

Interdisciplinarity Revealed by Transitive Reduction of Citation Networks

H. AlMuhanna^{1,2, ✉*}, V. Vasiliauskaite^{1,2,3†}, T.S. Evans^{1,2‡}

1. Centre for Complexity Science, Imperial College London, SW7 2AZ, U.K.
2. Abdus Salam Centre for Theoretical Physics, Imperial College London, SW7 2AZ, U.K.
3. Laboratory of Biosensors and Bioengineering, Institute for Biomedical Engineering, Gloriastrasse 37/39, 8092 Zurich, Switzerland

5th July 2026

Abstract

We investigate the impact of transitive reduction on citation networks. Our hypothesis is that documents which lose fewer citations under transitive reduction are likely to be multi-class (our definition of interdisciplinary), while a large loss of citations suggests a document is primarily cited within a single class. We test this hypothesis by using an artificial model of a citation network and by using data on citations from three sources: academic papers, court decisions and patents. Where needed, we applied modularity-based clustering techniques on a network defined using the undirected version of the citation network to classify documents by communities. A cluster-dependent measure was then used to classify the nodes as multi-class or single-class. Our results provide strong support for our hypothesis in all four cases.

1 Introduction

Citation data provide a useful way to study how knowledge develops across a collection of documents. Because citations leave an explicit record of intellectual connection [1, 2, 3], they allow one to examine not only the prominence of individual documents, but also broader structural patterns in how ideas spread, persist, and influence later work. This makes citation-based analysis a valuable tool for understanding the organisation and evolution of document collections across a range of domains [4, 5, 6, 7].

Previous research has used citation patterns to study breadth, diversity, and the extent to which documents connect otherwise distinct parts of the literature [8, 9, 10, 11]. Citation-based indicators such as centrality and diversity measures have also been used to study how documents or journals relate to multiple areas of research [12, 13, 14]. Related work has further shown that atypical combinations of prior knowledge can be associated with high scientific impact [15]. These studies show that citation data can reveal several different aspects of cross-area influence, although different measures capture different facets of that behaviour. In this paper, however, our aim is narrower; we study how the citation count of a document changes under the removal of redundant edges, and how that change relates to the distribution of the fields from which its citations originate.

Citation count is a commonly used measure to assess the influence of a document [16, 17, 18, 19, 20]. Documents with high citation counts are often considered significant, especially in the case of academic papers. However, relying on citation count alone may compromise capturing broad cross-class influence [21]. Citations may arise due to various motivations, including social or habitual factors, rather than genuine intellectual influence [22, 23]. In addition, many

*ORCID: 0009-0004-2564-0140; ✉ email: haya.almuhanna22@imperial.ac.uk

†ORCID: 0000-0003-2039-6236

‡ORCID: 0000-0003-3501-6486

citations may be redundant, since authors frequently draw information from more recent and up-to-date documents even when they also cite the older sources [24, 25].

Therefore, the influence of a document is not defined solely by its direct citations. Rather, a document may have a wider, often indirect, effect on how a subject evolves. These aspects are naturally represented by a network [1, 2, 3], and citation networks provide a particularly useful framework for capturing them. In citation networks, nodes represent documents, such as academic papers, patents, or legal decisions, while edges denote citations. Here we adopt the convention that edges point backward in time from newer documents to the older ones they cite. Consequently, the in-degree of a node represents its citation count, while the out-degree reflects the number of references in its bibliography. The time ordering of the edges means that a citation network is a special type of network known as a directed acyclic graph (DAG), whose unique properties we will exploit. More broadly, citation networks provide a valuable framework for analysing the structural flow of information [1, 2, 3, 4, 5, 6, 7].

Our focus is on the idea that many citations from a given document to older sources are redundant and do not necessarily imply a direct flow of information. In terms of the citation network, there are edges that can be removed in a process known as TRANSITIVE REDUCTION, illustrated in Fig. 1, without changing the connectivity of nodes. Transitive reduction gives the minimal structure needed to preserve the flow of information. One of the special properties of a directed acyclic graph, such as a citation network, is that there is no ambiguity in the transitive reduction process [26]: edges are either essential for maintaining connectivity or are redundant and can be removed. This is a distinctive feature of DAGs, in contrast to more general directed graphs, where cycles complicate the transitive reduction process. Transitive reduction is well known in computer science and discrete mathematics, with long-established algorithms that can be implemented in just a few lines [26].




Figure 1: Example of performing transitive reduction on a Directed Acyclic Graph. Redundant edges are removed but the overall connectivity of the network is preserved. The result of transitive reduction is always unique and well defined [26].

Our proposal is that documents exhibiting minimal changes in their citation count after transitive reduction tend to be multi-class, in the sense that they are cited across multiple network communities rather than predominantly within a single one. These communities are clusters of nodes which are our proxies for disciplines, fields or topics used to classify documents.

We test our hypothesis using citation datasets and an artificial model, employing richness, an information-theoretic measure of diversity, to quantify multi-class citation reach. If

our proposition holds, nodes with high diversity values should exhibit low citation loss after transitive reduction.

2 Datasets and Models

To validate our approach, we apply our methodology to citation networks generated by an artificial model: the Preferential Directed Acyclic Graph (PDAG) model. This is a generalisation of previous models and numerical implementations [27, 28, 29, 30, 31]. The nodes $n_i \in \mathcal{N}$ represent documents and are labelled by integers $i \in \{0, 1, 2, \dots, N - 1\}$, where the label of the i -th node n_i is also the time of publication of a document represented by that node. Each document n_i is assigned uniformly at random to one cluster, denoted by the label $c_i \in \mathcal{C}$ drawn from $\mathcal{C}$, the set of labels of the different clusters. Each cluster represents one type or class of document and these are our proxies for the single field, discipline or topic best associated with each document. The directed edges (n_i, n_j) from node n_i to node n_j represent links from the bibliography of the newer document n_i to the older document n_j . As a result, for any directed edge (n_i, n_j) , we have that $i > j$, ensuring an acyclic structure.

Nodes are also assigned to be one of two types; with probability p_{mc} nodes are a multi-class document, otherwise they are a single-class document. Single-class documents in cluster f can only cite single-class documents within their own cluster f . However, part of their bibliography will consist of multi-class documents n_j from cluster c_j , provided document n_j has cited at least one document in cluster f . Multi-class documents can cite documents from any cluster. The number of entries in the bibliography of any document n_i is fixed and each entry n_j is chosen with a probability roughly proportional to $(\mu + k_j^{(in)})$ where $k_j^{(in)}$ is the in-degree and μ is a model parameter which we choose to be 12.0. In addition, a fraction $p = 0.45$ of the bibliography comes from recent documents, here coming from the most recent 2200 or so documents of each cluster, i.e. when citing a recent document, we have that $2200 \gtrsim (i - j) > 0$. The remaining 55% of documents come from documents of any age. This is to simulate the well-known feature [29, 30] that recent documents appear more often in bibliographies than found in simple models using attachment probabilities of the form $(\mu + k_j^{(in)})$. The parameter values used were chosen to produce a citation network with characteristics similar to the hep-th arXiv data set discussed below including fat-tailed degree distributions [27, 28, 29]. A more detailed description of the model is given in Appendix A.

We also study three datasets to give us the citation network of three different types of documents. The first dataset gives us the citation relationships between papers posted on the High Energy Physics Theory (hep-th) category on the arXiv preprint repository between 1992 and 2003. This dataset, sourced from the Cornell KDD Cup website [32, 33], contains 27,770 nodes representing academic papers and 352,807 edges representing citations.

The second dataset represents the legal citations found in majority opinions of the United States Supreme Court between the years 1754 and 2002. The dataset, collected and developed by Fowler and Jeon as described in [34, 35], consists of 25,417 nodes representing opinions and 216,738 edges representing citations.

The last dataset gives us a citation network of patents. For this we use data from the National Bureau of Economic Research (NBER) [36], consisting of patents registered in the U.S. between 1975 and 1999, with 3,774,768 nodes representing granted U.S. patents and 16,518,949 edges representing citations. This dataset also assigns each patent a single subject category using the NBER technological classification. Further information on the source of the data and the input files used in our analysis is given in Appendix B.

