

# Assessing Network Representations for Identifying Interdisciplinarity

Eoghan Cunningham

School of Computer Science, University College Dublin  
Dublin, Ireland  
eoghan.cunningham@ucdconnect.ie

Derek Greene

School of Computer Science, University College Dublin  
Dublin, Ireland  
derek.greene@ucd.ie

## ABSTRACT

Many studies have sought to identify interdisciplinary research as a function of the diversity of disciplines identified in an article’s references or citations. However, given the constant evolution of the scientific landscape, disciplinary boundaries are shifting and blurring, making it increasingly difficult to describe research within a strict taxonomy. In this work, we explore the potential for graph learning methods to learn embedded representations for research papers that encode their ‘interdisciplinarity’ in a citation network. This facilitates the identification of interdisciplinary research without the use of disciplinary categories. We evaluate these representations and their ability to identify interdisciplinary research, according to their utility in interdisciplinary citation prediction. We find that those representations which preserve structural equivalence in the citation graph are best able to predict distant, interdisciplinary interactions in the network, according to multiple definitions of citation distance.

## CCS CONCEPTS

• **Computing methodologies** → Learning latent representations.

## KEYWORDS

datasets, scientific knowledge representation, interdisciplinarity

### ACM Reference Format:

Eoghan Cunningham and Derek Greene. 2022. Assessing Network Representations for Identifying Interdisciplinarity. In *Companion Proceedings of the Web Conference 2022 (WWW ’22 Companion)*, April 25–29, 2022, Virtual Event, Lyon, France. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3487553.3524653>

## 1 INTRODUCTION

*Interdisciplinary research* is most commonly defined as research activity which integrates expertise, data or methodologies from two or more distinct disciplines. In light of its perceived importance, many studies have been conducted that quantify article and author interdisciplinarity in an effort to identify relevant research trends and to explore their impact. The most widely-accepted methods for measuring interdisciplinarity assess *knowledge integration* using citation information, attempting to measure interdisciplinarity as

a function of the balance, diversity, and dissimilarity of the topics identified in an article’s cited papers [15, 18, 20]. Using these methods, a trend towards increased interdisciplinarity has been identified [19], which has been found to correlate positively with research impact [8, 15] and productivity [9].

Despite some agreement across measures of interdisciplinarity (with many studies based on work by [18]), the literature lacks convergence towards a single metric [24]. Most candidate metrics are reliant on explicit research topic or subject category information, for which sources are numerous, diverse, and sometimes unavailable. Moreover, the scientific community structure has been shown to be evolving rapidly [21]. This evolution may explain the lack of agreement between subject taxonomies and hierarchies, which vary considerably [13]. The limitations of prescribed, static, subject classifications necessitate a modern approach that identifies research disciplines – and indeed interdisciplinarity – according to research network structure. Such an approach may offer a dynamic view of interdisciplinarity, and may help to map emerging developments that do not fit any existing schema [24].

In this work, we begin to investigate appropriate graph representation learning methods which allow us to identify and quantify research interdisciplinarity in citation networks, so that we can monitor interdisciplinary interactions and track the development of research topics. We propose to learn unsupervised node (or paper)-level representations in research citation networks that encode research ‘interdisciplinarity’. In pursuit of such a representation, we explore the different network characteristics that embeddings must preserve in order to encode interdisciplinarity and we evaluate these embeddings according to their utility in an interdisciplinary citation prediction task. Specifically, we investigate how the performance of link prediction models (based on different node representations) is affected by the distance of citation, according to the hypothesis that representations that encode article interdisciplinarity will fare better when predicting “far-reaching” or “long-distance” interdisciplinary interactions/citations [8], compared to those that do not. Given the acknowledged limitations of journal-level subject categories, we explore multiple definitions of citation distance: both according to prescribed categories with predefined topic distances, and to network defined distances according to graph structure.

## 2 BACKGROUND

### 2.1 Measuring Interdisciplinarity

Numerous studies have developed metrics for quantifying research interdisciplinarity (IDR). Generally, these methods are designed to evaluate interdisciplinarity in pursuit of three goals: (i) to explore trends in IDR over time [19], (ii) to measure the benefits of IDR [8, 9, 15], (iii) to measure changes in interdisciplinarity in response

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*WWW ’22 Companion*, April 25–29, 2022, Virtual Event, Lyon, France

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9130-6/22/04.

