

Large-scale structure of time evolving citation networks

E.A. Leicht^{1,a}, G. Clarkson², K. Shedden³, and M.E.J. Newman^{1,4}

¹ Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA

² School of Information, University of Michigan, Ann Arbor, MI 48109, USA

³ Department of Statistics, University of Michigan, Ann Arbor, MI 48109, USA

⁴ Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109, USA

Received 31 May 2007

Published online 5 October 2007 – © EDP Sciences, Società Italiana di Fisica, Springer-Verlag 2007

Abstract. In this paper we examine a number of methods for probing and understanding the large-scale structure of networks that evolve over time. We focus in particular on citation networks, networks of references between documents such as papers, patents, or court cases. We describe three different methods of analysis, one based on an expectation-maximization algorithm, one based on modularity optimization, and one based on eigenvector centrality. Using the network of citations between opinions of the United States Supreme Court as an example, we demonstrate how each of these methods can reveal significant structural divisions in the network and how, ultimately, the combination of all three can help us develop a coherent overall picture of the network's shape.

PACS. 89.75.Hc Networks and genealogical trees – 89.75.Fb Structures and organization in complex systems – 89.70.+c Information theory and communication theory

1 Introduction

The physics community has in recent years devoted considerable attention to the study of networks, including social networks, biological networks, information networks, and others [1–3]. Many of these networks also have long histories of study in other fields. Citation networks, which are the principal focus of this paper, have been studied quantitatively almost from the moment citation databases first became available, perhaps most famously by physicist-turned-science-historian Derek de Solla Price, who authored two celebrated papers in the 1960s and 1970s highlighting the power-law degree distributions in networks of scientific papers and developing models to explain their origin [4, 5].

A citation network is an information network in which the vertices represent documents of some kind and the edges between them represent citation of one document by another. Citation networks differ from other networks in a number of important ways. First, they are directed: citations go from one document to another and hence constitute an inherently asymmetric relationship between the vertices involved. Mathematically, the network can be represented by an adjacency matrix \mathbf{A} , with elements

$$A_{ij} = \begin{cases} 1 & \text{if there is an edge from } j \text{ to } i, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

In a directed network the adjacency matrix is, in general, asymmetric.

A second feature of citation networks is that they evolve over time as new documents are created. The evolution of the network takes a special form, in that vertices and edges are added to the network at a specific time and cannot be removed later. This permanence of vertices and edges means that the structure of the network is mostly static: it changes only at the “leading edge” of the network, the current time at which new documents are being added. Citation networks differ in this respect from other information networks such as the world wide web, in which vertices and edges can be removed as well as added and edges can be repositioned after they are added. The limited form of time evolution found in citation networks makes them, in some ways, a simpler and cleaner laboratory for the study of network growth than the web.

The combination of the two features of citation networks described above leads to a third: citation networks are acyclic, meaning there are no closed loops of citations of the form A cites B cites C cites A, or longer. When a new vertex is added to a citation network it can cite any of the previously existing vertices, but it cannot cite vertices that have not yet been created. This gives the network a clear “arrow of time”, with all edges pointing backwards in time as shown in Figure 1. As a result it is typically possible, starting from a given vertex, to find a path of citations that takes us back in time through the network, but it is not possible to find one that takes us forward again, so no

^a e-mail: eleicht@umich.edu




Fig. 1. Citations run from vertices created at later times to those created at earlier times — in the opposite direction to the “arrow of time”.

closed loops exist. (Real citation networks are often not perfectly acyclic. For example, a scientific paper can sometimes cite work that is forthcoming but not yet published, resulting in a closed loop in the network. However, such loops are rare and necessarily short, being limited by the narrow span of time over which it is possible to predict future publications. In practice, therefore, it is usually a good approximation to assume the network to be acyclic.)

Citation networks arise in a variety of different areas. We have mentioned networks of scientific citations, which have been studied by many authors since the classic work of Price mentioned above. (See, for instance, the book by Egghe and Rousseau [6] or any volume of the journal *Scientometrics*, which is entirely devoted to the quantitative analysis of scholarly authorship and citation patterns.) Citation networks of patents have, to a lesser extent, also been studied. Patents cite other patents for a variety of reasons, but most often to establish their originality and distinction from previous work. Extensive data on patent citations have become available in recent years, allowing the construction of very large citation networks [7, 8]. For example, the time evolution of the United States patent citation network has been studied recently by Csárdi et al. [9], although their approach to the analysis of temporal patterns, while both interesting and useful, is quite different from the one adopted in this paper. Very recently, there has also been some interest in legal citation networks, networks of legal opinions written by judges and others, which cite one another to establish precedent [10–12]. We make extensive use of one particular legal citation network, the network of opinions of the United States Supreme Court, as an example in this paper, although the techniques we will be considering are certainly applicable to other networks as well.

Given the wide interest in and unique structure of citation networks, it is instructive to investigate what can be learned from an analysis of the statistical patterns present in these networks. A variety of studies have been presented in the past focusing on relatively standard network measures such as degree distributions [4, 13, 14]. To investigate the time-dependent structure that is the special property of citation networks, however, other methods are needed. In this paper we present several techniques that, as we will show, are — both individually and collectively — capable of revealing interesting new structure in these networks.