The advantage of using these three datasets is that the motivation for adding a reference to a document is not necessarily the same in these three cases, giving us a broader test of

our hypothesis and methods. This choice also places our analysis in the broader context of citation network studies spanning academic, legal, and patent domains [37, 38]. For academic papers represented by the hep-th dataset, there is little formal requirement for a rigorous standard in the choice of quoted papers. So while references to key earlier developments are often included, it is also clear that citations frequently reflect other pressures, such as personal relationships or popularity, as part of a so-called “citation culture” [39]. On the other hand, legal citations in the United States Supreme Court follow much stricter guidelines; references are made based on the legal principle of “precedent”, which restricts citations to relevant court rulings, maintaining consistency within the justice system [40]. Finally, patents follow a different citation logic. The main aim of citations in patents is to prove an invention’s novelty, and there is a legal requirement to cite “prior art”; existing patents relevant to the current one. Still, many references in a patent can be to highly influential patents from companies such as Google or IBM [41]. Additionally, self-citation is a common phenomena among large organisations; it broadens the inventor’s patent portfolio and acts as a sort of protection from legal issues by ensuring the robustness of their intellectual property [42].

3 Methods

We start by identifying which documents are multi-class and which are not. To do this, we need to identify a unique cluster label associated with each document. Single-class and multi-class documents are then identified from the diversity of cluster labels of their neighbours. Finally, we measure the change in the in-degree of each document after transitive reduction to see if it is smaller for multi-class documents than it is for single class documents.

3.1 Cluster Labels

A multi-class document will have influence across multiple clusters. Before we can even quantify cross-community reach, we must first define what constitutes a cluster. Each cluster is our proxy for one research field of academic papers and patents or one area of law for USSC judgements. We will define the clusters by partitioning the set of documents, i.e. to each node n_i we assign a label c_i where the number of distinct cluster labels is $N_c = |\mathcal{C}|$.

One option would be to use any subject labels provided in the metadata for the documents. For the patent data, we have such labels, and due to the size of this dataset, we use these labels to define our clusters. However, for the hep-th and USSC datasets, the documents are not assigned to a specific subject label. Therefore, we have a bottom-up strategy, creating our own cluster label for each document based on the network structure by using a community detection method. This can be applied to all datasets to ensure that we can make a fairer comparison across our citation networks by using the same clustering method.

In the case of our artificial model, every paper is already assigned to be either a multi-class or a single-class paper, as well as having an explicit cluster or community assignment. However, the stochastic nature of the model means that there is a chance that, in particular instances, a node may not behave exactly as the predefined label suggests. So we choose not to use the intrinsic node labels for data produced from the model. Instead, we use the same bottom-up technique as used on the hep-th and USSC datasets.

In network analysis, clustering of nodes is known as “Community Detection”. In the context of a network, a community or cluster is assumed to have a much larger number of edges between nodes in the same cluster than there are edges between nodes from different clusters [43]. Various community detection methods exist (for example see chapter 31 of [43]), but we cluster nodes using the well established Leiden algorithm [44], applied to the undirected version of the

citation network [45]. We discuss some of these choices and show results for a different method in Appendix D.

The Leiden community detection algorithm [44] finds a partition of the nodes into clusters $\{c_i\}$ by approximately maximizing the modularity Q [46, 43], where

$$Q(\{c_i\}) = \frac{1}{W} \sum_{i,j} \left(A_{ij} - \frac{s_i s_j}{W} \right) \delta(c_i, c_j), \quad s_i = \sum_j A_{ij}, \quad W = \sum_i s_i. \quad (3.1)$$

Here $\delta(c_i, c_j)$ is one if nodes n_i and n_j belong to the same cluster ($c_i = c_j$), and zero otherwise.

The matrix entry A_{ij} we use is one if document n_j cites n_i or vice-versa so this is the adjacency matrix of the symmetrised citation network. We note there are other undirected networks which capture the relationships between documents in a citation network such as the bibliographic coupling discussed in Appendix D.4. There appears to be no consensus as to which community detection method is best [45, 47]. We report an alternative approach using the Louvain community detection method and bibliographic-coupling in Appendix D.4 for comparison.

To check the quality of the clusters found, we will give the sizes of the ten largest communities and the value of modularity Q in each case. A modularity close to one indicates a good set of communities while $Q = 0$ is a community structure that is no more clustered than a random null model. We will also give the effective number of large clusters using an h-index h and the Shannon diversity $N^{(\text{eff})}$. The h-index h is the largest number of clusters which contain at least h documents [48]. The Shannon diversity $N^{(\text{eff})}$ is the exponential of the Shannon entropy calculated using the fractional size N_f/N of each cluster f where N_f is the number of documents in cluster f and $N = \sum_f N_f$.

These types of clusters found using a community detection algorithm applied to a citation network could reflect groups within the data that are not associated with different topics or fields. For instance, a local citation culture within a group of authors might produce a cluster in a citation network. For that reason we prefer to talk about multi-class rather than interdisciplinary documents. However, our results are sensitive to relatively large clusters of the type we expect for different subject areas and so smaller clusters due to some local effects will not alter our results. In addition, we have also used other approaches where possible, externally provided field labels for the patents or clusters in the hep-th data based on text analysis of their abstracts. By checking our methods on an artificial model, and by using different data sets and different clustering methods, we will demonstrate that our results are robust and we are confident that our clusters do represent one possible discipline-based division in our data.

3.2 Cross-Community Influence Measures

Once we have a unique cluster label c_i for each document n_i , we can quantify its cross-community influence. To do this, we look at the distribution of the cluster labels of the neighbours of node n_i in the citation network, using measures of the diversity of their cluster labels [49, 50, 51].

We use the simplest diversity measure, RICHNESS, which is simply the number of different cluster labels associated with a document n_i , for example see Fig. 2. In this case, we will use the cluster label c_j associated with each of the neighbours n_j of a node n_i in the original citation network. Since our focus is primarily on the in-degree, it makes sense to use the in-richness $R_i^{(\text{in})}$ of a node n_i as our diversity measure, where

$$R_i^{(\text{in})} = |\{c_j | (n_j, n_i) \in \mathcal{E}\}|. \quad (3.2)$$

Here (n_j, n_i) is a directed edge from node n_j to n_i in the set of edges $\mathcal{E}$ of the citation network.




Figure 2: The node n_i outlined in black has an in-Richness of $R_i^{(\text{in})} = 3$, as it is cited by 3 distinct clusters.

We then need to identify the multi-class documents in our data. Based on the diversity measure for each document, we split the documents into three classes: single-class (SC) documents, multi-class documents (MC) or an intermediate class (IC). This last class is needed because cross-community influence is not a simple binary concept in practice, and most documents should really lie on a spectrum between pure single-class documents and maximally multi-class documents.

In our case, we will use the in-richness $R_i^{(\text{in})}$ of each node n_i to determine single-class and multi-class documents. It is not so clear where to draw the boundary in richness values between single- and multi-class documents, reflecting the reality that some documents may just be “mostly” single-class documents or “slightly” multi-class. So we set the criteria that *only* documents with in-richness of one will be considered single-class documents, i.e. if $R_i^{(\text{in})} = 1$, we will say that document n_i is a single-class document. While only documents with in-richness of three or more $R_i^{(\text{in})} \geq 3$, will be considered to be multi-class documents. This means we will not study documents with in-richness of two as we assume their true nature is unclear from our measurements.

Our artificial model provides a good example of our motivation for these boundaries. In principle, in the PDAG model, each single-class document only cites documents in the same community. That is true for citations from single-class to single-class documents. However, a multi-class document n_j , even though it is assigned in the PDAG model to one cluster given by the value c_j , behaves in the model as if it is a member of many clusters. In particular, a multi-class document n_j can cite and be cited by a single-class document n_i regardless of the value of c_i . This means a single-class document might be cited by just one multi-class document from a different cluster, giving that single-class document a richness of two. In the PDAG model, this occurs reasonably often, and this mimics what might happen with real data. The authors may not consider their document to be multi-class, but later authors in another field may find value in it.