<https://doi.org/10.1145/3487553.3524653>

to IDR initiatives [18]. While the proposed approaches are diverse, interdisciplinarity is most commonly measured on the basis of citation analysis. Using citations to model the flow of information, *knowledge integration* is calculated as a function of the balance, diversity and/or distance between the disciplines identified in *referenced* articles [1, 15, 19, 20]. Conversely, *knowledge diffusion* is measured according to the disciplines of *citing* articles [17, 23]. These definition of knowledge integration, knowledge diffusion, or a combination of both have been used to enumerate interdisciplinarity for research papers. Invariably, such methods are reliant on explicit topic or discipline categories for research papers, such as those provided by Web of Science, Scopus, or Microsoft Academic. Such explicit categorisations for research papers, especially those assigned at a journal level, are problematic [1, 13]. Moreover, the inconsistencies evident across many of these subject taxonomies [22] may confirm that no singular, correct categorisation exists. However, recent graph learning methods may be capable of learning article representations that encode interdisciplinarity, without any knowledge of explicit discipline categorisation.

## 2.2 Graph Learning - Node Embeddings

Motivated by extensive work on embedded representations in natural language processing [12], recently there has been considerable focus in the field of network analysis on developing analogous approaches for transforming nodes into low-dimensional vector representations [3, 5]. Methods for producing such representations, or node embeddings, are designed to preserve important relationships between nodes on the graph. In particular, we focus on approaches which preserve *proximity* (where similar embeddings are learned for nodes which belong to the same community in the graph) and *structural equivalence* (where similar embeddings are learned for nodes which have the same structural role in the graph).

The *DeepWalk* method [16], which extends the natural language processing-inspired SkipGram model [12], generates fixed-length random walks from all nodes on the graph. The model learns to predict the nodes that will occur on walks given the starting node as input. As such, the model preserves proximities between nodes by learning similar embeddings for nodes which occur in similar contents, i.e. nodes that co-occur on random walks. *node2vec*, proposed by [6], extends *DeepWalk*, introducing two parameters which can be used to bias the random walks in order to explore different aspects of network structure. In our experiments, we implement both *DeepWalk* and *node2vec* as examples of proximity preserving node embeddings. We also explore the structural equivalence preserving ‘role embeddings’ learned via *role2vec* [2]. The *role2vec* model further generalises the *DeepWalk* model attributing structural features (e.g. graph motif counts) to nodes before mapping those nodes to ‘roles’ and conducting role-based random walks.

## 3 DATA

We conduct experiments using five citation networks, where each paper is represented as a node and citations are represented as undirected edges between them. Three of these datasets are well-known benchmark networks used to evaluate graph learning algorithms. We also contribute two new datasets which we have collected.

## 3.1 Benchmark Datasets

As our three benchmark bibliographic datasets we consider *Cora* [11], *CiteSeer* [10], and *PubMed* [14]. These datasets range from 2,708 to 19,717 nodes (papers), with an average node degree (citations per paper) ranging from 2.7 – 4.5. Each paper is assigned a label, representing a topic or field within a discipline. For example, the *Cora* dataset contains 7 topics within the field of Computer Science: ‘*case based*’, ‘*genetic algorithms*’, ‘*neural networks*’, ‘*probabilistic methods*’, ‘*reinforcement learning*’, ‘*rule learning*’, and ‘*theory*’. While these datasets offer an accessible medium for exploring research networks, the insights gained from them may be limited by their sparsity (low node degree) and the scale of interdisciplinarity, i.e. they represent interactions between topics within a discipline, rather than macro interdisciplinary interactions. As such, we also provide two additional citation graphs that are both more complete and more interdisciplinary in nature.