2 A mixture model of citation patterns

The first analysis we describe makes use of a stochastic mixture model of the citation process, which is fitted to the observed network data using the likelihood optimization technique known as the expectation-maximization algorithm.

A crucial property affecting the structure of citation networks is the pattern over time of the citation of documents following their publication. It is interesting for instance to ask if there are typical patterns that documents follow. Are there more citations immediately after publication than later, or do they grow in frequency over time? Are documents more likely to cite recent precedents or older better-established ones? Do documents tend to cite others published during a particular time period? There could also be more than one common pattern with different documents following different patterns. If so, how can we determine those patterns, and how can we tell which pattern particular documents follow, given that citation data are inherently noisy?

As an example, we consider the network of legal citations between cases handed down by the Supreme Court of the United States, from its inception in 1789 until the present day. We will use this example throughout this paper; it is well documented, shows clear and interesting structural signatures, and has been studied much less than other types of citation networks in the past, so that, although we use the network primarily as an illustrative device, the results we derive are in many cases of interest in their own right and not just as a demonstration of our methods.

Consider the following table, which gives the dates of the citations received so far by a single example opinion handed down by the Supreme Court in the year 1900: we will take citation profiles such as this as the basic inputs in our analysis.

year	cites	year	cites	year	cites
1900	1	1907	2	1925	1
1901	4	1910	1	1936	1
1902	3	1912	2	1947	1
1904	1	1920	1		

One interesting question (there are many) is whether there are distinct eras of citation in the history of this (or any) citation network. Are there, for instance, eras in which a certain set of documents are well cited, followed perhaps by another era or eras in which that set falls out of favor to be replaced by a different one? Many readers can probably think of anecdotal cases of behavior like this in scientific citation networks. Here we place these observations on a firm analytic foundation.

We will attempt to divide the vertices in a citation network into groups by identifying similarities in their citation profiles. Our method will be to define a set of citation profiles and then self-consistently assign each case to the profile it best fits while at the same time adjusting the shape of the profiles to best fit the cases assigned to

them. The means by which we accomplish this task is the expectation-maximization (EM) algorithm [15, 16].

The EM algorithm is an established tool of statistics, but one that is relatively new to network analysis. In a previous paper we described an application of the method to the classification of vertices in static networks, both directed and undirected [17]. Here we describe a different application to the analysis of the temporal profiles of citations.

In essence the EM algorithm is a method for fitting a model to observed data by likelihood maximization, but differs from the maximum likelihood methods most often encountered in the physics literature in that it does not rely upon Markov chain Monte Carlo sampling of model parameters. Instead, by judicious use of “hidden” variables, the maximization is performed analytically, resulting in a self-consistent solution for the best-fit parameters that can be evaluated using a relatively simple iteration scheme.

Suppose we have a network of n vertices representing our documents and we believe that they can be divided into c groups, each of which is characterized by a particular probability distribution of citations over time. (Ultimately, we will vary c to find the best description for our data, but for the moment let us assume it to be fixed.) Our approach to finding the groups will be to fit the network to a model consisting of two parts: (1) a set of *time profiles* $\{\theta_r(t)\}$, one for each group, such that $\theta_r(t)$ is the probability that a particular citation received by a document in group r is made during year t ; (2) a set of *probabilities* π_r , such that π_r is the probability that a randomly chosen document belongs to group r (i.e., π_r is the expected fraction of documents belonging to group r). We fit this model to the observed data by maximizing the probability of the observed set of citations given the model — the so-called likelihood function.

Suppose that document i belongs to group g_i and let $z_i(t)$ be the number of citations that the document receives in year t . Then the probability that document i received the particular citations it did and is in group g_i , given the model parameters, is

$$\Pr(z_i, g_i | \pi, \theta) = \Pr(z_i | g_i, \pi, \theta) \Pr(g_i | \pi, \theta), \quad (2)$$

where for convenience we use π, θ to denote the entire set $\{\pi_r, \theta_r\}$. Assuming random and uncorrelated citations drawn from the time profile $\theta_{g_i}(t)$, the terms on the right-hand side are given by

$$\Pr(z_i | g_i, \pi, \theta) = k_i! \prod_{t=t_1}^{t_2} \frac{[\theta_{g_i}(t)]^{z_i(t)}}{z_i(t)!}, \quad (3)$$

$$\Pr(g_i | \pi, \theta) = \pi_{g_i}, \quad (4)$$

where $k_i = \sum_t z_i(t)$ is the in-degree of document i , i.e., the total number of citations it receives, and t_1 and t_2 are the first and last years of data in our dataset.

Now taking the product over all vertices, the likelihood of the entire data set is $L = \prod_{i=1}^n \Pr(z_i, g_i | \pi, \theta)$. In fact, we will work with the logarithm \mathcal{L} of the likelihood, which

has its maximum in the same place:

$$\mathcal{L} = \ln L = \sum_{i=1}^n [\ln \Pr(g_i | \pi, \theta) + \ln \Pr(z_i | g_i, \pi, \theta)]. \quad (5)$$