3.3 Analysis of Degree Change

Our proposition states that a minimal change in the in-degree (or citation count) of a node after transitive reduction is an indicator of its cross community influence. Since citation counts vary widely across documents, we bin nodes by their original in-degree to control for degree-

dependent effects. Bin b is defined to contain all nodes with in-degree between $(k_{b-1} + 1)$ and k_b . For display purposes, results for each bin are shown at a location given by the degree $\bar{k}_b$ corresponding to the midpoint of the bin, defined as

$$\bar{k}_b = (k_{b-1} + k_b)/2. \quad (3.3)$$

Formally, we define

$$B_b = \{n \mid n \in \mathcal{N}, \quad k_{b-1} < k_n^{(\text{in})} \leq k_b\} \quad (3.4)$$

where $\mathcal{N}$ is the set of nodes and $k_n^{(\text{in})}$ is the in-degree of node $n \in \mathcal{N}$. We define k_b to be the smallest in-degree such that the bin B_b contains at least $n_{\min}$ documents, choosing $n_{\min} = 100$ in all cases except for the NBER patents, where we use $n_{\min} = 500$. That is

$$k_b = \min \mathcal{K}_b \quad \text{where} \quad \mathcal{K}_b = \left\{ \kappa \mid \sum_{k=k_{b-1}+1}^{\kappa} n(k^{(\text{in})}) \geq n_{\min} \right\}. \quad (3.5)$$

We are comparing the difference in the behaviour of single-class documents, which are represented by the set of nodes in $\mathcal{N}^{(\text{SC})}$, with multi-class documents, represented by the set¹ of nodes in $\mathcal{N}^{(\text{MC})}$. To make our comparison, we split each bin into two subsets, $B_b^{(t)}$, where $B_b^{(t)}$ contains all the documents of type t in bin B_b , that is

$$B_b^{(t)} = B_b \cap \mathcal{N}^{(t)}, \quad t \in \{\text{SC}, \text{MC}\}. \quad (3.6)$$

We find the in-degree of a node n before and after transitive reduction, $k_n^{(\text{in})}$ and $k_n^{(\text{in}, \text{TR})}$ respectively. We then measure the change Δ_n in the in-degree of a node n before and after transitive reduction, $\Delta_n = k_n^{(\text{in})} - k_n^{(\text{in}, \text{TR})}$. For a given bin b and a given subset of documents of type t , we estimate the mean in-degree change $\bar{\Delta}_b^{(t)}$ and associated standard deviation $\sigma_b^{(t)}$ from the average in-degree change Δ_n of nodes of type t within bin b , that is

$$\bar{\Delta}_b^{(t)} = \frac{1}{|B_b^{(t)}|} \sum_{n \in B_b^{(t)}} \Delta_n^{(t)}, \quad (3.7)$$

$$(\sigma_b^{(t)})^2 = \frac{1}{|B_b^{(t)}| - 1} \sum_{n \in B_b^{(t)}} (\Delta_n^{(t)} - \bar{\Delta}_b^{(t)})^2. \quad (3.8)$$

Finally, to compare the effect of transitive reduction on single-class (SC) and multi-class (MC) nodes, we compute the z-score z_b of the in-degree change per bin b

$$z_b = \frac{\bar{\Delta}_b^{(\text{SC})} - \bar{\Delta}_b^{(\text{MC})}}{\sqrt{(\bar{\sigma}_b^{(\text{SC})})^2 + (\bar{\sigma}_b^{(\text{MC})})^2}} \quad (3.9)$$

where $\bar{\sigma}_b^{(t)} = \sigma_b^{(t)} / |B_b^{(t)}|^{1/2}$ is the estimate of the standard error of the mean for nodes in bin b of type t .

This z-score quantifies how many standard deviations the citation loss of multi-class nodes differs from that of single-class nodes, helping to assess whether cross community influence correlates with lower citation loss after transitive reduction. A z-score of order one or more for each bin supports our hypothesis.

¹Note that documents cannot be of both types so $\mathcal{N}^{(\text{SC})} \cap \mathcal{N}^{(\text{MC})} = \emptyset$ but we also allow for a situation where not all documents are classified in this way, so $\mathcal{N}^{(\text{SC})} \cup \mathcal{N}^{(\text{MC})}$ may not contain all the nodes.

4 Results

For each dataset, we choose not to consider nodes with in-degree $k^{(\text{in})}$ less than the median in degree $\tilde{k}^{(\text{in})}$ of the citation network when assessing changes in the in-degree after transitive reduction. Nodes with low in-degree are very limited in the way they can change their few citations under transitive reduction, and so these documents have very noisy statistics. This cutoff of low-degree nodes is common in citation network analysis, for example see [52]. Some of the plots here display results only up to a certain bin size as there is a strong difference between number of single-class and multi-class nodes for higher bin values. Specifically, we do not consider bins where the standard error of the mean $\bar{\sigma}_b^{(t)} \geq 1$ for either type of nodes.

The results from the PDAG model in Fig. 3 show positive z-scores, suggesting that there is in fact a difference in in-degree count before and after TR, especially for bins containing lower in-degree nodes. For this partition, we have no singleton or discarded clusters. The clusters are also highly balanced in size (largest 10: [2724, 2716, 2715, 2706, 2699, 2698, 2688, 2687, 2685, 2682]), with high modularity ($Q = 0.79$) and a strong predominance of within-cluster edges (0.89 versus 0.11 cross-cluster), supporting the quality of the clustering.




Figure 3: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the PDAG model with $N = 27000$ nodes, average degree $\langle k \rangle = 11.95$ and median in-degree $\tilde{k}^{(\text{in})} = 5$. Using Leiden community detection on the undirected version of the citation network, we found $N_c = |\mathcal{C}| = 10$ clusters with an h-index of $h = 10$ and an effective number of clusters $N^{(\text{eff})} = 10.0$.

The hep-th results in Fig. 4 show a positive z-score greater than 1 across the majority of bins, indicating a clear difference in in-degree before and after transitive reduction throughout the full in-degree range. For this partition, we obtain no singleton clusters and only a very small discarded fraction (0.0043). The largest clusters contain [2624, 2610, 2090, 1909, 1861, 1641, 1637, 1557, 1439, 1435] nodes, while the partition has modularity $Q = 0.66$, with 75% of edges lying within clusters and 25% between clusters, again indicating substantial community structure. We emphasize that, particularly in narrow domains such as hep-th, these communities should not necessarily be interpreted as distinct disciplines. They may instead correspond to methodological traditions or author-level research groups.

By way of comparison, the hep-th data also comes with titles and abstracts and this allows us to use a non-network based method to cluster the documents for this data set. Our text-clustering methods and results are discussed in Appendix C.




Figure 4: A plot of z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the hep-th DAG with median in-degree $\tilde{k}^{(\text{in})} = 4$. Using Leiden community detection on the undirected version of the citation network, we found $N_c = |\mathcal{C}| = 35$ clusters. The h-index is $h = 22$ and the effective number of clusters is $N^{(\text{eff})} = 18.75$.

The USSC results shown in Fig. 5 exhibit positive z -scores most of which are greater than 1, with lower z -scores in bins corresponding to larger in-degrees. In this dataset, the z -scores are substantially higher than those observed in the other datasets. For this partition, we obtain no singleton or discarded clusters. The largest clusters contain [4268, 3114, 2378, 2335, 2154, 2151, 1769, 1232, 1174, 1152] nodes. The partition has modularity $Q = 0.61$, with 72% of edges lying within clusters and 28% between clusters, indicating clear but somewhat weaker community structure than in the PDAG and hep-th cases.




Figure 5: A plot of z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the USSC DAG with median in-degree $\tilde{k}^{(\text{in})} = 4$. Using Leiden community detection on the undirected version of the citation network, we get $N_c = |\mathcal{C}| = 14$ clusters. The h-index is $h = 14$ and the effective number of clusters is $N^{(\text{eff})} = 12.3$.