## 3.2 Scopus Indexed Datasets

We construct new citation networks using Microsoft Academic Graph citation data from two different sets of journal papers. These sets consist of two samples of articles from Scopus indexed journals, stratified according to their All Science Journal Categories (ASJC). We define two networks containing different topics, so that we can investigate interactions between different subsets of disciplines. We refer to these networks as *Scopus1* and *Scopus2*. The former contains samples of 1,500 articles published between 2017 and 2018 in Scopus indexed journals with the ASJCs ‘*Computer Science*’, ‘*Mathematics*’, ‘*Medicine*’, ‘*Chemistry*’, and ‘*Social Sciences*’. For *Scopus2*, we use the seed set of categories: ‘*Computer Science*’, ‘*Mathematics*’, ‘*Neuroscience*’, ‘*Engineering*’, and ‘*Biochemistry, Genetics and Molecular Biology*’. For both datasets, we maximise the completeness of the graph by including all available referenced articles that are published in Scopus indexed journals. In this manner, we produce dense, multidisciplinary citation networks, such that each article can be categorised according to the ASJC of the journal in which it was published. *Scopus1* contains 27,213 nodes with average degree 6.9, while *Scopus2* contains 25,961 nodes with average degree 6.2. See statistics for all datasets in Table 1 of the Appendix.

## 4 METHODS

### 4.1 Citation Prediction

One established method for evaluating the quality of node embeddings is to assess their performance in a down-stream link prediction task [6, 7]. We introduce the following citation prediction task to evaluate the quality of our paper representations and their ability to encode the citation graph structure. Further, we explore their ability to identify and predict interdisciplinary interactions in the network through their performance in interdisciplinary citation prediction. We evaluate citation prediction performance with different node representations using 75:5:20 train:validation:test splits of the edges in each network. Three sets of node representations (*DeepWalk*, *node2vec*, and *role2vec*) are learned using the subgraph induced by the training edges. In the case of each model, we learn node embeddings with 128 dimensions. To train a Multi-Layer Perception (MLP) model for the citation prediction task, we supplement

each set of edges with a set of negative edges that were not present in the graph. We then represent each edge as the concatenation of the embedded representations of the associated nodes, and use train and validation edge sets to train an MLP model to classify positive and negative edges. We report performance on the test data as the area under the Receiver Operating Characteristic curve (AUC). Each experiment is repeated for 5 random splits of the data and the average AUC score is reported. Using this approach we evaluate the utility of different node representations in both overall citation prediction and in interdisciplinary citation prediction.

We use the following methods to identify interdisciplinary citations in the test data according to both binary and continuous definitions of interdisciplinary. We report the binary interdisciplinary citation prediction score as the AUC performance on the test data for only the edges where the associated nodes belong to different classifications. These scores appear in parentheses in Table 1. In Figure 1 we compare the relationship between citation prediction performance and citation distance according to our continuous definitions of interdisciplinarity, according to the intuition that representations that encode article interdisciplinarity will fare better when predicting “long-distance” interdisciplinary citations. To explore this relationship, we aggregate the edges in the test data (into bins of equal width) according to the chosen definition of citation distance, and plot the AUC for bins that contain at least 25 positive and negative edges.

## 4.2 Citation Distance Metrics

To allow us to define a continuous form of interdisciplinarity, we implement and evaluate five definitions of citation distance: *Scopus Topic Distance*, *Network Distance*, *DeepWalk Embedding Distance*, *Node2vec Embedding Distance*, and *Role2vec Embedding Distance*. Scopus Topic Distance is defined according the neighbourhood similarity of the ASJC. Using a network of 6,000,000 papers published between 2010 and 2020, we create a discipline-level weighted adjacency matrix  $\mathbf{W}$ , that encodes all of the citations between categories in the ASJC. We define the topic distance between two categories  $i$  and  $j$  as the value  $1 - S_{ij}$ , where  $S_{ij}$  is the cosine-similarity between rows  $i$  and  $j$  in  $\mathbf{W}$ . Thus, we can calculate the Scopus Topic Distance of a citation as the topic distance between the ASJC assigned to each of the relevant papers. We define the Network Distance of a citation as the length of the shortest path on the graph between to the two relevant papers [25]. Finally, each of the embedding distances is defined as the cosine distance between the two vectors learned for the pair of the papers involved in the citation.