Unfortunately, \mathcal{L} depends on the group memberships g_i , which we don't know. Given the observed citation patterns, however, we can make a good guess about the group memberships, or more precisely we can compute the probability distribution of their values, which in Bayesian fashion we regard as a statement about our knowledge of the world, rather than a statement about the actual values of the group memberships, which are in theory perfectly well-defined quantities. Writing the probability of a particular assignment of vertices to groups as $\Pr(\{g_i\} | z, \pi, \theta)$, we can then calculate the expected value of the log-likelihood as the average of equation (5) over all possible assignments thus:

$$\begin{aligned} \bar{\mathcal{L}} &= \sum_{g_1=1}^c \dots \sum_{g_n=1}^c \Pr(\{g_i\} | z, \pi, \theta) \mathcal{L} \\ &= \sum_{g_1=1}^c \dots \sum_{g_n=1}^c \Pr(\{g_i\} | z, \pi, \theta) \\ &\quad \times \sum_{i=1}^n [\ln \Pr(g_i | \pi, \theta) + \ln \Pr(z_i | g_i, \pi, \theta)] \\ &= \sum_{i=1}^n \sum_{r=1}^c \Pr(g_i = r | z_i, \pi, \theta) \\ &\quad \times [\ln \Pr(g_i = r | \pi, \theta) + \ln \Pr(z_i | g_i = r, \pi, \theta)] \\ &= \sum_{i=1}^n \sum_{r=1}^c q_{ir} \left\{ \ln \pi_r + \ln k_i! + \right. \\ &\quad \left. \sum_{t=t_1}^{t_2} [z_i(t) \ln \theta_r(t) - \ln z_i(t)!] \right\}, \quad (6) \end{aligned}$$

where we have introduced the shorthand notation

$$q_{ir} = \Pr(g_i = r | z_i, \pi, \theta) \quad (7)$$

for the probability that vertex i belongs to group r , given the model and the observed citation pattern.

This expected log-likelihood represents our best estimate of the value of the log-likelihood given what we know about the system. By maximizing it, we can now calculate a best estimate of the most likely values of the model parameters, a process that involves two steps: first, we estimate the group membership probabilities q_{ir} ; second, we use those probabilities in the maximization of $\bar{\mathcal{L}}$. We take these steps in turn.

To calculate the q_{ir} we observe that

$$q_{ir} = \Pr(g_i = r | z_i, \pi, \theta) = \frac{\Pr(z_i, g_i = r | \pi, \theta)}{\Pr(z_i | \pi, \theta)}. \quad (8)$$

The two factors on the right can be determined by summing equation (2) over the appropriate sets of variables

and making use of equations (3) and (4) to give

$$q_{ir} = \frac{\pi_r \prod_t [\theta_r(t)]^{z_i(t)}}{\sum_k \pi_k \prod_t [\theta_k(t)]^{z_i(t)}}. \quad (9)$$

Once we have this expression, we can use it to evaluate the log-likelihood, equation (6), and hence to find the values of the model parameters that maximize the likelihood, which is our ultimate goal. The maximization is helped by the fact that π_r and θ_r enter equation (6) in independent terms. Considering π_r first and noting that it must satisfy the normalization condition $\sum_r \pi_r = 1$, we introduce a Lagrange multiplier α and then differentiate, holding q_{ir} constant, to get

$$\begin{aligned} 0 &= \frac{\partial}{\partial \pi_r} \left\{ \sum_{ir} q_{ir} \ln \pi_r + \alpha \left[1 - \sum_r \pi_r \right] \right\} \\ &= \frac{1}{\pi_r} \sum_{i=1}^n q_{ir} - \alpha. \end{aligned} \quad (10)$$

Rearranging this expression gives

$$\pi_r = \frac{1}{\alpha} \sum_{i=1}^n q_{ir}. \quad (11)$$

The Lagrange multiplier α is then fixed by the condition $\sum_r \pi_r = 1$ thus:

$$\sum_{r=1}^c \pi_r = 1 = \frac{1}{\alpha} \sum_{ir} q_{ir} = \frac{n}{\alpha}, \quad (12)$$

where we have made use of $\sum_r q_{ir} = 1$. Thus π_r is given by

$$\pi_r = \frac{1}{n} \sum_i q_{ir}. \quad (13)$$

In other words, the prior probability of a vertex belonging to group r is just the average over all vertices of the conditional probability of belonging to group r .

Similarly, the θ_r satisfy the normalization condition $\sum_t \theta_r(t) = 1$ for all r , so we introduce a set of c Lagrange multipliers $\{\beta_r\}$ and write

$$\begin{aligned} \frac{\partial}{\partial \theta_r(t)} \left\{ \sum_{ir} q_{ir} \sum_{t=t_1}^{t_2} z_i(t) \ln \theta_r(t) \right. \\ \left. + \sum_r \beta_r \left[1 - \sum_t \theta_r(t) \right] \right\} = 0. \end{aligned} \quad (14)$$

Again holding q_{ir} constant and employing Eq. (3), we find

$$\sum_i q_{ir} \frac{z_i(t)}{\theta_r(t)} - \beta_r = 0, \quad (15)$$

or

$$\theta_r(t) = \frac{\sum_i q_{ir} z_i(t)}{\sum_i q_{ir} k_i}, \quad (16)$$




Fig. 2. Results of the application of the EM analysis with $c = 2$ to the network of citations between Supreme Court opinions. The two curves show the fraction of cases assigned to each of the two groups found, as a function of time.

where we have evaluated β_r using the normalization condition and the fact that $\sum_t z_i(t) = k_i$ by definition.