Finally, for the results of the patents in the NBER DAG shown in Fig. 6, all but four of the thirty four bins shown have a positive z-score. However, there are a larger number of bins with z-scores between zero and one than can be seen in previous examples, and more of these values are not significantly different from zero. There is also a strong dependence on the degree of the bin $\bar{k}_b$. As before, the bins with small degrees, where all documents also have the same in-degree and where there are larger numbers of documents in each bin, exhibit very high z-scores. For higher-degree bins, which tend to have the minimum number of papers in each bin and a wider range of in-degree values for each document, we observe that the z-scores are noisier and show a weaker effect compared to the previous datasets. Nevertheless, the vast majority of bins display at least a weak positive z-score, with many showing a strong effect. So even with patents, there is good support for the hypothesis that multi-class patents lose fewer citations under transitive reduction than single-class patents.




Figure 6: Plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the NBER DAG, with median in-degree $\tilde{k}^{(\text{in})} = 5$. The dataset included a pre-assigned subcategory for each patent, which we treated as a cluster, giving $|\mathcal{C}| = 36$ clusters. Note that the bins with degree $\bar{k}_b$ between 10.0 and 18.0 inclusive have z-scores above 8.0 and are not shown here for clarity.

To provide a simple summary of these results, we show in Table 1 the median of the z-scores for all the bins shown for each dataset. These median z-scores show that there is indeed a significant difference between the in-degree change for multi-class nodes in comparison to their single-class counterpart. The range of median z-scores across DAGs could be ascribed to attributes inherent to the dataset, such as the different citation styles for different types of documents, size of the dataset and effects of the clustering method. We observe that the mean of the multi-class difference is at least 2.2 standard deviations above the mean of the single-class difference across the bins, again validating our proposition.

Citation Network	Median z-Score $\tilde{z}$
PDAG	3.0
hep-th text clustering (Appendix C)	3.4
hep-th Leiden clustering	2.2
USSC	3.4
NBER	2.7

Table 1: The median z-score for every DAG. We only include bins with $k^{(\text{in})} \geq \tilde{k}^{(\text{in})}$ and where the z-score is well defined (some very high degree bins have significantly less multi-class documents than their single-class counterpart).

5 Discussion and Conclusions

We started from the observation that there is often a substantial decrease in the citation count of documents after transitive reduction [53, 54], something confirmed in the statistics here (see Appendix B). However, the loss of citations after transitive reduction is very uneven; some documents can lose a very large fraction of their incoming edges, some very few.

Our proposition is that this difference can be explained by the existence of multi-class documents in these datasets [53]. Transitive reduction removes “inessential edges”, edges that are not required to maintain the connectivity in the citation network. If document A is only of interest to documents from one cluster, then it is likely many of the documents citing A will also cite each other. So there is a high chance that most of the edges are inessential and will be removed by transitive reduction. If, on the other hand, A is a multi-class document, then that means it picks up citations from many different clusters. It is unlikely that those documents citing A will cite each other as they come from different clusters, so this means that the citations to A are much less likely to be inessential. We illustrate this idea in Fig. 7.

In our approach, we first use modularity based clustering to identify documents in the same community and hence to identify multi-class documents. We can then look at collections of documents with a similar citation count, our bins of (3.4), and measure the change in their in-degree after transitive reduction. Finally, we estimate the mean and standard error of the mean for the change in citation count for each type of document in each bin. A z-score (3.9) is then used to show if there is a statistically significant difference in the change in the citation count for the multi-class and single-class documents.

We first explored the impact of transitive reduction on citation networks by using an artificial model, the PDAG model, described in Section 2 (see also Appendix A) to validate our approach. We then looked at three datasets: academic papers posted between 1992 and 2003 on the hep-th repository of arXiv [32, 33], judgements from the U.S. Supreme Court (USSC) between 1754 and 2002 [34, 35] and data on U.S patents from the National Bureau of Economic Research (NBER) from 1975 to 1999 [36]. We first assigned each document to a cluster using a modularity-based community detection method on the undirected version of each citation network. The cluster labels were then used as proxies for the communities of each document. To see if a document was multi-class (cited by documents from many different clusters, many communities) or single-class (cited by documents from one other cluster, a single community), we used the diversity of the cluster labels held by its neighbours.

Our statistical analysis through the use of z-scores provides clear evidence that single-class nodes tend to lose more citations than multi-class nodes after transitive reduction, thus supporting our proposition.

The “reduced in-degree centrality”, that is, the citation count after transitive reduction, is a useful structural statistic for studying how citation-network topology and community structure shape the retention of incoming links under transitive reduction.




Figure 7: A simple example of how transitive reduction removes more citations between documents in the same cluster than from multi-class documents. In each of the four citation networks, documents are nodes with older (youngest) at the top (bottom) and documents in the same cluster are shown in the same colour and with the same symbol in one column. The two citation networks on the left represent two possible examples shown before transitive reduction, while on the right the same citation network is shown but after transitive reduction. The top row is an example where the documents are all single-class documents. Here, transitive reduction removes several edges, reducing the citation count for half the documents. On the bottom row, the nodes now represent multi-class documents, so they often cite documents in another cluster. No edges are removed under transitive reduction in this case, illustrating the robustness of cross-community citations under transitive reduction.

We do not present our approach using transitive reduction to find multi-class documents as a replacement for existing indicators. Prior work has used betweenness centrality in normalized local citation environments as an indicator of interdisciplinarity [12], and has benchmarked interdisciplinarity using integration scores together with science mapping [8]. Other studies have compared diversity-, centrality-, and citation-based indicators [13], and have further developed diversity-based measures of citation patterns [14].

Our focus is different. We wanted to try to understand how the topology of a citation network changes under transitive reduction. Specifically, why some documents with a similar number of citations lose far more citations than others under transitive reduction [53]. Our results show that there is a strong correlation between documents which see a high loss of edges and documents that are cited by a single class of documents. Our approach should be seen as complementary to, rather than competing with, existing metrics of cross-class influence.

It is widely acknowledged that the presence of a citation can indicate many different things [22, 23], none of which are evident from the existence of that single edge in a network. What

we show here is that there is additional information implicit in an edge of a DAG that comes from the wider network topology. We give one example of how this extra information can be leveraged to gain more insight into the meaning of each citation.

More broadly, our results show that transitive reduction can reveal meaningful differences in the structural role of documents within citation networks. This suggests that reduced in-degree may be useful in other settings where directed acyclic networks represent the propagation of influence or dependency. For instance, in ecological food webs, it could help identify keystone species whose extinction would have a significant impact on the ecosystem.

Data Accessibility

The data for the three citation networks is available online as cited in the main text and discussed in Appendix B: hep-th academic preprints on arXiv [32, 33], USSC opinions [34] and NBER patents [36]. Copies of the [data and code](#) used in this work are available online [55].

Funding

AH is supported by a PhD scholarship from the Saudi Arabian Ministry of Education. VV would like to acknowledge financial support from the Swiss National Science Foundation and from EPSRC, grant EP-R512540-1.

Appendices

A PDAG Model Definition

This model is an extension of models and numerical implementations used in [29, 30].

We start with an initial graph $\mathcal{G}(t = N_{\text{init}}) = (\mathcal{N}(t), \mathcal{E}(t))$ of $N_{\text{init}} = 120$ nodes but no edges. These nodes are labelled by distinct integers i from 0 to $(N_{\text{init}} - 1)$, so $n_i \in \mathcal{N}$ is the node with index i . Each node n_i also carries a cluster label c_i given by the node's index i modulo the number of clusters $N_c = |\mathcal{C}| = 10$, so $c_i = i \bmod N_c$. The set of possible cluster labels is therefore $\mathcal{C} = \{0, 1, 2, \dots, (N_c - 1)\}$.

We also initialise attachment lists, multisets $\mathcal{A}_f$ of references to nodes from cluster f . So if $n_i \in \mathcal{A}_f$ then $c_i = c_f$. In this case, we initialise these lists with one copy of each node in $\mathcal{N}(t)$ from the appropriate cluster f .