## 5 RESULTS

### 5.1 Citation Prediction and Binary Interdisciplinarity

**5.1.1 Benchmark Datasets.** Table 1 shows citation results for 3 link prediction models, evaluated on 5 citation datasets. In each case, the AUC score is reported for both the entire test set and the interdisciplinary portion of the test set (IDR AUC). For each dataset, the structural representations learned by role2vec outperform those learned by DeepWalk-based methods. Although predicting interdisciplinary citations may be slightly more difficult than predicting

**Table 1: Citation prediction results: AUC (IDR AUC).**

Dataset	node2vec [6]	DeepWalk [16]	role2vec [2]
<i>Cora</i>	0.785 (0.789)	0.782 (0.769)	0.836 (0.817)
<i>CiteSeer</i>	0.700 (0.670)	0.692 (0.644)	0.766 (0.750)
<i>PubMed</i>	0.800 (0.788)	0.800 (0.778)	0.823 (0.818)
<i>Scopus1</i>	0.909 (0.903)	0.906 (0.901)	0.900 (0.896)
<i>Scopus2</i>	0.879 (0.875)	0.879 (0.874)	0.868 (0.869)

intradisciplinary citations, performance in each task seems to correlate well. That is, the models which perform best at predicting citations within topics are also the most performant when predicting citations between topics. However, due to the limitations in scale and completeness of the benchmark datasets (see Section 3.1), we continue experiments using two new citation graphs collected from journals indexed by Scopus.

**5.1.2 Scopus Indexed Datasets.** Table 1 reports citation prediction performance in relation to the Scopus datasets. In the case of these larger, more diverse and denser networks, we find link prediction models are more successful, likely due to the greater completeness of the graph. Moreover, the DeepWalk-based methods are no longer less effective than role2vec embeddings, and the gap between intradisciplinary and interdisciplinary scores has reduced further. To understand better the ability of different representations to encode the interdisciplinary role of research articles (by identifying those which can that can predict citations at greater distances), we now move to a continuous definition of citation distance.

### 5.2 Citation Distance Metrics

In our Scopus indexed networks, we explore 5 definitions of citation distance: 1 according to journal-level ASJC classifications and their citation neighbourhoods, 3 according to the distance between different unsupervised embedding representations, and 1 according to shortest path distance between papers. We compare these definitions of citation distance below, where we evaluate their relationship with network modularity and edge betweenness.

**5.2.1 Edge Distance and Modularity.** We compare frequency distributions for distances of positive and negative edges to examine the extent to which different metrics of edge distance capture modularity in the Scopus networks. See the supplemental figures in the Appendix. Crucially, we find that the DeepWalk and node2vec defined citation distances offer the best separation between distributions of positive and negative edges.

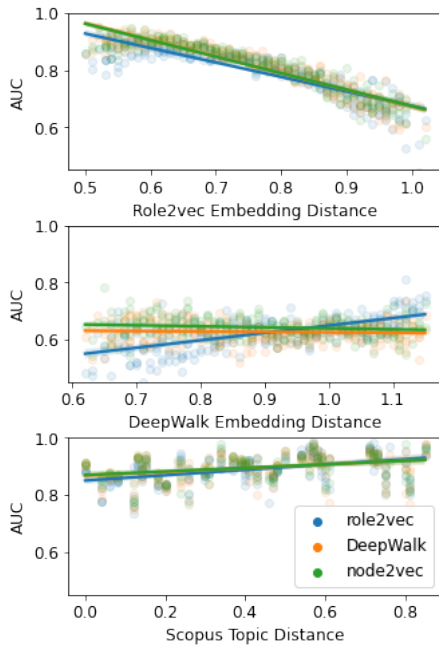
**5.2.2 Edge Distance and Edge Betweenness.** Table 2 shows Spearman rank correlations between citation distance and edge betweenness centrality, for each metric of citation distance in the Scopus indexed graph. The edge betweenness centrality for an edge (or citation) is a measure of the proportion of shortest paths on the graph which traverse that edge [4]. So citations which bridge or connect communities will have high edge betweenness centrality. As such, we expect a definition of citation distance which encodes the interdisciplinarity of citation, to correlate positively with edge betweenness. The distance measures defined according to the distance between DeepWalk-based embeddings are shown to have the

**Table 2: Citation distances correlated with edge betweenness.**

Distance	Scopus1	Scopus2
Scopus Topic	0.029	-0.020
DeepWalk	0.331	0.300
Node2vec	0.329	0.304
Role2vec	0.191	0.193

strongest correlation with edge betweenness centrality. Conversely, the distances between prescribed topics based on the ASJC from Scopus show no correlation with edge betweenness, indicating that these topic assignments, or at least the defined distances between them, do not reflect the community structure in the graph.