To calculate the optimal values of the model parameters, as well as the group membership variables q_{ir} , we now need to solve equation (9) simultaneously with equations (13) and (16). The simplest way to do this is numerical iteration. Starting from an initial guess about the values of $\{\pi_r, \theta_r(t)\}$, we evaluate equation (9) and then use the results to make an improved estimate of the model parameters from equations (13) and (16). Under reasonable conditions this process is known to converge upon iteration to a self-consistent solution.

2.1 Example

As a demonstration of the EM method we have applied it to the citation network of Supreme Court cases described in Section 2. Applied to this network, the algorithm will divide the network into any requested number c of groups, such that each group is characterized by a distinctive pattern of citations to cases in that group. We have performed the analysis for a variety of different values of c . We begin with the simplest case, $c = 2$, of division into two groups. Starting with random initial values for $\{\pi_r, \theta_r\}$ and applying the EM iteration, equations (9), (13), and (16), the parameters rapidly converge to a clear split of the network into two groups. Figure 2 shows the fraction of cases assigned by the algorithm to each of the groups as a function of time. Cases are assigned in proportion to their probability of membership in each of the groups so that, for instance, a case belonging to group 1 with probability 0.7 and to group 2 with probability 0.3 contributes 0.7 of a case to the first group and 0.3 of a case to the second.

Figure 2 reveals a dramatic split between the two groups: the best fit, in the maximum likelihood sense, of the mixture model with two groups to these data produces one group containing practically all cases before 1937 and

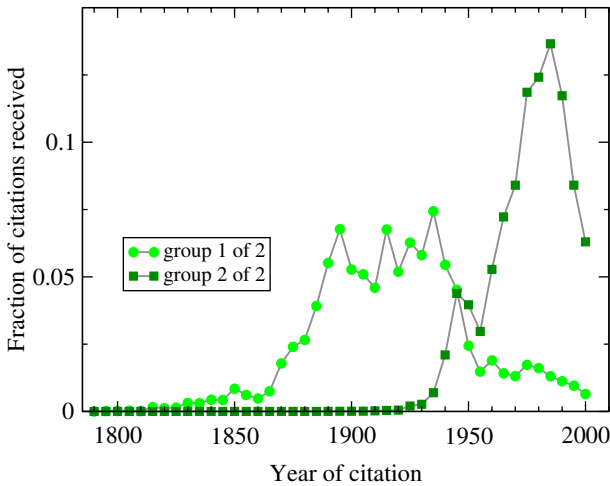


Fig. 3. The citation profiles $\theta_r(t)$ generated by the EM algorithm with $c = 2$ for the Supreme Court citation network.

another containing practically all cases after. This breakpoint coincides with a significant constitutional crisis for the Supreme Court. For the interested reader we give some further analysis in Section 5.

The EM algorithm tells us in this case that the Supreme Court’s rulings split quite cleanly into groups with distinct citation profiles. That is, the opinions of the court can be distinguished sharply by the cases that later cited them. The citation profiles themselves, meaning the temporal citation patterns represented by the parameters $\{\theta_r\}$ in the model, are shown in Figure 3. As we can see, they also divide into two time periods, which correspond closely to those of the group memberships depicted in Figure 2. This implies that the opinions that cite cases in each of our groups were handed down during roughly the same eras as the cited cases. This is not surprising if one assumes that the group divisions reflect different legal ideologies, but it is important to bear in mind that our analysis does not require it: it would be perfectly possible to detect groups that were distinguished by citations received during some entirely different era of the court arbitrarily later in its history, or even in no era at all but scattered widely over time.

We can also ask about best fits to the model for numbers of groups c greater than two. It is always the case that larger values of c will give better fits to the data, since larger values give us more parameters to fit with, but we must be wary of overfitting. In practice, we have been able to extract useful formation about networks by comparing the results for a variety of small values of c . Rigorous methods for deciding optimal values of c , such as minimum description length, methods based on approximations to the marginal likelihood, or information theoretic measures, have been developed for other applications of the EM algorithm [18,19] and we discuss these approaches elsewhere. For the moment we simply describe the results for various values of c .

Figure 4 shows results for the Supreme Court network with $c = 4$. The method again finds clear groups of cases,

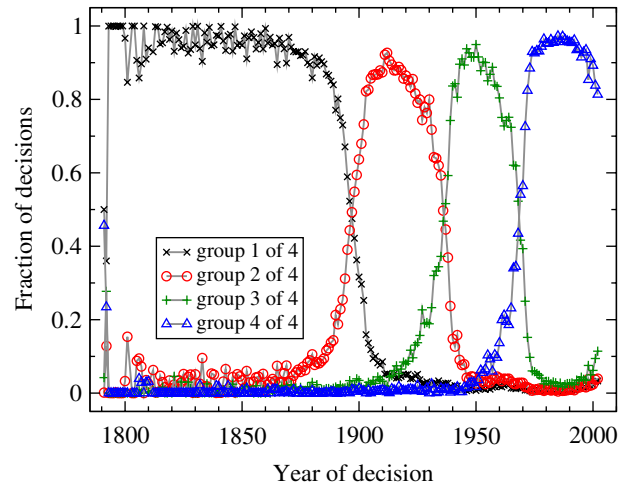


Fig. 4. Results of the application of the EM analysis with $c = 4$ to the network of citations between Supreme Court opinions.