We now add the remaining $(N - N_{\text{init}})$ new nodes to the graph, one at a time. At each step we create a new graph $\mathcal{G}(t)$ with (t) nodes which starts with a copy of the graph from the previous step, $\mathcal{G}(t - 1) = (\mathcal{N}(t - 1), \mathcal{E}(t - 1))$. We add one new node n_t , so with index t , to the existing nodes so that $\mathcal{N}(t) = \mathcal{N}(t - 1) \cup \{n_t\}$. The cluster label of the new node, c_t is chosen uniformly at random from the set $\mathcal{C}$ of possible cluster values. The new node is also classified as being of one of two types; with probability $p_{\text{mc}} = 1/(N_c + 1) \approx 0.0909$ the node is set to be a (multi-class document), otherwise it is set to be a single-class document.

Next, we add the edges by creating a bibliography of length $m = 12$ for the new node n_t . The bibliography is made from older documents, target nodes n_j where $j < t$. If the node t is single-class SC, then the bibliography is chosen only from older target nodes j (where $j < t$) of the same cluster type $c_t = c_j$, otherwise we have a multi-class document which can cite any older node. In either case, we choose each older target document j for the bibliography of node t as follows. First, the cluster f of the cited document j is chosen to be $f = c_t$ if the new node t is single-class, otherwise f is chosen uniformly at random from the set of cluster labels $\mathcal{C}$. Then, with probability $p = 0.45$, the cited document j is chosen from the last T_r entries

in the multi-set $\mathcal{A}_f$, otherwise it is chosen from the full set $\mathcal{A}_f$. Here, T_r is defined as “recent documents”. This means that single-class documents n_t can only cite documents n_j with the same cluster f , $c_t = c_j$.

In principle, it is possible for the method to fail at finding m older documents to cite in a reasonable period of time, particularly for single-class nodes with a low index where the number of possible older documents it can cite is close to m . So in practice, a few earlier nodes have a little short of m documents in their bibliography, but this is almost impossible for nodes n_t where $t \gg mN_c$.

Finally, for every edge added, we update the attachment lists by adding the new node $\mu = 12$ times. If the new node n_t is a single-class document, then it is added μ times to the attachment list of its own cluster $\mathcal{A}_f$ where $f = c_t$. Otherwise, if the new node is a multi-class node, the new node n_t is again added to attachments lists μ times but this is spread across the attachments lists of different clusters, where the clusters chosen reflect the clusters of the entries in the bibliography². Finally, every node n_j in the bibliography of n_t is added once to the attachment list $\mathcal{A}_g$ of the old document’s own cluster, so $g = c_j$.

The purpose of these attachment lists is that, with the exception of the initial nodes, every node appears a total of $(\mu + k_j^{(\text{in})})$ times in these lists. So choosing a node uniformly at random from these lists means we are selecting documents n_j to cite with a probability proportional to $(\mu + k_j^{(\text{in})})$. This is a form that leads to fat-tailed in-degree distributions for the nodes as shown in simple generalisations of the Price model for bibliometrics [27, 28]. The point of selecting from the more recent documents is that it is well known that bibliographies tend to have more entries from recent documents than the simple Price model and its variants suggest [29]. By having this sharp cut-off in time for a fraction p of a bibliography is a crude but simple way to capture this aspect. Given that we add $(m + \mu) = 24$ documents to each attachment list $\mathcal{A}_f$ for every new document n_i in cluster f , the recent time scale of $T_r = 52800$ corresponds to $T_r/(m + \mu) = 52800/24 = 2200$ new papers in that cluster.

Note that in some ways, multi-class documents behave in the citation network as if they are in many clusters, not just one, even though these documents are also assigned a unique cluster label c_i . This is because they are added to attachment multi-sets $\mathcal{A}_f$ that reflect the clusters of their bibliographies, not the label c_i . The label c_i of a multi-class document n_i plays a role in the statistics. That is, when we assess the diversity of documents in the artificial network, multi-class documents n_i are treated as coming from one cluster since they carry one label c_i . On the other hand, each attachment multi-set $\mathcal{A}_f$ contains single-class documents from cluster f and multi-class documents carrying from any cluster. This means that a single-class document can cite multi-class documents from other clusters. Thus, the out-richness of a single-class document can be more than one as the bibliography will contain on average about $m.p_{\text{mc}}(1 - 1/N_c)$ documents (always multi-class documents) from another cluster. For the values quoted here, we have $m.p_{\text{mc}}(1 - 1/N_c) \approx 12 \times 0.0909 \times 0.9 \approx 0.982$ documents from another cluster in the bibliography of a single-class document.

²The cluster of each document in the bibliography has a separate entry in a list $\mathbf{B}$, where $\mathbf{B}[i]$ is the cluster of the i -th entry in the bibliography. If the new node is a multi-class document, then this list of the clusters of bibliometric entries is cycled around μ times to give a list of μ clusters, that is $f = \mathbf{B}[j \bmod \mathbf{m}]$ is the j -th cluster chosen with j running over all integers from 0 to $(\mu - 1)$. For each cluster f chosen, the new node n_t is added to the associated attachment list $\mathcal{A}_f$. If μ is not an integer multiple of m , then the new multi-class node will not use the clusters of the documents in the bibliography equally when choosing which attachment lists to be added to.

B Data

Copies of the [data and code](#) used in this work are available online at <http://doi.org/10.6084/m9.figshare.29663918> [55].

B.1 Artificial Data

The artificial data is produced by `python` code implementing the model described in Appendix A. For the file used in Fig. D.1, the parameter values are given in Table B.1. These values were chosen to reflect the properties of the hep-th dataset as much as possible and are inspired by the values used in the models of [29].

parameter		value
number of nodes	$N = \mathcal{N} $	27000
number of edges	$E = \mathcal{E} $	322560
average degree	$\langle k \rangle$	11.95
initial nodes	N_{init}	120
recent document time scale	T_r	52800
length of bibliography	m	12
fitness	μ	12
fraction of recent documents in bibliography	p	0.45
fraction of multi-class documents	p_{mc}	0.0909
number clusters	$N_c = \mathcal{C} $	10
random seed		0

Table B.1: A list of the parameter values used to generate the example data from the PDAG model. The model is defined in Appendix A.

B.2 Real World Data

The data used for the citation network of academic papers posted on the hep-th section of the arXiv preprint repository between 1992 and 2003 comes from the Cornell KDD Cup website [32, 33].

The citation network of majority opinions of the United States Supreme Court (USSC) between the years 1754 and 2002 came from links on a 2016 version of a web page `James H. Fowler's Supreme Court Citation Network Data Page.html`, which cites Fowler and Jeon as authors and asks for citation to two published papers [34, 35].

The dataset used to construct the citation network of patents came from the National Bureau of Economic Research (NBER) [36], consisting of patents registered in the U.S. between 1975 and 1999, with 3,774,768 nodes representing granted U.S. patents and 16,518,949 edges representing citations. This also has a single subject category assigned to each patent.

In Table B.2 we list some of the basic properties of each network at each stage of our processing.

Edge correction is where we reverse the direction of any edge that cites a document in the future, at least according to the node data we have. We also remove all self-citations. The PDAG model gives a perfect DAG so this correction is not needed. For the real data sets, we only have to remove a very small number of edges, 0.15% or less.

We always analyse the largest weakly connected component (LWCC), so the next stage is to select this component. Again, the PDAG model gives a single weakly connected component

network	stage	nodes	edges	node%	edge %
hep-th [32, 33]	Original	27770	352807	100.0	100.0
	After Edge Correction	27770	352285	100.0	99.85
	After LWCC	27400	352021	98.67	99.78
	After TR	27400	63247	98.67	17.93
USSC [34]	Original	25417	216738	100.0	100.0
	After Edge Correction	25417	216456	100.0	99.87
	After LWCC	25389	216436	99.89	99.86
	After TR	25389	58829	99.89	27.14
NBER [36]	Original	3774768	16518949	100.0	100.0
	After Edge Correction	3774768	16518947	100.0	100.00
	After LWCC	3764117	16511740	99.72	99.96
	After Degree Filtering	1236699	8594651	32.76	52.03
	After TR	1233221	5386866	32.67	32.61
PDAG	Before TR	27000	322560	100.0	100.0
	After TR	27000	124855	100.0	38.71

Table B.2: Node and edge counts for each dataset at different preprocessing stages. In edge correction, we remove cycles and reverse backward edges to create a DAG. LWCC refers to the extraction of the largest weakly connected component. In degree filtering, we reduce the size of the NBER patent data by removing nodes with degree below ten. Finally we apply transitive reduction (TR).

so this is not required. For the other datasets, the change is small, at most a one percent loss, often much less.