We can further investigate the validity of the ASJC paper categorisations in the Scopus datasets, by comparing the edge betweenness scores of inter- and intra-topic edges. If the ASJC assignments represent the communities in the network, we would expect edges that are *between* ASJC topics (i.e. citations linking papers of different categories) to have greater betweenness than those edges that are *within* ASJC topics (i.e. linking papers with the same category). A Kolmogorov-Smirnov test comparing the edge betweenness distributions of inter- and intra-disciplinary edges finds interdisciplinary edges to have greater betweenness at 5% significance. This indicates that, while the predefined distances between ASJC topics do not correlate with citation distance, the ASJC categorisations are somewhat reflective of the communities in the network.



**Figure 1: Scopus1 prediction performance (AUC) plotted against citation distance, for different node representations and different metrics of citation distance.**

**Table 3: Link prediction AUC correlated with citation distance.**

Distance	role2vec	DeepWalk	node2vec
Scopus Topic	0.444*	0.327*	0.323*
DeepWalk Embedding	0.642*	-0.061	-0.139*
Node2vec Embedding	0.632*	-0.126*	-0.129*
Role2vec Embedding	-0.875*	-0.929*	-0.914*
Network Distance	0.245	-0.173	-0.191

**5.2.3 Citation Distance and Link Prediction.** Table 3 reports the linear regression coefficients comparing citation prediction performance with citation distance in the *Scopus1* network. The table include stars to highlight regression coefficients significant at 5% significance. In particular, we note that the models employing role2vec structural representations of the nodes perform better as citation distance increases according to most definitions of citation distance. Table 5 in the Appendix plots similar results for *Scopus2*, where role2vec performance again increases with those distances that agree with the underlying network structure.

## 6 DISCUSSION AND CONCLUSIONS

In our investigations of the potential for graph learning methods to identify interdisciplinary research, we explored three different methods of node representation and evaluated their utility in predicting interdisciplinary citations. According to a binary definition of interdisciplinary citation, we found that predicting links between research topics is indeed more difficult than predicting within-topic interactions. This effect is more evident in sparser networks, such as the benchmark citation graphs *Cora*, *CiteSeer*, and *PubMed*. In the more complete citation networks that we have collected (*Scopus1* and *Scopus2*), the distinction between citation prediction and interdisciplinary citation prediction performances are less apparent.

We explored a continuous definition of interdisciplinary citation, as some interactions are viewed as being more distant than others [8]. We evaluated five different definitions of citation distance. Although it remains difficult to determine the most accurate means of measuring the distance or interdisciplinarity of citation, we highlighted a number of issues with the predefined ASJC distances, which do not appear to agree with the underlying network structures. Further, we found that the role2vec-based model achieved citation prediction AUC that correlates positively with all metrics of citation distance that do appear to encode the underlying network structure. This provides evidence that the role2vec-based paper representations are capable of encoding the structural signatures associated with interdisciplinary interactions. We suggest that the embeddings which preserve structural equivalence on the graph are capable of encoding the ‘interdisciplinary role’ of different articles.

We plan to extend this work to learning more advanced paper representations by including text features on nodes in the graph. There is also potential to develop means of better explaining these embeddings to quantify interdisciplinarity and interpret the different interdisciplinary roles of research papers.

**Acknowledgements.** This research was supported by Science Foundation Ireland (SFI) under Grant Number SFI/12/RC/2289\_P2.