Fig. 5. Results of the application of the EM algorithm with $c = 4$ to data for citations *made* (rather than received) by opinions in our Supreme Court dataset. The groups found are quite similar to those for the analysis based on citations received.

and as in the $c = 2$ case they are strongly delineated according to the dates of the opinions and thus appear to offer evidence for the presence of distinct eras in the court’s history. In particular, the analysis finds a clear grouping of cases between 1897 and 1937, corresponding approximately to the so-called *Lochner* era of Supreme Court jurisprudence, the significance of which is described in Section 5.

In these analyses we have characterized our documents by the pattern of citations they receive. However, one can equally well look at the pattern of citations that documents *make* and this also, at least in some cases, can be a useful cue for detecting patterns in the network. The EM algorithm can be applied to this analysis as well. The developments are identical and the same computer code can be used — one simply takes the transpose of the adjacency matrix. Figure 5, for example, shows the results of the application of the method to citations made by

the opinions in our Supreme Court dataset, with $c = 4$. As the figure shows, the results are remarkably similar to those for citations received: it appears that, in this case at least, there is a high degree of agreement about how cases should be classified into eras. This could indicate agreement between the opinions' writers and those that came after them, about the position staked out by individual opinions within the larger body of literature represented in our data set.

3 Clustering in citation networks

The general problem of the division of networks into groups of related vertices has been extensively studied in the past. The classic problem of “clustering” or “community detection” is to find groups of vertices within networks that have a higher than average density of internal edges and relatively few connections to the rest of the network [20,21]. The second analysis technique we investigate for citation networks is a clustering method of this kind. As we will see, it is instructive to compare the results with those of our EM analysis in the previous section. The two methods do not do the same thing: the EM analysis groups together vertices that have similar time profiles to their citations, while the community analysis groups together vertices that are specifically linked to one another by edges. Nonetheless, as we will show, the two approaches can produce similar outcomes, for instance in the example of the Supreme Court data set.

Considerable effort has been devoted to the development of methods to find community structure within networks. The authors are aware of dozens of different methods (at least) published within the last few years. Here we make use a method recently proposed by Newman [22] based on the maximization of the benefit function known as “modularity”. Although many competing methods appear to give excellent results, we focus on this particular method for two reasons: first, it is based on firm statistical principles that make its operation transparent to the user; second, it has been shown in recent head-to-head comparisons to give better results on standardized tests than competing methods [21].

Briefly the method works as follows. Given a network and a particular division of the vertices of that network into nonoverlapping groups or communities, the modularity is defined as the number of edges that lie within those groups minus the expected number of such edges if edges are placed at random between the vertices (but respecting vertex degree) [23]. In essence, the modularity measures whether a larger than expected number of edges fall within the groups defined. In principle, the task of finding the best division of the network into groups is then one of maximizing the modularity over all possible divisions [24]. In practice, this maximization problem is known to be NP-complete [25], so approximate solution methods must be used for all but the smallest networks. Newman's method works by rewriting the modularity in the language of linear algebra as a quadratic form involving an index vector and a characteristic matrix dubbed




Fig. 6. A histogram of the number of decisions versus the year of the decision for cases assigned to each group in the two-way split produced by the modularity maximization algorithm.

the “modularity matrix.” It can then be shown that the signs of the elements of the leading eigenvector of this modularity matrix give an approximation to the division of the network into two parts that maximizes the modularity. This approximate maximum can optionally be further refined by, for instance, applying a greedy algorithm that moves vertices between groups as described in [22]. By repeatedly dividing the network in two in this way, a network can be divided into any number of communities, although typically one stops dividing when no divisions exist that will increase the modularity any further.

This repeated subdivision of the network into smaller and smaller groups is particularly attractive for the purposes of our present analysis, because it allows us to observe the major divisions in the network first, followed by more minor ones, and to stop the process at any point to compare with our other analyses. A limitation of the method is that it is designed for use with undirected rather than directed networks. This however is not a great hindrance. It seems reasonable to consider edges in a citation network to be a sign of connection between documents, and that connection exists regardless of the direction the edge runs in. So we simply ignore the directions in our analysis and apply the eigenvector calculation to the undirected network. This approach has been taken before by other authors and appears to work well — see, for example, reference [26].

We can visualize the results of our clustering analysis in a manner similar to our visualizations of the output of the EM algorithm, as a histogram over time. The results for the leading split of the Supreme Court network into two clusters are depicted in this way in Figure 6. The results are similar to those for the EM algorithm, with a significant break around 1937. This appears to bolster the conclusions of our EM analysis, that there have been separate periods in the court's history that leave identifiable signatures in the citation record. There are some differences between the two sets of results, particularly the early

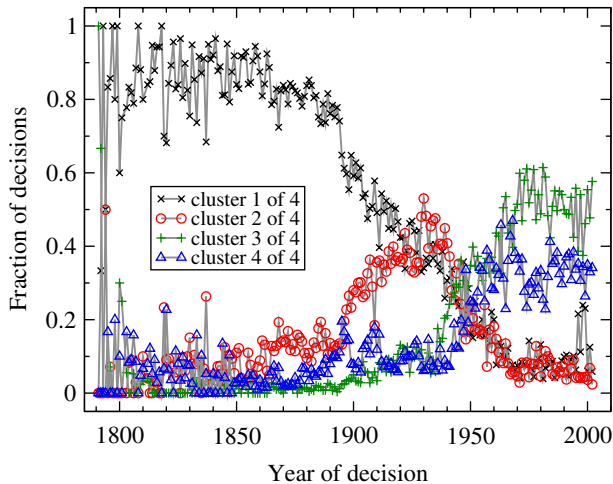


Fig. 7. A histogram of the number of decisions versus the year of the decision for cases assigned to each group in the four-way split produced by the modularity maximization algorithm.

