1(1), 27–64.
- Serrano, M. A., Boguñá, M., & Vespignani, A. (2009). Extracting the multiscale backbone of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 106(16), 6483–6488.
- Shibata, N., Kajikawa, Y., Takeda, Y., Sakata, I., & Matsushima, K. (2011). Detecting emerging research fronts in regenerative medicine by the citation network analysis of scientific publications. *Technological Forecasting and Social Change*, 78(2), 274–282.
- Silva, F., Rodrigues, F., Oliveira, O., da, F., & Costa, L. (2013). Quantifying the interdisciplinarity of scientific journals and fields. *Journal of Informetrics*, 7(2), 469–477.
- Small, H. (1973). Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science*, 24(4), 265–269.
- Small, H. (1999). Visualizing science by citation mapping. *Journal of the American Society for Information Science*, 50(9), 799–813.
- Small, H., & Greenlee, E. (1989). A co-citation study of aids research. *Communication Research*, 16(5), 642–666.
- Small, H., & Griffith, B. C. (1974). The structure of scientific literatures I: Identifying and graphing specialties. *Social Studies of Science*, 4(1), 17–40.

- Small, H., Boyack, K. W., & Klavans, R. (2014). Identifying emerging topics in science and technology. *Research Policy*, 43(8), 1450–1467.
- Traud, A. L., Kelsic, E. D., Mucha, P. J., & Porter, M. (2011). Comparing community structure to characteristics in online collegiate social networks. *SIAM Review*, 53(3), 526–543.
- Velden, T., Yan, S., & Lagoze, C. (2017). Mapping the cognitive structure of astrophysics by infomap clustering of the citation network and topic affinity analysis. *Scientometrics*, 111(2), 1033–1051.
- Waltman, L., van Eck, N. J., & Noyons, E. C. M. (2010). A unified approach to mapping and clustering of bibliometric networks. *Journal of Informetrics*, 4(4), 629–635.
- Wang, L., & Hopcroft, J. (2010). Community structure in large complex networks. In J. Kratochvíl, A. Li, J. Fiala, & P. Kolman (Eds.), *Proceedings of Theory and Applications of Models of Computation: 7th Annual Conference, TAMC 2010* (pp. 455–466).
- Wang, Q., & Waltman, L. (2016). Large-scale analysis of the accuracy of the journal classification systems of Web of Science and Scopus. *Journal of Informetrics*, 10(2), 347–364.
- Wang, S., & Koopman, R. (2017). Clustering articles based on semantic similarity. *Scientometrics*, 111(2), 1017–1031.
- Wasserman, S., & Anderson, C. (1987). Stochastic a posteriori blockmodels: Construction and assessment. *Social Networks*, 9(1), 1–36.
- Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications* (Vol. 8) Cambridge university press.
- Weinberg, B. H. (1974). Bibliographic coupling: A review. *Information Storage and Retrieval*, 10(5–6), 189–196.
- Xie, J., Kelley, S., & Szymanski, B. K. (2013). Overlapping community detection in networks. *ACM Computing Surveys*, 45(4), 1–35.
- Yang, J., & Leskovec, J. (2012). Community-affiliation graph model for overlapping network community detection. *IEEE 12th International Conference on Data Mining (ICDM)*, 1170–1175.
- Yang, J., & Leskovec, J. (2012). *Structure and overlaps of communities in networks*. , arXiv:12056228 [physics].
- Yang, J., & Leskovec, J. (2015). Defining and evaluating network communities based on ground-truth. *Knowledge and Information Systems*, 42(1), 181–213.
- Zhang, L., Liu, X., Janssens, F., Liang, L., & Glänzel, W. (2010). Subject clustering analysis based on ISI category classification. *Journal of Informetrics*, 4(2), 185–193.
- Zhang, P., Krzakala, F., Reichardt, J., & Zdeborová, L. (2012). Comparative study for inference of hidden classes in stochastic block models. *Journal of Statistical Mechanics: Theory and Experiment*, 2012(12), P12021.