## REFERENCES

- [1] Giovanni Abramo, Ciriaco Andrea D'Angelo, and Lin Zhang. 2018. A comparison of two approaches for measuring interdisciplinary research output: The disciplinary diversity of authors vs the disciplinary diversity of the reference list. *Journal of Informetrics* 12, 4 (2018), 1182–1193.
- [2] Nesreen K. Ahmed, Ryan A. Rossi, John Boaz Lee, Theodore L. Willke, Rong Zhou, Xiangnan Kong, and Hoda Eldardiry. 2019. role2vec: Role-based network embeddings. In *Proc. DLG KDD*. 1–7.
- [3] Hongyun Cai, Vincent W Zheng, and Kevin Chen-Chuan Chang. 2018. A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering* 30, 9 (2018), 1616–1637.
- [4] Michelle Girvan and Mark E.J. Newman. 2002. Community structure in social and biological networks. *Proceedings of the national academy of sciences* 99, 12 (2002), 7821–7826.
- [5] Palash Goyal and Emilio Ferrara. 2018. Graph embedding techniques, applications, and performance: A survey. *Knowledge-Based Systems* 151 (2018), 78–94.
- [6] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proc. 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 855–864.
- [7] Thomas N Kipf and Max Welling. 2016. Variational graph auto-encoders. *arXiv preprint arXiv:1611.07308* (2016).
- [8] Vincent Larivière, Stefanie Haustein, and Katy Börner. 2015. Long-distance interdisciplinarity leads to higher scientific impact. *PloS one* 10, 3 (2015), e0122565–e0122565.
- [9] Erin Leahey, Christine M. Beckman, and Taryn L Stanko. 2017. Prominent but less productive: The impact of interdisciplinarity on scientists' research. *Administrative Science Quarterly* 62, 1 (2017), 105–139.
- [10] Qing Lu and Lise Getoor. 2003. Link-Based Classification. In *Proc. 20th International Conference on International Conference on Machine Learning* (Washington, DC, USA) (*ICML '03*). AAAI Press, 496–503.
- [11] Andrew Kachites McCallum, Kamal Nigam, Jason Rennie, and Kristie Seymore. 2000. Automating the construction of internet portals with machine learning. *Information Retrieval* 3, 2 (2000), 127–163.
- [12] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013).
- [13] Staša Milojević. 2020. Practical method to reclassify Web of Science articles into unique subject categories and broad disciplines. *Quantitative Science Studies* 1, 1 (02 2020), 183–206.
- [14] Galileo Namata, Ben London, Lise Getoor, Bert Huang, and UMD EDU. 2012. Query-driven active surveying for collective classification. In *Proc. 10th International Workshop on Mining and Learning with Graphs*, Vol. 8. 1.
- [15] Keisuke Okamura. 2019. Interdisciplinarity revisited: evidence for research impact and dynamism. *Palgrave Communications* 5, 1 (2019), 1–9.
- [16] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. Deepwalk: Online learning of social representations. In *Proc. 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. 701–710.
- [17] ALDE Porter and D Chubin. 1985. An indicator of cross-disciplinary research. *Scientometrics* 8, 3-4 (1985), 161–176.
- [18] Alan Porter, Alex Cohen, J. David Roessner, and Marty Perreault. 2007. Measuring researcher interdisciplinarity. *Scientometrics* 72, 1 (2007), 117–147.
- [19] Alan Porter and Ismael Rafols. 2009. Is science becoming more interdisciplinary? Measuring and mapping six research fields over time. *Scientometrics* 81, 3 (2009), 719–745.
- [20] Ismael Rafols and Martin Meyer. 2010. Diversity and network coherence as indicators of interdisciplinarity: case studies in bionanoscience. *Scientometrics* 82, 2 (2010), 263–287.
- [21] Martin Rosvall and Carl T Bergstrom. 2008. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences* 105, 4 (2008), 1118–1123.
- [22] Zhesi Shen, Fuyou Chen, Liying Yang, and Jinshan Wu. 2019. Node2vec representation for clustering journals and as a possible measure of diversity. *Journal of Data and Information Science* 4, 2 (2019), 79.
- [23] Richard Van Noorden et al. 2015. Interdisciplinary research by the numbers. *Nature* 525, 7569 (2015), 306–307.
- [24] Caroline S. Wagner, J. David Roessner, Kamau Bobb, Julie Thompson Klein, Kevin W. Boyack, Joann Keyton, Ismael Rafols, and Katy Börner. 2011. Approaches to understanding and measuring interdisciplinary scientific research (IDR): A review of the literature. *Journal of Informetrics* 5, 1 (2011), 14–26.
- [25] Stanley Wasserman and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. Cambridge University Press.