The NBER patents dataset is much larger than the other datasets, so we apply a further procedure in this case to extract a smaller relevant network. We choose to remove nodes and all their associated edges with degree of nine or less from the data. This is not done recursively, so it is still possible to have nodes with degree less than ten. This leaves us with most of the data to analyse, about one third of the original nodes and half of the edges.

Finally, we apply transitive reduction. As reported in [53], for the academic papers in hep-th and the court citations of USSC, we get large drops in the number of edges, as we lose 70% to 80% of the edges. The NBER patents however lose far fewer edges, only 37% of edges were left after the degree filtering. By way of comparison, the artificial PDAG model sits in between, losing 61% under transitive reduction.

C Text-Based Clustering

One other clustering approach we employ is to use text associated with each document to generate topic labels. The titles of documents, available for all our data, provide too little text to work on. However, the arXiv data set [32, 33] comes with abstracts, and so we group those based on the similarity of their words.

To do this, we take the words from the abstracts of the hep-th documents, we tokenize and stem the words, excluding all common tokens, those in more than 80% of documents, and rare tokens, those that are in just 20% of documents. We then compute a “Term Frequency-Inverse Document Frequency” matrix $T_{\alpha i}$ which is the number of times that token α occurs in document i divided by the logarithm of the number of documents where token α appears at least once. We then use cosine similarity to find the distance D_{ij} between two documents i and j

$$D_{ij} = 1 - \frac{1}{|T_i||T_j|} \sum_{\alpha} T_{\alpha i} T_{\alpha j}, \quad |T_i|^2 = \sum_{\alpha} T_{\alpha i} T_{\alpha i}. \quad (\text{C.1})$$

Finally, the distance matrix D is used to find cluster labels for each document i using the K-Means algorithm³ [59].




Figure C.1: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the hep-th DAG using k-means clustering on abstracts with $N_c = |\mathcal{C}| = 7$ clusters. The majority of z-score are positive with a few negative points in the centre of the plot.

Testing this measure on the hep-th dataset, the results shown in Fig. C.1 suggest a trend with z-scores greater than 1 for most bins and decreasing z-scores for bins with larger in-degrees. This supports what we found for the hep-th data using community detection in the main text. However, there are some limitations inherent to text-based clustering. In our dataset, we observed significant keyword repetition across multiple clusters, which may indicate that text-based clustering is not effectively distinguishing between the citation network communities. We suspect this issue arises due to the lack of variability in hep-th abstracts, as papers in this domain often revolve around closely related topics. Prior research suggests that abstracts alone may not contain enough distinguishing terms to serve as a reliable basis for clustering, particularly when compared to full-text analysis [60]. This limitation likely contributes to the weaker trend observed in the z-score results.

D Community Detection

Community detection is a method in network science which gives a partition of the set of nodes into clusters, known as “communities” in this context. While there can be many ways to define a community both qualitatively and quantitatively, a general guide is that in a good clustering the edges tend to be between nodes in the same cluster while nodes in different clusters are connected less often [46, 43].

Each community detection method gives a precise definition of what that method calls a cluster. It is not clear if the result from one method is any “better” as there may be many

³We tried a variety of other clustering methods, including Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [56], Hierarchical DBSCAN (HDBSCAN) [57] and Ordering Points to Identify the Clustering Structure (OPTICS) [58]. However, we found K-means offered the best trade-off between granularity and cohesion—avoiding both excessive fragmentation (over-clustering) and overly broad groupings (under-clustering).

different types of clusters at different scales in any realistic network. That is different methods may reveal different ‘truths’. Even if we have some external “ground truth” provided, perhaps given by experts labelling nodes, it is not obvious that such a ground truth is some absolute truth or just another reasonable view of a possible clustering of the nodes in a network. Accordingly, in our work we do not interpret the detected communities as uniquely correct ground-truth groups. We just accept the clusters we find as a reasonable description of the subgroups within our data.

D.1 Quality of Clustering

To check the quality of the clusters found, we give the sizes of the ten largest communities. The value of modularity Q of (D.3) is another useful indicator since a modularity value close to one indicates a good community structure while $Q = 0$ shows our communities are no better clustered than we would expect in a random null model.

When interpreting partitions, it is useful to summarise how many communities contain a substantial number of documents, while discounting communities with negligible size. This is relevant because relatively large clusters dominate our multi-class statistics, while smaller clusters do not contribute much and can be uninteresting artefacts of the data or the community detection method.

One measure we use is the h-index h defined over the cluster size distribution $\{N_f\}$, where N_f is the number of documents in the cluster labelled $f \in \mathcal{C}$. The h-index h is the largest number of clusters which contain at least h documents [48]. Formally, if $N(\nu)$ is the number of clusters that contain at least ν documents, then

$$h = \max \{ \nu \in \mathbb{N} \mid N(\nu) \geq \nu \}, \quad N(\nu) := |\{f \mid N_f \geq \nu, f \in \mathcal{C}\}|. \quad (\text{D.1})$$

Here $\mathcal{C}$ is the set of labels of the different clusters.

We also employ Shannon diversity $N^{(\text{eff})}$ as a second estimate of the effective number of clusters. We use the fraction of nodes N_f/N in each cluster where N_f is the number of documents in cluster f and $N = \sum_f N_f$ is the total number of nodes. This is used to define the Shannon entropy of the cluster size distribution. The Shannon diversity $N^{(\text{eff})}$ is then given by the exponential of the Shannon entropy, where

$$N^{(\text{eff})} = \exp \left(- \sum_{f \in \mathcal{C}} \frac{N_f}{N} \ln \left(\frac{N_f}{N} \right) \right). \quad (\text{D.2})$$

D.2 Modularity Maximisation

The main community detection method we use is the Leiden community detection algorithm [44]. This is an example of a method which finds a partition of the nodes into clusters $\{c_i\}$ by approximately maximizing modularity $Q(\{c_i\})$ [46, 43] where

$$Q(\{c_i\}) = \frac{1}{W} \sum_{i,j} \left(A_{ij} - \frac{s_i s_j}{W} \right) \delta(c_i, c_j), \quad s_i = \sum_j A_{ij}, \quad W = \sum_i s_i. \quad (\text{D.3})$$

Here, A_{ij} is the weight of an edge from node n_j to node n_i in an undirected network while $\delta(c_i, c_j)$ is one if nodes n_i and n_j belong to the same cluster ($c_i = c_j$), and zero otherwise.

Modularity Q is larger when there are more edges are between nodes in the same cluster than we might reasonably expect. Hence, larger values of Q indicate higher-quality cluster assignments $\{c_i\}$ and a more meaningful partitioning of the network. The value of modularity $Q(\{c_i\})$ measures how well the network is partitioned into clusters.

We note that modularity-based clustering has known limitations, including a resolution limit that can force smaller communities to remain part of larger ones [61, 62]. There are many different algorithms which find partitions $\{c_i\}$ with large modularities and some have algorithms have their own issues. However, as we are interested in the broad statistical properties of our network, any community detection method that has been tested on shown to be effective could be used.

D.3 Leiden Clustering on the Symmetrised Citation Network

In our main text, we use the Leiden community detection algorithm [44] as it is practical and well-established method for identifying robust large-scale structure in an undirected network. For instance, compared with the older Louvain method [63], the Leiden algorithm [44] avoids disconnected communities by construction and yields better-connected partitions.

It is simplest to work with modularity defined for undirected networks. So for the results in the main text, we take the undirected citation network. That is suppose the adjacency matrix for the original directed citation network is the matrix $\mathbf{C}$, where $C_{ij} = 1$ if there is an edge from node n_j to node n_i , i.e. document n_j cites the older document n_i . To perform community detection, we construct a new undirected network on the same set of nodes with the adjacency matrix $\mathbf{A}$ where

$$A_{ij} = C_{ij} + C_{ji}, \quad A_{ii} = 0. \quad (\text{D.4})$$

Thus, two nodes n_i and n_j are connected by an undirected edge whenever one cites the other. Since the original citation network is a directed acyclic graph, at most one of C_{ij} and C_{ji} can be nonzero, so the undirected network is unweighted.