“tail” to the second group in the clustering analysis and an overall difference in the number of cases assigned to each group. A possible explanation for these differences is that the EM analysis makes use only of citations received by cases, whereas the clustering analysis, which ignores edge direction, takes into account both citations received and citations made. This allows the classification into groups of some vertices that were unclassifiable with the EM algorithm by virtue of never receiving any citations. (About 10% of cases were never cited.) It could also be responsible for the tail in the second group because citations made, which are necessarily to cases in the past, connect vertices to earlier times, perhaps pulling them from the second group into the first in the clustering analysis.

As with the EM analysis, we can go further and look at splits into larger numbers of groups. For instance, Figure 7 shows the best split into four groups according to the modularity-based approach. Again the split is similar in overall form to the split found by the EM algorithm with $c = 4$, although the results are not as clean as those for the EM algorithm. As before, a new split point appears around 1900, which could be associated with the start of the *Lochner* era.

4 Vertex authority score and time evolution

For our third analysis, we turn away from studies of groups or clusters and focus on another class of network measures, centrality scores, which quantify the importance or influence of individual vertices in a network. As we will see, the pattern of centrality scores as a function of time in our evolving citation networks can reveal interesting patterns.

The simplest of centrality scores is the degree of a vertex. In a directed network such as a citation network, there are two degrees, the in-degree and the out-degree. It is reasonable, for instance, to imagine that important or influential vertices in a citation network will receive many

citations and therefore have high in-degree. A more sophisticated version of the same idea is eigenvector centrality [27], in which, rather than merely counting the number of citations a vertex gets, we award a higher score when the citing vertices are themselves influential. The simplest way to do this is to define the centrality to be proportional to the sum of the centralities of the citing vertices, which makes the centralities proportional to the elements of the leading eigenvector of the adjacency matrix. Unfortunately, this method does not work for acyclic directed networks, such as citation networks, for which all such centralities turn out to be zero.

An interesting variant of eigenvector centrality has been proposed by Kleinberg [28] that works well for acyclic networks. In this variant each vertex has two centralities, known as the *authority score* and the *hub score*, the first derived from the incoming links and the second from the outgoing links. In this view a “hub” is a vertex that points to many important authorities — a review paper in a scientific citation network, for instance — while an authority is a vertex pointed to by many important hubs — such as an important or authoritative research article on a particular subject. In the simplest version of the method the authority score x_i of vertex i is simply proportional to the sum of the hub scores y_j of the vertices citing it:

$$x_i = \frac{1}{\lambda} \sum_j A_{ij} y_j, \quad (17)$$

for some constant λ , while the hub score is proportional to the sum of the authority scores of the vertices it cites:

$$y_i = \frac{1}{\mu} \sum_j A_{ji} x_j. \quad (18)$$

In matrix form, these equations can be written

$$\mathbf{A}\mathbf{y} = \lambda\mathbf{x}, \quad \mathbf{A}^T\mathbf{x} = \mu\mathbf{y}. \quad (19)$$

Or, eliminating either \mathbf{x} or \mathbf{y} ,

$$\mathbf{A}\mathbf{A}^T\mathbf{x} = \lambda\mu\mathbf{x}, \quad (20)$$

$$\mathbf{A}^T\mathbf{A}\mathbf{y} = \lambda\mu\mathbf{y}. \quad (21)$$

Thus \mathbf{x} and \mathbf{y} are eigenvectors of the symmetric matrices $\mathbf{A}\mathbf{A}^T$ and $\mathbf{A}^T\mathbf{A}$ (also known as the *cocitation* and *bibliographic coupling* matrices respectively). In Kleinberg’s formulation of the problem one focuses on the leading eigenvector of each of the matrices, although in principle there could be useful information to be gleaned from other eigenvectors too.

Taking the Supreme Court network as an example again, we note that hub and authority scores have been applied to Supreme Court cases previously by Fowler and Jeon [12], who showed that it can be revealing to calculate scores not only for a complete citation network but also for subsets of the network containing only opinions handed down on or before a certain date, effectively recreating the citation network as it existed at that date. We

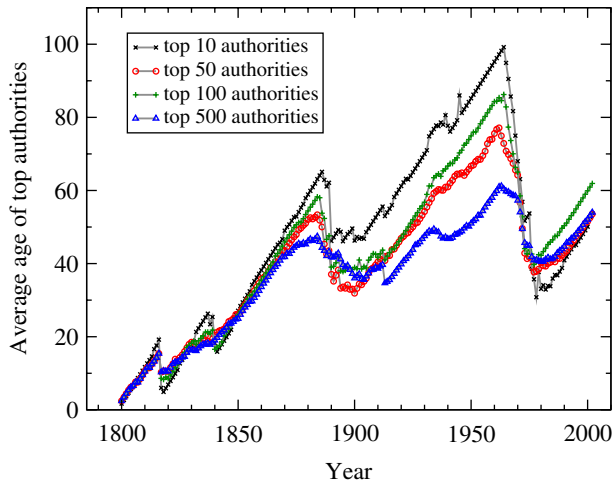


Fig. 8. The average age of the highest-authority cases in the Supreme Court citation network as a function of time.

have adopted a similar approach in our calculations, calculating authority scores for the network as it existed at some time t , as a function of t . We focus primarily on the most central cases, those with the highest scores.