D.4 Louvain Clustering on the Bibliographic Coupling Network

In this section, we show the results obtained with nodes cluster labels found by using the modularity-based Louvain method [63] applied to the network defined by bibliometric coupling but otherwise we use the much same approach as before.

In this approach, we construct a new undirected network using bibliographic coupling [64, 65, 45]. Let the adjacency matrix for the citation network be the matrix $\mathbf{C}$ where $C_{ij} = 1$ if there is an edge from node n_j to n_i , i.e. document n_j lists the older document n_i in its bibliography. The new network retains the same nodes as the original citation network, so we will denote the weight of an edge between n_i and n_j by the entry A_{ij} in the adjacency matrix $\mathbf{A}$ of the bibliographic-coupling network. Each pair of nodes n_i and n_j is connected by an undirected but weighted edge if have at least one document in common in their bibliographies. Suppose node n_k is cited by both n_i and n_j , so that $C_{ik} = C_{jk} = 1$. This node n_k will contribute the inverse of its in-degree in the original citation network, $1/k_k^{(\text{in})}$, to the weight of the bibliographic-coupling link A_{ij} . This avoids undue contributions from highly-cited documents. So we have that

$$A_{ij} = \sum_k \frac{1}{k_k^{(\text{in})}} C_{ki} C_{kj}, \quad k_k^{(\text{in})} = \sum_i C_{ki}. \quad (\text{D.5})$$

We then apply the Louvain community detection algorithm [63]. The algorithm finds a partition of the nodes into communities by finding an approximate maximum for modularity Q of (D.3).

In this case we choose not to consider nodes with in-degree $k^{(\text{in})} < 10$ (rather than using the median value as a cutoff) when assessing changes in the in-degree after transitive reduction. Additionally, we do not consider bins where the standard error of the mean $\bar{\sigma}_b^{(t)} \geq 1$ for either type of nodes.

The results from our artificial PDAG model, shown in Fig. D.1, indicate that all but two of the bins have positive z-scores significantly greater than zero. This suggests that in our PDAG model, multi-class nodes lose fewer citations after transitive reduction. However, there is a clear trend toward lower z-scores for bins containing higher-degree documents. These tend to have a smaller number of documents in each bin, but also a wider range of in-degrees for each document, making these statistics inherently noisier than the lower degree bins.




Figure D.1: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the PDAG model with $N = 27000$ nodes and average degree $\langle k \rangle = 11.95$. Using Louvain clustering on the network of bibliographic coupling, we found $N_c = |\mathcal{C}| = 130$ clusters with an h-index of $h = 10$ and an effective number of clusters $N^{(\text{eff})} = 10.4$.

We now turn to analyse the hep-th papers from arXiv. The resulting z-score plot is shown in Fig. D.2. All but three of the forty-two bins have positive z-scores, and all but one of those have a z-score of at least 1.0. Again, this is strong evidence that multi-class papers in hep-th lose significantly fewer citations under transitive reduction than single-class papers.




Figure D.2: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the hep-th DAG using Louvain clustering on a bibliographic-coupling network, with $N_c = |\mathcal{C}| = 3091$ clusters. The h-index is $h = 28$ and the effective number of clusters is $N^{(\text{eff})} = 48.8$.

All the z-scores observed for the court documents in the USSC DAG in Fig. D.3 are positive. Further, all but three of the twenty-five bins have z-scores of 1.0 or more, so the positive values are statistically significant. This gives strong support for our hypothesis in the context of these legal opinions.




Figure D.3: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the USSC DAG using Louvain clustering on a bibliographic-coupling network with $N_c = |\mathcal{C}| = 6338$ clusters. The h-index is $h = 19$ and the effective number of clusters is $N^{(\text{eff})} = 103.8$.

Finally, for the patents in the NBER DAG, the results in Fig. D.4 match those in Section 4. However, as our cut-off of $k^{(\text{in})} \geq 10$ is larger than the median in degree $\tilde{k}^{(\text{in})} = 5$, our median z-score shown in Table D.1 is lower as we have removed the bins with $5 \leq k^{(\text{in})} < 10$ which contained high positive z-scores.




Figure D.4: A plot of z-score z_b (3.9) against bin centre $\bar{k}_b$ (3.3) for the NBER DAG. The dataset included a pre-assigned subcategory for each patent, which we treated as a cluster, giving $N_c = |\mathcal{C}| = 36$ clusters. Note that the bins with degree $\bar{k}_b$ between 10.0 and 18.0 inclusive have z-scores above 8.0 and are not shown here for clarity.

To provide a simple summary of these results, we show in Table D.1 the median of the z-scores for all the bins shown for each dataset. These median z-scores show that there is indeed a significant difference between the in-degree change for multi-class nodes in comparison to their single-class counterpart, even when using a different clustering method (Louvain community-detection) on a bibliographic coupling network. We observe that the mean of the multi-class difference is at least 1.7 standard deviations above the mean of the single-class difference across the bins, again validating our proposition.

Citation Network	Median z-Score $\tilde{z}$
PDAG	2.7
hep-th Louvain clustering	3.4
USSC	2.2
NBER	1.7

Table D.1: The median z-score for every DAG. We only include bins with $k^{(\text{in})} \geq 10$ and where the z-score is well defined (some very high degree bins have significantly less multi-class documents than their single-class counterpart).

References

- [1] Garfield E, Sher IH, Torpie RJ. 1964 *The use of citation data in writing the history of science*. Institute for Scientific Information Inc. Philadelphia. Pennsylvania, USA.
- [2] Price DS. 1965 Networks of Scientific Papers. *Science* **149**, 510–515.
- [3] Zhao D, Strotmann A. 2015 *Analysis and Visualization of Citation Networks*. Morgan & Claypool Publishers.
- [4] Golosovsky M. 2019 *Citation Analysis and Dynamics of Citation Networks*. Springer International Publishing.

- [5] Šubelj L, Waltman L, Traag V, van Eck NJ. 2020 Intermediacy of publications. *Royal Society Open Science* **7**, 190207.
- [6] Ho M, Price HC, Evans TS, O’Sullivan E. 2024 Who made the mRNA vaccine? Measuring Division of Labour in Therapeutic Innovation. *Nature Biotechnology* **42**, 1643–1648.
- [7] Zhou Y, Gong F, Wang Y, Wang R, Zeng A. 2025 Fusing structural and temporal information in citation networks for identifying milestone works. *Chaos, Solitons & Fractals* **192**, 116004.
- [8] Porter A, Rafols I. 2009 Is science becoming more interdisciplinary? Measuring and mapping six research fields over time. *Scientometrics* **81**, 719–745.
- [9] Huutoniemi K, Klein JT, Bruun H, Hukkinen J. 2010 Analyzing interdisciplinarity: typology and indicators. *Research Policy* **39**, 79–88.
- [10] Rafols I, Meyer M. 2010 Diversity and network coherence as indicators of interdisciplinarity: case studies in bionanoscience. *Scientometrics* **82**, 263–287.
- [11] Bu Y, Waltman L, Huang Y. 2021 A multidimensional framework for characterizing the citation impact of scientific publications. *Quantitative Science Studies* **2**, 155–183.
- [12] Leydesdorff L. 2007 Betweenness centrality as an indicator of the interdisciplinarity of scientific journals. *Journal of the American Society for Information Science and Technology* **58**, 1303–1319.
- [13] Leydesdorff L, Rafols I. 2011 Indicators of the interdisciplinarity of journals: diversity, centrality, and citations. *Journal of Informetrics* **5**, 87–100.
- [14] Leydesdorff L, Wagner CS, Bornmann L. 2019 Interdisciplinarity as diversity in citation patterns among journals: Rao-Stirling diversity, relative variety, and the Gini coefficient. *Journal of Informetrics* **13**, 255–269.
- [15] Uzzi B, Mukherjee S, Stringer M, Jones B. 2013 Atypical combinations and scientific impact. *Science* **342**, 468–472.
- [16] Garfield E. 1955 Citation indexes to science: a new dimension in documentation through association of ideas. *Science* **122**, 108–111.
- [17] Cole JR, Cole S. 1973 *Social Stratification in Science*. University of Chicago Press.
- [18] Martin BR, Irvine J. 1983 Assessing basic research. *Research Policy* **12**, 61–90.
- [19] Bornmann L, Daniel HD. 2008 What do citation counts measure? a review of studies on citing behavior. *Journal of Documentation* **64**, 45–80.
- [20] Leydesdorff L, Bornmann L, Comins JA, Milojević S. 2016 Citations: indicators of quality? the impact fallacy. *Frontiers in Research Metrics and Analytics* **1**.
- [21] Cai X, Lyu X, Zhou P. 2023 The relationship between interdisciplinarity and citation impact—a novel perspective on citation accumulation. *Humanities and Social Sciences Communications* **10**, 1–12.
- [22] Brooks TA. 1986 Evidence of Complex Citer Motivations. *Journal of the American Society for Information Science* **37**, 34–36.

- [23] MacRoberts MH, MacRoberts BR. 1989 Problems of Citation Analysis: A Critical Review. *Journal of the American Society for Information Science* **40**, 342–349.
- [24] Simkin MV, Roychowdhury VP. 2003 Read Before You Cite!. *Complex Systems* **14**, 269–274.
- [25] Simkin MV, Roychowdhury VP. 2005 Stochastic Modeling of Citation Slips. *Scientometrics* **62**, 367–384.
- [26] Aho A, Garey MR, Ullman JD. 1972 The Transitive Reduction of a Directed Graph. *SIAM Journal on Computing*, **1**, 131–137.
- [27] Price DS. 1976 A general theory of bibliometric and other cumulative advantage processes. *J.Amer.Soc.Inform.Sci.* **27**, 292–306.
- [28] Newman M. 2010 *Networks: an introduction*. Oxford University Press first edition.
- [29] Goldberg S, Anthony H, Evans T. 2015 Modelling Citation Networks. *Scientometrics* **105**, 1577–1604.
- [30] Evans T, Calmon L, Vasiliauskaite V. 2020 The Longest Path in the Price Model. *Scientific Reports* **10**, 10503.
- [31] Vasiliauskaite V, Evans TS. 2020 Making Communities Show Respect for Order *Applied Network Science* **5**, 15.
- [32] Gehrke J, Ginsparg P, Kleinberg J. 2003 Overview of the 2003 KDD Cup. *ACM SIGKDD Explorations Newsletter* **5**, 149–151.
- [33] KDD cup. 2003 KDD cup. <https://www.cs.cornell.edu/projects/kddcup/datasets.html>.
- [34] Fowler JH, Jeon S. 2008 The Authority of Supreme Court Precedent. *Social Networks* **30**, 16–30.
- [35] Fowler JH, Johnson TR, Spriggs JFI, Jeon S, Wahlbeck PJ. 2007 Network Analysis and the Law: Measuring the Legal Importance of Precedents at the U.S. Supreme Court. *Political Analysis* **15**, 324–346.
- [36] Hall BH, Jaffe AB, Trajtenberg M. 2001 The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools. Technical Report 8498 National Bureau of Economic Research (NBER).
- [37] Barnett GA, Huh C, Kim Y, Park HW. 2011 Citations among communication journals and other disciplines: a network analysis. *Scientometrics* **88**, 449–469.
- [38] Ji J, Barnett GA, Chu J. 2019 Global networks of genetically modified crops technology: a patent citation network analysis. *Scientometrics* **118**, 737–762.
- [39] Wouters P. 1999 *The Citation Culture*. Phd thesis University of Amsterdam Amsterdam, Netherlands. Unpublished doctoral dissertation, Centre for Science and Technology Studies (CWTS).
- [40] Cross FB, Spriggs JEI, Johnson TR, Wahlbeck PJ. 2010 Citations in the U.S. Supreme Court: An Empirical Study of Their Use and Significance. *University of Illinois Law Review* **2010**, 489–576.

- [41] Messeni Petruzzelli A, Rotolo D, Albino V. 2014 Determinants of Patent Citations in Biotechnology: An Analysis of Patent Influence Across the Industrial and Organizational Boundaries. *Technovation* **34**, 1–13.
- [42] Blind K, Cremers K, Mueller E. 2009 The Influence of Strategic Patenting on Companies' Patent Portfolios. *Research Policy* **38**, 428–436.
- [43] Coscia M. 2021 *The Atlas for the Aspiring Network Scientist*. Michele Coscia.
- [44] Traag VA, Waltman L, van Eck NJ. 2019 From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports* **9**.
- [45] Boyack KW, Klavans R. 2010 Co-citation analysis, bibliographic coupling, and direct citation: Which citation approach represents the research front most accurately?. *Journal of the American Society for Information Science and Technology* **61**, 2389–2404.
- [46] Newman MEJ, Girvan M. 2004 Finding and evaluating community structure in networks. *Phys. Rev. E* **69**, 026113.
- [47] Klavans R, Boyack KW. 2017 Which type of citation analysis generates the most accurate taxonomy of scientific and technical knowledge? *Journal of the Association for Information Science and Technology* **68**, 984–998.
- [48] Hirsch JE. 2005 An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences of the United States of America* **102**, 16569–16572.
- [49] MacArthur RH. 1965 Patterns of species diversity. *Biological Reviews* **40**, 510–533.
- [50] Jost L. 2006 Entropy and diversity. *Oikos* **113**, 363–375.
- [51] Gheorghiadu P, Vasiliauskaite V, Diachenko A, Price H, Evans T, Rivers R. 2023 Entropol-ogy: an Information-Theoretic Approach to Understanding Archaeological Data. *Journal of Archaeological Method and Theory* **30**, 1109–1141.
- [52] Chen P, Redner S. 2009 Community Structure of the Physical Review Citation Network. *Journal of Informetrics* **3**, 111–123.
- [53] Clough JR, Gollings J, Loach TV, Evans TS. 2014 Transitive reduction of citation net-works. *Journal of Complex Networks* **3**, 189–203.
- [54] Vasiliauskaite V. 2020 *Paths and Directed Acyclic Graphs*. PhD thesis Imperial College London.
- [55] AlMuhanna H, Vasiliauskaite V, Evans T. 2025 Data and Code for 'Interdisci-plinarity Revealed by Transitive Reduction of Citation Networks'. figshare DOI [10.6084/m9.figshare.29663918](https://doi.org/10.6084/m9.figshare.29663918).
- [56] Schubert E, Sander J, Ester M, Kriegel HP, Xu X. 2017 DBSCAN revisited, revisited: why and how you should (still) use DBSCAN. *ACM Transactions on Database Systems (TODS)* **42**, 1–21.
- [57] McInnes L, Healy J. 2017 Accelerated hierarchical density based clustering. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)* pp. 33–42. IEEE.
- [58] Ankerst M, Breunig MM, Kriegel HP, Sander J. 1999 OPTICS: Ordering points to identify the clustering structure. *ACM Sigmod Record* **28**, 49–60.

- [59] Sculley D. 2010 Web-scale k-means clustering. In *Proceedings of the 19th International Conference on World Wide Web* pp. 1177–1178. ACM.
- [60] Alexandrov M, Gelbukh A, Rosso P. 2005 An Approach to Clustering Abstracts. In *Natural Language Processing and Information Systems* vol. 3513 pp. 275–285. Springer-Verlag.
- [61] Fortunato S, Barthelemy M. 2007 Resolution limit in community detection. *PNAS* **104**, 36–41.
- [62] Good BH, de Montjoye YA, Clauset A. 2010 Performance of modularity maximization in practical contexts. *Physical Review E* **81**, 046106.
- [63] Blondel VD, Guillaume JL, Lambiotte R, Lefebvre E. 2008 Fast unfolding of community hierarchies in large networks. *J. Stat. Mech.* p. P10008.
- [64] Kessler MM. 1963a Bibliographic coupling between scientific papers. *American documentation* **14**, 10–25.
- [65] Kessler M. 1963b An experimental study of bibliographic coupling between technical papers (Corresp.). *IEEE Transactions on Information Theory* **9**, 49–51.