Figure 8 shows one particularly informative statistic, the average age of the ten highest-ranked cases in our data set as a function of the year at which the network is cut off. As the plot shows, there is a marked trend for the average age to increase in step with the passage of time. This is precisely the behavior one would expect if the top authorities in the network are remaining the same as time goes by. Every once in a while, however, the plot shows a sudden and precipitous drop in the average age, indicating that a much younger set of vertices have, in a short space of time, taken over as the new leaders in the authority score rankings. Thus the plot indicates a repeated pattern in the evolution of the network in which a certain set of vertices — certain cases considered by the Supreme Court — remain the top authorities for substantial periods of time before being swiftly replaced by a different set. One example of such a turnover can be seen in Figure 8 around 1900 and a smaller one around 1940, dates that, as we have seen, correspond roughly to the beginning and end of the *Lochner* era. Another very large dip in the curve occurs around 1970. (Our four-group EM analysis also found a group division at approximately the same point — see Fig. 4.) The large size of this dip may be due in part to the much larger number of cases decided per year by the Supreme Court in more recent decades than in its earlier history, which makes it easier for newly appearing cases to quickly become top authorities. The results of the centrality analysis are thus compatible with but different from those of previous sections. Such variations are one reason why a variety of different analytic techniques are useful in studies of network structure.

The behavior described is clearest in the age of the top ten vertices, but persists if a different number is used. Figure 8 shows the results of the same calculation for the top 50, 100, and 500 authorities, and in each case a similar pattern of maturation followed by swift renewal is visible.

5 Discussion

Although the purpose of this paper is primarily to highlight new methods for the analysis of network data, the ultimate goal of these methods is of course to give researchers insight into the structure and meaning of their data. Thus it is interesting to ask whether the analyses described here do indeed shed light on the system studied — in this case the network of citations between Supreme Court cases. In fact the results do appear to offer an interesting perspective on the workings of the Supreme Court as we argue briefly in this section.

The United States underwent a transition from an agricultural economy to an industrial economy in the latter part of the nineteenth century. Federal and state legislators adapted to the new economic environment by passing laws that regulated emerging industries. These regulations, however, were not without opposition from those who preferred a *laissez-faire* or hands-off approach. Among those outspoken in opposition were several members of the Supreme Court and, beginning in 1897, the court began invalidating a number of cases that imposed regulations on industry and business, starting with *Allgeyer v. Louisiana*. The legal doctrines of *substantive due process* and *freedom of contract* were merged together into a significant limitation on the police power of the state. After *Allgeyer*, any statute, ordinance, or administrative act that imposed any kind of limitation upon the right of private property or freedom of contract became suspect, even if the regulation was intended to promote safety and general welfare [29].

The most famous (or infamous) of the cases to use substantive due process to invalidate state regulation was *Lochner v. New York* in 1905, a case that became so notorious that this entire era of jurisprudence, between 1897 and 1937, came to be known as the *Lochner* era. During the *Lochner* era the Supreme Court struck down nearly 200 regulations [30]. The *Lochner* era is clearly visible, for example, in our EM analysis with $c = 4$ (Fig. 4) — the analysis picks out one group of cases with start and end dates that correspond closely to the accepted dates of the era.

Ultimately, the Supreme Court’s hostility to state and federal regulation began to interfere with the “New Deal” programs instituted by US President Franklin Roosevelt to combat the Great Depression. Between 1934 and 1936, the court invalidated more federal statutes than during any other two-year period in its history and by 1936 nearly all of the statutes passed as part of the New Deal had been struck down. In response, Roosevelt launched in early 1937 a counteroffensive against the Supreme Court in which he proposed to appoint to the court up to six additional justices more receptive to the New Deal. This “court packing” plan was, to say the least, highly controversial, but Roosevelt had the support of significant majorities in both houses of Congress, and the nation as a whole, still in the throes of the depression, was eager for something new.

Following Roosevelt’s proposal, the court abruptly reversed course and, beginning in March of 1937, validated a series of state and federal measures. Contemporary

commentators have humorously dubbed this change the “switch in time that saved nine,” but whether the switch was substantive or illusory has been the subject of much debate. Some scholars believe that the court responded to political pressure while others have suggested that the court already contained a majority of justices who would have been inclined to sustain the New Deal if legislation had been drafted better or if certain unanswered questions had been appropriately posed to the court.

Our EM analysis shows a clear break around 1937, corresponding closely to the end of the *Lochner* era. It is important to appreciate that the analysis takes into account only citations received by cases, and thus that the opinions of the Supreme Court appear to have taken a substantial change of direction not merely in their conclusions but also in their arguments: later cases cited the new opinions rather than those coming before them because, presumably, their arguments better supported the decisions of the post-1937 court. Thus our analysis appears to indicate not merely a change in case outcomes that was a natural, if novel, result of positions long held by the sitting justices, but a more fundamental change in legal thinking itself — or at least its expression in the written opinions of the court and the later citation of those opinions.

6 Conclusions

In this paper we have described several methods for the analysis of citation networks, which are acyclic directed graphs of citations between documents. Using the network of citations between opinions handed down by the US Supreme Court as an example, we have described and demonstrated three analysis techniques. The first makes use of a probabilistic mixture model fitted to the observed network structure using an expectation–maximization algorithm. The second is a network clustering method making use of the recently introduced method of modularity maximization. The third is an analysis of the patterns of time variation in eigenvector centrality scores, particularly the “authority” score introduced by Kleinberg [28].

When applied to the Supreme Court network, each of these analyses reveals interesting structure, particularly highlighting qualitative changes in citation patterns that may be associated with specific eras of legal thought in the Supreme Court. However, it is in combination that the methods become most effective. Features that appear clearly in analyses performed using several different techniques possess correspondingly greater persuasive force. In the case of the Supreme Court, there emerges quite a clear picture of the eras of the court as marked by shifts in citation patterns, particularly around the time of the so-called *Lochner* era in the early 20th century.

This work was funded in part by the National Science Foundation under grant number DMS-0405348 and by the James S. McDonnell Foundation.

References

1. R. Albert, A.L. Barabási, *Reviews of Modern Physics* **74**, 47 (2002)
2. S.N. Dorogovtsev, J.F.F. Mendes, *Advances in Physics* **51**, 1079 (2002)
3. M.E.J. Newman, *SIAM Rev.* **45**, 167 (2003)
4. D.J.de S. Price, *Science* **149**, 510 (1965)
5. D.J.de S. Price, *J. Am. Soc. Inf. Sci.* **27**, 292 (1976)
6. L. Egghe, R. Rousseau, *Introduction to Informetrics* (Elsevier, Amsterdam, 1990)
7. A. Jaffe, M. Trajtenberg, *Patents, Citations and Innovations: A Window on the Knowledge Economy* (MIT Press, Cambridge, MA, 2002)
8. G. Clarkson, D. DeKorte, in *Progress in Convergence: Technologies for Human Wellbeing*, edited by W.S. Bainbridge, M.C. Roco (New York Academy of Sciences, 2006), Vol. 1093 of *Annals of the New York Academy of Science*, pp. 180–200
9. G. Csárdi, K.J. Strandburg, L. Zalanyi, J. Tobochnik, P. Erdi, *Physica A* **374**, 783 (2007)
10. F.B. Cross, T.A. Smith, A. Tomarchio, *Determinants of cohesion in the supreme court's network of precedents* (2006), <http://ssrn.com/abstract=924110>
11. J.H. Fowler, T.R. Johnson, J.F. Spriggs II, S. Jeon, P.J. Wahlbeck, *Political Analysis* doi:10.1093/pan/mpm011 (in-press)
12. J.H. Fowler, S. Jeon, *Social Networks* doi:10.1016/j.socnet.2007.05.001 (2007)
13. P.O. Seglen, *J. Am. Soc. Inf. Sci.* **43**, 628 (1992)
14. S. Redner, *Eur. Phys. J. B* **4**, 131 (1998)
15. A.P. Dempster, N.M. Laird, D.B. Rubin, *J. Roy. Stat. Soc. B* **39**, 1 (1977)
16. G.J. McLachlan, T. Krishnan, *The EM Algorithm and Extensions* (John Wiley & Sons, Inc., New York, 1997)
17. M.E.J. Newman, E.A. Leicht, *Proc. Nat. Acad. Sci. USA* **104**, 9564 (2007)
18. H. Akaike, *IEEE Trans. Auto. Control* **19**, 716 (1974)
19. G. Schwarz, *Annals of Statistics* **6**, 461 (1978)
20. M.E.J. Newman, *Eur. Phys. J. B* **38**, 321 (2004)
21. L. Danon, J. Duch, A. Diaz-Guilera, A. Arenas, *J. Stat. Mech.* P09008 (2005)
22. M.E.J. Newman, *Proc. Nat. Acad. Sci. USA* **103**, 8577 (2006)
23. M.E.J. Newman, M. Girvan, *Phys. Rev. E* **69**, 026113 (2004)
24. M.E.J. Newman, *Phys. Rev. E* **69**, 066133 (2004)
25. U. Brandes, D. Delling, M. Gaertler, R. Görke, M. Hofer, Z. Nikoloski, D. Wagner, *On Finding Graph Clusterings with Maximum Modularity*, in *Proceedings of the 33rd International Workshop on Graph-Theoretic Concepts in Computer Science* (Springer, Berlin, in press), Lecture Notes in Computer Science (in press)
26. A.E. Krause, K.A. Frank, D.M. Mason, R.E. Ulanowicz, W.W. Taylor, *Nature* **426**, 282 (2003)
27. P.F. Bonacich, *Am. J. Sociol.* **92**, 1170 (1987)
28. J.M. Kleinberg, *J. ACM* **46**, 604 (1999)
29. A.H. Kelly, W.A. Harbison, H. Belz, *The American Constitution: Its origins and development*, 7th edn. (Norton, New York, 1991)
30. M.I. Urofsky, P. Finkelman, *A March of Liberty: A Constitutional History of the United States*, 2nd edn. (Oxford University Press, New York, 2002)